A Dynamic Structural Analysis of the PC Microprocessor Industry

by

Brett Russell Gordon

Submitted in Partial Fulfillment of the

Requirements for the Degree of Doctor of Philosophy in

Economics

Tepper School of Business

Carnegie Mellon University

May 2007

Dissertation Committee:
Professor Ronald L. Goettler (co-chair)
Professor Kannan Srinivasan (co-chair)
Professor Dennis Epple
Professor Holger Sieg
Professor Vishal Singh

Abstract

In durable goods markets, sellers face a dynamic trade-off: more units sold today come at the expense of selling more units tomorrow. Buyers face a similar dynamic trade-off: should they purchase a new product today or retain their existing product and purchase a potentially better product tomorrow? These issues lay at the heart of durable goods markets - especially those involving high-tech products - yet relatively little research addresses them from an empirical perspective. To that end, this dissertation provides a dynamic structural analysis of demand and competition within the context of the PC microprocessor industry. This industry is particularly interesting because it is a duopoly that has experienced intense technological and price competition.

First, I estimate a model of dynamic demand that allows for both product adoption and replacement decisions when consumers are uncertain about future product price and quality. In the absence of panel data, I show how to infer replacement from a combination of aggregate data. The results show that heterogeneity in replacement behavior provides an opportunity for firms to tailor their product introduction and pricing strategies by targeting specific consumer replacement segments.

Second, I extend this analysis to construct an equilibrium model of dynamic oligopoly with durable goods and endogenous innovation. Firms make dynamic pricing and investment decisions while taking into account competitors' strategies and the aggregate dynamic behavior of consumers. I examine the role of product durability on firms' optimal policy functions and the nature of competition in the industry. I find that industry profits are 24 and 41 percent lower in the duopoly and monopoly settings, respectively, when the firms ignore the durable nature of the product while setting prices. This demonstrates the strong link between optimal firm behavior and accounting for durability and dynamic demand. Welfare outcomes also differ significantly: compared to a socially benevolent monopolist, consumer welfare is 22 percent lower in a duopoly and 54 percent lower in a monopoly. While investment is higher in the duopoly than in the monopoly, a counterfactual analysis suggests that most of the welfare loss associated with monopoly comes from higher margins and not slower innovation.

Acknowledgments

I am immeasurably indebted to a great number of people. Steve Spear helped to spark and nurture my interest in economics when I was an undergraduate. Alan Montgomery deserves special thanks for having continued to guide me as both an undergraduate and graduate student. Mike Peress and Fallaw Sowell served as sounding boards for research questions on many occasions, and Rick Green provided advice that greatly facilitated my progress through the Ph.D. program.

I have been extraordinarily lucky in my dissertation committee: Ron Goettler, Kannan Srinivasan, Dennis Epple, Holger Sieg, and Vishal Singh. Their knowledge and wisdom is only exceeded by their patience and common sense. I am grateful to each individually, and to my committee collectively, for challenging me intellectually and helping me grow as a researcher. Ron deserves my special gratitude for tolerating years of questions relating to practically everything. Kannan helped smooth my transition into a new field and spent a vast amount of time with me, discussing a range of topics. Dennis gave me outstanding research direction and unwavering encouragement. Holger encouraged me to pursue my research projects and gave me excellent advice. Vishal introduced me to a new field and generously gave his time and counsel as I made the transition.

I am grateful for support from the William Larimer Mellon fund and a dissertation fellowship from the Center for Analytical Research in Technology (CART), both at Carnegie

Mellon. I am also grateful to Margaret Kyle for providing me with funding needed to purchase data, and to Alan, Dennis, and Holger for providing additional financial support throughout my time as a graduate student.

I would like to thank my family for their support, especially my father and stepmother, who gave me continuous encouragement and advice, and my grandmother Blanche Gordon, whom I know is happy to welcome the second *real* doctor into the family.

Finally, I would like to thank Sue Yeh, for her love, patience, and understanding. I cannot imagine having gone through this experience without her. She has been my best friend and companion since long before this work started, and I know the same will be true long afterwards.

Dedication

This dissertation is dedicated to my father Barry Gordon and my late grandfather Bernard Gordon. They taught me to love learning and to seek out knowledge. While I was resistant at times, I am grateful for their persistence. I would not have completed this dissertation without the lessons they instilled in me.

Table of Contents

Chapter 1

Introduction

In durable goods markets, sellers face a dynamic trade-off: more units sold today come at the expense of selling more units tomorrow. Buyers face a similar dynamic trade-off: purchase a new product today or retain their existing product and purchase a potentially better product tomorrow? These issues lay at the heart of durable goods markets, especially those involving high-tech products. Yet, while there is a well-developed literature on the microeconomic theory of durable goods, relatively little work has empirically analyzed these markets from a structural perspective. This imbalance is not due to a lack of intrinsic empirical interest; numerous questions exist that are specific to such markets concerning the inter-related roles of product replacement cycles, R&D investment decisions, dynamic pricing strategies, and industry evolution. Instead, the lack of empirical analysis is primarily due to the difficulty of finding suitably rich data sets and overcoming modeling and computational obstacles.

High-tech durable goods markets represent an especially interesting class of industries for several reasons. First, these industries play an increasingly central role in the economy and their influence extends beyond their formal boundaries. This makes it important for researchers to be able to accurately model such markets to be able to answer policy questions and conduct counterfactual analyses, such as evaluating potential mergers or measuring the welfare impact of new technologies. Second, from a modeling perspective, these markets frequently undergo both rapid improvements in quality and falling prices. This implies that a consumer's decision to replacement a high-tech product, such as a digital camera or an MP3 player, is most often due to obsolescence. As a result, the replacement decision is dynamic because the consumer that forgoes a purchase today might buy a potentially better product tomorrow for less money. This is in contrast to"low-tech" durables, such as washing machines, which are usually replaced due to wear and tear and do not require consumers to posses forward-looking expectations.

These unique features of high-tech durable goods make it problematic to apply standard demand estimation techniques and equilibrium modeling approaches. First, the standard static discrete choice model does not account for the forward-looking behavior of consumers, and using such a model generally leads to biased parameter estimates and potentially flawed counterfactual analysis. Second, consumers' replacement decisions depend on the value they place on the product they already own (if any). This suggests that it is important to model the distribution of consumers over their currently owned products, and that is important to model how this distribution of product ownership changes over time in response to price changes and new product introductions.

This dissertation seeks to address these issues by providing a dynamic structural analysis of demand and competition in a high-tech durable goods market. I develop techniques for dealing with the data and modeling problems, and investigate the substantive issues within the context of the PC processor industry. This industry is particularly interesting because it is a duopoly, with Intel and Advanced Micro Devices (AMD) controlling about 95 percent of the market, and sales have been driven by intense technological innovation and price competition. While I focus on the processor industry, the analysis should be relevant for any industry where innovation and obsolescence drive product replacement. To conduct the empirical analysis, I have constructed a unique and comprehensive data set including prices, characteristics, sales, and consumer ownership of PC processors from 1993 to 2004.

I briefly summarize the other two chapters below.

Chapter 2. In this chapter, I focus on the critical role of product replacement in high-tech markets. Despite the seeming importance of the replacement decision, most of the empirical work to date has only focused on the initial product adoption decision (Melnikov, 2001, Song and Chintagunta, 2003). While such a restriction may be reasonable in industries where replacement sales are negligible, many high-tech markets have matured to the point where replacement purchases are a significant portion of overall sales volume. In these cases, firms know that consumers follow replacement cycles, but little is known about how and why these replacement cycles change over time. And recognizing that consumers do follow such cycles, how might firms alter their own strategies to take advantage of this?

To address these issues, I develop and estimate a consumer model of dynamic demand for PC processors. The model allows for both product adoption and replacement decisions when consumers are uncertain about future product price and quality. In the absence of panel data, I show how to infer replacement behavior from a combination of aggregate data sources and examine the role of technological innovation and pricing on product replacement over time.

After applying the model to the PC processor industry, the results reveal substantial variation in replacement behavior over time. First, I demonstrate that this heterogeneity in consumer replacement behavior provides an opportunity for managers to tailor their product introduction and pricing strategies to target the particular segment of consumers that is most likely to replace in the near future. Managers can alter their product introduction and pricing strategies to correspond to the preferences of each particular consumer segment. Second, the model suggests that an"averaging effec" may provide an alternative explanation for the observed increase in replacement cycle length in the PC market: more consumers with inherently longer replacement cycles have entered the market over time, producing a natural increase in the aggregate replacement cycle length. Third, I find that a myopic model of replacement underestimates price elasticities by approximately 30 to 40 percent and overestimates the frequency of replacement by 50 percent. Lastly, using counterfactual simulations, I also examine the impact of alternative rates of innovation on replacement cycle length and consumer welfare. I find that the marginal effect of innovation on the replacement cycle length has decreased over time, implying a decline in the ability of quality innovations to generate replacement sales.

Chapter 3. Dynamic demand is a key characteristic of durable goods markets. An event in the current period, such as a price cut, may cause some consumers to shift future consumption into the present, leading to lower sales in the following period. Depending on the magnitude of such events, their impact on future sales may persist for many periods, and thus affect the optimal firm policies as well. Despite the importance of dynamic demand in durable goods markets, the equilibrium implications on firms' and consumers' strategies in an imperfectly competitive market remain unclear.

To this end, the third chapter of this thesis constructs a model of dynamic oligopoly with durable goods and endogenous innovation. Firms make both dynamic pricing and investment decisions while taking into account competitors' strategies and the aggregate dynamic behavior of consumers. A consumer must decide whether to keep their existing product (if any) or to buy a new product, given her expectations about future product characteristics. The distribution of currently owned products affects current demand and evolves endogenously as consumers make replacement purchases. The equilibrium model allows us to understand the role of forward-looking consumer behavior on firms' optimal policy functions and the nature of competition in the industry and to conduct policy and counterfactual simulations in equilibrium.

Our work extends the framework developed by Ericson and Pakes (1995) and Pakes and McGuire (1994) to incorporate durable goods. Our work is also related to recent empirical models of dynamic demand that take the firms' behavior as exogenous, whereas this chapter endogenizes firms' actions through a dynamic multi-agent game setting. Finally, our work is connected to the large theoretical literature on durable goods, such as the literatures on optimal durability, starting with Swan (1970, 1971), and Sieper and Swan (1973), and the literature on the monopolist's time inconsistency problem, beginning with Coase (1972) and followed by Stokey (1981) and Bulow (1982).

Our model primarily differs along two key dimensions. First, we consider a dynamic oligopoly where competition in pricing and innovation play critical roles. Second, we allow for endogenous innovation and account for the endogenous evolution of the distribution of currently owned (used) products. Product quality increases over time as the firms invest in innovation. While the product a consumer owns never deteriorates in an absolute sense, consumers make replacement purchases as the quality of the product they own becomes worse relative to the frontier product in the market.

We show that accounting for product durability and the distribution of consumer ownership have significant implications for firms' profits and consumer surplus. In a duopoly, we find that industry profits are 24 percent lower when the firms ignore the durable nature of the product when setting prices. In the monopoly case, the firm's profits are 41 percent lower. This demonstrates the strong link between optimal firm behavior, with and without competition, and accounting for product durability and dynamic demand. Margins are 48 percent lower in the duopoly and 68 percent lower in the monopoly, confirming the intuition that prices are higher under dynamic demand.

Welfare outcomes also differ significantly: compared to a socially benevolent monopolist, consumer welfare is 22 percent lower in a duopoly and 54 percent lower in a monopoly. While investment is higher in the duopoly than in the monopoly, a counterfactual analysis suggests that most of the welfare loss associated with monopoly comes from higher margins and not a slower rate of innovation.

Chapter 2

A Dynamic Model of Consumer Replacement Cycles

2.1 Introduction

Product replacement plays a critical role in many high-tech durable goods markets, from cell phones to camcorders to computers. Since these markets frequently undergo both rapid improvements in quality and falling prices, the consumer's replacement decision is most often due to product obsolescence, as opposed to wear and tear. As a result, the replacement decision is dynamic because the consumer that forgoes a purchase today might buy a potentially better product tomorrow for less money. However, from an empirical perspective, this issue has received insufficient attention in both marketing and economics.[1] Managers know that consumers follow replacement cycles, but little is known about how and why these replacement cycles change over time. And recognizing that consumers do follow such cycles,

[1]This relates to the literature on planned obsolescence, which has mostly addressed these issues from a theoretical perspective. See, for example, Levinthal and Purohit (1989) and Waldman (1993).

how should managers alter their own strategies to take advantage of this?

To address these issues, I develop and estimate a consumer model of dynamic demand for PC processors. The model allows for both product adoption and replacement decisions when consumers are uncertain about future product price and quality. The PC processor industry is particularly interesting because it is composed of two major players, Intel and Advanced Micro Devices (AMD). Intense technology and price competition between them has helped increase market penetration of PC's from 28% in 1993 to 74% in 2004. As penetration has grown, the industry has increasingly relied on product replacement as a vehicle for sales growth.

Using a unique data set that combines different aggregate data sources, I examine the role of technological innovation and pricing on product replacement over time. I show that heterogeneity in both consumer preferences and replacement behavior has important implications for understanding consumer demand. This heterogeneity provides an opportunity for managers to tailor their product introduction and pricing strategies to target the particular segment of consumers that is most likely to replace in the near future.

Using a structural model of dynamic demand for differentiated durable goods, I allow the consumer's replacement decision to depend on the quality of the product they currently own. Each period, consumers in the model choose whether to keep their existing product (if any) or to replace it with one of the new products available. Consumers form expectations over future product qualities and prices. The uncertainty in these expectations creates an inherent trade-off. If consumers expect little change in product qualities or prices, they are more likely to make a purchase in the current period. But if they expect a near-term decrease in quality-adjusted prices, they are more likely to postpone replacement. Heterogeneity enters the model in two ways. First, consumers belong to latent preference segments. Second, the value of a consumer's outside option (no purchase) is determined endogenously; the product a consumer already owns determines the value of not making a purchase in the future. Thus, both a consumer's replacement decision and the exogenous evolution of the product market affect the composition of their future choice sets. Note that the product adoption decision is implicitly nested within this specification.

Models of demand for durable goods have already been examined in a variety of contexts, but relatively few studies have considered the replacement of durable goods in a structural empirical model.[2]

In marketing, much of the existing work on product adoption follows in the tradition of the Bass diffusion model (Bass, 1969). Some relevant extensions consider the diffusion of successive generations of a technology product (Norton and Bass, 1987) and the optimal introduction timing of new product generations (Wilson and Norton, 1989). This research methodology has successfully described the empirical diffusion pattern of new consumer durables ranging from air conditioners to color TVs to clothes dryers. A few papers incorporate product replacement into forecasting models of durable goods sales (Bayus, 1988; Bayus, Hong, and Labe, 1989; Steffens, 2003).[3]

Building off the literature on single-agent dynamic models, researchers in both marketing and economics have also constructed structural models of durable goods adoption.[4] A common feature of most of this work is that the consumer's decision problem is formulated as an optimal stopping problem.[5] In this formulation, the consumer decides on the optimal time

---

[2]The discussion here is limited to infrequently purchased durable goods. The replacement decision for frequently purchased products, such as paper towels or cereal, is fundamentally different, because it is usually motivated by stockout effects and the characteristics of the products are constant. See, for example, Gonul and Srinivasan (1996), Sun *et al* (2003), Mehta *et al* (2004), and Sun (2005).

[3]Ratchford, Balasubramanian, and Kamakura (2000) provide an excellent review of the literature on diffusion models with replacement and multiple purchases. The author thanks a reviewer for bringing this paper to his attention.

[4]See Rust (1994) and Dubé *et al* (2005) for reviews of the literature on dynamic models in economics and marketing, respectively.

[5]Horsky (1990) and Chatterjee and Eliashberg (1990) take a different approach by constructing an aggregate diffusion curve based on the individual level adoption decisions derived from a consumer's maximization problem.

period to enter the market, makes a single purchase from among the available products, and then exits the market permanently. An appealing consequence of this assumption is that the policy function describing the optimal entry period takes the form of a cutoff rule, which can facilitate the estimation process.

A number of researchers use the optimal stopping framework to study the adoption of new consumer technology durables. Melnikov (2001) uses it in a novel model of demand for differentiated durable goods that he applies to the adoption of computer printers. Carranza (2004) and Gowrisankaran and Rysman (2005) extend Melnikov's model to include consumer heterogeneity and examine the introduction of digital cameras and DVD players, respectively. Song and Chintagunta (2003) develop a similar empirical model of digital camera adoption that incorporates unobserved consumer segments. Erdem *et al* (2005) use an optimal stopping problem to model a the process of consumer learning and information search about the choice of a PC technology. Using panel data, they track the various information sources used by consumers to help inform their decisions about whether to buy an IBM-PC compatible computer or an Apple computer.[6]

All of these papers successfully apply the optimal stopping formulation to the durable goods adoption problem. However, incorporating replacement into the optimal stopping framework is difficult because it is not designed to model multiple, interdependent decisions. Rust (1987) and Prince (2005) are the only other papers to consider the replacement of durable goods, but both differ significantly from the present work. Rust (1987) studies the replacement of a durable good due to wear and tear, as opposed to technological innovation. The decision maker solves a regenerative optimal stopping problem, where the choice set is fixed and the decision environment is stationary.[7] Prince (2005) creates a structural model

[6]Nair (2007) studies the optimal pricing problem for a monopolist selling video game consoles to forward-looking consumers, who are modeled as solving an optimal stopping problem.

[7]The model is *regenerative* because the replacement decision resets the current state to the initial state.

of PC demand to quantify the effects of subsidies to first-time buyers, but his data does not permit him to study replacement cycles. Also, consumers in his model do not face any product market uncertainty: they have perfect foresight about future product quality and price.

To the best of my knowledge, there is no published work involving durable technology markets that incorporates the replacement decision into the consumer's dynamic problem. The modeling approach developed here contributes to the literature on estimating dynamic demand models of durable goods in three significant ways. First, I show how to incorporate the consumer's replacement decision into a dynamic estimation problem that allows for a sequence of purchases. This represents an important contribution over previous work that solely considers the adoption decision. Second, I demonstrate how to combine aggregate data on sales and ownership to infer replacement behavior. The estimation results show that the additional information provided by the ownership data are necessary for capturing the dynamics of replacement. This is the first attempt to study replacement behavior using aggregate data, motivated in part because of the lack of adequate panel-level data on durable goods replacement.[8] Third, I show that modeling heterogeneity in product ownership has important implications for managers' product introduction and pricing strategies. The model must not only distinguish between consumers that are non-owners versus owners, but it also tracks the distribution of product ownership by quality among the owners.

To estimate the model, I construct a comprehensive data set of prices, characteristics, sales, and ownership of PC processors manufactured by Intel and AMD from January 1993 to June 2004. These two firms controlled roughly 95% of the market during the sample period. The data set includes unit shipments from an industry research firm and proprietary survey data on ownership by processor type from a consumer research company. The survey

[8]The ideal panel data set would allow me to observe a sequence of PC purchase decisions conditional on the existing PC. To the best of my knowledge, such a data set does not exist.

data is only available at the aggregate level, but this allows me to estimate the distribution of ownership by processor quality at any given point in time. Changes in the ownership distribution from one period to the next and the current period's sales make it possible to infer replacement behavior at the aggregate level. In addition, I collect detailed information on the history of processor prices and characteristics from old press releases, news reports, and industry periodicals. Finally, since processors across firms are not always comparable based solely on processor frequency (e.g. 1.8 gigahertz), I supplement the data with speed benchmark information that takes multiple processor attributes into account in determining processor quality and allows for the accurate comparison of processors both within and across firms. Using these data, I estimate and solve the model using generalized method of moments (GMM) as part of a nested fixed-point algorithm to match a set of simulated moment conditions to their empirical counterparts.

The results reveal the important role that the replacement decision plays in shaping demand within this technologically dynamic industry. First, estimates from the structural model imply there is significant variation in the distribution of replacement cycles across consumers both within a period and over time. Not surprisingly, consumers with a lower marginal value for PC processor quality have longer average replacement cycle lengths. However, the length of the replacement cycle for this segment experienced the largest shift over time. Overall, the replacement cycle for all consumer segments increased over time. Second, changes in the relative competitive structure of the industry show that while AMD has made some competitive gains, it is still unable to compete head on with Intel. I find that AMD's premium set of products appear to compete most closely with Intel's value line of processors. Third, the marginal effect of innovation on the length of the replacement cycle has decreased over time, implying that PC hardware manufacturers may not always be able to rely on quality improvements to generate replacement sales in the future.

A comparison of the benchmark dynamic model to a version with myopic consumers yields substantially different results. The myopic model underestimates the length of consumer replacement cycles by roughly one year, implying that consumers replace their products more frequently than the dynamic model implies. The myopic model also underestimates price elasticities by up to 45%. The intuition behind this result is that forward-looking consumers take into account the long-run benefits and costs of purchasing a product today.

From a managerial perspective, I find that a firm's product introduction and pricing strategies can be adjusted to take into account the particular consumer segment that is most likely to replace at a point in time. For example, if more price sensitive consumers are likely to replace their products in the near future, a firm could release a more value-oriented product to coincide with this event. The firm could increase the prices of their non-premium products and decrease the prices of their premium products. In addition, consumer replacement cycles could be integrated into a firm's customer relationship management (CRM) system (Lewis, 2005). Understanding and managing a consumer's replacement cycle is certainly a key factor in determining the consumer's life-time value.

These findings also have some implications at the industry level, and are likely to apply to other high-tech markets. For the PC industry, if current technology trends continue, the replacement cycle is likely to continue to increase over time. Price competition may increase as consumers' demand for PC processor speed becomes increasingly satiated.[9] Some PC and software makers have already responded to these changes by implementing more detailed diversification strategies, but it is still too early to know whether these strategies will ultimately find success in the marketplace.[10] In general, firms in high-tech markets may want to track the distribution of product ownership and understand how consumer replacement cycles change over time in their particular industry. This could help the firms

[9] Gartner Global PC Forecast Q1 Update, March 2006, and IDC Worldwide Quarterly PC Tracker, March 2006.

[10] FIND/SVP, PC Trend Report, 2004.

improve the long-term planning of their product introduction and design strategies.

The paper proceeds as follows. Section 2.2 discusses the data set and provides an overview of the PC processor industry, as related to Intel and AMD. Section 2.3 presents the model and Section 2.4 describes the estimation approach and presents the parameter estimates. Section 2.6 examines replacement behavior. Section 2.7 discusses the managerial implications of the paper's findings and concludes.

2.2 The PC Processor Industry and The Data Set

This section begins with a brief history of the PC processor industry, focusing on the relationship between Intel and AMD. A description of the data set and some discussion follows.

2.2.1 The PC Processor Industry: Intel and AMD

The relationship between Intel and AMD dates back to the early 1980's.[11] Intel developed the first microprocessor in 1974. IBM helped them become the market leader after IBM chose Intel's processor design to be the standard for PCs. However, not wanting to depend on a single supply source, IBM demanded that Intel contract with another company and license it to manufacture Intel's x86 chips. AMD agreed to abandon its own competing architecture and began producing x86 chips as a second source. Relations between the two firms later turned sour, and AMD sued Intel in 1987 over the alleged use of anticompetitive tactics that breached the good faith of the original licensing agreement.

AMD continued to produce Intel's chip designs under the disputed contract until the lawsuit was completely settled in January 1996. This marked an important turning point in the industry because the resolution of the dispute allowed each company's strategy to evolve

[11]See Langlois (1992, 2002) for excellent histories of the semiconductor and PC industries and Bresnahan and Greenstein (1999) for a historical examination of computer industry market structure.

in its own way. Intel concentrated on the Pentium chip, which AMD had no legal right to produce. In response, AMD purchased NexGen in an attempt to upgrade its microprocessor design capabilities and to establish itself as a credible alternative to Intel. From 1995 to 1999, AMD reduced the lag time between Intel's release of a new design and AMD's release of a competing chip from over 18 months to almost nothing. In mid-1999, AMD introduced the Athlon processor, its first x86-based chip that did not depend on any previously licensed technology from Intel. According to McKinsey, this evidence of stronger competition from AMD prompted Intel to increase the frequency of new chip releases.[12] Older products became obsolete more rapidly as both firms increased the pace of innovation. These actions reduced the average market lifespan of a PC processor from about three years to one and half years (Stevens, 1994).

Despite AMD's efforts, Intel has always been the recognized market leader: its market share has fluctuated between 70% and 92% since the early 1990's. AMD's market share has been less stable, hovering around 15% for most of the early 1990's, then dropping to as low as 6% in 1997, and later rising to nearly 23% in 2001.

2.2.2 Data Description

The data set focuses on the desktop PC market, and consists of PC processor unit shipments, consumer PC ownership, manufacturer prices, and quality measurements, by processor, over the period January 1993 to June 2004.[13] The market of interest is consumers and businesses in the U.S, which I assume follow similar purchasing patterns. According to the Computer Industry Almanac, a trade publication, businesses historically tend to own roughly two thirds

[12]McKinsey Global Institute (2001).

[13]Historically, server and mobile processors have occupied a small share of the overall PC market, though laptop sales have significantly increased in the last few years.

of all PC's. Thus, the market size is set to three times the number of households.

Shipments. Quarterly unit shipment data were obtained from In-Stat/MDR, an industry research firm that specializes in the microprocessor industry. In-Stat/MDR uses a combination of contacts in various distribution channels and a detailed knowledge of manufacturing costs to back out estimates of shipment data by processor type. Data directly from the manufacturers are not available. I obtained data from the Computer Industry Almanac on the U.S. portion of global chip sales, the portion of sales for replacement machines, and the market size. This information is used to convert the global shipment figures from In-Stat/MDR to U.S. quantities.

PC manufacturers have strong incentives to minimize their inventories, so the delay between the shipment of a PC processor and the subsequent purchase of a PC containing that processor by an end-user should be negligible. Note that the data presumably include a component of unit sales that go to consumers who purchase multiple computers.[14]

Ownership. Aggregate information on consumer PC ownership and penetration rates comes from the Homefront study created by Odyssey, a consumer research firm that specializes in technology products. The firm conducts semi-annual telephone surveys using a nationally representative sample ranging from 1,500 to 2,500 households. The households do not belong to a pre-chosen panel and a new sample is drawn for each wave of the study. The survey data is neither available at the household level nor in panel form. To increase the accuracy of the relevant sample, approximately 500 additional households that own a PC are oversampled. The survey gathers basic information on PC ownership, including details such as the CPU manufacturer, architecture, and processor speed. This allows me to estimate the percent of consumers in the population who own a CPU from a particular speed range,

[14] While the issue of multi-PC households is interesting, the present data set does not provide the ability to study such a problem.

such as Pentium III processors operating between 500 megahertz and 800 megahertz.

Figure 2.1 displays a select sample of the ownership data, aggregated to the processor generation level. The figure covers the period from 2000 to 2004 and shows the share of households that own an Intel 486, Pentium, Pentium II, Pentium III, or Pentium 4 processor. The Pentium III was released in late 1999 and the Pentium 4 in late 2001. The other processors were no longer being sold. The graph reveals several interesting points. First, note that despite the fact that Intel has not sold 486 processors since 1995, a significant portion of households still owned 486-based PC's in 2000. Second, more households own a PC with a Pentium chip compared to a Pentium II, even though the Pentium II is the more advanced technology and neither chip is available for sale. Third, the rate of decline in the ownership shares for the Pentium and Pentium II appears to have leveled off in recent years, suggesting perhaps that the remaining owners are less likely to replace their product's in the near future.

Prices. The complete price history for the processors was obtained from a large number of sources. The manufacturer prices are quoted in chip quantities of one thousand units. These quotes represent the "official" prices; the actual prices Intel and AMD charge to PC manufacturers are likely to differ. Nevertheless, Intel and AMD adjust their official prices frequently (between five and ten times per year), which implies that the posted prices can still serve as adequate indicators. The primary data sources for prices were news websites, In-Stat/MDR, historical Intel and AMD press releases, technology newsgroups, and other sources. Data from these sites was also supplemented with information from historical issues of PC-related magazines and periodicals. Whenever possible, prices are checked against multiple sources. In the case of a contradiction, the more reliable source (such as a company's press release) is used.

Quality. Processor frequency does not adequately capture the computational power of a CPU because of differences in chip architecture and characteristics.[15] To account for such differences, I use a CPU speed benchmark to generate composite quality ratings for each chip.[16] The CPU Scorecard (www.cpuscorecard.com) provides a comprehensive list of benchmarks that adequately covers the sample period. The list of processor speed ratings from the CPU Scorecard does not contain all the processors in the data set: 74 of 217 processors did not have benchmarks (38 from AMD and 36 from Intel). To fill in the missing values, I impute the missing benchmark based on the available ratings. Regressing the existing speed ratings against processor frequency and brand dummies produces an R^2 of 97.4%. Adding other processor characteristics, such as bus frequency, cache size, and dummies for the processor architecture, increases the R^2 to 99.8%.

2.2.3 Discussion

Since Intel and AMD sell many processors, I aggregate the choice set and create composite *frontier* and *non-frontier* processors based on the current period product offerings of each firm. This allows each firm to sell multiple products, while keeping the model tractable. For a given period and firm, I divide the set of available processors into groups above and below the median quality processor. The frontier product is formed by taking the average price and quality of the upper processor group and the non-frontier product is assigned the average price and quality of the lower group.[17] As the prices of the underlying products change, the prices of the composites change. Changes in the set of actual products available leads to changes in the quality of the composite products. Composite market shares are calculated

[15]For example, the AMD Athlon XP 3000+ processor has a frequency of 2.16 gigahertz, but it performs comparably to an Intel Pentium 4 at 3.0 gigahertz.

[16]A benchmark measures a processor's speed based on its actual computational performance on a common set of tasks, facilitating speed comparisons between different processors.

[17]One alternative would be to use the top K processors in each sub-group, using the average price and quality of these to form the composite products. The model produced qualitatively similar results with $K = 3$.

in a similar fashion.

Figures 2.2 and 2.3 plot the quality and price of each product over time. The quality plots show how Intel started as the dominant technology provider, but when AMD released its Athlon processor in mid-1999, the two firms entered into close technological competition. A similar story exists in the price graphs. Intel started as the high-cost and high-quality provider, while AMD served as the lower-cost and lower-quality provider. But after mid-1999, the price differences for each of the firms' products narrows significantly.

Finally, the frequency of the data varies: prices are available continuously, shipments quarterly, and ownership information semi-annually. Since the shipment and ownership data are less volatile than prices, I convert all the series to monthly observations, for a total of $T = 138$ time periods. Prices are set to the mean price observed in a month, shipments are distributed evenly over the quarter, and ownership shares are interpolated using cubic splines.[18]

2.3 The Model

In this section the basic structure of the model and the solution and estimation methods are discussed. The first subsection describes the product market and a consumer's period utility function. The second subsection describes the stochastic processes that consumers use to form expectations over future product qualities and prices. The third subsection presents the dynamic version of the consumer's decision problem. Finally, the fourth subsection shows how to calculate demand and provides the laws of motion for the distribution of product ownership.

[18]Tests using other conversion methods produce qualitatively similar results.

2.3.1 Basic Setup

Products are represented using a single, composite quality attribute $q_{jk} \in Q = \{1, 2, \ldots, \bar{q}\}$ for product j of firm k, with $p_{jk} \in \mathbb{R}_+$ the associated price. The market contains two firms, both of which sell two products. For each firm $k = 1, 2$, I refer to the higher quality product as the frontier product and the lower quality product as the non-frontier product. The vector $\mathbf{q}_t \in Q^4$ denotes the set of products available in a period and the vector $\mathbf{p}_t \in \mathbb{R}^4_+$ is their associated set of prices. Thus, the state of the product market can be summarized by the pair of vectors $\mathbf{s}_t = \{\mathbf{q}_t, \mathbf{p}_t\}$. Product qualities and prices evolve according to exogenous stochastic processes. Either or both firms may have the highest quality product in a given period. The frontier product for one firm may be lower in quality than the non-frontier product of the other firm. The time subscript will be dropped when possible.[19]

The basic problem for a consumer is whether to purchase a product today or to wait. For consumers who do not own a product, this represents a technology adoption decision: enter the market now or stay out. For consumers who own a product, this represents a replacement decision: purchase a more advanced product today or keep the existing product. Uncertainty enters into the decision because consumer are unsure about future product qualities and prices, while they know the utility they get from existing products. I incorporate heterogeneity by allowing consumers to be segmented according to their preferences for quality and price.

A consumer owns an existing product from firm k at time t from the set:

$$\tilde{Q}_t = \left\{\tilde{q}_k : \tilde{q}_k \in \left\{\cup_{\tau=1}^{t} \mathbf{q}_\tau\right\} \cup \{0\}\right\}$$

where $\tilde{q} = 0$ represents a consumer who owns no product. The set $\tilde{Q}_t$ is the set of products

[19]This specification is a generalization of Song (2003a).

that have been sold up until time t, such that $\tilde{Q}_t \subseteq \tilde{Q}_{t+1}$. Several restrictions are required for tractability. First, a consumer is constrained to own at most one product. Second, old products are either kept or discarded, that is, a second-hand market does not exist. Second-hand markets play an important role in some industries, such as the used automobile market, but the used PC market still represents a small share of overall demand.[20] Third, products retain their original quality indefinitely – there are no depreciation or upkeep costs.

The period utility function for a consumer in segment i, for $i = 1, \ldots, I$, who purchases some product $q_{jk} \in \mathbf{q}$ is:

$$u_i(q_{jk}, \mathbf{s}) = \gamma_i q_{jk} - \alpha_i p_{jk} + \xi_k + \varepsilon_{ijk}$$

where γ_i is consumer i's taste for quality, α_i is the marginal utility of income, ξ_k is a brand fixed-effect, and ε_{ijk} represents unobservable factors that influence the consumer's utility. In this case, the consumer pays for the cost of purchasing the new product. No restrictions are placed on which product the consumer can purchase: a consumer with a frontier product may choose to replace it with a non-frontier product.

On the other hand, a consumer may retain her existing product. Then the period utility for a consumer who owns $\tilde{q}_k$ and does not make a purchase is:

$$\underline{u}_i(\tilde{q}_k, \mathbf{s}) = \begin{cases} \gamma_i \tilde{q}_k + \xi_k + \varepsilon_{ik} & \text{if } \tilde{q}_k > 0 \\ \varepsilon_{i0} & \text{if } \tilde{q}_k = 0 \end{cases}$$

If the consumer owns a product ($\tilde{q}_k > 0$), they receive the utility associated with it without having to pay any additional cost. I normalize the utility to zero in the case that a consumer does not own a product ($\tilde{q}_k = 0$). This specification demonstrates the fact that a consumer's

[20]See Esteban and Shum (2004) for a model of the automobile industry with second-hand markets.

existing product represents her state-specific outside option. The utility of the outside option is a function of both a consumer's exogenously determined segment i and the endogenously determined past choice $\tilde{q}_k$.

2.3.2 Price and Quality Expectations

I assume that consumers possess rational expectations about the stochastic processes governing the evolution of prices and qualities. Let the stochastic process generating market state transitions follow a regular Markov transition kernel $\Pi(\mathbf{s}'|\mathbf{s})$. The processes are modeled independently, such that $\Pi(\mathbf{s}'|\mathbf{s}) = \Pi_{\mathbf{q}}(\mathbf{q}'|\mathbf{q})\Pi_{\mathbf{p}}(\mathbf{p}'|\mathbf{p})$. The impact of this assumption should be minimal because, in effect, prices are quality-adjusted. This is due to the fact that prices are measured with respect to the composite product qualities, which also change over time.

Price Expectations. The highly competitive nature of the CPU industry creates a significant amount of interdependence between processor prices. To capture these complicated relationships, prices follow a first-order vector autoregressive (VAR) process:

$$\log(\mathbf{p}_t) = A_0 + A_1 \log(\mathbf{p}_{t-1}) + z_t, \qquad z_t \sim N(0, \Sigma)$$

where A_0 and z_t are (4×1) vectors, and A_1 and Σ are (4×4) matrices. The cross-terms in each regression equation account for price competition between Intel and AMD and the off-diagonal elements of Σ capture the covariance between different product prices. Allowing for correlation in the random shocks further captures the co-movement of prices of the competing firm. While it is somewhat difficult to imagine a supply-side model that would generate such a price process, such a process is a reasonable assumption about consumers' expectations and memories. This approach is similar to Adda and Cooper (2000), who use a VAR to model the stochastic evolution of prices, and to Erdem, Imai, and Keane (2003), who model the price process for multiple products as contemporaneous functions of the price of a base product.[21]

Quality Expectations. Since product quality in the model is assumed to lie on a discrete grid, I discretize the continuous measure of quality obtained from the speed benchmark. Changes in product quality are modeled as discrete increases, or 'jumps', on a quality ladder. Let

$$\Delta(q_{jkt}) = q_{jkt} - q_{jk,t-1}$$

be the integer change in a product's quality from one period to the next, where $\Delta(q_{jkt}) \geq 0$ for all j, k, t. One possible approach is to use a simple Poisson count process to model the quality jumps. However, we observe under-dispersion in the number of zeros – periods in which the product quality does not change – and this can lead to inconsistent parameter estimates (Cameron and Trivedi, 1999). To account for this, I use a modified version of a zero-inflated Poisson (ZIP) process, and separately model the probability of no quality change and some positive change. The probability that a product's quality remains constant from one period to the next is allowed to depend on the current product quality:

$$\Pr(\Delta(q_{jkt}) = 0 | q_{jk,t-1}) = \Phi(\kappa_0 + \kappa_1 q_{jk,t-1})$$

where Φ is the standard normal distribution. The probability of a positive quality change is

$$\Pr(\Delta(q_{jkt}) = z | q_{jk,t-1}) = \frac{1 - \Delta_0(q_{jkt})}{1 - e^{-\lambda_{jk}}} \frac{e^{-\lambda_{jk}} \lambda_{jk}^z}{z!} \qquad \text{for } z > 0$$

where $\Delta_0(q_{jkt}) \equiv \Pr(\Delta(q_{jkt}) = 0 | q_{jk,t-1})$ and the first fraction is required as a normalization. Let $I_0 = I_{\Delta(q_{jkt})=0}$ be the indicator function representing no innovation. Assuming the innovation processes are independent across products, the transition kernel for all product

[21] See also Hall and Rust (1999), Song and Chintagunta (2003), and Sun (2005) for alternative mechanisms to model price expectations.

qualities is

$$
\begin{aligned}
\Pi_q(\mathbf{q_{t+1}}|\mathbf{q}_t) &= \Pr(\Delta(\mathbf{q}_t)|\mathbf{q}_t) \\
&= \prod_{q_{jkt}\in\mathbf{q}_t} \Delta_0(q_{jk,t+1}|q_{jkt})^{I_0} \Pr(\Delta(q_{jk,t+1}) = q_{jk,t+1} - q_{jkt}|q_{jkt})^{1-I_0}
\end{aligned}
$$

2.3.3 Dynamic Consumer Problem

I model the consumer's decision to purchase a new CPU as a dynamic optimization problem under price and quality uncertainty. A consumer's task is to decide whether to keep her existing product (if any) or to purchase one of the new products. Consumers are uncertain about future product qualities and prices, but possess rational expectations about the future state of the product market. The consumer is endowed with an initial product $\tilde{q}_0 \in \tilde{Q}_0$ at $t = 0$. The consumer then chooses a sequence of product purchases that maximize the sum of discounted expected future utility over the infinite horizon:

$$
\max_{\{q_t\in\mathbf{q}_t\cup\tilde{q}_{kt}\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{\infty}\beta^t\left[\underline{u}_i(\tilde{q}_{kt},\mathbf{s}_t)\cdot I_{\{q_t=\tilde{q}_{kt}\}} + u_i(q_t,\mathbf{s_t})\cdot(1-I_{\{q_t=\tilde{q}_{kt}\}})\right]\Big|\tilde{q_{kt}},\mathbf{s_t}\right\}
$$

where $E \equiv E_{\mathbf{s},\varepsilon}$ and the discount factor β is fixed. The recursive form of the consumer's optimization problem is written as:

$$
V_i(\tilde{q}_k,\mathbf{s},\varepsilon) = \max\left\{\underline{u}_i(\tilde{q}_k,\mathbf{s}) + \beta E[V_i(\tilde{q}_k,\mathbf{s}',\varepsilon'_{i,0})|\mathbf{s}], \max_{q_{j\ell}\in\mathbf{q}}\{u_i(q_{j\ell},\mathbf{s}) + \beta E[V_i(q_{j\ell},\mathbf{s}',\varepsilon'_{ij\ell})|\mathbf{s}]\}\right\}
$$

Alternatively, write a consumer's period utility function as

$$
U_i(q,\tilde{q}_k,\mathbf{s}) = \underline{u}_i(\tilde{q}_k,\mathbf{s})\cdot I_{\{q=\tilde{q}_k\}} + u_i(q,\mathbf{s})\cdot(1-I_{\{q=\tilde{q}_k\}})
$$

Then rewrite the dynamic problem for the consumer as

$$V_i(\tilde{q}_k, \mathbf{s}, \varepsilon) = \max_{q \in \mathbf{q} \cup \tilde{q}_k} \{U_i(q, \tilde{q}_k, \mathbf{s}) + \beta E[V_i(q, \mathbf{s}', \varepsilon')|\mathbf{s}]\} \tag{2.1}$$

where

$$E[V_i(q, \mathbf{s}', \varepsilon')|\mathbf{s}] = \int_{\varepsilon'} \int_{\mathbf{p}'} \sum_{\mathbf{q}'} V_i(q, \mathbf{s} + \Delta(\mathbf{q}'), \varepsilon')\nu(\varepsilon')\Pi_{\mathbf{q}}(\mathbf{s} + \Delta(\mathbf{q}')|\mathbf{q})\Pi_{\mathbf{p}}(\mathbf{p}'|\mathbf{p})$$

2.3.4 Demand

The aggregate demand for a product is determined by the solution of the consumer's decision problem for each type $(i, \tilde{q}_k)$ and the distribution of product ownership over consumer types. The distribution of product ownership in the next period is determined by consumer demand in this period. Denote the $(|\tilde{Q}_t| \times I)$ matrix F_t as the discrete ownership distribution, where $F_t(i, \tilde{q}_k)$ is the proportion of consumers who are of type i and own $\tilde{q}_k$ at time t. As time passes and product qualities increase, $\tilde{Q}_t \subseteq \tilde{Q}_{t+1}$, which implies that F_t also grows in the $\tilde{Q}_t$ dimension. Similarly, define $F_t(\tilde{q}_k) = \sum_i F_t(i, \tilde{q}_k)$ as the discrete marginal distribution of ownership across existing products and $F_t(i|\tilde{q}_k) = F_t(i, \tilde{q}_k)/F_t(\tilde{q}_k)$ as the discrete conditional distribution of consumer types given product ownership.[22]

Following Rust (1987), I assume that $\{\varepsilon_{ijk}\}$ are drawn from a multivariate extreme-value distribution. This produces the standard multinomial logit formula for product demand from consumers of type $(i, \tilde{q}_\ell)$, for $\ell = 1, 2$, who purchase some $q_{jk} \in \mathbf{q}_t$:

$$d_{jkt}(\tilde{q}_\ell, i) = \frac{\exp\{\bar{V}_i(q_{jk}, \tilde{q}_\ell, \mathbf{s}_t)\}}{\sum\limits_{q' \in \mathbf{q}_t \cup \tilde{q}_\ell} \exp\{\bar{V}_i(q', \tilde{q}_\ell, \mathbf{s}_t)\}} \tag{2.2}$$

[22] The specification of a finite number of types is the aggregate analogue to Kamakura and Russell's (1989) consumer level latent-class models.

where $\bar{V}_i(q_{jk}, \tilde{q}_\ell, \mathbf{s}_t)$ is the product-specific value function obtained after integrating out the unobserved consumer heterogeneity:

$$\bar{V}_i(q_{jk}, \tilde{q}_\ell, \mathbf{s}_t) = u_i(q_{jk}, \tilde{q}_\ell, \mathbf{s}_t) + \beta \int_{s_{t+1}} \log \left(\sum_{q' \in \mathbf{q}_{t+1} \cup q_{jk}} \exp \left\{ \bar{V}_i(q', q_{jk}, \mathbf{s}_{t+1}) \right\} \right) \Pi(\mathbf{s}_{t+1}|\mathbf{s}) \quad (2.3)$$

Let $\tilde{d}_{kt}(\tilde{q}_k, i)$ denote the set of consumers of type $(i, \tilde{q}_k)$ who choose not to make a purchase and retain their existing product.

The market size M_t is observed and evolves deterministically. Demand for a product is determined by integrating over consumer preferences and summing over all other existing products, which yields:

$$x_{jkt} = M_t \sum_{\substack{\tilde{q}_\ell \in \tilde{Q}_t \\ \tilde{q}_\ell \neq q_{jk}}} F_t(\tilde{q}_\ell) \sum_{i \in I} d_{jkt}(\tilde{q}_\ell, i) F_t(i|\tilde{q}_\ell) \quad (2.4)$$

Note that this represents the total new demand for a product, but not the proportion of consumers who will own the product in the following period. Consumers who already own this product and choose not to purchase anything must be accounted for in the next period distribution of product ownership. The number of consumers who do not purchase a new product and retain their existing product is

$$\tilde{x}_t = M_t \sum_{\tilde{q}_k \in \tilde{Q}_t} F_t(\tilde{q}_k) \sum_{i \in I} \tilde{d}_{kt}(\tilde{q}_k, i) F_t(i|\tilde{q}_k)$$

Market shares for current products are

$$\mu_{jkt} = \frac{x_{jkt}}{\tilde{x}_t + \sum_{q_{j'k'} \in \mathbf{q}_t} x_{j'k't}} \quad (2.5)$$

The proportion of consumers who own a product in the following period is the sum of those who purchased the product in the previous period, plus those who already owned the product and did not make a new purchase. For all $q_k \in \mathbf{q}_t \cap \mathbf{q}_{t+1}$, this is given by

$$F_{t+1}(\tilde{q}_k) = M_t^{-1} x_{jkt} + F_t(\tilde{q}_k) \sum_{i \in I} \tilde{d}_{kt}(\tilde{q}_k, i) F_t(i|\tilde{q}_k) \tag{2.6}$$

The law of motion for the marginal distribution over existing products that are no longer sold by either firm, such that $\tilde{q}_k \notin \mathbf{q}_{t+1}$, is

$$F_{t+1}(\tilde{q}_k) = F_t(\tilde{q}_k) \sum_{i \in I} \tilde{d}_{k,t}(\tilde{q}_k, i) F_t(i|\tilde{q}_k) \tag{2.7}$$

The law of motion for the conditional distribution of consumer tastes over existing products that are in the product market in the next period is defined as follows. For all $q_k \in \mathbf{q}_t \cap \mathbf{q}_{t+1}$, let

$$F_{t+1}(i|\tilde{q}_k) = \frac{\tilde{d}_{kt}(\tilde{q}_k, i) F_t(\tilde{q}_k) F_t(i|\tilde{q}_\ell) + F_t(\tilde{q}_\ell) \sum_{\tilde{q}_\ell \in \tilde{Q}_t} d_{jkt}(\tilde{q}_\ell, i) F_t(i|\tilde{q}_\ell)}{F_{t+1}(\tilde{q}_\ell)} \tag{2.8}$$

This expression captures consumers of type i who retained q_j who entered the market and those substituting ownership of this product for a more advanced product. The conditional distribution must also be updated differently for products that are not currently being marketed. For any $\tilde{q} \notin \mathbf{q}_t$

$$F_{t+1}(i|\tilde{q}_k) = \frac{\tilde{d}_{k,t}(\tilde{q}_k, i) F_t(i|\tilde{q}_k)}{\sum_{i' \in I} \tilde{d}_{k,t}(\tilde{q}_k, i') F_t(i'|\tilde{q}_k)} \tag{2.9}$$

2.4 Estimation

I estimate the model using GMM as part of a nested fixed-point. This procedure sets parameters that make the moments of the simulated model as close as possible to their empirical counterparts. First, the price and quality processes are estimated using maximum likelihood. These estimates are treated as known and substituted into the consumer's dynamic optimization problem. Second, given a parameter vector $\theta \in \mathbb{R}^d$, a nested-fixed point procedure minimizes a GMM objective function in the outer loop and computes the value function in the inner loop. This procedure is similar to others in the literature (Pakes, 1986, Rust, 1987, and Erdem and Keane, 1996) except that it is not feasible to estimate the model in a maximum likelihood setting, so instead I use a method of moments estimator.[23]

2.4.1 Overview of Estimation

The data contain information on the marginal distribution of product ownership $F_t(\tilde{q}_k)$, but not on the conditional distribution of consumer segments given product ownership $F_t(i|\tilde{q}_k)$. The next period values for the conditional distribution can be calculated using the laws of motion defined in the previous section, but I require estimates of the conditional distribution at $t = 0$ to serve as initial conditions. To address this issue, for each product, I need to estimate the initial proportion of consumers who belong to a particular segment. To reduce the number of parameters, I assume that the initial distribution of consumers across segments is identical for both firms for a given product. That is, I assume that $F_0(i|\tilde{q}_f^{Intel}) = F_0(i|\tilde{q}_f^{AMD}) \equiv F_0(i|f)$ and $F_0(i|\tilde{q}_{nf}^{Intel}) = F_0(i|\tilde{q}_{nf}^{AMD}) \equiv F_0(i|nf)$, where $\tilde{q}_f^k$ and $\tilde{q}_{nf}^k$ are the frontier and non-frontier products for firm k. I also need to estimate the proportion of each segment who do not own any product, $F_0(i|\tilde{q} = 0)$.

2.4.2 Identification

With the discount factor fixed, the dynamic parameters are identified through the combination of demand and ownership data and standard arguments found in Rust (1996). Replacement behavior is inferred by the relationship between changes in the distribution of

[23] Luan (2005) and Lee and Wolpin (2006), among others, have taken similar approaches.

ownership and period sales. Given the distribution of ownership and sales in a given period, these two quantities uniquely determine the distribution of ownership in the following period. Thus, the time series of ownership and demand helps me uncover consumer replacement behavior. Consumer heterogeneity is identified through differences in replacement cycle length and purchasing behavior. For example, a consumer who owns an old product and upgrades to a non-frontier product probably follows different preferences than a consumer who owns a slightly outdated product and purchases a frontier product. Variation in price and quality identifies the consumer's sensitivity to money and product quality. Combined variation in demand and ownership, as well the overall market penetration rate, helps identifies consumer heterogeneity through the rate at which a new product is purchased by different consumer segments.

2.4.3 Implementation

Combining the dynamic parameters of interest with the required initial conditions, the parameter vector $\theta \in \mathbb{R}^d$ contains the set of quality coefficients γ, price coefficients α, firm fixed-effects ξ, and the initial conditions.

Estimation proceeds as follows. First, for a given parameter vector, I solve the consumer's fixed-point problem for all possible consumer types $(i, \tilde{q}_k)$. Second, starting at $t = 1$, I use equations (2.2) and (2.5) to compute aggregate consumer demand, followed by equations (2.6) through (2.9) to calculate the implied distribution of ownership at $t = 2$. This is repeated until $t = T$ over the entire sequence of *observed* states. Third, I form moments based on the simulated and empirical values and minimize the GMM objective function.

The vector of moments consists of the period market shares $\mu_{jkt}(\theta)$ and the period ownership shares $F_t(\tilde{q}_k; \theta)$. I also include to additional moments that are useful in estimating the model, namely the penetration rate $\mu_{pen,t}(\theta)$ and the share of sales due to replacement purchases $\mu_{rep,t}(\theta)$. These are compactly written as $\boldsymbol{\mu}_t(\theta) = [\mu_{jkt}(\theta), \mu_{pen,t}(\theta), \mu_{rep,t}(\theta)]'$ and

$\mathbf{F}_t(\theta) = [F_t(\tilde{q}_k;\theta)]', \forall \tilde{q}_k \in \tilde{Q}_t$. To estimate parameters of the model, I assume the following moment condition holds:

$$E[m_t(\theta_0)] = 0$$

where $m_t(\cdot) \in \mathbb{R}^p$ with $p \geq d$ is a vector of moment functions that specifies the differences between the observed quantities and those predicted by the model. The moment conditions are:

$$m_t(\theta) = \begin{bmatrix} \boldsymbol{\mu}_t(\theta) - \hat{\boldsymbol{\mu}}_t \\ \mathbf{F}_t(\theta) - \hat{\mathbf{F}}_t \end{bmatrix}$$

A generalized method of moments estimator, $\hat{\theta}$, minimizes the weighted quadratic form:

$$Q_T(\theta) = \min_{\theta \in \mathbb{R}^d} \frac{1}{2} \left[\frac{1}{T} \sum_{t=1}^{T} m_t(\theta) \right]' \Omega \left[T^{-1} \sum_{t=1}^{T} m_t(\theta) \right] \tag{2.10}$$

where Ω is an $p \times p$ positive semidefinite weighting matrix. Under the assumption that $\Omega \xrightarrow{p} \Omega_0$, define the $p \times d$ matrix $\mathbf{M}_0 = E[\nabla_\theta m_t(\theta_0)]$. Let $\Lambda_0 = E[m_t(\theta_0)m_t(\theta_0)']$ and substitute a consistent estimator for Λ_0^{-1} into the weighting matrix. Under the standard assumption that $m_t(\theta)$ is independent across t, we have that

$$\sqrt{T}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, (\mathbf{M}_0' \Lambda_0^{-1} \mathbf{M}_0)^{-1}/T) \tag{2.11}$$

The first obstacle in using this standard GMM method is that errors in the moment conditions are likely to be autocorrelated. The GMM estimator remains consistent under dependent data, but the covariance matrix needs to be corrected to take this dependence into consideration. In particular, I replace the asymptotic covariance matrix of the moment functions in equation (2.11) using the estimator in Newey and West (1987), written as

$$\tilde{\Lambda}_0 = \hat{\Gamma}_0 + \sum_{v=1}^{q(T)} \left[1 - \frac{v}{q+1} \right] (\hat{\Gamma}_v + \hat{\Gamma}_v')$$

where $\hat{\Gamma}_v = T^{-1}\sum_{t=v+1}^{T} m_t(\hat{\theta})m_{t-v}(\hat{\theta})'$ and $q(T)$ grows at a sufficiently slow rate.

The full two-step GMM procedure is used to produce efficient standard error estimates. Note that the primary source of error in the estimation processes arises from sampling error, as opposed to simulation error.[24] I use the non-derivative based Nelder-Meade algorithm to get within a neighborhood of the optimal parameters and then switch to a quasi-Newton method.[25]

Computing the fixed-point requires integrating over the continuous four-dimensional price vector. I use the randomization technique developed by Rust (1997a) to solve for an approximate value function.[26] After integrating out the idiosyncratic components of the utility function, the continuous form of the value function is replaced with a discretized equivalent. Let $P \subset \mathbb{R}^4_+$ be the compact space of price vectors. Take N_p i.i.d. uniform random draws from P to produce the set of random grid points $\{\tilde{p}_1, \ldots, \tilde{p}_{N_p}\}$.[27] To guarantee the discretized probabilities are sufficiently smooth and sum to one, the original continuous transition function is replaced with the discrete probability densities $\Pi_{N_p}(\tilde{\mathbf{p}}'|\tilde{\mathbf{p}})$ constructed using the normalization:

$$\Pi_{N_p}(\tilde{\mathbf{p}}'_a|\tilde{\mathbf{p}}) = \frac{\Pi(\tilde{\mathbf{p}}'_a|\tilde{\mathbf{p}})}{\sum_{i=1}^{N_p}\Pi(\tilde{\mathbf{p}}'_i|\tilde{\mathbf{p}})}$$

Denote $\hat{V}_{i,N_p}(q_{jk}, \tilde{q}_\ell, \mathbf{q}, \mathbf{p})$ as the discretized version of the product-specific value function for

[24]Errors due to the randomized approximation of the integral enter nonlinearly into the moment conditions. See Benitez-Silva *et al* (2000) for an evaluation of different approximation techniques.

[25]I also perform a check of the numerical condition for local identification. Let $\hat{m}_t^s(\theta)$ be a subvector of $m_t(\theta)$ such that $\dim(\hat{m}_t^s) = \dim(\theta)$. Then a local identification condition requires that $\det\left(\frac{\partial \hat{m}_t^s}{\partial \theta}\right) \neq 0$. Roughly interpreted, if the determinant of the Jacobian is non-zero then the moments m_t are informative about the structural parameters θ.

[26]Standard Monte Carlo integration techniques are not appropriate because the value function has to be evaluated at arbitrary random points, which may lie off the predefined state space grid.

[27]See the literature on minimum discrepancy grids cited in Rust (1997b) for alternatives to uniform grid distributions.

$q_{jk} \in \mathbf{q} \cup \tilde{q}_\ell$. Since the approximation is computed over the price vector, the product market state variable, $\mathbf{s}$, has been split up into its two components, $\{\mathbf{q}, \mathbf{p}\}$. The discretized version of the product-specific value function can be computed based on the expectation over the product-specific value function:

$$\hat{V}_{i,N_p}(q_{jk}, \tilde{q}_\ell, \mathbf{q}, \mathbf{p}) = U_i(q_{jk}, \tilde{q}_\ell, \mathbf{s}) + \beta \sum_{\mathbf{q}'} \sum_{\tilde{p}_a}^{N_p} \hat{v}_{i,N_p}(q_{jk}, \tilde{q}_\ell, \mathbf{q}', \tilde{\mathbf{p}}'_a) \Pi_q(\mathbf{q}'|\mathbf{q}) \Pi_{N_p}(\tilde{\mathbf{p}}'_a|\tilde{\mathbf{p}}) \quad (2.12)$$

where

$$\hat{v}_{i,N_p}(q_{jk}, \tilde{q}_\ell, \mathbf{q}', \tilde{\mathbf{p}}'_a) = \log \left(\sum_{q' \in \mathbf{q}' \cup q_{jk}} \exp \left\{ \bar{V}_i(q', q_{jk}, \mathbf{q}', \tilde{\mathbf{p}}'_a) \right\} \right)$$

The contraction mapping defined in (2.12) converges stochastically to the true value function under certain regularity conditions, which are satisfied in this case. The randomized Bellman approach is appealing because it does not require interpolation and the random grid points are drawn once and then remain fixed at successive iterations.[28]

2.5 Parameter Estimates

This section presents the parameter estimates from the first-stage estimation of the price and quality expectations processes, an evaluation of model fit and comparison, the estimates from the dynamic structural model, and a comparison of the price elasticities from a permanent price change under the benchmark and myopic models.

[28] In practice, I use a random multigrid algorithm to solve for the value function. More details can be found in the Appendix.

2.5.1 Price and Quality Expectations.

The parameter estimates for the price and quality processes are reported in Tables 2.1 and 2.2. Overall, both performed well. For the price VAR, all the own-price coefficients are significant at the 99% level. Also, several of the cross-price coefficients are significant, suggesting that competitor's prices do have some affect on a firm's own price decisions. Price competition appears to be asymmetric: Intel's prices have a larger effect on AMD's prices than the converse. All the roots lie inside the unit circle, indicating that the VAR is stable.[29]

Similarly, the parameters of the quality process were estimated with a high degree of precision. The significance of the κ_1 parameters indicates that there is some non-stationarity in the probability of a product's quality changing, though given the magnitude of the parameter, this effect is not very large. The values of κ_1 for AMD's products are larger because AMD produced more quality innovations later in the sample period, even though these innovations tended to be smaller than the average innovation from Intel.

2.5.2 Model Fit and Comparison.

I estimated both myopic and dynamic specifications of the model while varying the numbers of segments. The models are evaluated using several measures. Table 2.3 reports the mean squared errors (MSE) for each set of moments, the objective function value and J-statistic, and the Distance Metric (DM) statistic.

The MSE's for the moments show that the two-segment dynamic model fits best. The myopic model performs the worst, particularly on fitting the replacement share and ownership share moments. This is not surprising because one would not expect a static demand model to adequately capture replacement behavior, which is inherently dynamic. The p-

[29]Higher-order lags do not add significant predictive power and including the set of product qualities as exogenous regressors only increases the R^2 by roughly 0.2%.

values for the J-statistics show that none of the models are rejected by the data, although the p-value for the myopic model comes close. To compare the models to one another, I use the DM statistic, $DM = 2T[Q_T(\hat{\theta}_1) - Q_T(\hat{\theta}_2)] \sim \chi^2$, which is the GMM counterpart of the likelihood ratio test. I compare the homogeneous and myopic models against the benchmark model.[30] This produces test statistics of 24.04 and 7.98 for the myopic and homogeneous cases, respectively, both of which are significant at the 99% level, indicating that the forward-looking model with heterogeneity should be preferred to the alternatives and is most consistent with the data. The two-segment dynamic model produces similar results under alternative values of the discount factor (0.95 and 0.99).

Figure 2.4 shows the market penetration rate as predicted by the model versus the actual penetration. The predicted rates were calculated at each time period given the observed level. The model does a good job of fitting the penetration rate, though the predictive accuracy does suffer a little during a few periods. One concern in the model is that the value of the absolute outside option – not owning a PC processor at all or owning a Mac-based chip – might be changing over time. Fortunately, given the accuracy of the model in predicting the share of non-owners in the population, this issue does not appear to be significant.

Figures 2.4 shows the market penetration rate generated by the model versus the empirical quantities. The predicted penetration rate is calculated in each time period given the observed quantities. The plot shows that the model does a good job of fitting the penetration rate, though the predictive accuracy does deteriorate during a few periods. One concern in the model is that the value of the absolute outside option – not owning a PC processor at all or owning a Mac-based chip – might be changing over time. Fortunately, given the accuracy of the model in predicting the share of non-owners in the population, this issue does not appear to be significant.

[30] See Newey and West (1987).

2.5.3 Structural Parameters.

Parameter estimates for the different specifications are presented in Table 2.4, all of which are statistically significant.

The parameter estimates in Table 2.4 show that a similar segment structure exists between the myopic and dynamic models: each has one segment that is more price sensitive and less quality sensitive and another segment that is less price sensitive and more quality sensitive. Henceforth, I refer to the first segment as the low valuation segment and the second segment as the high valuation segment. One difference between the myopic and dynamic estimates is that myopic values are smaller in absolute value. In all the specifications, the difference in the firm fixed-effects indicates that consumers place a significant premium on Intel processors. This is expected because a high value is required for the model to rationalize the difference in market shares and prices for the two firms.

I calculate the size of each consumer segment by combining the estimates of the initial conditions with the initial distribution of ownership. The initial conditions imply that at the beginning of 1993, 94% of non-owners were low segment consumers. High segment consumers, although only 14% of the population, owned disproportionate shares of both frontier and non-frontier processors. Overall, the benchmark dynamic model estimates that the low segment consumers make up approximately 85% of the population.

One distinguishing feature of the model is that, while consumer segments are static, the mix of segments among owners varies over time. Figure 2.5 displays the proportion of each consumer segment from the dynamic model for the set of owners. The plot reveals a familiar story: high segment consumers made up the majority of owners early in the market's history, declined as a portion of all owners over time. At the beginning of the sample period, the high segment consumers represented slightly more than half of all owners, despite the fact that they only represent 14.6% of the population. As the market penetration increased, the share of owners who belonged to the high segment declined. Figure 2.6 breaks down the segments for the frontier and non-frontier products, and shows that more low segment consumers own the non-frontier than the frontier.

2.5.4 Price Elasticities

The elasticity estimates are based on *permanent* changes in the price of a product to capture the long-term effects of the change on the consumer's expectations.[31] The elasticities are generated as follows. First, I use the observed quantities to solve the consumer problem and estimate a baseline for demand. I then create the new price time series by adjusting the values from one period to the end of the sample. This is repeated for each time period. The price process is re-estimated using the new time series (though the changes in parameter estimates were negligible). All other prices remain fixed. Finally, I solve for the optimal consumer behavior given these alternate time series and compare the new demand estimates to the baseline. The reported estimates are the average of the elasticities calculated in each period using the observed quantities.

Table 2.5 provides summaries of the price elasticities in the benchmark dynamic model with two segments and the myopic model with two segments based on price increases of 10 percent. Table 2.6 decomposes these elasticity values according to each potential consumer choice. The last two columns of Table 2.6 contain the cross elasticities for each product with respect to a consumer's no-purchase option.

First, there is an asymmetry in the market structure: Intel's products have a larger impact on AMD's products than the converse. One somewhat counterintuitive result is that the more established brand, Intel, has higher own-price elasticities for the non-frontier products compared to the lesser-known brand, AMD. Second, non-owners are more sensitive to price changes than owners, because non-owners must have a larger marginal return for product

[31]Erdem, Imai, and Keane (2003) and Hendel and Nevo (2005) estimate elasticities in a similar fashion.

adoption than an owner does for product replacement. For example, a 10% increase in the price of Intel's frontier product raises demand by 2.3% from owners versus 3.1% from non-owners. Third, I find that the myopic version of the model underestimates price elasticities by 30 percent to 40 percent. Myopic consumers do not consider the future utility associated with owning a product (leading to downwardly biased price and quality coefficients), and thus under-react to a permanent price change.

I also used clout and vulnerability, two common measures in the marketing literature, to analyze the evolution of market structure over time for each composite product (Kamakura and Russell, 1989). The clout at time t for product j of brand k is defined as the total impact of this product on the demand for *all* other products, and the vulnerability is the total impact of all other products on the demand for this product. I included both products from each brand because the products are close substitutes and each firm is concerned about product line cannibalization.[32]

Figure 2.7 plots the time path of the clout and vulnerability for each product over time. The plot shows that Intel's frontier product has generally led the market, while AMD's non-frontier product has lagged the market. AMD's frontier product and Intel's non-frontier product have become increasingly close competitors, suggesting that AMD's premium chips are competing more closely with Intel's Celeron chips, its value-line of processors. Overall, these asymmetries suggest that while AMD has made some competitive gains in the industry over time, it is still unable to compete head on with Intel.

[32]The clout and vulnerability for a product is defined as

$$\text{clout}_{jkt} = \sum_{q' \in \mathbf{q}_t, q' \neq q_{jk}} \eta(q_{jk}, q') \qquad \text{vulnerability}_{jkt} = \sum_{q' \in \mathbf{q}_t, q' \neq q_{jk}} \eta(q', q_{jk})$$

where $\eta(q_{jk}, q')$ is the cross-price elasticity of q_{jk} with respect to q'.

2.6 Innovation and Replacement Cycles

This section uses the parameter estimates to analyze consumer behavior in the model. First, I characterize the replacement behavior of consumers implied by the benchmark model and conduct policy simulations of replacement behavior under alternative rates of technological innovation. Second, I show how consumer heterogeneity in both preferences and replacement behavior could provide some insights into how firms might tailor their strategies according to consumer replacement cycles.

2.6.1 Product Replacement

I begin by documenting the features of replacement cycles implied by the model, and then conduct a policy simulation to examine the effects of alternative rates of innovation on replacement behavior.

Characterizing Replacement Cycles. Table 2.7 reports the average replacement cycle length by consumer segment over the entire sample, in the benchmark and myopic cases. The results from the benchmark case indicate that, on average, consumers replace their existing processors approximately every 3.29 years. This figure is consistent with estimates from the market researcher firms IDC and Gartner. As expected, consumers in the high segment have a shorter replacement cycle than those in the low segment. The myopic consumer model underestimates replacement cycle length, suggesting that consumers replace their products more frequently than implied by the dynamic model. Table 2.7 also compares the length of the replacement cycle in the first half of the sample to the second half, showing that the average increased 2.7 years before 1999 to 3.6 years after 1999.[33]

[33] These figures are generated by tracking product replacement within a given time period until every consumer who owned a product in the beginning had replaced their product. Thus, a consumer who purchased a product in 1998 and replaced it in 2002 was counted under the pre-1999 period.

To demonstrate the degree of variation across consumers, Figure 2.8 displays the distribution of replacement cycle lengths in the first and second half of the sample, for each consumer segment. The mean and variance of each replacement distribution has increased over time, with both distributions shifting to the right. There is less variation in the replacement behavior of the high segment consumers because they represent a smaller part of the consumer population. The distribution for the low segment consumers is particularly skewed to the right, illustrating the fact that some consumers rarely replace their products.

Figure 2.9 shows that the mean replacement cycle length for each consumer segment has steadily increased over the sample period. The rate of change appears to slow during slower periods of technological innovation and adoption.

Alternative Innovation Rates. Empirically, the average product quality tends to double about every two years. For this policy simulation, I alter the parameters of the product quality process so that product qualities are expected to double every 1.5 years and one year, representing increases in the rate of innovation by 25% and 50%, respectively. All prices are fixed.[34]

Table 2.8 provides a comparison of the benchmark replacement cycle estimates to those generated by the policy simulations. A 25% increase in the rate of technological innovation lowered the mean replacement time by 10.7%, or about 4 months, and a 50% increase lowered the mean by 25.5%, or about 10 months. In both cases, the effect of the faster innovation rates was greater for the high segment consumers than the low segment consumers. For the high segment consumers, the alternative innovation rates reduced the mean replacement cycle length by 14.5% and 34.0%, compared to reductions of 10.2% and 24.5% for the low segment consumers. These results suggest that technology innovations can have a significant

[34]One issue with this policy simulation is that we would expect the firms to alter their pricing strategies given the new rates of innovation and consumer demand. An equilibrium model of competition in this industry would be required to account for the endogeneity of prices in policy simulations.

impact on the replacement cycle, especially for consumers who place a premium on processor quality.

Table 2.9 examines how the length of the replacement cycle responds over time to changes in the innovation rate. I fix the innovation parameters at their default values until a particular period. After this period, I increase the rate of innovation for all subsequent periods and examine the resulting change in the length of the replacement cycle. One column in Table 2.9 shows the percentage change in replacement cycle length between the benchmark and alternative models. The declining percentages in the column show that the marginal effect of innovation on the length of the replacement cycle has decreased over time. This implies that in the future hardware manufacturers may have to rely less on quality improvements to generate replacement sales. The effect is more pronounced in the case of a 50% increase in the rate of innovation compared to a 25% increase.

2.6.2 Replacement Cycle Segmentation

There is significant variation in replacement behavior across consumer segments and over time. This variation arises due to differences in the marginal distribution of ownership across the two segments and differences in consumers' replacement cycles lengths.

Figure 2.10 plots the distribution of ownership by product age, as implied by the model, for the years 2000 and 2004. The product age is defined as the number of years since the consumer purchased the product. For example, in 2000, Figure 2.10 shows that roughly 30 percent of owners had purchased within the last year, roughly 30 percent in the last one to two years, roughly 16 percent in the last two to three years, and so on. The key observation is that the mode of the distribution of ownership shifted during these years. In 2000, a large number of consumers had bought in the past two years, but in 2004, more consumers owned a product that was about three years old. Given that the mean replacement cycle length in this period was close to 3.5 years, we would have expected a significant number of consumers to replace their products starting in the beginning of 2005.

However, Figure 2.10 ignores the variation across consumer segments in product ownership and replacement behavior. To take advantage of the differences in the consumer segments, Figure 2.11 decomposes the distribution of ownership into the segments implied by the model.

A brief comparison of the upper and lower graphs demonstrates the asymmetry in each segments' distribution of product ownership and replacement. In 2000, a large number of high segment consumers (lower graph) had recently replaced their products, but in 2004, a larger number had replaced their products in the past two years. Given the estimates of the high segment's replacement cycle length, it would have been unlikely for many high segment consumers to replace their products in early 2005, because so many of them had purchased new products so recently.

This should be contrasted with the results for the low segment (upper graph). In 2000, a large portion of low segment consumers had recently purchased in 2000, but in 2004, many of the low segment consumers had products that were between three and four years old. This implies that there was a disproportionately large number of low segment consumers who, at the end of 2004, were increasingly likely to replace their products.

2.7 Conclusions and Future Research

As a technology market matures, replacement sales must eventually surpass adoption sales. A model of product replacement is required to help understand the dynamics of demand when multiple segments of consumers are simultaneously making adoption and replacement decisions. This paper presents a structural model of dynamic demand for PC processors that explicitly accounts for the replacement decision. The model helps provide an understanding of the impact of price and quality changes on consumer replacement behavior. Taking into account the forward-looking behavior of consumers is necessary in high-tech markets with rapid innovation, and modeling consumer heterogeneity in both preferences and product ownership has important implications for managers' strategies.

At the firm level, the results show that a manager can take advantage of variation in consumer replacement cycles across segments to adjust their product introduction and pricing strategies. First, the distribution plots in Figure 2.11 suggest that more value-oriented product could have been released in 2005 to target the replacement cycle of the low-segment consumers. Second, the menu of prices in the firms' product lines could have been adjusted: non-frontier product prices might have increased, and frontier product prices could have decreased. These actions could potentially increase profit on the non-frontier set of products, while encouraging some low-segment consumers to upgrade to the frontier products.

Firms must also recognize that their pricing strategies have an impact on consumers' aggregate replacement cycles. For example, a period of significant price competition will encourage more price sensitive consumers to replace their products in the near-term. This segment of consumers probably has a longer replacement cycle than less price sensitive segments. Given this, it may be useful for firms' to collect additional data to track consumer replacement behavior to gain a better understanding of the preferences of users who are replacing at different points in time. This information could be integrated into a CRM system (Lewis, 2005). Understanding and managing consumer replacement cycles is a largely unexplored facet of CRM, but certainly, a consumer's replacement cycle length plays a large role in determining their lifetime consumer value.

At the industry level, the results from Section 2.6.1 support observations that the replacement cycle length is increasing.[35] If processor speed continues to lose importance as a purchasing characteristic, processor manufacturers should search for other ways to differen-

[35]See CNET News.com, "Software: The End of Forced Upgrades?", October 23, 2001 and FIND/SVP, PC Trend Report, 2004.

tiate their products. To be clear, this does not imply that processor manufacturers should reduce their investments in R&D. Processor speed will always matter, to some degree, but manufacturers must attempt to differentiate their products along dimensions other than speed. There is already some evidence that Intel has begun to follow such a strategy with its Intel Centrino™ and Intel Viiv™ brands. Each product contains a processor as part of a larger system, with the first targeted at providing wireless capabilities in laptops, and the second within so-called "media-center" PC's. The implications for retailers are similar, whose interests are firmly linked to those of the manufacturers.

The results also provide a possible explanation for the apparent increase in consumers' replacement cycles from three years in the 1990's to four years after 2000.[36] A number of factors have been offered to explain this observation, such as evolving software requirements on the part of consumers. The results in this paper suggest one non-exclusive, yet intuitive, explanation that has been overlooked: the average replacement cycle length has increased because, as market penetration has risen, more consumers with inherently longer replacement cycles have entered the market.[37]

The analysis is not without its limitations, and these suggest possible avenues for future research. First, the model does not consider the role of software. A model that links a consumer's decision to purchase a computer with their software demand might be able to shed some light on the historical relationship between the hardware and software industries. Second, the number of households with multiple PC's has increased over the years, as has consumer ownership of laptop computers. If suitable data were available, a model could examine the effects of existing PC's on replacement behavior, and the potential substitution effects between desktop PC usage and laptop usage. Third, while the survey data shows

[36] Gartner Inc., Global PC Forecast Q2 Update, 2004 and FIND/SVP, PC Trend Report, 2004.

[37] To be clear, the paper is unable to rule out alternative explanations of this phenomenon. Multiple factors are likely the source of the change. However, most analysis of the industry appears to have ignored this averaging effect as a possible explanation.

that consumers view the processor as the dominant characteristic of a PC, they also care about the size of the hard drive, the amount of memory, and other features. Modeling the consumer's choice of the entire PC could reveal additional insights for the firm in terms of product line design and cannibalization. Fourth, alternative specifications could be used to model consumer expectations about price and quality, such as allowing innovations across products to be correlated.

Finally, the current model assumes that price and quality are exogenous. To address this issue requires an equilibrium model with both dynamic consumers and dynamic firms making decisions simultaneously. This would allow us to understand precisely how forward-looking behavior on the consumers' side affects the firms' policy functions. I develop such an equilibrium model in the next chapter.

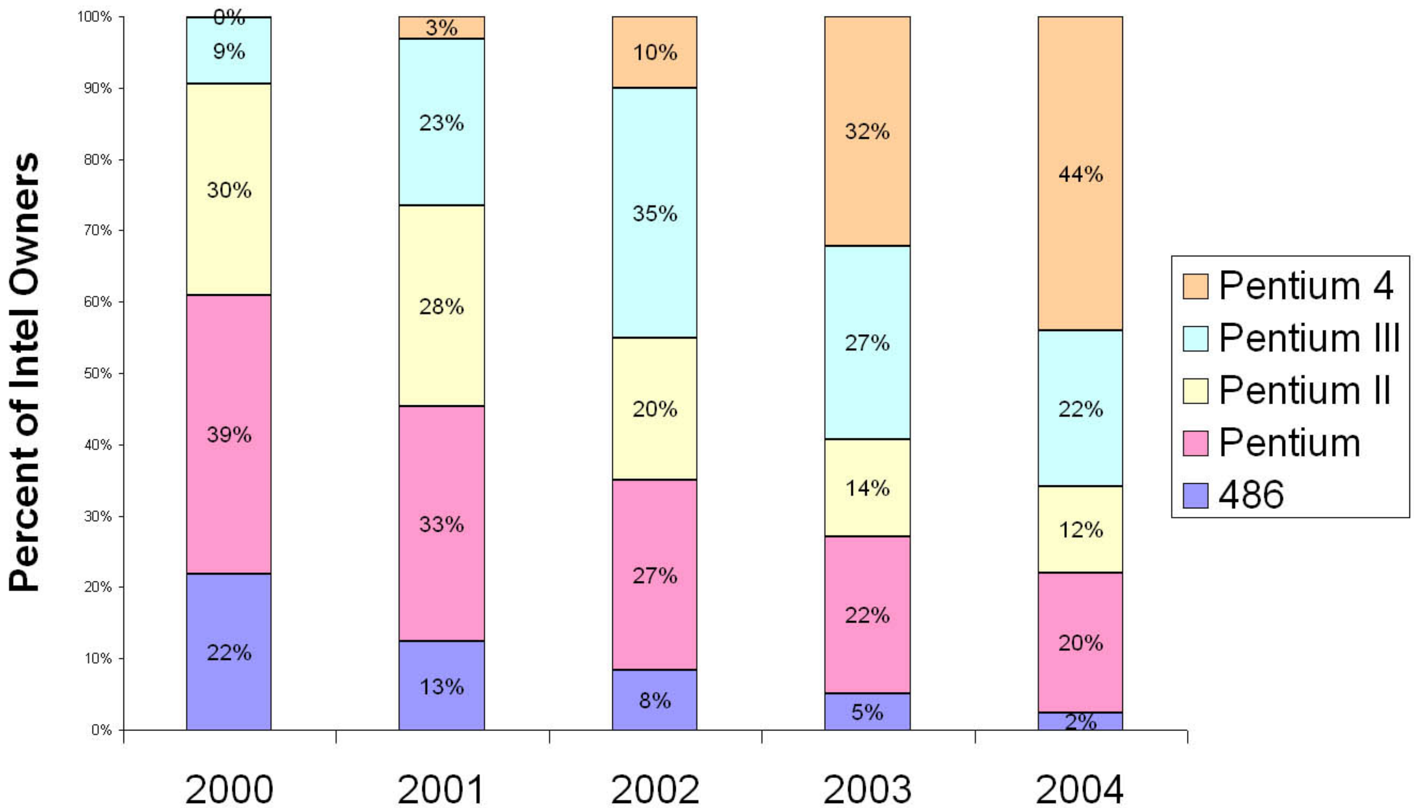

Figure 2.1: PC processor ownership by Intel processor architecture from 2000 to 2004. Note that the Pentium 4 processor was introduced in April 2001.

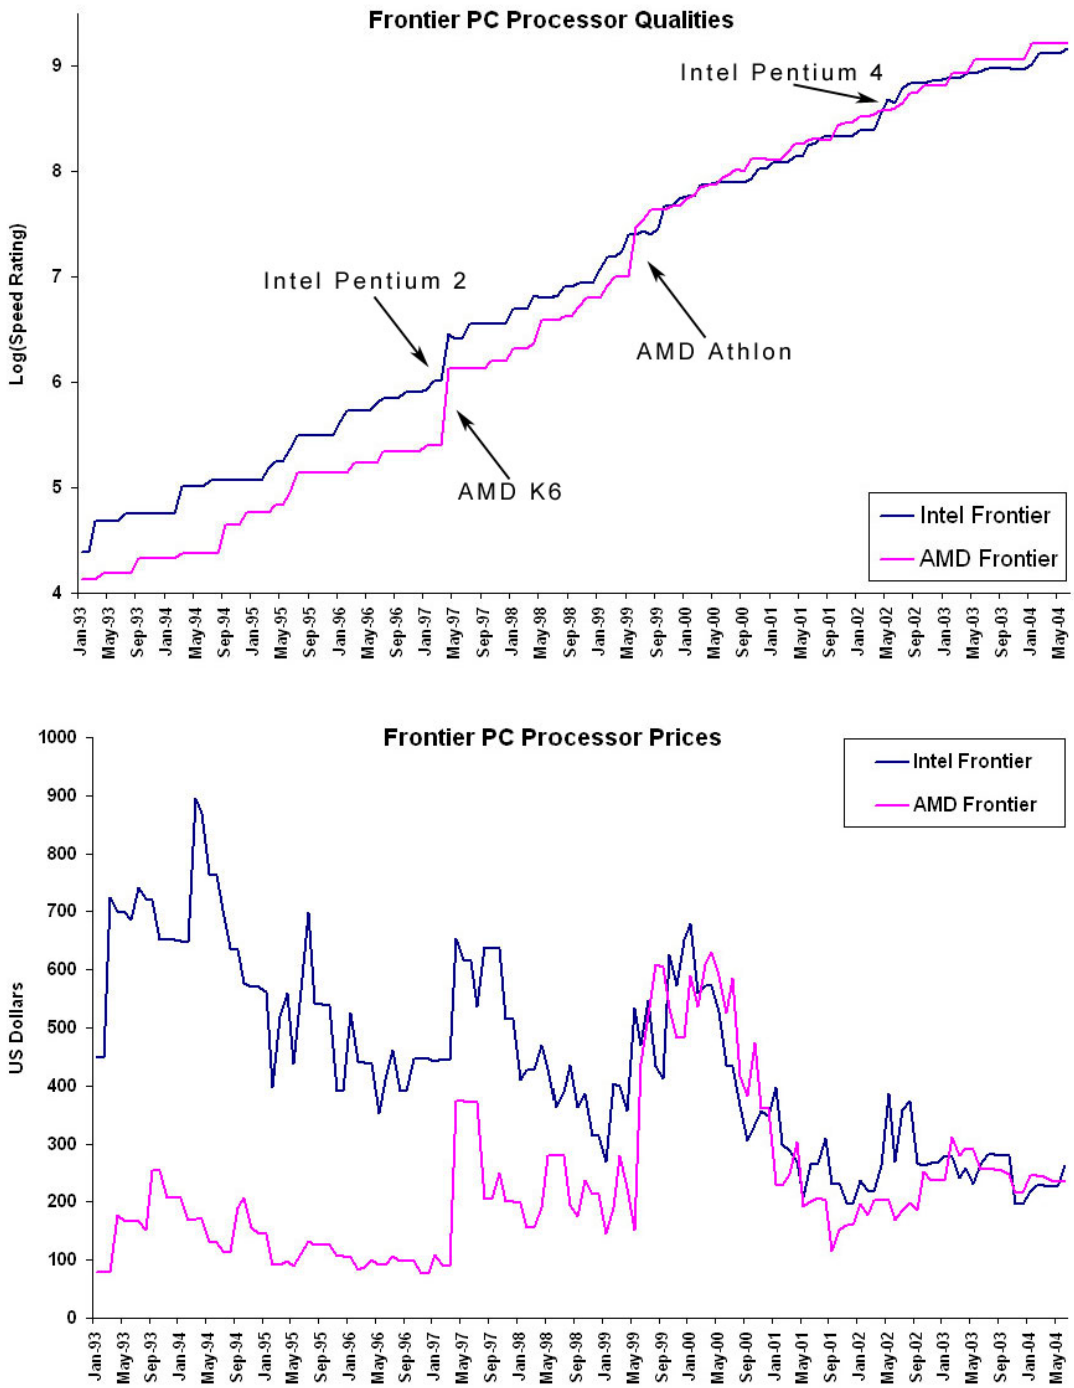

Figure 2.2: Composite frontier qualities and prices for Intel and AMD. Prices are in constant January 2000 dollars.

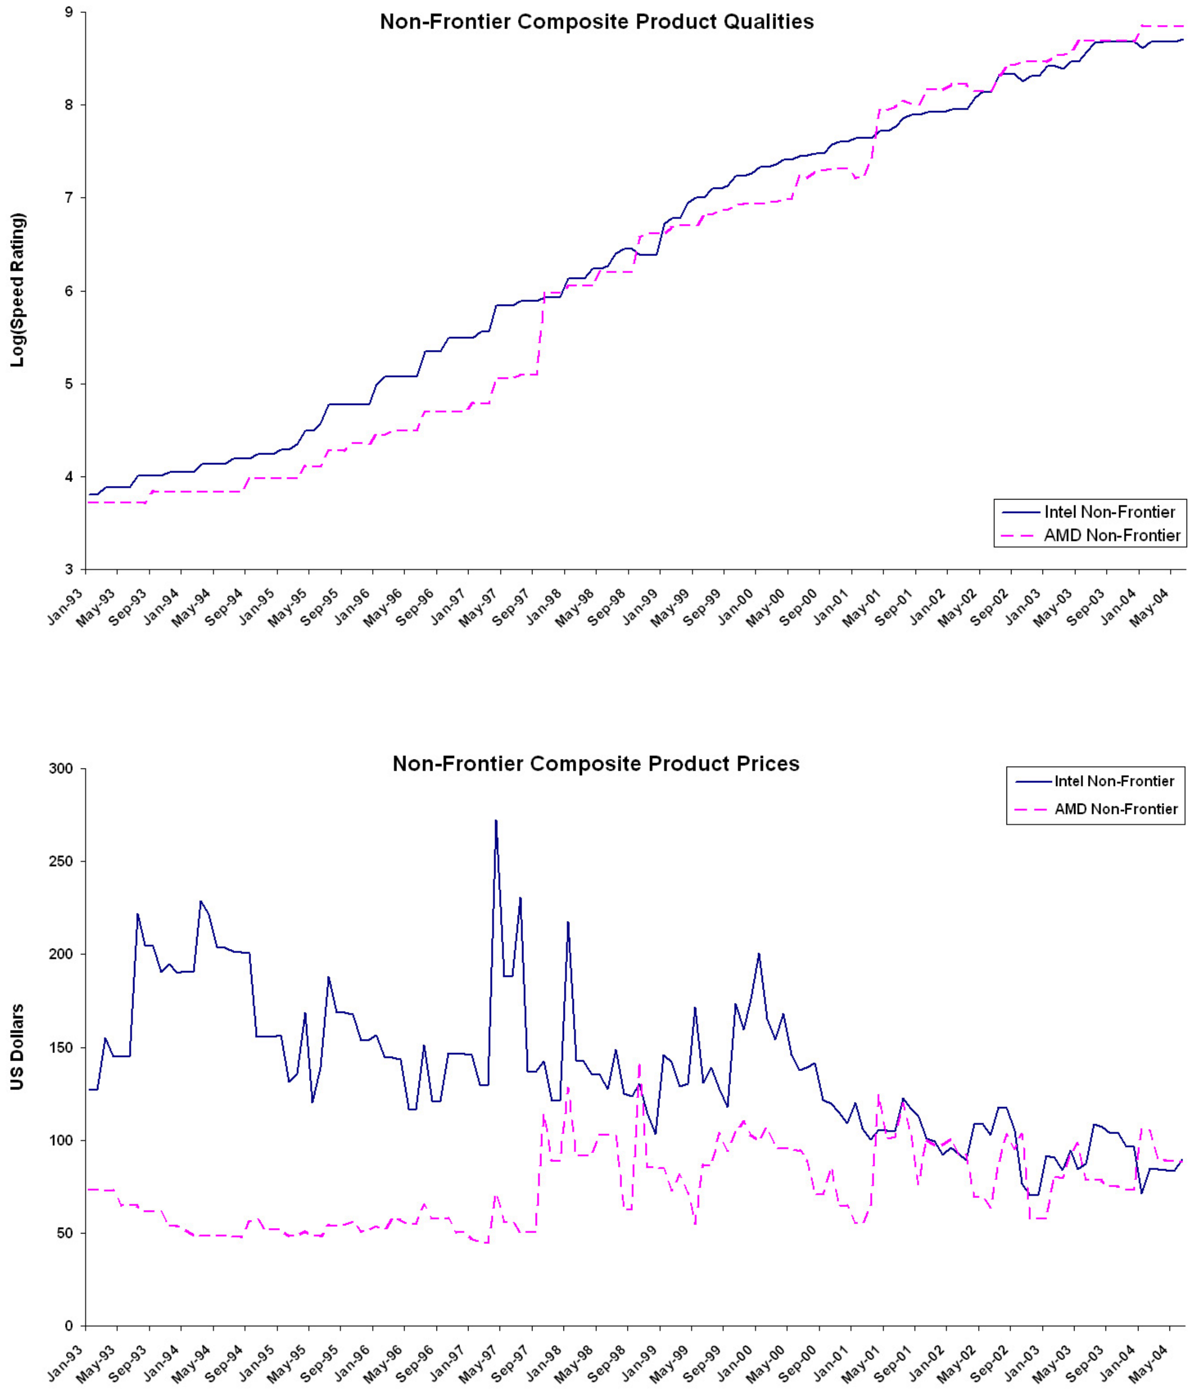

Figure 2.3: Composite non-frontier qualities and prices for Intel and AMD. Prices are in constant January 2000 dollars.

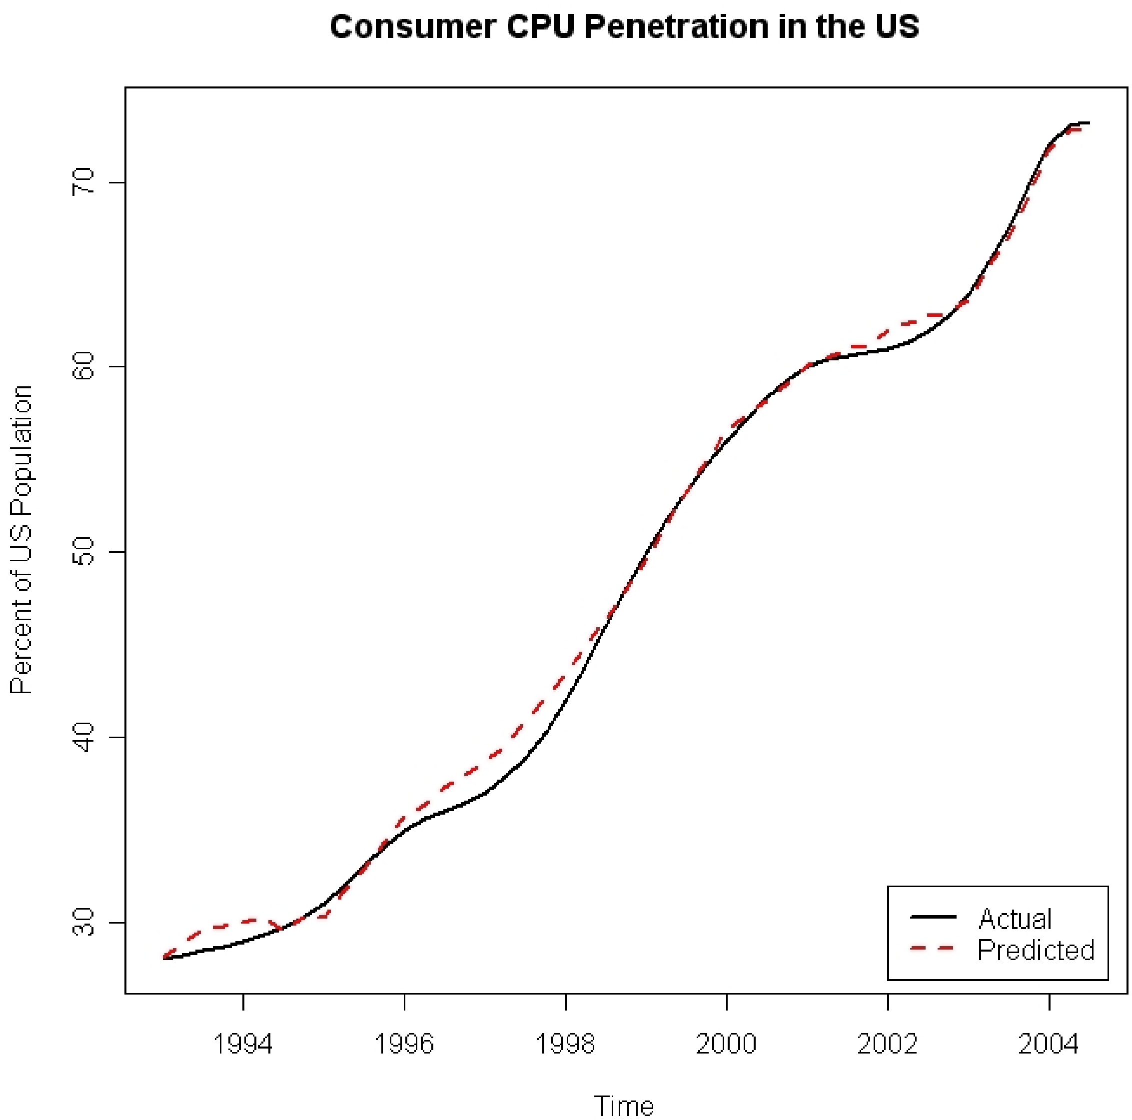

Figure 2.4: Actual versus predicted market penetration rate.

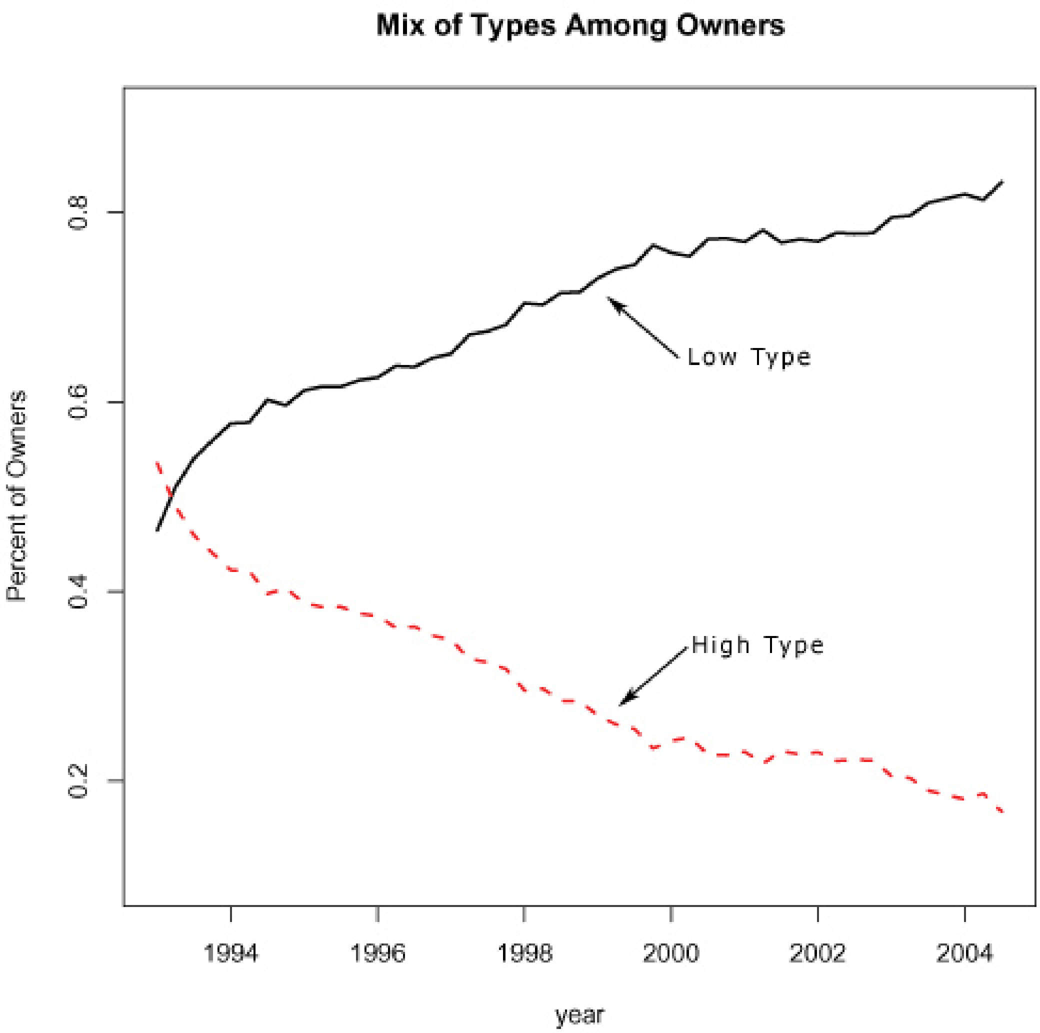

Figure 2.5: Distribution of consumer segments among owners.

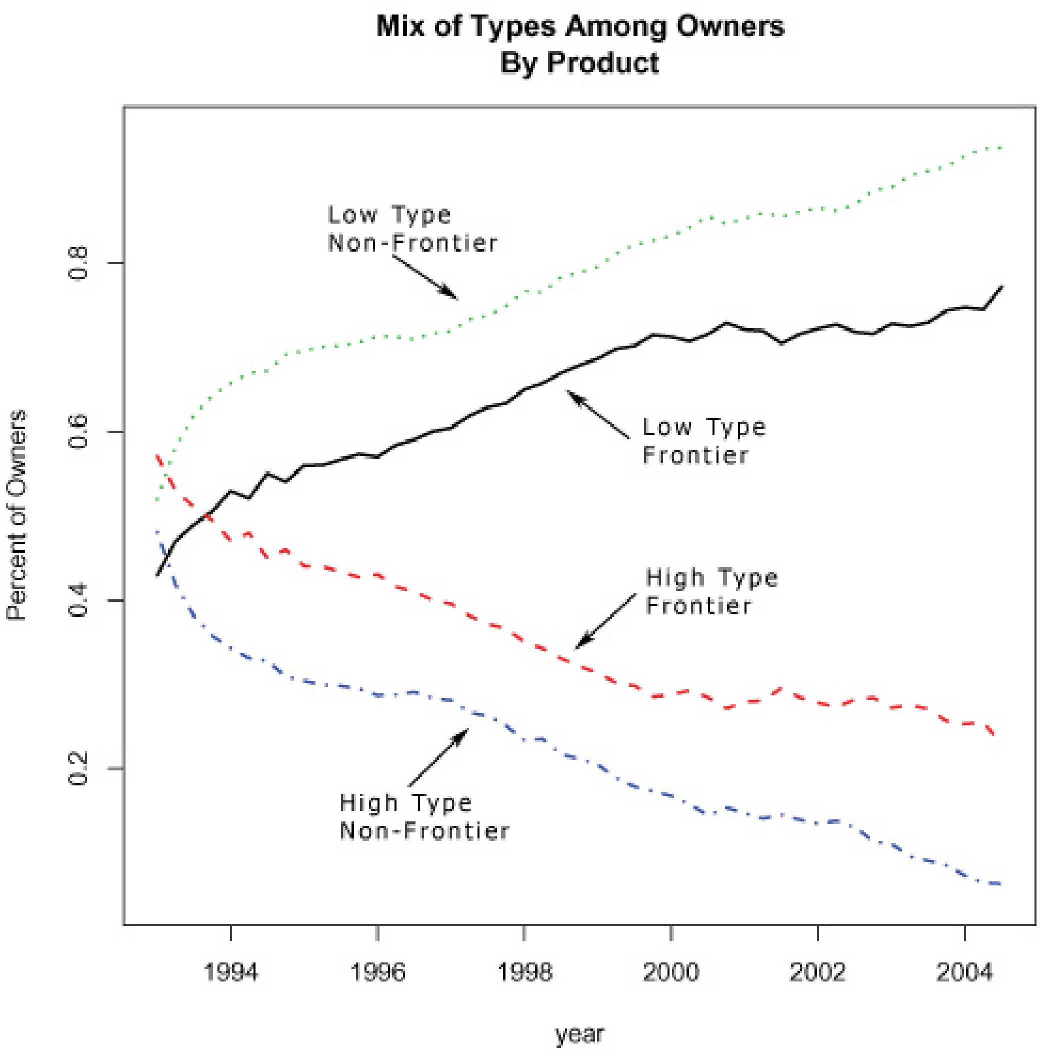

Figure 2.6: Distribution of consumer segments among owners for the frontier and non-frontier products.

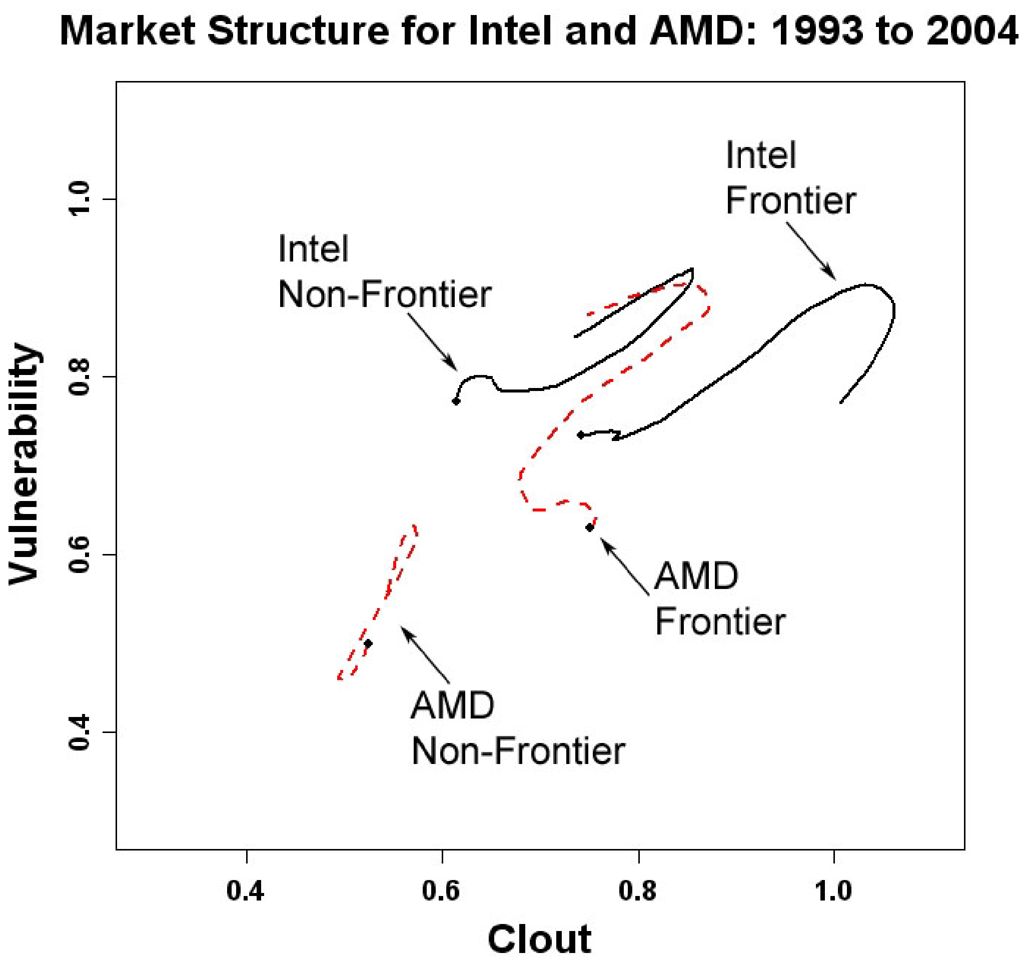

Figure 2.7: Clout and vulnerability for each product over time, with the time period starting at the black dot for each product.

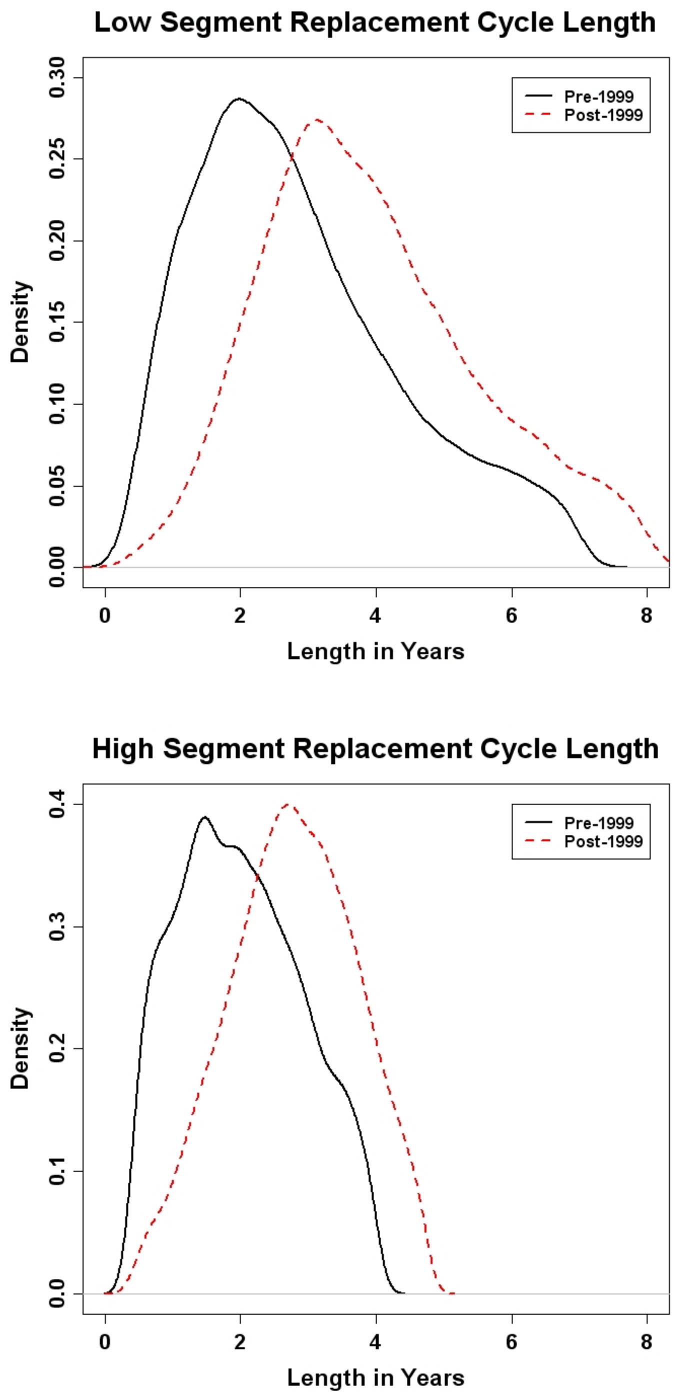

Figure 2.8: Smoothed plots of the distribution of CPU replacement length for low and high segment consumers, before and after 1999.

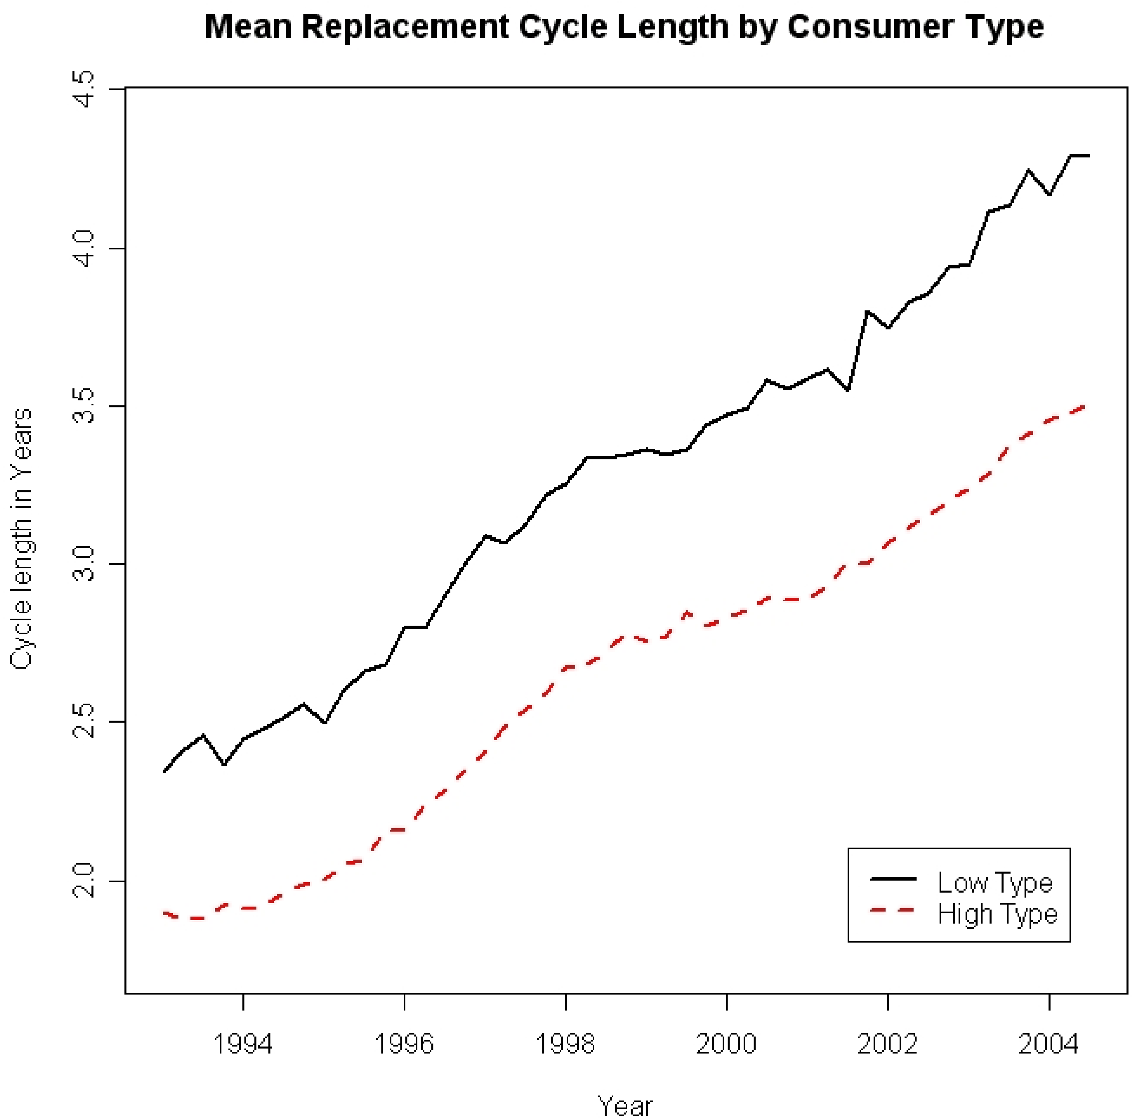

Figure 2.9: Mean replacement cycle length by consumer segment from 1993 to 2004.

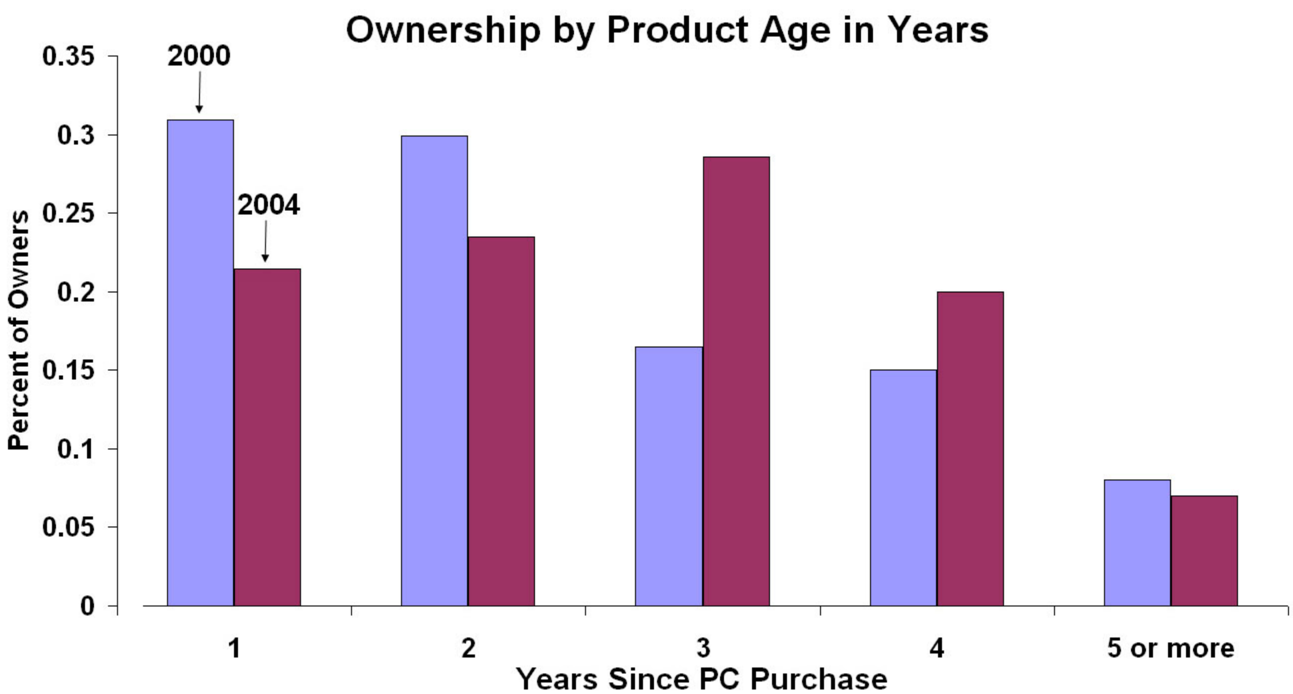

Figure 2.10: The distribution of product ownership over all consumers for the years 2000 and 2004.

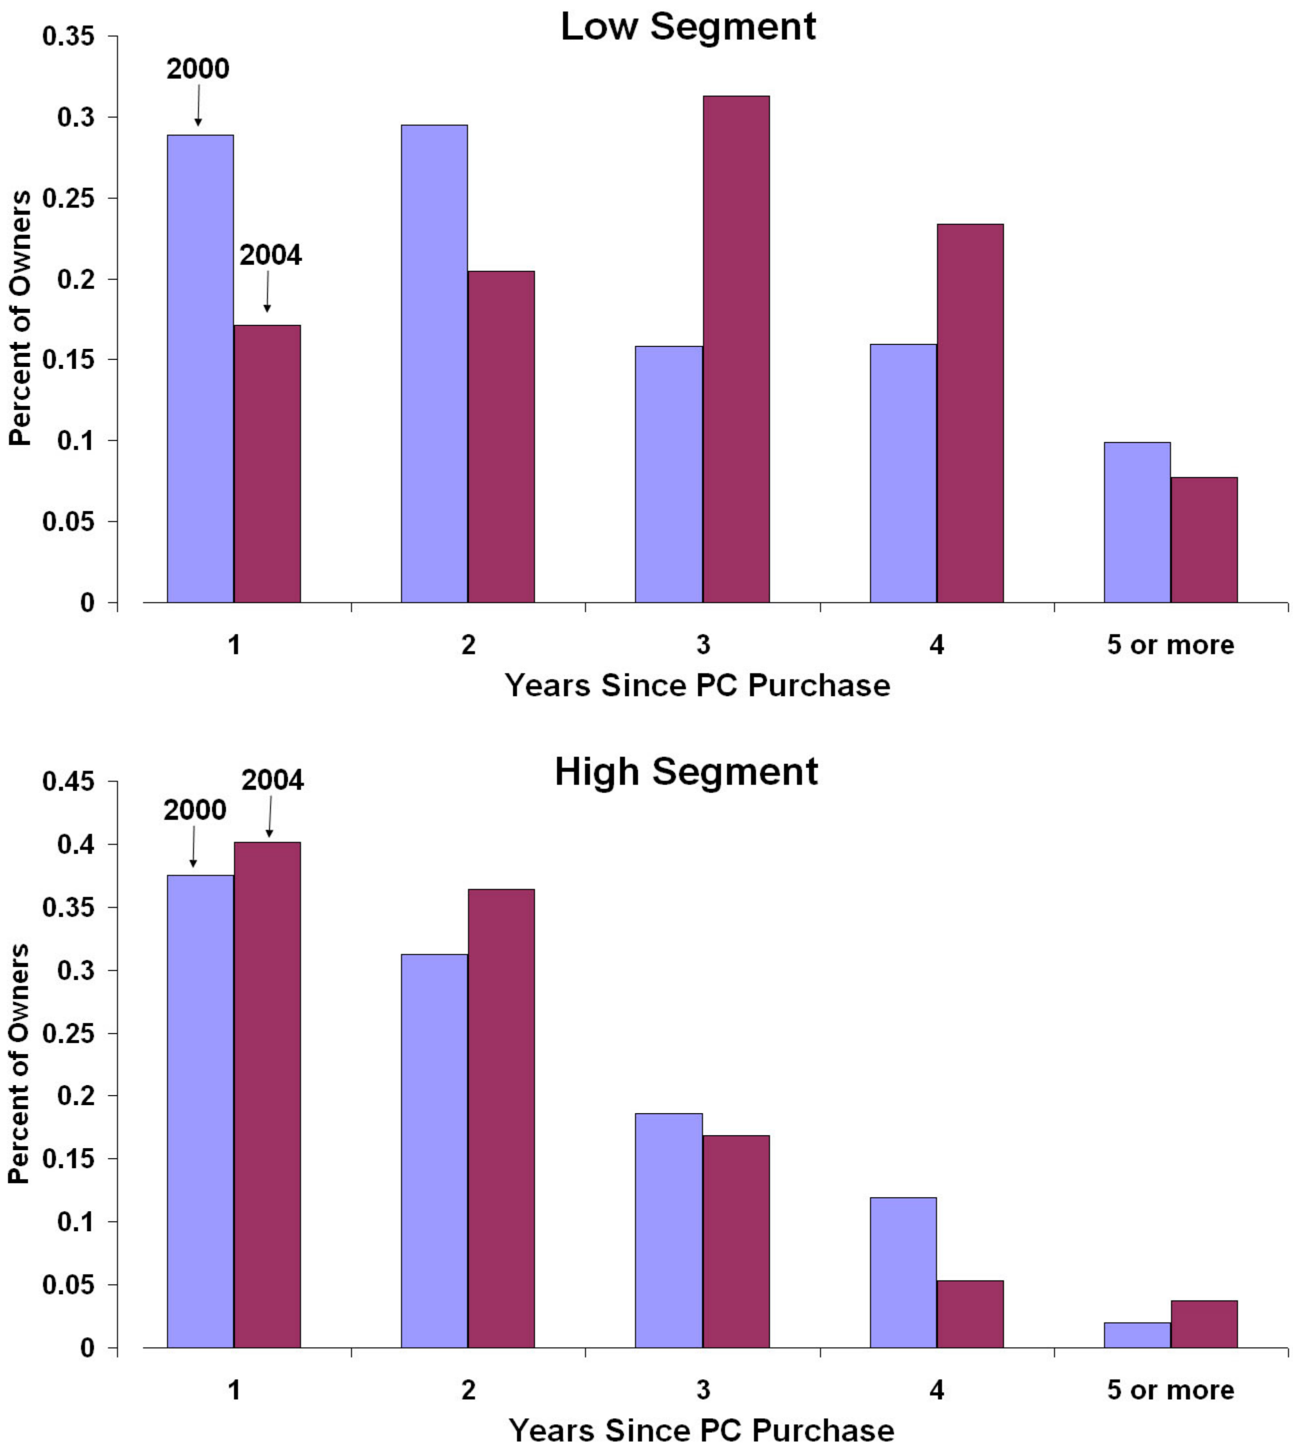

Figure 2.11: The distribution of product ownership by each consumer segment for the years 2000 and 2004.

Table 2.1: Vector Autoregression Process for Prices

	Intel Frontier(t)	Intel Non-frontier(t)	AMD Frontier(t)	AMD Non-frontier(t)
Intel Frontier($t-1$)	0.8151	0.3249	-0.1342	0.0187
	(0.0752)	(0.0689)	(0.0387)	(0.0472)
Intel Non-frontier($t-1$)	0.1072	0.4179	-0.1397	-0.1836
	(0.0425)	(0.0916)	(0.0302)	(0.1117)
AMD Frontier($t-1$)	-0.0150	-0.0485	0.8853	0.0414
	(0.0316)	(0.0226)	(0.0486)	(0.0353)
AMD Non-frontier($t-1$)	-0.1016	-0.0470	-0.0161	0.6836
	(0.0654)	(0.0599)	(0.1007)	(0.0731)
Constant	1.0984	1.1952	0.8620	1.9313
	(0.4065)	(0.3724)	(0.6258)	(0.4544)
R^2	0.8367	0.7361	0.7953	0.6359
-LL	61.3607	73.3306	2.2381	46.0780

Standard errors in parentheses.

Table 2.2: Product Quality Process Estimates

	Intel Frontier	Intel Non-frontier	AMD Frontier	AMD Non-frontier
κ_0	0.4912	0.4877	0.4501	0.4537
	(0.0289)	(0.0272)	(0.0190)	(0.0203)
κ_1	0.0218	0.0250	0.0334	0.0329
	(0.0053)	(0.0069)	(0.0041)	(0.0048)
λ	1.5656	1.4406	1.5248	1.3717
	(0.0572)	(0.5320)	(0.4798)	(0.0412)
-LL	15.3830	16.2184	15.5011	16.8874

Standard errors in parentheses.

Table 2.3: Model Fit

Model	Myopic	One Segment	Two Segment
Moments		Mean Squared Error	
Penetration Rate	1.298	0.622	0.446
Replacement Share	5.894	2.838	1.577
Market Shares	10.027	8.655	7.120
Ownership Shares	18.703	9.630	6.359
Objective Value	0.174	0.115	0.086
J Statistic	23.943	15.911	11.923
p-value	0.120	0.530	0.805
DM Statistic	24.040	7.976	-
p-value	0.000	0.019	-

The last line reports the DM statistics in a comparison of the first two models against the two-segment dynamic model.

Table 2.4: Structural Parameter Estimates

	Myopic		One Segment		Two Segment	
	Estimate	Std Err	Estimate	Std Err	Estimate	Std Err
Utility Parameters						
Quality (γ_l)	1.313	0.388	2.236	0.522	2.039	0.329
Quality (γ_h)	2.165	0.483			2.848	0.385
Price (α_l)	−1.542	0.427	−1.873	0.390	−2.299	0.316
Price (α_h)	−0.954	0.304			−1.708	0.142
Intel (ξ_I)	63.252	14.725	48.113	8.750	38.027	6.413
AMD (ξ_A)	14.585	3.551	11.910	1.005	7.857	1.098
Discount factor (β)	0		0.98		0.98	
Initial Conditions						
$F_0(i\|0)$	0.837	0.177			0.942	0.205
$F_0(i\|f)$	0.520	0.152			0.671	0.121
$F_0(i\|nf)$	0.825	0.191			0.781	0.166
Low segment size	0.780	0.084			0.854	0.071
High segment Size	0.220	0.026			0.146	0.039

Parameter estimates for different model specifications.

Table 2.5: Summary of Price Elasticities

	Dynamic model		Myopic model	
	Mean	Std Dev	Mean	Std Dev
Intel				
Own Elasticities	−5.68	1.03	−3.74	0.83
Cross Elasticities	2.73	0.39	1.65	0.30
AMD				
Own Elasticities	−5.25	1.34	−3.41	0.92
Cross Elasticities	2.18	0.17	1.56	0.18

Average value of elasticities for a permanent 10% change in prices.

Table 2.6: Average Price Elasticities

	Intel Frontier	Intel Non-frontier	AMD Frontier	AMD Non-frontier	No Purchase Owners	No Purchase Non-Owners
Intel Frontier	−5.015	3.592	3.405	1.985	2.317	3.121
	[1.285]	[0.365]	[0.592]	[0.317]	[0.439]	[0.558]
Intel Non-frontier	2.935	−6.345	2.621	1.814	1.970	2.429
	[0.330]	[0.748]	[0.485]	[0.195]	[0.301]	[0.388]
AMD Frontier	3.332	2.797	−5.711	1.602	2.628	3.515
	[0.407]	[0.126]	[0.419]	[0.179]	[0.321]	[0.413]
AMD Non-frontier	1.723	2.021	1.608	−4.788	1.623	2.006
	[0.098]	[0.168]	[0.051]	[0.247]	[0.258]	[0.287]

Standard deviations appear in square brackets. All estimates are statistically significant; standard errors are not reported. Each entry is the average percentage change in demand for the column product given a permanent 10% change in the price of the row product. The last two columns on the right are the change in demand for the outside option (no purchase) given a 10% change in the price of the row product, conditional on whether the consumer owns a product or does not. These elasticities are the average value of the elasticity calculated in each sample period.

Table 2.7: Comparison of Replacement Cycle Length by Period

	Benchmark model Mean	Benchmark model Std Dev	Myopic model Mean	Myopic model Std Dev
Full Sample				
Low segment	3.395	1.622	2.478	1.361
High segment	2.560	1.036	2.105	0.918
All	3.287	1.575	2.429	1.318
Pre-1999				
Low segment	2.878	1.506	1.962	1.386
High segment	2.001	0.925	1.728	0.893
All	2.694	1.384	1.905	1.342
Post-1999				
Low segment	3.896	1.574	2.941	1.477
High segment	2.862	0.919	2.570	0.907
All	3.624	1.436	2.893	1.438

Mean replacement cycle lengths, in years, for the dynamic two-segment model and the myopic model.

Table 2.8: Effects of Different Innovation Rates on Replacement Cycle Length

	Benchmark model		+25% Innovation		+50% Innovation	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Low segment	3.395	1.622	3.046	1.492	2.562	1.427
High segment	2.560	1.036	2.193	0.980	1.689	0.935
All	3.287	1.546	2.935	1.425	2.449	1.370

Mean replacement cycle lengths, in years, over the entire sample for the benchmark model and counterfactual cases.

Table 2.9: Alternative Innovation Rates and Replacement Cycle Length, by Time

Year	Benchmark	+25% Innovation		+50% Innovation	
	Mean	Mean	% Change	Mean	% Change
1995	2.14	1.89	−11.83	1.32	−38.24
1997	2.47	2.21	−10.77	1.63	−34.23
1999	2.86	2.62	−9.01	1.99	−30.74
2001	3.30	3.05	−7.66	2.38	−27.72
2003	3.68	3.41	−7.43	2.83	−23.00

Mean replacement cycle lengths, in years, in the benchmark and counterfactual cases.

Chapter 3

Dynamic Duopoly

In durable goods markets, sellers face a dynamic trade-off: more units sold today come at the expense of reduced demand tomorrow. Despite the importance of dynamic demand in durable goods markets, the equilibrium implications on firms' and consumers' strategies in an imperfectly competitive market remain unclear. To this end, we construct a model of dynamic oligopoly with durable goods and endogenous innovation. Firms make both dynamic pricing and investment decisions while taking into account the aggregate dynamic behavior of consumers. In such markets, the consumer's purchase decision is inherently dynamic because many technology durables undergo both rapid improvements in quality and falling prices. A consumer must decide whether to keep their existing product (if any) or to buy a new product, given her expectations about future product characteristics. This implies that the distribution of currently owned products affects current demand, and we model the endogenous evolution of this distribution. We show that accounting for product durability and the distribution of consumer ownership has significant implications for firms' profits and policy functions and consumer surplus.

Our work extends the framework developed by Ericson and Pakes (1995) and Pakes and McGuire (1994) to incorporate durable goods. To solve for equilibrium, we introduce a normalization that transforms the state space from one that is unbounded to one that is finite. Rather than measuring quality on an absolute scale, we measure all qualities relative to the current period's maximum quality. A number of papers present similar transformations of an infinite state space into a stationary finite state space. In the Pakes-McGuire dynamic oligopoly with differentiated products, firms' product qualities are measured relative to the "no purchase" alternative. The outside alternative improves in quality according to an exogenous process, and provides incentives for the "inside" firms to invest. However, these transformations carry implicit assumptions about the nature of the outside alternative.

A peculiar implication of their assumption is that the long-run (steady-state) rate of innovation is determined solely by the outside good's exogenous rate of quality improvement. A related point is that investment decreases to zero as a firm's quality improves, regardless of whether there are competitors with products of similar quality. These implications seem unlikely to apply to many markets and are a by-product of the assumptions Pakes and McGuire use to generate a finite state space.

Two advantages of the approach taken in this paper are that our assumptions have less influence on the behavior of the frontier firms, and these are the firms responsible for generating most of the sales, profits, and surplus. Moreover, a firm's innovation is driven by responses to other firms' innovations (and pricing policies), as opposed to some exogenous competing technology. Hence, the steady-state rate of innovation is endogenously determined.

We calibrate the model to the PC microprocessor industry. This industry is well-suited for the analysis because it is a duopoly, with Intel and Advanced Micro Devices (AMD) controlling about 95 percent of the market, and sales have been driven by intense technological innovation and price competition. While our quantitative results are specific to this industry, we believe the insights derived from them will be relevant for any durable goods market where innovation and obsolescence drive product replacement.

We find that industry profits are 24 percent lower when the duopolists ignore the durable nature of the product when setting prices. In the monopoly case, the firm's profits are 41 percent lower when the firm ignores the dynamic aspect of demand. Margins are 48 percent lower in the duopoly and 68 percent lower in the monopoly, confirming the intuition that prices are higher under dynamic demand. This demonstrates the strong link between optimal firm behavior, with and without competition, and accounting for product durability and dynamic demand.

Welfare outcomes also differ significantly: compared to a socially benevolent monopolist, consumer welfare is 22 percent lower in a duopoly and 54 percent lower in a monopoly. There is excessive innovation in the duopoly compared to the social monopolist case. While investment is higher in the duopoly than in the monopoly, a counterfactual analysis suggests that most of the welfare loss associated with monopoly comes from higher margins and not the slower pace of innovation.

The difference in profits and prices highlights the importance of explicitly modeling durability for policy analysis and counterfactual simulations, such as for the evaluation of mergers and antitrust. We consider a similar exercise by varying the degree of monopoly power for one firm. This policy simulation is motivated by the recent antitrust lawsuit filed by AMD alleging that Intel has engaged in anti-competitive practices that effectively shutout AMD from part of the market.[1] While any degree of monopoly power for one firm decreases consumer welfare, we find that the decrease in welfare did not become significant until one firm had greater than a 40 percent market monopoly. In a pure monopoly, consumer welfare was about 41 percent lower than in the pure duopoly, but with a 40 percent restriction, consumer welfare only decreased by 2.8 percent. In general, the increase in market power for one firm had a minimal impact on investment; the decrease in consumer welfare was primarily due to

[1] See Singer, M. and D. Kawamoto, "AMD Files Antitrust Suite Against Intel," CNet News.com, June 28, 2005.

increased margins and not a slower pace of innovation. This result is due to the fact that in a durable goods market, the monopolist still must compete with themselves in future periods, and so must continue to investment to spur product replacement.

The Ericson-Pakes framework has previously been applied to durable goods markets, though the forward-looking nature of consumers has generally been ignored. Benkard (2004) accounts for dynamic demand for airplanes by allowing the market size to change stochastically based on current period sales to mimic forward-looking behavior. Markovich (2004) and Markovich and Moenius (2005) have consumers that look two periods into the future in an oligopoly model, and Nair (2006) has consumers who solve an optimal stopping problem in a monopoly model. The most closely related paper is Esteban and Shum (2005), who consider the effects of durability and secondary markets on equilibrium firm behavior in the automobile industry.[2] They assume that consumers' current decisions are independent of past and future choices and derive a deterministic linear-quadratic model, relying on the certainty equivalence property. They assume a stationary environment by ignoring the entry and exit of car models from the market.[3]

Finally, our work is connected to the large theoretical literature on durable goods.[4] Two strands of this literature are most relevant. The first area, starting with the works of Kleiman and Ophir (1966), Swan (1970, 1971), and Sieper and Swan (1973), asks whether a durable goods monopolist would provide the same level of durability as competitive firms and whether such a firm would choose the socially optimal level of durability. The second area, beginning

[2]There is also a small existing literature on durable goods oligopoly, such as Sobel (1984) and Gul (1987).

[3]Our work is also related to recent empirical models of dynamic demand that take the firms' behavior as exogenous. Most of these papers consider high-tech durables, but restrict their attention to the initial product adoption decision. Melnikov (2001) develops a model of demand for differentiated durable goods that he applies to the adoption of computer printers. Carranza (2005) extends his model and examines the introduction of digital cameras. Song and Chintagunta (2003) develop a similar empirical model of digital camera adoption that incorporates unobserved consumer segments. More recently, Gordon (2006) and Gowrisankaran and Rysman (2006) have developed models that allow consumers to replace their products over time.

[4]For an excellent review, see Waldman (2003).

with Coase (1972) and followed by Stokey (1981) and Bulow (1982), among others, considers the time inconsistency problem faced by a durable goods monopolist.[5]

Early work in both areas focused on obtaining analytical results that often required various assumptions. The most prominent of these assumptions was that old and new goods were perfect substitutes, that the infinite durability of the goods eliminated the need for replacement, and that the markets were either monopolies or perfectly competitive. Later work typically investigated the robustness of the original conclusions to the relaxation of certain assumptions.

In the optimal durability literature, the so-called "Swan Independence Result" states that a monopolist provides the socially optimal level of durability. This result hinged on the assumption that the lifetime distribution of the durable is exogenous. Rust (1986) relaxed this assumption by allowing for endogenous scrappage rates based on the equilibration of supply and demand in the secondary market. Another common assumption was that some number of old units could substitute perfectly for a new unit. Waldman (1996) and Hendel and Lizzeri (1999) assumed that new and used units were imperfect substitutes that differed in terms of quality. They showed that Swan's independence results failed to hold under this weaker assumption. In the time inconsistency literature, Bond and Samuelson (1984) show that depreciation and replacement sales reduce the monopolist's tendency to cut prices.

Despite these recent efforts, the theoretical literature still has been hindered by the need to make strong assumptions. By turning to numerical methods, we are able to relax some of these assumptions and provide us with a model we can calibrate to data to generate quantitative results.

Our model primarily differs along two key dimensions. First, we consider a dynamic

[5]A related area studies the problem of a monopolist pricing a new product, such as the next generation of a durable good. See, for example, Levinthal and Purohit (1989), Fudenberg and Tirole (1998), and Lee and Lee (1998).

oligopoly where competition in pricing and innovation play critical roles. Second, we allow for endogenous innovation and account for the endogenous evolution of the distribution of currently owned (used) products. While the product a consumer owns never deteriorates in an absolute sense, consumers make replacement purchases as the quality of the product they own becomes worse relative to the frontier product in the market. This leads to an endogenous scrappage rate and hence an endogenous lifetime distribution for the durable good, but whereas Rust (1986) relies on secondary markets to endogenize scrappage, consumers naturally upgrade as quality increases over time.

The rest of the paper is organized as follows. Section 3.1 defines the model and equilibrium. Section 3.2 discusses the method we use to compute the equilibrium. Section 3.3 characterizes the firms' equilibrium policy functions and presents results from simulating the calibrated model. Section 3.4 offers some concluding remarks and future directions for research.

3.1 Model

In this section we present a dynamic model of differentiated products oligopoly for a durable good. Time, indexed by t, is discrete with an infinite horizon. Each firm $j \in \{1, \ldots, J\}$ sells a single product with time-varying quality denoted $q_{jt} \in \{0, \delta, 2\delta, \ldots\}$.[6] In each period, firms simultaneously choose their prices p_{jt} and investment x_{jt}.[7] Price is a dynamic control since lowering price in period t increases current sales, but reduces future demand. Investment is a dynamic control since future quality is stochastically increasing in investment. Consumers decide each period whether to buy a new product or to continue using

[6] We could normalize q_{jt} to be positive integers, but the calibrated model is easier to interpret if the quality grid (and the implied innovation process) matches the data.

[7] The model does not allow entry or exit, primarily because of the lack of significant entry in the CPU industry.

their currently owned product (if any).[8] Hence, the distribution of currently owned products affects current demand. We denote this endogenous distribution Δ_t.

Firms and consumers are forward-looking and take into account the optimal dynamic behavior of the other agents (firms and consumers) when choosing their respective actions. We assume the vector of firms' qualities $q_t = (q_{1t}, \ldots, q_{Jt})$ and the ownership distribution Δ_t is observed by all agents. These two state variables comprise the state space of payoff relevant variables for firms. The consumer's state space consists of the quality of her currently owned product $\tilde{q}_t$, the firms' current offerings q_t, and the ownership distribution Δ_t. This latter state variable is relevant to the consumer since it affects firms' current and future prices and investment levels (i.e., innovation rates).

3.1.1 Consumers

Utility for a consumer from firm j's new product with quality q_{jt} is given by

$$u_{jt} = \gamma q_{jt} - \alpha p_{jt} + \xi_j + \varepsilon_{jt} \tag{3.1}$$

where γ is the taste for quality, α is the constant marginal utility of money, ξ_j is a brand preference for firm j, and ε_{jt} captures idiosyncratic variation in utility which is i.i.d. across consumers, products, and periods.

Utility for the outside alternative (i.e., no purchase option) is

$$u_{0t} = \gamma \max(\tilde{q}_t, max(q_t) - \bar{q}) + \varepsilon_{0t} \tag{3.2}$$

where $\tilde{q}_t$ is zero if a consumer does not yet own the product and $\bar{q}$ is the maximum difference

[8]Some durable good markets, such as automobiles, have established used good markets. Only a small fraction of purchases of durable goods with rapid innovation, such as CPUs and consumer electronics, transact in used markets. As such, our model does not allow for resale of used goods.

between the quality of the best new product and the no purchase option.

A key feature of this demand model is that the value of a consumer's outside option is endogenous, since it depends on past choices. This feature generates the dynamic trade-off for firms' pricing decisions: selling more in the current period reduces demand in future periods since recent buyers are unlikely to buy again in the near future. Dynamic demand also has an impact on firms' investment decisions because the potential marginal gain from a successful innovation depends on the future distribution of consumer product ownership. The potential gain from an innovation will be larger if many consumers own older products, and the gain will be smaller if many consumers had recently upgraded to a product near the frontier.

One can think of our model as having two outside alternatives – one for consumers who have purchased at least once in the past, and one for "non-owners" who have never purchased the good. In the context of CPUs, the outside good for non-owners may consist of using computers at schools and libraries or using very old computers given to them by family or friends who have upgraded. Such an outside alternative surely improves as the quality frontier improves. The lower bound in the quality term of u_{0t} due to $\bar{q}$ captures this notion that the outside alternative, even for non-owners, improves as quality improves.

Given the lower bound $\max(q_t) - \bar{q}$, the ownership distribution can treat all consumers with $\tilde{q} \leq \max(q_t) - \bar{q}$ as owning the lower bound itself. Hence, $\Delta_t = (\Delta_{\max(q_t)-\bar{q},t}, \ldots, \Delta_{\max(q_t),t})$, where Δ_{kt} is the fraction of consumers in the population whose outside option (i.e., current product) has quality q_{kt}.[9]

Each consumer maximizes her expected discounted utility, which can be formulated using

[9]Here, we use the subscript k instead of j because these subscripts do not necessarily refer to products currently offered by any of the J firms. Furthermore, the dimension of Δ_t is $\max(q_t) - \bar{q}$ which has no relation to J.

Bellman's equation as the following recursive decision problem:

$$V(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) = \max_{y_t \in (0,1,\ldots,J)} u_{y_t,t} + \beta \mathrm{E}\left[V(q_{t+1}, \Delta_{t+1}, \tilde{q}_{t+1}, \varepsilon_{t+1}) | y_t, q_t, \Delta_t, \varepsilon_t\right] \tag{3.3}$$

where y_t denotes the optimal choice in period t. The expected continuation value depends on the consumer's expectations about future products' qualities and on the law of motion for Δ_t. With an appropriate distributional assumption on $\{\varepsilon_{jt}\}$, we can derive an expression for the demand for each product based on the consumers' value function. The implied demand system governs the law of motion for Δ_t and is used below in the model of firm behavior.

Note that once a consumer purchases a product at some quality level, the brand of the product no longer matters. That is, the consumer receives a one-time utility payoff of ξ_j from purchasing a product from firm j. This payoff does not occur in future periods since the outside option depends only on $\tilde{q}_t$. Relaxing this assumption would require Δ_t to be brand specific, which would substantially increase the state space.

Each consumer is small relative to the size of the market so that their individual actions do not affect the evolution of the aggregate Δ_t. We also assume consumers are ex-ante identical. Relaxing this assumption to allow γ and α to vary across consumers would require expanding the state space to include separate ownership distributions for each consumer type. While such an extension may be worth pursuing in future research, the current specification is sufficient for capturing the most relevant feature of durable goods demand – current sales affect future demand.

3.1.2 Firms

Each period firms make dynamic pricing and investment decisions. Each firm has access to an R&D process that governs their ability to introduce higher quality products into the market. Firms choose a level $x_j \in \mathbb{R}_+$ to invest in the R&D process. The outcome of this process, denoted $q_{j,t+1} - q_{j,t} = \tau_{j,t}$, is probabilistic, and stochastically increasing in the level of investment. We restrict $\tau_{j,t} \in \{0, \delta\}$ and denote its probability distribution $f(\cdot|x)$.[10]

The period profit function, excluding investment costs, for firm j is

$$\pi_j(p_t, q_t, \Delta_t) = M s_{jt}(p_t, q_t, \Delta_t)(p_{jt} - mc_j)$$

where M is the (fixed) market size, $s_{jt}(\cdot)$ is the market share for firm j, p_t is the vector of J prices, and mc_j is firm j's constant marginal cost of production. The Bellman equation for firm j's maximization problem is

$$W_j(q_{jt}, q_{-j,t}, \Delta_t) = \max_{p_{jt}, x_{jt}} c x_{jt} + \beta \sum_{\tau_{jt},\, q_{-j,t+1}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}(\cdot)) h_f(q_{-j,t+1}|q_t, \Delta_t) f(\tau_{jt}|x_{jt}) \tag{3.4}$$

where c is the unit cost of investment, $h_f(\cdot|\cdot)$ is firm j's beliefs about its competitors future product quality levels, and $\Delta_{t+1}(\cdot)$ is the transition kernel for Δ_t which depends on prices and investment as discussed below.

Following Rust(1987) we assume the consumers' $\{\varepsilon_{ijk}\}$ are multivariate extreme-value so that we can obtain the standard multinomial logit formula for product demand for $q_{jt} \in q_t$ by consumers who currently own $\tilde{q}$. In particular, we can integrate over the future ε_{ijk} to obtain the product-specific value function

$$\hat{V}_j(q_t, \Delta_t, \tilde{q}_t) = u_{jt} - \varepsilon_{ijt} + \beta \int_{q_{t+1}} \log \left(\sum_{j' \in \{0,\ldots,J\}} \exp \left\{ \hat{V}_{j'}(q_{t+1}, \Delta_{t+1}(\cdot), \tilde{q}_{t+1}) \right\} \right) h_c(q_{t+1}|q_t, \Delta_t) \tag{3.5}$$

where $h_c(\cdot|\cdot)$ is the consumer's beliefs about future product qualities. The conditional choice

[10] We actually think of quality as being in logs, so that improvements are proportional increases in quality. Using a log scale makes more sense than a linear scale when calibrating the model to the CPU industry.

probabilities for a consumer owning product $\tilde{q}$ are therefore

$$s_{jt|\tilde{q}} = \frac{\exp\{\hat{V}_j(q_t, \Delta_t, \tilde{q}_t\}}{\sum_{k\in\{0,\dots,J\}} \exp\{\hat{V}_k(q_t, \Delta_t, \tilde{q}_t)\}} \ . \tag{3.6}$$

Using Δ_t to integrate over the distribution of $\tilde{q}_t$ yields the market share of product j

$$s_{jt} = \sum_{\tilde{q}\in\{\max(q_t)-\bar{q},\dots,\max(q_t)\}} s_{jt|\tilde{q}} \ \Delta_{\tilde{q},t} \ . \tag{3.7}$$

These market shares translate directly into the law of motion for the distribution of ownership.[11] Recall that Δ_t only tracks ownership of products within $\bar{q}$ quality units of the highest quality offering. Assuming this highest quality is unchanged between t and $t+1$, the share of consumers owning a product of quality k at the start of period $t+1$ is

$$\Delta_{k,t+1}(\cdot) = s_{0t|k}\Delta_{kt} + \sum_{j=1,\dots,J} s_{jt} I(q_{jt} = k) \tag{3.8}$$

where the summation accounts for the possibility that multiple firms may have quality k. For quality levels not offered in period t, this summation is simply zero. If a firm advances the quality frontier with a successful R&D outcome pushing its $q_{j,t+1}$ beyond $\max(q_t)$ then Δ_{t+1} shifts: the second element of Δ_{t+1} is added to its first element, the third element becomes the new second element (and so on), and the new last element is initialized to zero. The transition kernel $\Delta_{t+1}(\cdot)$ is therefore a deterministic function of prices, except for the potential shift due to the stochastic innovation of frontier products.

[11] For conciseness our notation suppresses the dependence of market shares on prices.

3.1.3 Optimal Prices and Investments

Each firm chooses price and investment simultaneously, fixing other firms' price and investment choices. Fortunately, we can reduce the computational burden of this two-dimensional optimization using a sequential approach. The outer search is a line optimization over prices which contains a closed-form solution for investment given price.

Consider the first-order condition for investment $\frac{\partial W_j}{\partial x_{jt}} = 0$ evaluated at some arbitrary price p_{jt}:

$$-c + \beta \sum_{\tau_{jt},\, q_{-j,t+1}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}(\cdot)) h_f(q_{-j,t+1} | q_t, \Delta_t) f(\tau_{jt}|x_{jt}) \frac{\partial f(\tau_{jt}|x_{jt})}{\partial x_{jt}} = 0 \ .$$

Given outcomes $(\tau_{jt},\ q_{-j,t+1})$ the transition for Δ_{t+1} is a deterministic function of the control p_{jt}. Thus, with a suitable choice for f we can analytically compute the optimal investment as a function of price, $x^*_{jt}(p_{jt})$.[12] Let

$$\begin{aligned} EW^+(p_{jt}) &= \textstyle\sum_{q_{-j,t+1}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}(\cdot))\ h_f(q_{-j,t+1}|q_t, \Delta_t)\ f(\tau_{jt} = \delta|x_{jt}) \\ EW^-(p_{jt}) &= \textstyle\sum_{q_{-j,t+1}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}(\cdot))\ h_f(q_{-j,t+1}|q_t, \Delta_t)\ f(\tau_{jt} = 0|x_{jt}) \end{aligned}$$

be the expected continuation values conditional on, positive and negative innovation outcomes, respectively. The dependence of these expectations on p_{jt} is through the effect of price on the ownership transition to Δ_{t+1}. For an arbitrary price p_{jt}, the optimal investment is

$$x^*_{jt}(p_{jt}) = \frac{1}{a_j} \left(\frac{c}{\beta a_j (EW^+(p_{jt}) - EW^-(p_{jt}))} \right)^{-1/2} - 1 \ . \tag{3.9}$$

To determine the optimal price, consider the derivative of the firm's value function with

[12] For the details of deriving optimal investments, see Pakes, Gowrisankaran, and McGuire (1993).

respect to price, $\frac{\partial W}{\partial p_{jt}} = 0$, which implies

$$\frac{\partial \pi_j(p_t, q_t, \Delta_t)}{\partial p_{jt}} + \beta \sum_{\tau_{jt},\, q_{-j,t+1}} \frac{\partial W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}(\cdot))}{\partial \Delta_{t+1}} \frac{\partial \Delta_{t+1}}{\partial p_{jt}} h_f(q_{-j,t+1}|q_t, \Delta_t) f(\tau_{jt}|x^*_{jt}(p_{jt})) = 0$$

A higher price today implies that more people will be available in the next period to purchase the product. The second term captures this benefit to raising price, and leads to forward-looking firms pricing higher than myopic firms who ignore this dynamic aspect of demand.

We use Brent's method to solve for the optimal price. For each candidate price we use $x^*(p_{jt})$, the optimal investment level given this price, to evaluate the probability of a successful innovation. While we have yet to prove that the optimal price is uniquely determined, inspection of the first-order-condition as a function of p_{jt} at many states indicates that this appears to be the case. The pair $(p^*_{jt}, x^*_{jt}(p^*_{jt}))$ is the optimal set of controls at this state.

3.1.4 Equilibrium

We consider pure-strategy Markov-Perfect Nash Equilibrium (MPNE) of this dynamic oligopoly game. Our definition of a MPNE extends that found in Ericson and Pakes (1995) to account for the forward-looking expectations of consumers. In brief, the equilibrium fixed point has the additional requirement that consumers possess consistent expectations on the probability of future firm states. The firms must choose their optimal policies based on consistent expectations on the distribution of future consumer states.

The equilibrium specifies that (1) firms' and consumers' equilibrium strategies must only depend on the current state variables (which comprise all payoff relevant variables), (2) consumers possess rational expectations about the evolution of the firms' product qualities, (3) each firm possesses rational expectations about its competitor's price and investment policy functions, and (4) the perceived distribution for the industry's evolution corresponds to the actual Markov transition kernel defined by the previous points.

Formally, a MPNE in this model is the set $\left\{V^*, h_c^*, \left\{W_j^*, x_j^*, p_j^*, h_{fj}^*\right\}_{j=1}^J\right\}$, which contains the equilibrium value functions for the consumers and their beliefs h_c^* about future product qualities, and the firms' value functions, policy functions, and beliefs h_{fj}^* over their $J-1$ rivals' future qualities. The expectations are rational in that $h_c^*(q_{t+1}|q_t, \Delta_t, \tilde{q}) = \prod_{j=1}^J f(q_{j,t+1}|x_{jt}^*)$ and $h_{fj}^*(q_{-j,t+1}|q_t, \Delta_t) = \prod_{j'\neq j}^J f(q_{j',t+1}|x_{j't}^*)$.[13]

The functional form of the investment transition function satisfies the UIC admissibility criterion in Doraszelski and Satterthwaite (2005). To guarantee existence of equilibrium requires us to show that there is a pure-strategy equilibrium in both investment choices and prices. Ericson and Pakes (1995) and the various extensions found in Doraszelski and Satterthwaite (2005) do not consider dynamic demand. As such, they are able to construct a unique equilibrium in the product market in terms of prices or quantities (depending on the specific model of product market competition). We are working on a proof to show that pure strategies in prices also exist.

3.2 Computation

This section discusses the details behind the computation of the Markov-perfect equilibrium defined above. First, we present a normalization that converts the non-stationary state space into a finite stationary environment. Second, we introduce an approximation to the distribution over ownership that significantly reduces the size of the state space. Third, we present an overview of the steps required to compute the equilibrium. Lastly, we discuss the model calibration and evaluation.

[13]Symmetry corresponds to $W_j^* = W^*$, $x_j^* = x^*$, $p_j^* = p^*$, and $h_{fj}^* = h_f^*$ for all j.Symmetry obviously requires that firm specific parameters, such as brand intercepts ξ_j, are the same across firms.

3.2.1 Normalization

The state space in the model presented in section 3.1 is unbounded since product qualities increase without bound. To solve for equilibrium, we transform the state space to one that is finite. Rather than measuring qualities on an absolute scale, we measure all qualities relative to the current period's maximum quality (i.e., $\max(q_t)$). This transformation has no effect on the firm's value function since discrete choice demand models yield identical choice probabilities if a constant is added to each alternative. The only subtlety of this transformation is in the consumer's continuation value in equation (3.3). In the event that the frontier is increased (always by δ units), the consumer's continuation value for that particular outcome is $\gamma\delta/(1-\beta) + V(q_{t+1} - \delta, \Delta_{t+1}^{shift}, \tilde{q}_{t+1} - \delta, \varepsilon_{t+1})$ instead of $V(q_{t+1}, \Delta_{t+1}, \tilde{q}_{t+1}, \varepsilon_{t+1})$, where Δ_{t+1}^{shift} shifts the ownership distribution as described after equation (3.8).[14] Other firms' qualities are also shifted down by δ to account for the advance of the frontier. We invoke a knowledge spillover argument to bound the difference between each firm's own quality and the frontier quality. We denote the maximal difference in firms' qualities $\bar{\delta}$.[15] We choose $\bar{\delta}$ to be sufficiently large that it has minimal effect on equilibrium strategies.

In our extension of the Ericson-Pakes framework to durable goods we endogenize the utility offered by the "no purchase" option. As discussed in Section 3.1, in a sense we have two outside alternatives—one for current owners considering an upgrade and one for non-owners considering an initial purchase. Our truncation of the ownership distribution Δ_t to only track ownership of products within $\bar{q}$ quality steps of the industry frontier product is consistent with the lower bound of the outside good's utility in equation (3.2). Essentially,

[14]The adjustment of the continuation value can be derived using the fact that $V(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) = \gamma\delta/(1-\beta) + V(q_t - \delta, \Delta_t^{shift}, \tilde{q}_t - \delta, \varepsilon_t)$. This equation holds since a consumer at state $(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t)$ will behave exactly the same as a consumer at state $(q_t - \delta, \Delta_t^{shift}, \tilde{q}_t - \delta, \varepsilon_t)$ since all relative values are the same. The only difference in expected discounted utility is the present discounted value of the extra $\gamma\delta$ utils enjoyed in the higher $\tilde{q}$ state in every period of every possible future path.

[15]The role of spillovers could be implemented in a smooth fashion by specifying the probability of successful innovation $f(\tau_{tj}|x_{jt})$ to be increasing in the degree to which firm j is behind the frontier.

we group non-owners with owners of the lowest quality tracked by Δ_t. Since utility of these non-owners increases when the frontier expands, we do allow the "never purchased" utility to change over time. However, the change in this utility is driven by the endogenous innovation processes of the "inside" firms.

A number of papers present similar transformations of an infinite state-space into a stationary finite state-space (e.g. Ericson and Pakes (1995) and extensions by numerous researchers, Goettler, Parlour, and Rajan (2005, 2006), and Gordon (2006)). In the case of dynamic oligopoly models, such transformations carry implicit assumptions about the nature of the outside alternative. For example, in the Ericson-Pakes framework with differentiated products (e.g. Pakes and McGuire (1994)), firms' product qualities are measured relative to the "no purchase" alternative. The quality of this outside alternative improves according to an exogenous process, which provides a continual need for the "inside" firms to invest to remain competitive with the outside alternative. Hence, the long run (steady-state) rate of innovation is determined solely by the outside good's exogenous rate of quality improvements. In particular, if the outside good never improves, then the steady-state equilibrium has no investment and no innovations.

This exogenous long run growth in quality is a by-product of the assumptions Pakes and McGuire use to generate a finite state space (i.e., lower and upper bounds for firms' relative qualities). The lower bound is generated by a salvage value that triggers exit if relative quality gets sufficiently low. The upper bound is generated by bounding the firm's profit function, which implies its value function eventually becomes concave in its own quality. In particular, they assume a consumer's mean utility from product j is an increasing, concave, and bounded function of the product's quality minus the outside good's quality. Standard discrete choice models often specify diminishing marginal utility for absolute levels of quality, but not for quality measured relative to an outside alternative. An implication of this specification is that investment decreases to zero as a firm's quality improves *even if its competitors' qualities* *are near its own* since the derivative of market share with respect to own quality goes to zero regardless of competitors' qualities. This implication that two competitors do not care about being the leader, as long as they are both well above the attractiveness of some outside alternative seems unlikely to apply to many markets. Another odd implication is that the relative market shares for two products with different qualities depends on the level of the outside good's quality. For example, a consumer indifferent between two products that differ by one quality level will strictly prefer the higher of the two products if the outside alternative improves.

Our alternative normalization – measuring all qualities relative to the frontier instead of relative to the outside good's quality – avoids the somewhat undesired implications of the Ericson-Pakes normalization. We obtain a finite state space without eliminating the incentive for two firms near the frontier to innovate. Our upper bound of quality is obtained mechanically by normalizing relative to the highest quality firm. Our lower bound of quality is obtained by directly assuming a firm's quality is never more than $\bar{\delta}$ below the frontier. Hence, the the finiteness of the state space is implied by assumptions at the low end of the quality grid, rather than at the high end. An advantage of this approach is that our assumptions have less influence on the behavior of the leading firms – that is, those firms responsible for generating most of the sales, profits, and surplus. Moreover, the innovation of firms is driven by responses to each other's innovations (and pricing policies), not by an exogenous competing technology. Hence, the long run (steady-state) rate of innovation is endogenously determined, not exogenously specified as in Ericson-Pakes.

Another advantage of our normalization is that it is truly a "normalization" in that it is an exact transformation of a dynamic game that is initially expressed in absolute terms. Our focus on durable goods forces us to develop a model that is, from the consumer's perspective, consistent with respect to shifts in relative qualities when the baseline good's quality changes. A consumer purchases a product with quality q_{jt} knowing that this product has expected utility γq_{jt} next period as well (since depreciation is zero). If derived utility were a non-linear function of a relative quality measure then the future utility would not be as expected when the normalization's baseline changes. Such consistency issues are best addressed by initially writing the model in absolute levels. Another requirement for the normalization to be exact is that the quality term must enter the utility function linearly – otherwise the normalization will shift the utilities of each discrete choice by different amounts, thereby affecting choices. Of course, the quality grid may be on a log scale, in which case the "quality" term in utility is actually log quality and innovations are proportional improvements. Indeed, this is the interpretation we find suitable for the CPU industry since Moore's Law refers to CPU speed doubling every 18 months (or so).

Finally, we note that while we developed this alternative normalization in the context of durable goods, its merits apply equally to the nondurable case. In future research, we will assess the effect of using our specification for the nondurable case explored in Pakes and McGuire (1994).

3.2.2 Approximation of Δ_t

The challenge in solving the model outlined above is that Δ_t is a high-dimensional simplex. To compute the equilibrium, we must discretize Δ_t over some grid. Suppose that $\dim(\Delta_t) =$ 10 and a grid with 20 points in each dimension is used. A naive discretization that ignores the fact that $\sum_{\tilde{q}} \Delta_{\tilde{q},t} = 1$ produces 10.2 trillion states. An efficient encoding method that takes advantage of this constraint leads to 7 million states. While a substantial reduction, the cost of computing the equilibrium is still prohibitive.

Our goal is to reduce the dimensionality of the distribution of current ownership. We do this through an approximation based on the cumulative density function (CDF) of product ownership across vintages. This approach is motivated by the observation that any CDF is monotonically increasing, which makes it easier to smoothly approximate compared to a probability density function.

Suppose the CDF of the ownership distribution can be described by the parametric function $F(\rho_t)$, where $\rho_t \in \mathbb{R}^m$. That is, given a value of ρ_t, we can compute the implied density $\Delta_{kt}(\rho_t) = F_k(\rho_t) - F_{k-1}(\rho_t)$, and given the density, we can compute the CDF as $F_k(\rho_t) = \sum_{k'} \Delta_{k't} I(\tilde{q}_{k'} \leq \tilde{q}_k)$. Note that this changes the representation of our state space. The model described in Section 3.1 included the complete distribution Δ as a state variable, whereas now the state space for a firm is $(q_{jt}, q_{-jt}, \Delta(\rho_t))$ and, similarly, for a consumer is $(q_t, \Delta(\rho_t), \tilde{q}_t)$.

While computing a firm's value function, we perform the following steps. First, given ρ_t, compute the distribution $\Delta(\rho_t)$. Note that this distribution is exact since we are solving the firm or consumer's maximization problem directly on a point located on the discretized state space. Second, consumers make their purchase decisions based on the firms' policy functions at the current iteration. Then we use the law of motion for the distribution of ownership in equation (3.8) to generate the next period distribution of ownership Δ_{t+1}. The next period distribution is a deterministic function of the current period's state variables and represents the exact distribution of consumer ownership in the next period. We use least-squares fitting of the ownership density through the CDF to map this distribution back into our state space representation, obtain ρ_{t+1}:

$$\rho_{t+1} = \underset{\rho \in \mathbb{R}^m}{\operatorname{argmin}} \sum_k \left(\Delta_{k,t+1} - \left(F_k(\rho) - F_{k-1}(\rho)\right)\right)^2$$

In general, the value of ρ_{t+1} obtained from this approximation will lay off the state space grid, so we interpolate the value function through ρ to calculate the continuation values.

3.2.3 Steps to Solving the Industry Equilibrium

We compute the equilibrium using a Gauss-Jacobi scheme to update the value and policy functions.[16] Starting at iteration $k = 0$, initialize the consumer value functions $\hat{V}_j^0$ and the firms' value functions W_j^0 and policy functions (x_j^{*0}, p_j^{*0}).

Then for iteration $k = 1, 2, \ldots$, follow these steps:

1. For each $\tilde{q} \in \Delta$, evaluate the consumer's value function $\hat{V}^k$ given the firms' policy functions from the previous iteration $\left\{x_j^{*k-1}, p_j^{*k-1}\right\}_{j=1}^J$.

2. For each $j \in J$, evaluate firm j's value function W_j^k given the other firms' policy functions from the previous iteration $\left\{x_{j'}^{*k-1}, p_{j'}^{*k-1}\right\}_{j' \neq j}^J$.

3. Update the consumer value functions $\hat{V}^{k+1} \leftarrow \hat{V}^k$, the firms value functions $W_j^{k+1} \leftarrow W_j^k, \forall j$, and their policy functions $\left\{x_j^{*k+1}, p_j^{*k+1}\right\} \leftarrow \left\{x_j^{*k}, p_j^{*k}\right\}$.

4. Check for convergence in the sup norm of the all agents' value functions. If convergence is not achieved, return to step (1).

3.2.4 Model Calibration and Evaluation

We calibrate the model to the PC processor industry using data over the period from 1993 to 2004. We base the consumer's parameters on estimates obtained in Gordon (2006), who estimates a model of dynamic demand in the PC processor market using a similar demand model, and provides a detailed description of the data. This leaves us with two parameters of interest for each firm: the innovative efficiency a_j and the marginal cost of production mc_j.

[16] We also tried a Gauss-Seidel scheme. While this generally converged to the same value and policy functions, we found that this approach sometimes produced large oscillations in the values and hindered convergence.

We obtain data for each company, but to simplify the analysis, we focus on an industry with symmetric firms. As others have done, we use quarterly R&D expenditures from each firm's annual reports and estimate a_j directly in a first-stage regression. We set the innovative efficiency for both firms in the model to the revenue-weighted average value of both firms. We obtain blended unit production costs from In-Stat/MDR, a market research firm that specializes in the microprocessor industry. There is variation over time in each firm's marginal cost, but the time series average for both firms did not differ significantly. Thus, we set the industry-wide marginal cost to the sales weighed averages of each firm's marginal cost.

To provide a sense of the model's fit, Table 3.4 compares some moments obtained from the data with those generated by simulating the model. We compute the sales weighed price margins, where margin is defined as $(p_{jt} - mc)/mc$, the ratio between the leading firm's market share and the lagging firm's market share, the industry-wide innovation rate, and the average percent quality difference between the leader and laggard. We find that the parameterized model produces reasonable behavior that is roughly consistent with the industry. Note that the results in Table 3.4 are generated without performing any type of optimization or estimation (apart from the first-stage regression for a_j).

Compared to the data, the simulated margins are lower and the market share ratios are lower. This is consistent with the intuition that if one firm had an even larger market share than the other, its margin would also be higher. Given the market share ratios, our parameterized model is able to generate significant differences in the firms' behavior over time, but the model is not yet able to generate the extremes observed in the data.[17]

One important point to note is that the firms in our model are symmetric, and thus we do not account for any perceived differences in brand value or pre-existing institutional

[17]In Section 3.3.2, we conduct a counterfactual analysis that varies the degree of market power for one firm. Under this configuration, we are able to generate much closer market share ratios and price cost margins when one firm acts as a monopolist in one portion of the market.

arrangements.[18]

3.3 Numerical Illustrations

In this section we first present findings that correspond to a "base case" which uses parameters chosen so that the model's implied equilibrium behavior matches features of the CPU industry over the years 1993 to 2004. We then consider various counterfactual experiments to further illustrate the properties of the model and its implications for policy analysis.

3.3.1 Base Case

Table 3.4 lists the parameter values for our base case. We use these parameters in five different scenarios, which correspond to the column headers in Table 3.4: 1) durable duopolist, 2) myopic pricing duopolist, 3) durable monopolist, 4) myopic pricing monopolist, and 5) social monopolist.[19] For each scenario we solve for optimal policies and simulate 1000 industries each for 100 periods, starting from a state in which all consumers are non-owners (i.e., $\tilde{q} = 0$) and each firm has $q_{jt} = \bar{q}$. We then analyze the simulated data to characterize the equilibrium behavior of firms and consumers and to identify observations of particular interest.

We start by characterizing the firms' optimal price and investment policy functions. The complexity of the state space, due to the distribution of ownership, prohibits a state-by-state inspection of the policy functions. Instead, we report in Table 3.4 regressions from the simulated data that fit the policy functions to select moments of the state variables.[20]

[18]It is straightforward to allow for asymmetries in the marginal costs of production or innovative efficiency. However, this would probably increase the likelihood of multiple equilibria.

[19]We leave for future research the computation of the social planner's optimal behavior when he controls two firms.

[20]We use the innovation rates because they are direct functions of the firms' investment decisions and are

Standard errors are not reported because they are so small. With separate regressions for the leader and laggard, the Table also illustrates how responses to state variables depend on whether a given firm is the leader. In the regressions we include a dummy indicating whether the lead firm had an innovation that extended his lead in the previous period, the percentage of consumers who own a product, the mean product quality of consumers who own a product, the percentage quality difference between the leader and laggard, and the square of this quality difference.[21]

The signs of all the coefficients are consistent with intuition. For the leader, the positive coefficient on whether they extended the frontier implies that they benefit from such an increase, while the negative coefficient for the laggard suggests that their margin suffers. The negative coefficient on mean owned quality shows that both firms have lower margins when, on average, consumers own newer products.

Figures 3.1 and 3.2 plot the relationship between the margins and innovation rates and quality differences based on the coefficients in the regressions. The leader's margin is an increasing function of the quality difference and the laggard's margin is a decreasing function of the quality difference, and both of these results are consistent with our expectations. Figure 3.2 reveals that the leader has an incentive to invest heavily when their quality lead is relatively small, but that this incentive disappears as the quality lead grows, to the point where the leader invests significantly less when their lead is quite large. Conversely, the laggard's innovation rate declines for small differences in quality, implying that the marginal gains to an innovation are insufficient when the leader's advantage is small. However, when the difference in quality is sufficiently large, the laggard begins to increase their investment. This suggests that the laggard does not want to fall too far behind in the quality competition,

easier to interpret than investment levels.

[21]We obtain similar regression results using dummies for each quality difference level, suggesting that our functional form assumption is reasonable.

though the laggard is also less willing to compete directly with the leader when the two firms are closer in quality.

Turning to consumers, Figure 3.3 presents the portion of purchasers that comes from each vintage of current ownership, for both the duopoly and monopoly cases. Most purchasers upgrade from products that are between 4 and 8 grid points (i.e., δ-sized steps) from the frontier. Since $\delta = .3$ this corresponds roughly to having a frontier that is 120 to 240 percent faster than their current product.

Firms' profits are calculated as the discounted sum of per period profits. To calculate consumer welfare, we must take into account the quality normalization discussed in Section 3.2.1.

Period level consumer surplus is calculated as

$$CS_{\tilde{q},t} = \frac{\log\left(\sum_{j' \in \{0,\dots,J\}} \exp\left\{\mathcal{I}(\tilde{q}_{jt} = \tilde{q}_t)u_{0t} + \mathcal{I}(\tilde{q}_{jt} \neq \tilde{q}_t)u_{jt}\right\}\right)}{\alpha}. \tag{3.10}$$

Total consumer surplus over a simulation run is the discounted sum of the per period surplus, integrated over the distribution of consumer product ownership:

$$CS = \sum_{t=0}^{T} \beta^t \sum_{\tilde{q}=\max_{q_t} - \bar{q}}^{\max_{q_t}} CS_{\tilde{q},t} \cdot \Delta_{\tilde{q},t}. \tag{3.11}$$

The simulations produced the following observations.

Observation 1 *Margins (i.e., prices) and profits are significantly higher when firms correctly account for the dynamic nature of demand. The differences are larger for monopoly than duopoly.*

In Table 3.4 we see that monopoly profits are 68 percent higher and margins are 210 percent higher when the monopolist accounts for the dynamic nature of demand, compared to “myopic pricing” which ignores the decline in future demand due to current sales. Duopoly (industry) profits are 32 percent higher and margins are 94 percent higher when the firms account for the dynamic nature of demand, compared to “myopic pricing.”

Clearly, one should account for the dynamic nature of demand when analyzing pricing behavior in durable goods markets. Often researchers observe prices and use first order conditions from a static profit maximization to infer marginal costs (e.g. Berry, Levinsohn, and Pakes (1995)). Observation 1 suggests that marginal cost estimates computed in this manner for durable goods will be too high. That is, prices are high because the firm does not want to reduce future demand, not because its marginal costs are high.

Observation 2 *The monopolist invests too little, relative to the socially optimal level, and duopolists invest too much.*

The investment levels reported in Table 3.4 for the duopoly, monopoly, and social planner (with 1 firm to control) are, respectively, 5.979, 3.693, and 5.682. Hence, the monopolist invests 54 percent less than the social planner, which results in a substantially lower rate of innovation. While we have not computed the socially optimal investment when the planner controls two firms, we can still conclude that (for this parameterization) the duopolists over invest since each firm invests more than the planner with one firm. If the planner had two firms, total investment would likely increase, but the increase would clearly be less than double since the marginal benefit (to the planner) of one firm’s innovation is reduced when the other firm may also successfully innovate.

Given that duopolists over invest, the possibility exists that a monopoly may yield higher social surplus (summing consumer surplus and industry profits). However, we find this to not be the case: social surplus with a duopoly achieves 81 percent of the planner’s surplus, whereas the monopoly social surplus is only 52 percent of the planner’s surplus.

Observation 3 *Consumers capture the lion's share of surplus generated in durable good markets with innovation.*

From Table 3.4 we compute that consumers enjoy 96 percent of the duopoly's social surplus and 89 percent of the monopoly's social surplus. In essence, consumers benefit substantially from the innovation required to induce upgrades.

As mentioned, the initial state for all simulations has all consumers being non-owners (i.e., with $\tilde{q} = 0$), which is clearly not in the steady-state distribution (i.e., ergodic set) of states. Hence, we split the simulations into two sub-periods. The "pre-penetration" sub-period consists of periods 1 to 30, at which point about 98 percent of the consumers have bought at least once. The remaining periods comprise the "post-penetration" sub-period.

In Table 3.4 we present the same measures that appear in Table 3.4 broken down by sub-period, in addition to a few other measures. This breakdown illustrates the effect of firms having to compete with their previously sold units.

Observation 4 *The monopolist innovates much more after full penetration occurs, whereas duopolists innovate slightly less after full penetration occurs.*

In Table 3.4 we report the monopolists innovation rate of .511 in the pre-penetration period compared to .897 after penetration. The duopolists have innovation rates of .926 and .921 in the two periods, respectively. These differences reflect competitive forces. The monopolist does not have a need to innovate until after most consumers have purchased the good. To induce consumers to upgrade the monopolist must innovate. In the duopoly case, competition for consumers, and hence the advantage of high quality, is present even when penetration is low. The duopolists therefore invest heavily in both sub-periods.

Thus far, the observations focus on comparing means. We also observe substantial variation in outcomes from period to period.

Observation 5 *Margins vary substantially from period to period, as firms adjust prices to market conditions.*

In Table 3.4 we present the distribution of margins across periods. The min and max duopoly margins are .626 and 2.598 and the monopolist's margins vary from 1.36 to 10.0.

3.3.2 Counterfactual Analysis

Recently Advanced Micro Devices (AMD) has filed a lawsuit contending that Intel has engaged in anti-competitive practices that deny AMD access to a share of the CPU market. We can use our model to study the effect of such practices on innovation and pricing, and ultimately consumer surplus and firms' profits. We perform a series of counterfactual simulations in which we vary the portion of the market to which one firm has exclusive access. The firm which has exclusive access to a portion of the market is restricted to offer the same price in both sub-markets.

In Figure 3.4 we plot the margins, innovation rates, and consumer surplus when the "access denied" portion of the market varies from zero to one (in .1 increments). Interestingly, the margins do not begin to increase appreciably until the restricted portion of the market reaches 30 percent. Consumer surplus does not decline appreciably until the restriction hits 0.7. This suggests that as the market moves from a duopoly to a monopoly, most of the welfare loss is due to higher margins and not the slower rate of innovation.

3.4 Conclusions and Future Research

This paper presents a dynamic model of durable goods oligopoly with endogenous innovation. We show that accounting for the durable nature of products in an equilibrium setting has important implications on firm level profits and consumer surplus. To incorporate the endogenous distribution of product ownership, we introduce an approximation to this distribution. Calibrating the model produces results that are reasonably consistent with data from the industry.

There are numerous possibilities for future research. One potential route concerns when it is optimal for a firm to lease versus sell its durable product. A firm could take a durable good and sell it as a non-durable good, and this may be an alternative interpretation of the lease/sell scenario. There are several nature questions that arise, such as under what conditions should a firm sell versus lease, is it more advantage for the leader or laggard to lease, which outcome is best from a social welfare standpoint, etc.

To make the model computationally tractable, we have used an approximation to the true distribution of consumers over product ownership. We would like to perform additional robustness checks on the quality of our approximation, such as using higher-order approximations or alternative functional forms. Ideally, we could solve the equilibrium of the model using a sufficiently coarse grid over the ownership distribution without the use of the approximation. If this grid could be made fine enough to minimize discretization error, this would provide us with the best picture of how accurately our approximation captures the effects of the ownership distribution on the firms' policy functions. If the policy functions from the discretized model and the approximated model were similar, we could conclude that the approximation is adequate.

Finally, our goal is to more carefully calibrate or estimate the model. This depends crucially on finding an appropriate set of moments to match and on addressing the issue of multiple equilibria. The computational burden of a full-blown estimation routine may currently be infeasible, given that it takes between four and seven hours to solve for a single equilibrium. We are currently examining ways to address these issues.

Table 3.1: Base Case Parameters

Parameter	Value
γ	1.1
α	0.05
ξ	0 (all firms)
β	0.95
$\bar{q}$	3.9
mc	43
δ	0.3
$f(\tau = \delta \| x)$	ax/(1+ax)
a	2
c	1
M	100

Base parameterization of variables and functional forms.

Table 3.2: Empirical vs. Simulated Moments

	Industry	Model
Margins	3.260	1.235
Leader/Laggard Market Share	6.244	2.632
Innovation Rate	0.873	0.914
Avg. Quality Difference	32.42%	45.20%

Compares the empirical versus simulated moments, which include the price margins, the ratio of the leader and laggard market shares, the mean innovation rate, and the mean difference in quality between the laggard and leader.

Table 3.3: Policy Regressions

	Price Margin		Innovation Rate	
	Leader	Laggard	Leader	Laggard
Constant	-0.148	1.589	0.959	0.935
Leader Extended Frontier ?	0.073	-0.023	-0.002	0.008
Mean Owned Product Quality	-0.216	-0.172	-0.011	-0.017
Quality Difference	0.961	-0.593	0.053	-0.045
Quality Difference Squared	-0.209	0.109	-0.046	0.025
R^2	0.652	0.639	0.967	0.934

Regression results using simulated steady-state data of the firms' choice variables against various transformations of the state variables. The explanatory variables are a dummy for whether the leading firm extended their quality lead, the share of consumers that own a product, the mean quality of the product owned by these consumers, the quality difference between the leader and laggard, and the square of the quality difference.

Table 3.4: Industry Measures under Five Scenarios

	Durable Duopolist	Myopic Pricing Duopolist	Durable Monopolist	Myopic Pricing Monopolist	Social Monopolist
Profits (industry)	11473	8699	19449	11562	0
Consumer Surplus	264703	259605	157125	191399	340145
Margins	1.235	0.637	3.023	0.974	0
Innovation Rate	0.922	0.885	0.782	0.831	0.903
Investment	5.979	3.931	3.693	4.197	5.682

Comparison of the simulated outcomes under five different scenarios: durable duopoly, myopic pricing duopoly, durable monopoly, myopic pricing monopoly, and social monopoly. In the myopic cases, the firm(s) choose price taking into account the current distribution of product ownership, but without regard for how this distribution will change tomorrow based on demand today.

Table 3.5: Pre- vs. Post- Penetration

	Durable Duopolist	Myopic Pricing Duopolist	Durable Monopolist	Myopic Pricing Monopolist	Social Monopolist
Pre-Penetration					
Profits (industry)	7535	6472	14271	9141	0
Consumer Surplus	138823	137457	64007	83735	174524
Margins	1.056	0.760	3.846	1.577	0.000
Investment	6.341	4.520	2.117	2.905	3.855
Innov Rate	0.926	0.897	0.511	0.658	0.836
Quality Diff	0.864	1.413			
Leader Turnover	0.029	0.037			
Post-Penetration					
Profits (industry)	3939	2227	5178	2421	0
Consumer Surplus	125880	122148	93118	107664	165621
Margins	1.311	0.584	2.670	0.716	0.000
Investment	5.825	3.678	4.369	4.750	6.465
Innov Rate	0.921	0.880	0.897	0.905	0.931
Quality Difference	0.276	0.428			
Leader Turnover	0.003	0.002			
Frequency of Equal Quality	0.170	0.083			
Period of Full Penetration	23	23	25	22	22

Comparison of simulated outcomes with the results broken down by sub-periods. “Quality difference” is the mean percent quality difference between the leader and laggard, “leader turnover” is the fraction of periods in which the leading firm changes identities.

Table 3.6: Margin Distribution

	Durable Duopolist	Myopic Pricing Duopolist	Durable Monopolist	Myopic Pricing Monopolist
Min	0.626	0.558	1.360	0.414
5%	0.815	0.598	1.996	0.693
10%	0.842	0.625	2.138	0.702
25%	1.170	0.665	2.404	0.720
50%	1.451	0.689	2.802	0.721
75%	1.641	0.710	3.312	0.722
90%	1.836	0.852	4.691	1.550
95%	1.922	1.044	5.212	3.244
Max	2.598	3.184	10.010	4.422
Mean	1.235	0.637	3.023	0.974
Stdev	0.359	0.127	1.142	0.208

Distribution of sales-weighted average price margins under the different market scenarios, where the margin is calculated as $(p_{jt} - mc_j)/mc_j$.

Table 3.7: Comparison of the Leader and Laggard

	Durable Duopolist	Myopic Pricing Duopolist
Leader Market Share	0.104	0.158
Laggard Market Share	0.051	0.021
Leader Profits	8523.020	7341.690
Laggard Profits	2950.140	1357.290
Leader Price	114.991	75.537
Laggard Price	76.881	65.202
Leader Investment	7.527	4.308
Laggard Investment	4.432	3.554
Leader Innovation Rate	0.938	0.896
Laggard Innovation Rate	0.899	0.877
Period of Full Penetration	23	23

Summary statistics according to whether the firm is the leader or laggard in the market.

Figure 3.1: Price Margin Regressions

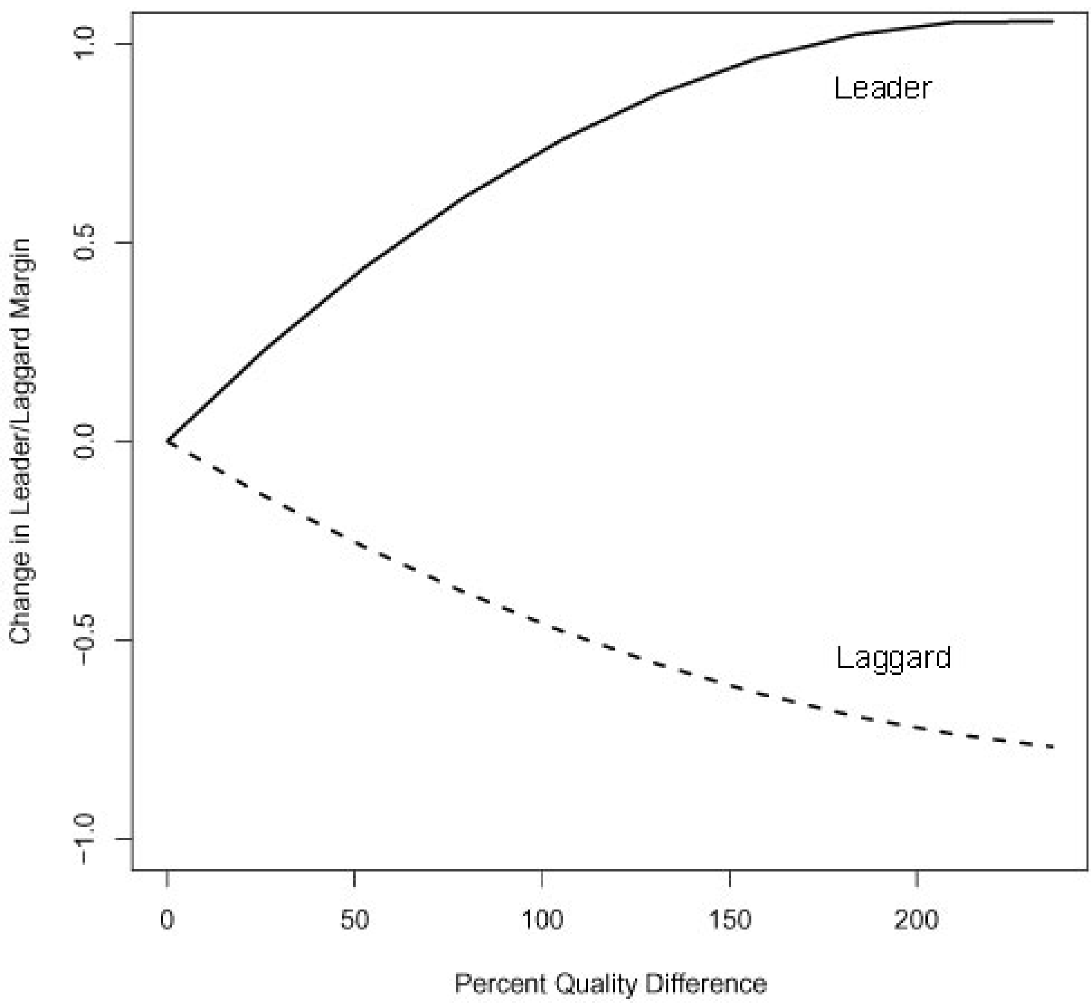

Figure 3.2: Innovation Rate Regressions

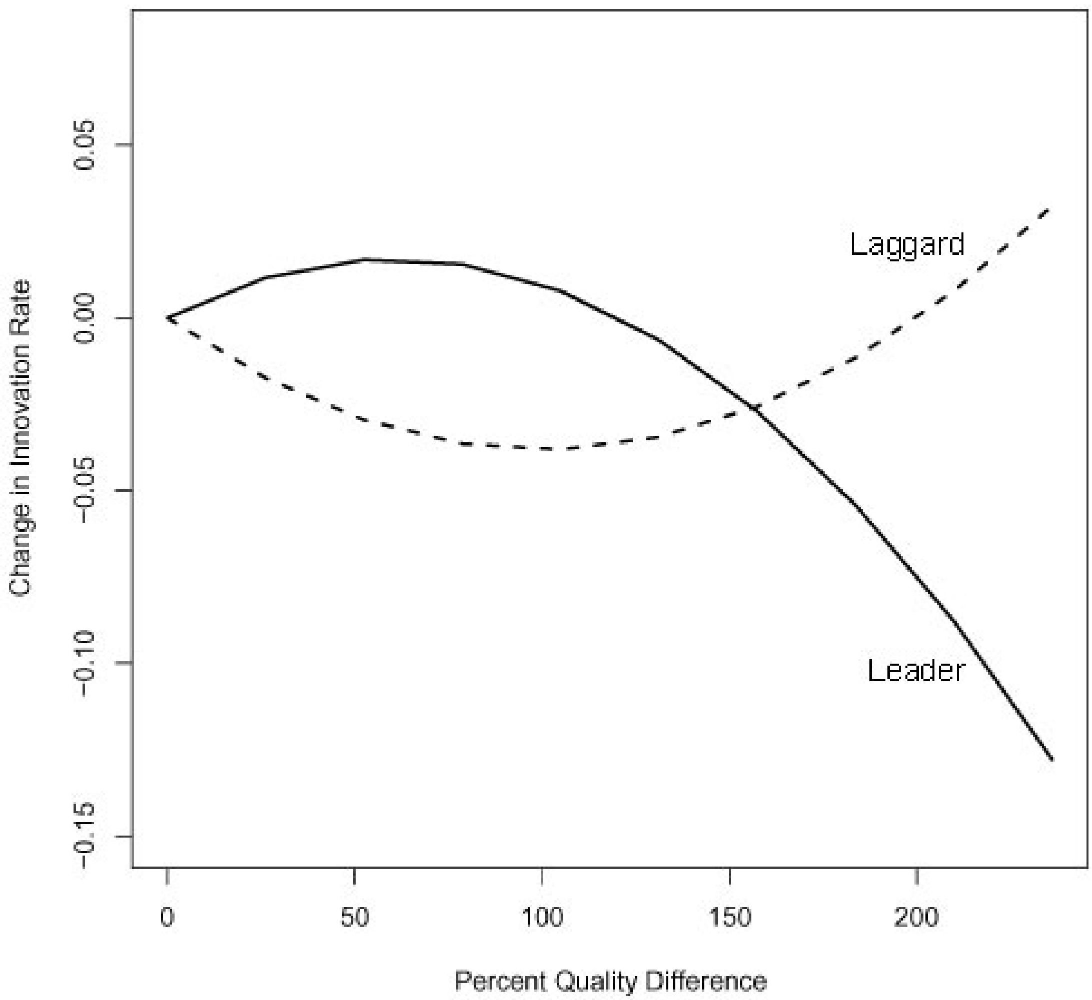

Figure 3.3: Demand by Vintage

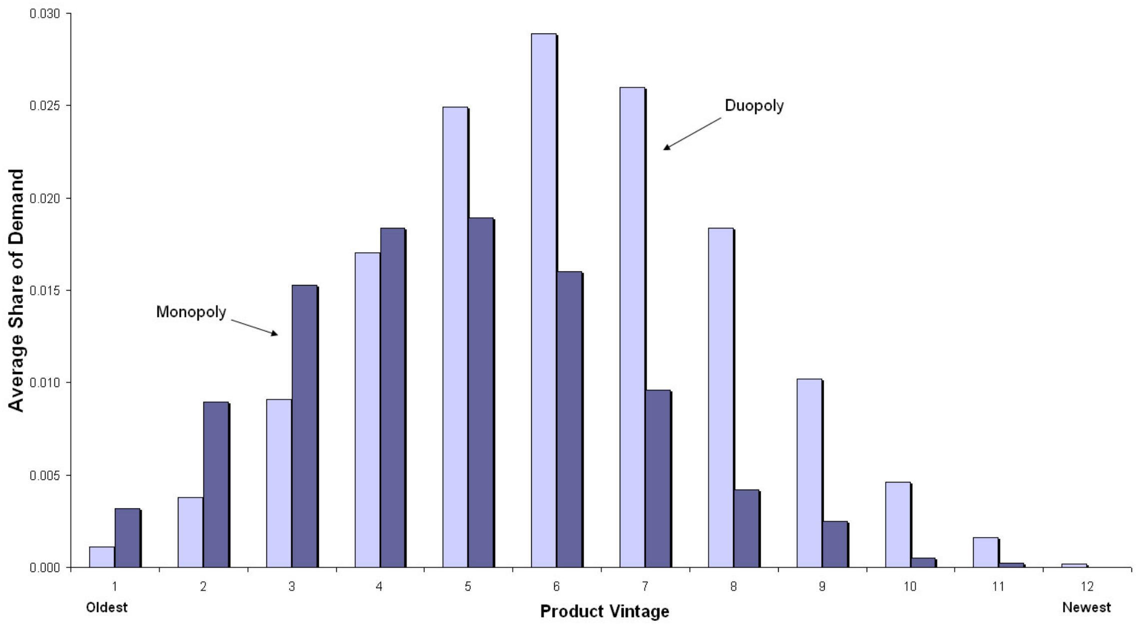

Figure 3.4: Market Restriction

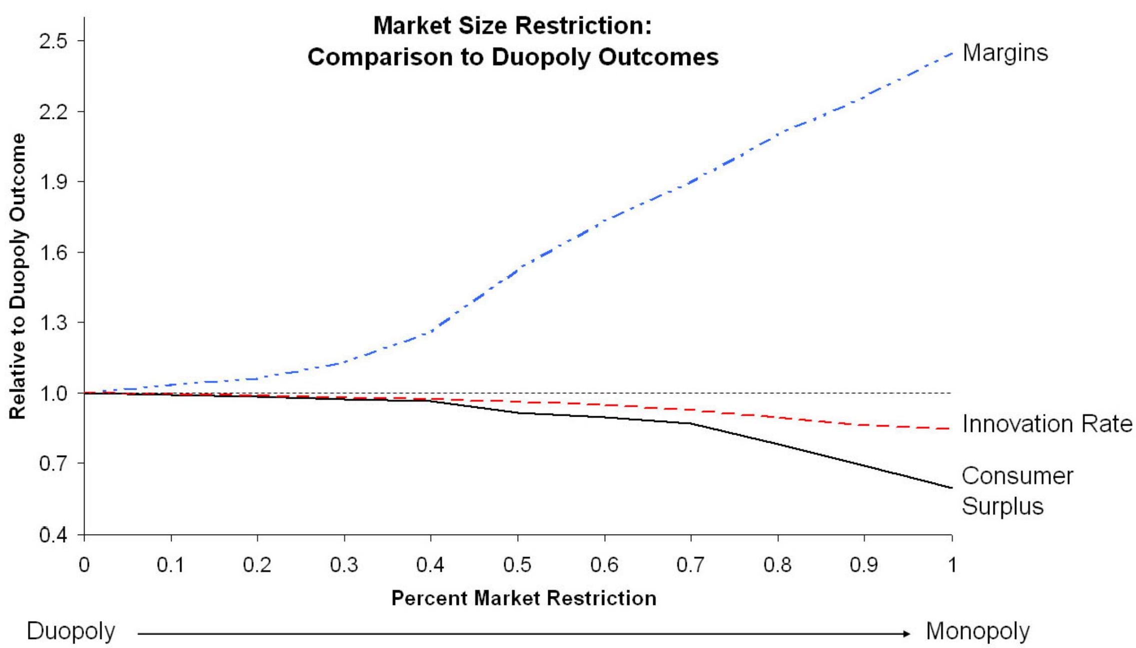

Bibliography

[1] Adda, J. and R. Cooper (2000). "The Dynamics of Car Sales: A Discrete Choice Approach," NBER Working Paper 7785.

[2] Adda, J. and R. Cooper (2000), "Balladurette and Juppette: A discrete analysis of scrapping subsidies", *Journal of Political Economy*, 108(4), 778–806.

[3] Bass, F. M. (1969), "A New Product Growth Model for Consumer Durables," *Management Science*, 15, 215–227.

[4] Bayus, B. L. (1988), "Accelerating the Durable Replacement Cycle with Marketing Mix Variables," *Journal of Product Innovation Management*, 5, 216–226.

[5] Bayus, B. L. (1991), "The Consumer Durable Replacement Buyer," *Journal of Marketing*, 55(1), 42–51.

[6] Bayus, B. L., S. Hong, and R. P. Labe Jr. (1989), "Developing and Using Forecasting Models of Consumer Durables," *Journal of Product Innovation Management*, 6, 5–19.

[7] Benitez-Silva, H., G. Hall, G. Hitsch, G. Pauletto, and J. Rust (2000), "A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems," Working paper, Yale University.

[8] Benkard, L. (2004), "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft," *Review of Economic Studies*, 71(3), 581–611.

[9] Bond, E. and L. Samuelson (1984), "Durable Good Monopolies with Rational Expectations and Replacement Sales," *RAND Journal of Economics*, 15(3), 336–345.

[10] Bresnahan, T. F. and S. Greenstein (1999), "Technological Competition and the Structure of the Computer Industry," *Journal of Industrial Economics*, 47(1), 1–40.

[11] Bulow, J. (1982), "Durable Goods Monopolists," *Journal of Political Economy*, 90, 314–332.

[12] Cameron, A. C. and P. K. Trivedi (1999), "Essentials of Count Data Regression," Working paper, University of California, Davis.

[13] Carranza, J. (2004a), "Demand for Durable Goods and the Dynamics of Quality," mimeo, University of Wisconsin-Madison.

[14] Coase, R. (1972), "Durability and Monopoly," *Journal of Law and Economics*, 15, 143–149.

[15] Dubé, J. P., K. Sudhir, A. Ching, G. Crawford, M. Draganska, J. T. Fox, W. Hartmann, G. J. Hitsch, V. B. Viard, M. Villas-Boas, and N. Vilcassim, (2005), "Recent Advances in Structural Econometric Modeling: Dynamics, Product Positioning, and Entry," *Marketing Letters*, 16(3–4), 209–224.

[16] Ericson, S. and A. Pakes (1995), "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62(1), 53–82.

[17] Erdem, T., S. Imai, and M. P. Keane (2003), "Brand and Quality Choice Dynamics Under Price Uncertainty," *Quantitative Marketing and Economics*, 1, 5–64.

[18] Erdem, T., M. P. Keane, and B. Sun (2005), "A Dynamic Model of Brand Choice when Price and Advertising Signal Product Quality," Working paper, Carnegie Mellon.

[19] Erdem, T., M. P. Keane, T. S. Öncü, and J. Strebel (2005), "Learning About Computers: An Analysis of Information Search and Technology Choice", *Quantitative Marketing and Economics*, 3, 207–246.

[20] Esteban, S. and M. Shum (forthcoming), "Durable-Goods Oligopoly with Secondary Markets: the Case of Automobiles," *RAND Journal of Economics*.

[21] Fudenberg, D. and J. Tirole (1998), "Upgrades, Trade-ins, and Buybacks," *RAND Journal of Economics*, 29(2), 237–252.

[22] Goettler, R., C. Parlour, and U. Rajan (2005), "Equilibrium in a Dynamic Limit Order Market," *Journal of Finance*, vol. 60, no. 5, 2149–2192.

[23] Goettler, R., C. Parlour, and U. Rajan (2006), "Microstructure Effects and Asset Pricing," Carnegie Mellon University, Tepper School of Business, working paper #2006-E97.

[24] Gonul, F. and K. Srinivasan (1996), "Impact of Consumer Expectations of Coupons on Purchase Behavior," *Marketing Science*, 15, 262–279.

[25] Gordon, B. R. (2006), "A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry," Working paper, Carnegie Mellon University.

[26] Gowrisankaran, G. and M. Rysman (2006), "Dynamics of Consumer Demand for New Durable Consumer Goods," Working paper, Boston University.

[27] Gowrisankaran, G. and R. Town (1997), "Dynamic Equilibrium in the Hospital Industry," *Journal of Economics and Management Strategy*, 6(1), 45–74.

[28] Gul, F. (1987), "Noncooperative Collusion in Durable Goods Oligopoly," *RAND Journal of Economics*, 18(2), 248–254.

[29] Hall, G. and J. Rust (1999), "An Empirical Model of Inventory Investment by Durable Commodity Intermediaries," Cowles Foundation Discussion Paper 1228, Cowles Foundation, Yale University.

[30] Hendel, I. and A. Lizzeri (1999), "Interfering with Secondary Markets," *RAND Journal of Economics*, 30(1), 1–21.

[31] Hendel, I. and A. Nevo (2005), "Measuring the Implications of Sales and Consumer Inventory Behavior," NBER Working Paper #11307.

[32] Horsky, D. (1990), "A Diffusion Model Incorporating Product Benefits, Price, Income, and Information," *Marketing Science*, 9(4), 342–365.

[33] Kamakura, W. and G. Russell (1989), "A Probabilistic Choice Model for Market Segmentation and Elasticity Structure," *Journal of Marketing Research*, 26, 370–390.

[34] Kleiman, E. and T. Ophir (1966), "The Durability of Durable Goods," *Review of Economic Studies*, 33(2), 165–178.

[35] Langlois, R. (1992), "External Economies and Economic Progress: The Case of the Microprocessor Industry," *Business History Review*, 66, 1–50.

[36] Langlois, R. (2002), "Digital Technology and Economic Growth: the History of Semiconductors and Computers," in B. Steil, D. Victor, and R. R. Nelson (eds.), *Technological Innovation and Economic Performance*. Princeton: Princeton University Press, 265–284.

[37] Lee, I. H. and J. Lee (1998), "A Theory of Economic Obsolescence," *Journal of Industrial Economics*, 46(3), 383–401.

[38] Lee, D. and K. Wolpin (2006), "Accounting for Wage and Employment Changes in the U.S. from 1968–2000: A Dynamic Model of Labor Market Equilibrium," PIER Working Paper 06-005.

[39] Levinthal, D. and D. Purohit (1989), "Durable Goods and Product Obsolescence," *Marketing Science*, 8(1), 35–56.

[40] Lewis, M. (2005), "A Dynamic Programming Approach to Customer Relationship Pricing," *Management Science*, 51(6), 986–994.

[41] Luan, J. Y. (2005), "Optimal Inter-Release Timing for Sequential Releases," Working paper, Yale School of Management.

[42] Markovich, S. (2004), "Snowball: A Dynamic Oligopoly Model with Network Externalities," Working paper, Northwestern University.

[43] Markovich, S. and Moenius, J. (2005), "Winning While Losing: Competition Dynamics in the Presence of Indirect Network Effects," Working paper, Northwestern University.

[44] McKinsey Global Institute (2001), "US Productivity Growth 1995-2000: Understanding the Contribution of Information Technology Relative to Other Factors," Washington, DC: McKinsey & Company.

[45] Mehta, N., S. Rajiv, and K. Srinivasan (2004), "Role of Forgetting in Memory-based Choice Decisions", *Quantitative Marketing and Economics*, 2(2), 107–140.

[46] Melnikov, O. (2001), "Demand for Differentiated Durable Products: The Case of the US Computer Printer Industry," mimeo, Yale University.

[47] Nair, H. (2007), "Intertemporal Price Discrimination in the Market for 32-bit Video-Games," Working paper, Stanford University.

[48] Newey, W. K., and K. D. West (1987), "Hypothesis Testing with Efficient Method of Moments Estimation," *International Economic Review*, 28(3), 777–787.

[49] Norton, J. A. and F. M. Bass (1987), "A Diffusion Theory Model of Adoption and Substitution for Successive Generations of High-Technology Products," *Management Science*, 33(9), 1069–1086.

[50] Pakes, A. (1986). "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54, 755–784.

[51] Pakes, A. and P. McGuire (1994), "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics*, 25(4), 555–589.

[52] Pakes, A., G. Gowrisankaran, and P. McGuire (1993), "Code for Implementing the Pakes-McGuire Algorithm for Computing Markov Perfect Equilibrium," Mimeo, Yale University.

[53] Prince, J. T. (2005), "Measuring the Digital Divide: Structural Estimation of the Demand for Personal Computers," Working paper, Cornell University.

[54] Ratchford, B. T., S. K. Balasubramanian, and W. A. Kamakura (2000), "Diffusion Models with Replacement and Multiple Purchases," in *New Product Diffusion Models*, V. Mahajan, E. Muller, and Y. Wind, Springer-Verlag.

[55] Rust, J. (1986), "When is it Optimal to Kill off the Market for Used Durable Goods?" *Econometrica*, 54(1), 65–86.

[56] Rust, J. (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999–1033.

[57] Rust, J. (1994), "Structural Estimation of Markov Decision Processes," in R. Engle and D. McFadden (eds.), *Handbook of Econometrics*, Vol. 4.

[58] Rust, J. (1996), Numerical Dynamic Programming in Economics. In *Handbook of Computational Economics*, Ch. 14. H. Amman, D. Kendrick, and J. Rust (eds).

[59] Rust, J. (1997a), "Using Randomization to Break the Curse of Dimensionality," *Econometrica*, 65(3), 487–516.

[60] Rust, J. (1997b), "A Comparison of Policy Iteration Methods for Solving Continuous-State, Infinite-Horizon Markovian Decision Problems using Random, Quasi-random, and Deterministic Discretizations," Mimeo, Yale University.

[61] Small, K. and H. Rosen. (1981), "Applied Welfare Economics with Discrete Choice Models," *Econometrica*, 49(1), 105–130.

[62] Sieper, E. and P. Swan (1973), "Monopoly and Competition in the Market for Durable Goods," *Review of Economic Studies*, 40(3), 333–351.

[63] Song, M. (2003a), "A Dynamic Model of Cooperative Research in the Semiconductor Industry," Working paper, Georgia Institute of Technology.

[64] Song, M. (2003b), "Measuring Consumer Welfare in the CPU Market: An Application of the Pure Characteristics Demand Model," Working paper, Georgia Institute of Technology.

[65] Song, I. and Chintagunta, P. (2003), "A Micromodel of New Production Adoption with Heterogeneous and Forward-Looking Consumers: Application to the Digital Camera Category," *Quantitative Marketing and Economics*, 1, 371–407.

[66] Sun, B. (2005), "Promotion Effect on Endogenous Consumption," *Marketing Science*, 24(3), 430–443.

[67] Sun, B., S. A. Neslin, and K. Srinivasan (2003), "Measuring the Impact of Promotions on Brand Switching Under Rational Consumer Behavior," *Journal of Marketing Research*, 4, 389–405.

[68] Steffens, P. R. (2003), "A Model of Multiple Ownership as a Diffusion Process," *Technological Forecasting and Social Change*, 70(9), 901–917.

[69] Stevens, J. (1994). *New Games: Strategic Competition in the PC Revolution*, Pergamon Press, New York.

[70] Stokey, N. (1981), "Rational Expectations and Durable Goods Pricing," *Bell Journal of Economics*, 12, 112–128.

[71] Swan, P. (1970), "Durability of Consumption Goods," *American Economic Review*, 60(5), 884–894.

[72] Swan, P. (1971), "The Durability of Goods and Regulation of Monopoly," *Bell Journal of Economics*, 2(1), 347–357.

[73] Waldman, M. (1993), "A New Perspective on Product Obsolescence," *Quarterly Journal of Economics*, 108, 273–283.

[74] Waldman, M. (1996), "Durable Goods Pricing When Quality Matters," *Journal of Business*, 69(4), 489–510.

[75] Waldman, M. (2003), "Durable Goods Theory for Real World Markets," *Journal of Economics Perspectives*, 17(1), 131–154.

[76] Wilson, L. O. and J. A. Norton, (1989), "Optimal Entry Time for a Product Line Extension," *Marketing Science*, 8, 1–17.

Appendix A

Computational Details

A.1 Chapter 2: The Value Function Approximation

Two necessary conditions are that the state space is compact and that the transition density is sufficiently smooth. To ensure compactness, it is only necessary to bound prices from below at zero and above at some arbitrarily large price where the density is sufficiently close to zero, and then to renormalize the density so that it integrates to one. Smoothness requires that the transition density does not contain too many spikes, which is guaranteed through the use of a multivariate normal. If the density contains any discontinuities, the approximate integral underestimates the true integral, requiring a higher number of random grid points to obtain the same degree of accuracy. With these conditions satisfied, the random Bellman operator possesses a bound that is a linear function of d. Thus, the randomization technique breaks the curse of dimensionality associated with approximating the integral as part of the contraction mapping.

More precisely, I use the random multigrid algorithm outlined in Rust (1997). This algorithm consists of a set of outer iterations $k = 1, 2, \ldots,$ where a number $N^{(k)}$ of uniform random sample points $\{\tilde{p}_1, \ldots, \tilde{p}_{N^{(k)}}\}$ is drawn at each iteration independently of the samples drawn at previous iterations $k-1, k-2, \ldots$. The basic idea is to start with relatively few sample points $N^{(0)}$ at $k = 0$ and to successively increase the number of sample points according to the rule

$$N^{(k)} = 2^{2k} N^{(0)}$$

Within each outer iteration k, a number $T(k)$ of successive approximation steps are taken using the random Bellman operator $\hat{\Gamma}_{N^{(k)}}$. Let $\hat{V}^{(k)}$ denote the value function produced after $T(k)$ steps at outer iteration k. The starting point for the value function at outer iteration $k+1$ is the value function $\hat{V}^{(k)}$ produced at iteration k. This is most easily done using nonparameteric regression. This leads to the recursion

$$\hat{V}_i^{(k+1)} = \hat{\Gamma}_{N^{(k)}}(\hat{V}_i^{(k)})$$

where the starting point for the value function at iteration 0 is maximum of the period utility function. This is made possible by the fact that $\hat{\Gamma}_{N_p^{(k)}}$ is self-approximating – it may be evaluated at arbitrary points in the state space without the need for interpolation.

In practice, I set $N^{(0)} = 100$ and find that the multigrid algorithm converges after three or four iterations. This yields significant computational savings because using the earlier values as starting values over the finer grid increases the rate of convergence.