

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Price Setting in an Innovative Market

Adam Copeland,
Federal Reserve Bank of New York

Adam Hale Shapiro,
Federal Reserve Bank of San Francisco

February 2013

Working Paper 2013-04

<http://www.frbsf.org/publications/economics/papers/2013/wp2013-04.pdf>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of New York, the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

Price Setting in an Innovative Market*

Adam Copeland[†] and Adam Hale Shapiro[‡]

February 28, 2013

Abstract

We examine how the confluence of competition and upstream innovation influences downstream firms' profit-maximizing strategies. In particular, we analyze how, in light of these forces, the downstream firm sets the price of the product over its life cycle. We focus on personal computers (PCs) and introduce two novel data sets that describe prices and sales in the industry. Our main result is that a vintage-capital model that combines a competitive market structure with a rapid rate of innovation is well able to explain the observed paths of prices, as well as sales and consumer income, over a typical PC's product cycle. The analysis implies that rapid price declines are not caused by upstream innovation alone, but rather by the combination of upstream innovation and a competitive environment.

Key words: innovation, market structure, computers

JEL classification: D40, L10, L63, O30

*We thank Ana Aizcorbe, Olivier Armantier, Steve Berry, Ron Borkovsky, Ben Bridgman, Ron Goettler, Phil Haile, Bronwyn Hall, Kyle Hood, David Mowery, Matt Osborne, Michael Ostrovsky, Ariel Pakes, Jeff Prince, Dave Rapson, James Roberts, and Marc Rysman for their comments and suggestions. An earlier version of this paper circulated under the title "The Impact of Competition on Technology Adoption: An Apples-to-PCs Analysis." Much of the work on this paper occurred while both authors worked at the Bureau of Economic Analysis. The views expressed here are those of the authors and do not necessarily reflect the position of the Bureau of Economic Analysis, the U.S. Department of Commerce, the Federal Reserve Bank of New York, the Federal Reserve Bank of San Francisco, or the Federal Reserve System.

[†]Federal Reserve Bank of New York; e-mail: adam.copeland@ny.frb.org; webpage: <http://www.copeland.marginalq.com> .

[‡]Federal Reserve Bank of San Francisco; e-mail: adam.shapiro@sf.frb.org; webpage: <http://sites.google.com/site/adamshap/> .

1 Introduction

We examine how the confluence of competition and upstream innovation influences downstream firms' profit-maximizing strategies. In particular, we analyze how, in light of these forces, the firm sets the price of the product over its life cycle. We focus on the personal computer (PC) industry and begin our work by describing the data and presenting stylized facts. Our main result is that a vintage-capital model which combines a competitive market structure with a rapid rate of innovation is well able to explain the observed paths of prices, sales, and consumer income over a typical PC's product cycle. The simplicity of the model leaves ample room for extensions to capture other important features of the PC market. Nevertheless, we argue the model provides a useful benchmark for comparison with more complicated models.

We use data from two sources, the NPD Group and MetaFacts. The NPD Group provides us with product-level data on monthly revenues and units sold from 2001 to 2009 as well as product characteristics (e.g., chip type and screen size). MetaFacts provides survey data allowing us to link income and the timing of a computer purchase. Using these data we present evidence that PC manufacturers set prices that decline rapidly over a short product cycle. A typical computer's product cycle lasts only four months and, over this time period prices fall 12 percent. Furthermore, we find that sales rapidly decline over the product cycle and firms frequently introduce new, higher-quality products. Finally, we show that average income of PC purchasers also falls over the product cycle. The exception are Apple's products, which have less frequent product introductions, roughly constant prices over their product cycle, and consumers with high and narrow income distributions.

The rapid decline in computer prices could be the result of a variety of forces. Process innovation, falling input costs, intertemporal price discrimination, and competition are the explanations that may be relevant for the retail computer industry. Given the short time frame of computer product cycles, falling input costs or process innovation can explain, at most, a small fraction of the 36 percent (annual rate) decline in PC prices. Indeed, prices for screens, batteries, and other components of PCs do not decline at such a rapid rate. Furthermore, Apple uses many of the same intermediate inputs used by PCs, yet its prices decline only negligibly over time. Consequently, we rule out process innovation and falling input costs as a main driver of declining PC prices over the product cycle.¹

¹Certainly, however, these factors may be important for explaining longer term price trends in this industry.

Intertemporal price discrimination, whereby the firm charges a high price early in the product cycle to those with the highest willingness to pay, seems like a plausible explanation at first glance. Indeed, we find that the average income of consumers who purchase PCs falls over the product cycle. Stokey (1979), however, showed that this type of price discrimination is profit maximizing only under very strict assumptions. First, the firm needs a considerable amount of market power; otherwise, competitive forces will determine the price. Second, consumers' reservation prices must be correlated with their time preferences; otherwise high willingness-to-pay consumers would prefer to wait for the price to fall. Finally, the firm must have the ability to commit to future prices or future production to avoid the time inconsistency dilemma posed by Coase.² Apple's pricing behavior casts doubt on price discrimination being the main force behind the rapid price declines. If the rapid price declines for PCs market are attributable to intertemporal price discrimination, then there must be some particular reason Stokey's conditions are met for all PCs except for Apple. We have no reason to believe that willingness to pay is more correlated with time preference for consumers of non-Apple PCs than for consumers of Apple computers. Furthermore, market power should be positively correlated with a firm's ability to commit to a price or production schedule. Given the industry wisdom that Apple has more market power than other PC manufacturers, intertemporal price discrimination seems like an unlikely explanation for the declining pricing patterns.

Competition, however, seems to be a plausible force behind declining PC prices over the product cycle. It is conceivable that the frequent adoption of higher-quality PCs can drive down the prices of PCs currently on the market. To see how well competition can explain this market dynamic we develop a vintage-capital model. While the model we develop is parsimonious, it captures the key features of the industry such as the joint behavior of rapid product introductions and the time series of prices, sales, and purchasers' income over the product cycle. On the demand side, we use the quality-ladder framework of Shaked and Sutton (1982, 1983). Consumers differ in their budgets for computers, and computers differ by quality (i.e., vintage). On the supply side, firms offer computers of different vintages and set the product's price. Firms face a constant marginal cost and pay a fixed cost to update the quality of their product. This fixed cost makes the firm's problem dynamic. Because firms need to account for the pricing and updating decisions of their competitors, their problem is also strategic.

We calibrate the vintage-capital model to fit the time series of prices and sales for a

²Bulow (1982) shows that a firm will use an inefficient production technology or produce goods that are less durable to skirt the commitment problem.

typical computer over its product cycle. Despite the model’s simple structure, we are able to closely match the data through the combination of competition and rapid innovation. The model rationalizes the frequent introduction of new products alongside the rapid price declines as market-stealing behavior. Given consumer preferences for quality, the firm with the highest-quality computer is able to capture a large market share and still charge a substantial mark-up. Consequently, there are large profits associated with having the highest-quality product on the market. These gains, however, are quickly eroded as competing manufacturers introduce higher quality computers. The introduction of a new computer obsolesces existing computers, generating rapid price declines over a computer’s product-cycle. Finally, the decline in unit sales over the product cycle is primarily driven by consumer heterogeneity. The combination of consumer heterogeneity and falling prices implies that consumers with smaller budgets purchase computers later in the product cycle, consistent with the implications of the Metafacts survey data on income and the timing of computer purchases.

The results of our calibration imply that the decline in prices, jointly with sales and consumer income, is due to the interaction between competitive forces and the rapid rate of upstream innovation in the market. To isolate the roles that innovation and competition play in generating rapid price declines, we use the model to explore pricing under alternative settings of innovation and competition. We find that a faster growth rate in upstream innovation implies a steeper price decline. Specifically, more rapid innovation implies that different vintages of computers are farther apart on the quality ladder. This greater vertical differentiation leads to a higher introductory price for a computer, followed by bigger price declines.

To assess the impact of competition on price setting, we consider our model in the case of monopoly. The monopoly case can be considered the case where we set the fixed cost of entry so high that only one firm is able to earn a positive profit. We find that under this scenario the firm’s pricing strategy radically changes—pricing is flat over the product cycle. This result implies that upstream innovation alone does not cause rapid price declines, rather it is the combination of upstream innovation and a competitive environment.

We then use the monopoly case considered above as an out-of-sample test of the model. Recall that Apple products have a different operating system which make them quite dissimilar to other personal computers on the market.³ Apple represents a firm whose

³In our calibration exercise we excluded Apple because of its differentiation along the horizontal dimension.

product is highly differentiated in the horizontal dimension; within the framework of our model it is a monopolist. We examine how well Apple’s price and sales decisions match the model’s predictions in the case of monopoly, but under the same calibrated parameters and upstream innovation rate as the competitive setting. Validating the model, we find that the model’s predictions for the monopoly case closely match the near constant prices and sales observed in the Apple data.

The paper is structured as follows: Section 2 reviews the related literature. Section 3 describes the data from the NPD Group and Technology User Profile survey and provides a description of the stylized facts for the PC market. In Section 4, we present the model, and in Section 5 we take the model to the data. In Section 6 we describe an out-of-sample exercise and then we conclude in Section 7.

2 Related Literature

This paper builds upon the literature analyzing the effect of competition on pricing behavior⁴ as well as studies of the prices for durable technological goods.⁵ It is closely tied to Aizcorbe and Kortum (2005), who use a vintage-capital model to analyze pricing and production in the semiconductor industry. They argue that the rapid price declines for semiconductor chips are driven by the introduction of better vintages. Similarly, we claim the incorporation of innovations into new computers drives down the price of existing computers. The novelty of our approach, however, is that we allow for competitive strategic interaction between firms and incorporate consumer heterogeneity. Our analysis emphasizes the role of competition in driving prices down over the product cycle, and so providing incentives for computer manufacturers to quickly incorporate innovations into new products. The results of Aizcorbe and Kortum (2005), in contrast, hold regardless of market structure.

Our work also touches upon a large literature commencing with Schumpeter (1934, 1942), and later Arrow (1962), who examined the impact of competition on research and development (R&D) activity. An ongoing line of research has been dedicated to the topic.⁶ Schumpeter conjectured that firms with larger market power would more aggressively

⁴See, for example, Borenstein and Rose (1994); and Gerardi and Shapiro (2009).

⁵See Erickson and Pakes (2008); Aizcorbe (2005); Berndt and Rappaport (2001); Gowrisankaran and Rysman (2009); Conlin (2010); and Pakes (2003).

⁶See Dasgupta and Stiglitz (1980); Gilbert and Newbery (1982); Aghion and Howitt (1992); Greenstein and Ramey (1998); Aghion et al. (2009); Biesebroeck and Hashmi (2009); and Goettler and Gordon (2009); Nosko (2010).

pursue R&D activity. Arrow, however, described a scenario in which a firm with less market power would have a higher incentive to undertake R&D since innovation provides a tool for escaping competition by differentiating itself from its competitors. While these studies referred to industries in which the innovating firm undertakes R&D directly, their question is also relevant for technology-adopting firms, such as PC manufacturers, which we study. Our result—that PC manufacturers seek to embed innovations into their retail products in order to leap frog their competitors and (temporarily) grab market share—is more in line with Arrow’s work.

Finally, our paper builds upon a large literature concerning product differentiation in the computer industry. Specifically, our model provides insight into the manner in which computer manufacturers are able to retain market share in such a highly competitive environment. The model highlights the importance of technology adoption as a means of gaining market share by allowing the firm to vertically differentiate its product. The nice fit with the data implicitly downplays the importance of certain types of horizontal differentiation, such as branding. This result contrasts with the findings of Bresnahan, Stern, and Trajtenberg (1997), who find that horizontal differentiation in the form of brand is needed in addition to vertical differentiation to make accurate predictions about sales. One difference between our study and theirs is that we examine a model that incorporates pricing and sales dynamics within an individual product cycle whereas Bresnahan, Stern, and Trajtenberg take a static cross-sectional approach. In particular, we find that a majority of the firm’s earnings are made in the short time frame following product introduction. A static analysis will inherently assume constant earnings over the course of the entire product cycle, which may downplay the importance of vertical differentiation and “racing to the frontier.” Another major difference between our study and theirs is that Bresnahan, Stern, and Trajtenberg analyzed the personal computer market in the late 1980s, before the introduction of the hugely successful Microsoft Windows 3.0 in 1990, as well as before the “Intel Inside” marketing program began in 1991. It is conceivable that as Microsoft and Intel cemented their dominance over the 1990s, consumers have come to pay closer attention to the operating system-CPU bundle and focused less on the manufacturer’s brand.

3 Data

Our study uses data from two sources: scanner data compiled by NPD Techworld and household survey data from the Technology User Profile (TUP) administered by MetaFacts.

The NPD data are point-of-sale⁷ transaction data (i.e., scanner data) sent to NPD Techworld weekly through automatic feeds from its participating outlets.⁸ The data cover the course of 90 months, November 2001 to April 2009, and consist of sales occurring at outlet stores.⁹ Thus, manufacturers such as Dell that sell primarily directly to the consumer are not included.¹⁰ Each observation consists of a model identification number, specifications for that model, the total units sold, and revenue. From units sold and revenue, we calculate a unit price of each PC sold. Table 1 displays the share of units sold in the data for the entire sample as well as for the notebook and desktop subsamples. Hewlett Packard (HP) and Compaq make up the bulk of computers sold in the data, at 29 and 15 percent, respectively.¹¹

Table 1: Market Share in NPD Sample

	Total	Desktops	Notebooks
Hewlett Packard	0.29	0.35	0.25
Compaq	0.15	0.19	0.11
Toshiba	0.13	0	0.22
Apple	0.12	0.09	0.14
Emachines	0.09	0.20	0
Gateway	0.07	0.06	0.08
Sony	0.07	0.05	0.08
Other	0.09	0.05	0.11

Notes: Market shares are based on units sold in each of three samples: all computers, desktop computers, and notebook computers. Source: NPD Group.

In the TUP survey data, we have access to four annual surveys conducted from 2001 to 2004. TUP is a detailed two-stage survey of households' use of information technology and consumer electronics products and services at home and in the workplace. The first stage is a screener, which asks for the characteristics of each head of household (such as income,

⁷Point-of-sale means that any rebates or other discounts (for example, coupons) that occur at the cash register are included in the price reported; mail-in rebates and other discounts that occur after the sale are not.

⁸The weekly data are organized into monthly data using the Atkins month definition, where the number of weeks assigned to the three months of each quarter are four, four, and five.

⁹This includes sales on outlet stores' websites.

¹⁰This would pose an issue for our analysis only if Dell were an outlier relative to the other PC manufacturers.

¹¹While HP and Compaq merged in 2003, we chose to keep the brands separate in our analysis.

education level, marital status, and presence of children). The second stage consists of the technology survey, which asks a multitude of questions ranging from brand to year of purchase to where the computer is used.¹²

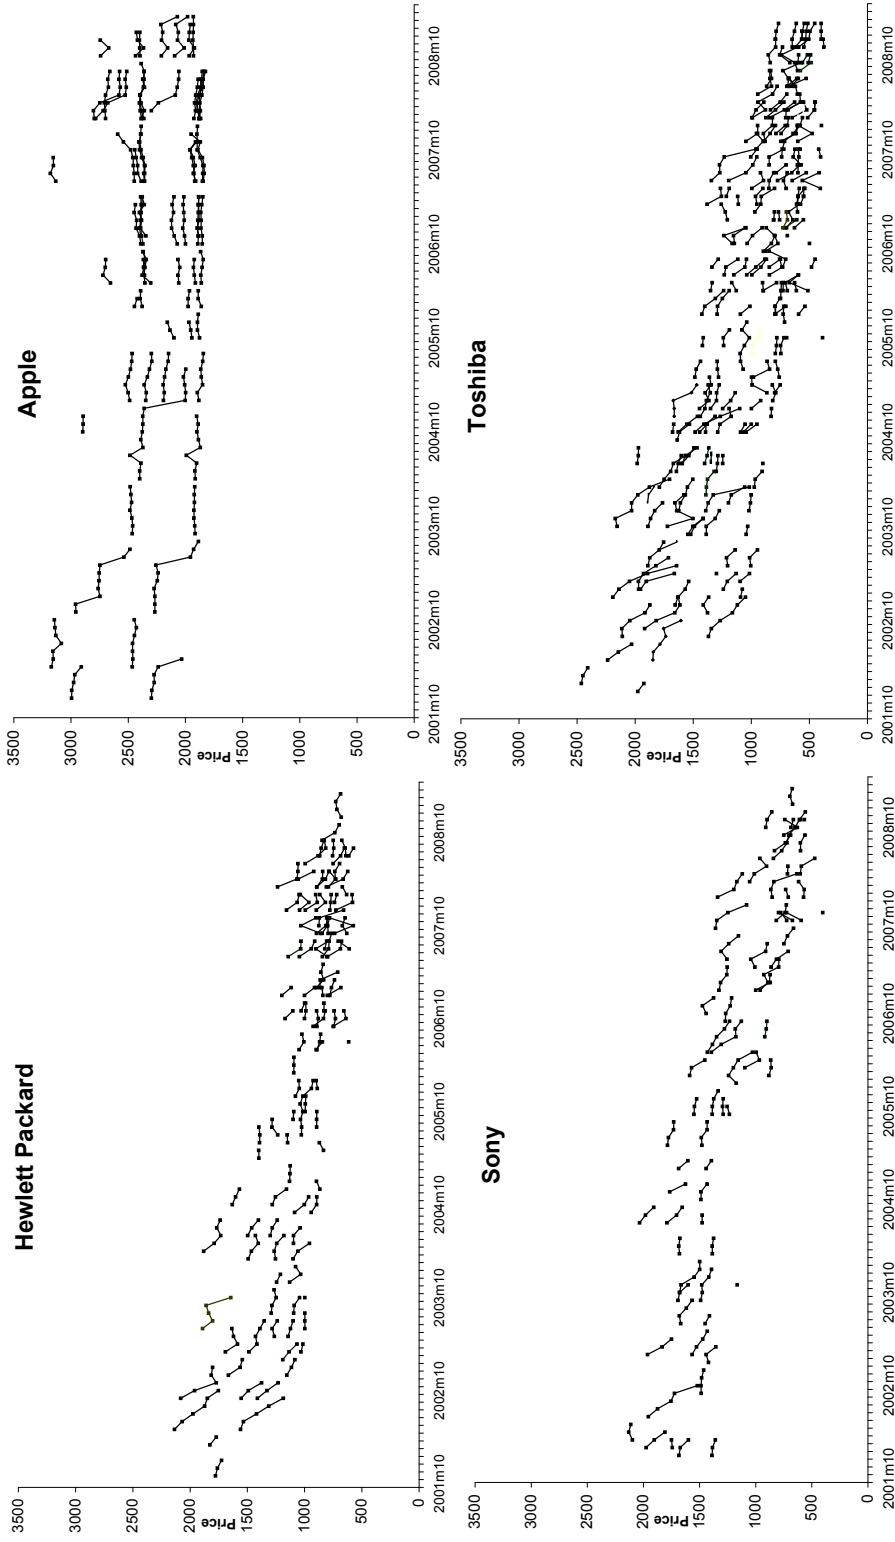
We use the NPD data to document descriptive statistics on price dynamics, product cycle length, and technology adoption. Figure 1 highlights many of the key aspects of these characteristics, where each point in the figure represents the unit price for a particular computer model in the sample of 15-inch notebook computers. The price time series for a given computer model is created by linking the model's prices over its life on the market.¹³ The three PC manufacturers (HP, Sony, and Toshiba) have short product cycles, frequent staggered entry, and declining prices over the life of the good. We show that these patterns are consistent with our entire data set.

The exception to these patterns are computers manufactured by Apple (see the upper right-hand corner of Figure 1). Apple products are characterized by long product cycles, less frequent and more uniform entry, and flatter price contours. Because of their unique operating system, Apple products are not close substitutes for other manufacturers' computers. Because HP, Sony, and Toshiba all offer Microsoft's Windows operating system and similar bundles of computer characteristics (e.g., Intel chips), we consider these products to be highly substitutable. For most of the analysis that follows, we focus on the personal computer market excluding Apple. We label this subset of the market PCs and note that over our sample period these computers account for 88 percent of all sales. As we explain later, however, the pricing and technology-adopting strategy pursued by Apple does help inform our analysis and provide an out-of-sample test for our model.

¹²All observations are reported on the user's "primary computer." An observation in this data consists of household demographics and computer specifications, including the price paid. We isolate observations where the PC is used at home, and we drop observations where the specification of the PC is not reported.

¹³Prices after the cumulative density function (CDF), in terms of units sold, reached 90 percent for each model were omitted in the analysis that follows, as these are generally stock-out sales. For ease of view, in Figure 1 we omitted depicting computer models with less than 20,000 total units sold for HP, 15,000 for Sony and Toshiba, and 4,000 for Apple.

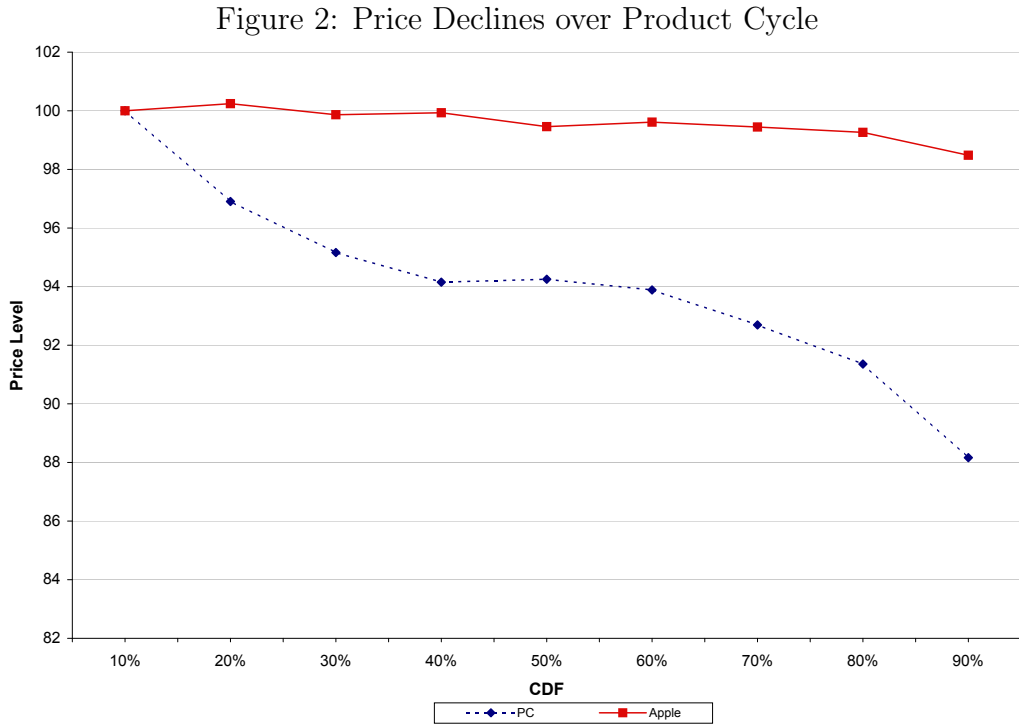
Figure 1: Prices: 15-Inch Notebook Computers



Notes: Depicted are the price contours of all 15 inch notebook computers sold by Hewlett Packard, Sony, Toshiba, and Apple computers over the course of the sample period. Prices after the sales CDF reached 90 percent for each model were omitted. Source: NPD Group.

3.1 Pricing Patterns

Figure 1 highlights some key features of price dynamics in the PC industry. Generally speaking, PC manufacturers introduce their products at a high price and then lower that price over the product cycle. We measure the rate at which prices fall over the life of the computer by estimating a fixed-effects regression of the logarithm of price on dummy variables representing deciles along the cumulative density function (CDF) of units sold. Depicted in Figure 2 are predicted values of the price level over the sales CDF, where we have normalized the price of the first decile to 100.¹⁴ It is clear from Figure 2 that PC prices fall quite rapidly over the product cycle. By mid-cycle, PC prices fall 6 percent, and by the end of the cycle, they fall by 12 percent. We performed separate regressions of each brand but, besides Apple, there were no significant differences between PC manufacturers. The figure shows that, strikingly, Apple maintains a flat price profile over its product cycle.



Notes: Depicted are the fitted values of a fixed effects (using model number as the fixed effect) regression of the logarithm of price on CDF decile dummy variables. Source: NPD Group.

¹⁴To calculate price changes over the product cycle seen in Figure 2, we run the following regression: $\ln P_{ik} = \alpha + \sum_{k=2}^9 \beta_k D_k + \gamma_i + \varepsilon_{ik}$ where $\ln P_{ik}$ is the logarithm of the price of model number i with CDF location k , D_k is a dummy variable indicating the location on the CDF, and γ_i is a model-number fixed effect. Specifically, $k = z$ indicates that product i lies between deciles $z - 1$ and z on product i 's CDF.

There are interesting dynamics between prices and product entry among PC manufacturers. In particular, PC manufacturers often leapfrog one another with the introduction of new, higher-quality computers. To display this feature in the data, in Figure 3 we isolate 512 MB RAM 15-inch notebooks where the entering PC happened to have the highest price in the product line.¹⁵ Due to the numerous innovative components, it is difficult to precisely assess the highest quality product in any given time period. This exercise attempts to isolate the computer models with both the newest and the highest-quality technology under the assumption that the computer with the highest quality is also the highest priced. Supporting our claim that newer products are of higher quality, we also report four computer characteristics that highlight in which dimension the newly introduced computer is of higher quality relative to existing computers. The manufacturer with the highest-quality 512 MB RAM 15-inch notebook rotates among HP, Compaq, Toshiba, and Sony. Introductory prices of these computers are quite high, around \$2,100, but then quickly fall to \$1,800.

Looking ahead, in Section 4 we develop a formal industry model of the personal computer industry. The model generates price declines over the product cycle through competitive effects, much like we observe for PC manufacturers. These price declines, along with decreasing sales, subsequently increase the incentives for adopting a new technology and ultimately drive the product off the market.

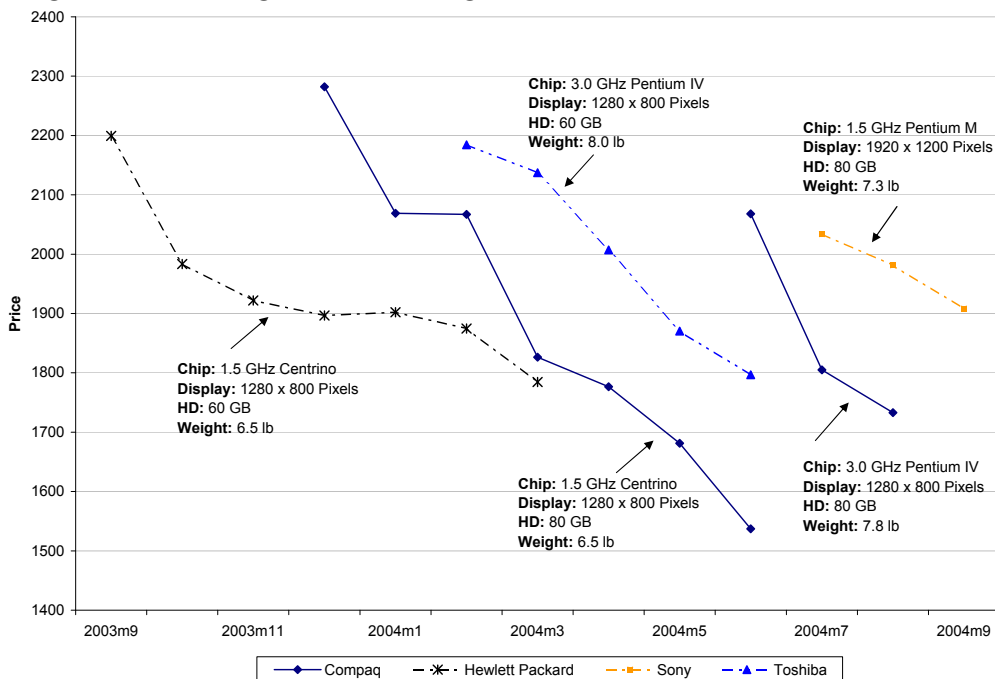
3.2 Product Cycle Length and Technology Adoption

To get a better sense of the timing of sales along the product cycle, we depict CDFs of units sold in Figure 4. Because computers do not necessarily enter the market at the beginning of the month, the first month of data will include anywhere between 1 day and 30 days worth of units sold. Thus we can create only upper and lower bands for the CDF, the lower band representing the case where the first month includes 30 days and the upper band representing the case where the first month is only one day.¹⁶ The CDFs demonstrate that PC manufacturers generally sell over half their units by the second month on the market, and that by the third month they have sold between 70 and 90 percent of their units. Apple, however, keeps its computers on the market about twice as long as the

¹⁵This line of computer represented 40 percent of all notebook units sold in our NPD sample during this time period.

¹⁶We measure the lower bound of the CDF by summing up total units by a variable indicating the age of the computer (in months). The upper bound is calculated in a similar manner; however, we assume that sales in month one are equal to total unit sales of computers one and two months old.

Figure 3: Entering PC is also Highest Priced: 15-Inch 512MB Notebooks



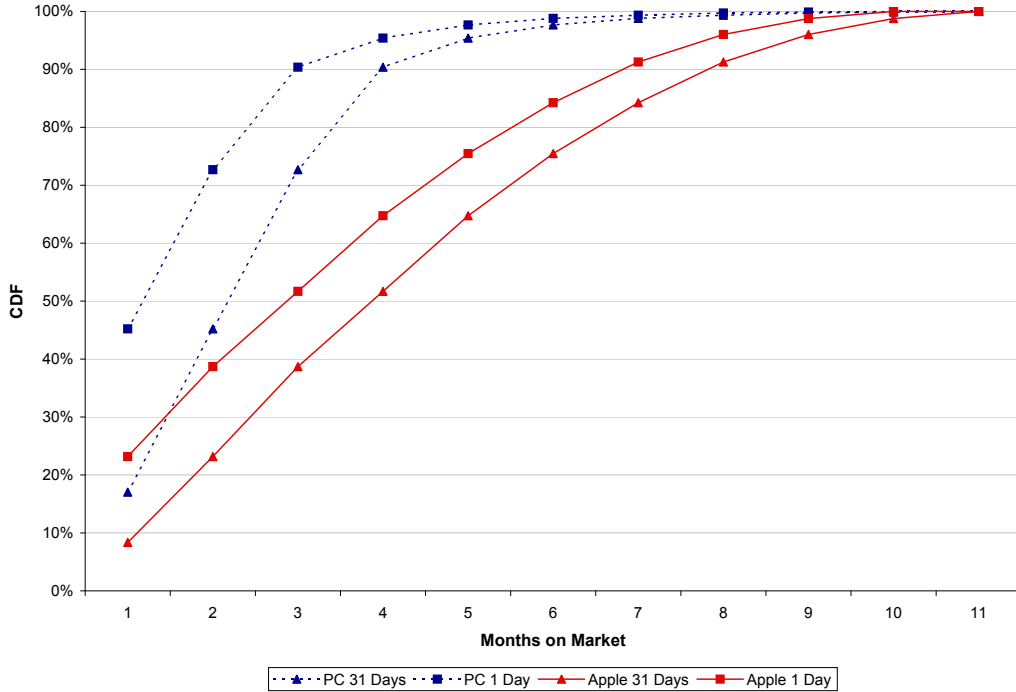
Notes: Computer models shown are models in which the entering PC happened to have the highest price in the category of 512 MB RAM 15-inch notebook computers. Prices after the units CDF reached 90 percent for each model were omitted. Source: NPD Group.

other PC manufacturers.

The fact that a firm holds its computers on the market for a longer period of time does not necessarily mean that it introduces new computers less frequently. For instance, a manufacturer could very well stagger the introduction of its new computers to release a new computer every period. Our data, however, show that this is not the case. Table 2 reports the fraction of months in the sample where no new computer was introduced. In effect, this table reports measures of the fraction of months in which the manufacturer’s entire product space is composed of computers that are at least one month old. Of the seven manufacturers, Apple has the largest proportion of months in which no new models are introduced (28 percent); by contrast, HP had only one month in the sample with no introduction of a new model.

Table 2 also depicts the maximum amount of time the manufacturer goes without introducing a new model. These numbers also show that Apple is relatively slow to introduce new computers. For instance, Apple underwent a period of nine months in which it did not introduce a new desktop computer and a period of six months without introducing a

Figure 4: CDF of Units Sold



Notes: The upper band is the estimated CDF under the assumption that the first month represents 1 day of sales, while the lower band is the estimated CDF under the assumption that the first month represents 30 days of sales. See Appendix A for details. Source: NPD Group.

new notebook computer—by far the longest periods in the sample.¹⁷

To gauge how frequently manufacturers adopt new CPUs, we plot the age of the newest

¹⁷The large number of components, as well as their complexity, makes it a nontrivial task to monitor and measure their adoption by computer manufacturers. For computer firms, however, we argue that new computer models (i.e., SKUs) usually incorporate an upstream innovation. Consequently, the rate at which a computer manufacturer adopts new computers is the rate at which the manufacturer is adopting new technologies and embedding them into its products. Our reasoning for equating product entry with technology adoption is based on the production technology for computers. Computers have many internal components that are produced by a diverse array of distinct upstream firms. Upstream firms undertake R&D in an attempt to increase the quality of the components they sell to the downstream computer manufacturers. Computer manufacturers, consequently, have ample opportunity to adopt new technologies when introducing a new computer to the retail market. For instance, one month Intel may introduce a new CPU, while the following month Samsung may introduce a new dynamic random access memory (DRAM) chip. While the assumption that the introduction of a new model equates to the adoption of a new technology could be flawed if, for instance, CPU manufacturers are frequently crimping their products, we believe that it is realistic to assume that newer computers generally embody more innovative, higher-quality components.

Table 2: Adoption of New Models

	Fraction of Months With No Model Adoption			Maximum Time Period Between Model Adoptions		
	All	Desktops	Notebooks	All	Desktops	Notebooks
Hewlett Packard	0.01	0.03	0.06	1	1	2
Compaq	0.09	0.22	0.11	2	3	2
Toshiba	0.12	-	0.12	2	-	2
Apple	0.28	0.53	0.56	4	9	6
Emachines	0.20	0.20	-	4	4	-
Gateway	0.02	0.18	0.11	1	1	2
Sony	0.14	0.41	0.19	3	6	3

Notes: This table reports the fraction of months in the sample period when no new model (columns 1 to 3) was introduced as well as the maximum time period (in months) for which no new model was introduced (columns 4 to 6). The sample is taken over all computers (columns 1 and 4), desktop computers (columns 2 and 5) and notebook computers (columns 3 and 6). Source: NPD Group.

Intel CPU by month for the post-PowerPC period for Hewlett Packard, Toshiba, and Apple notebook computers in Figure 5.¹⁸ Two features of this figure are striking. First, Toshiba and Hewlett Packard are twice as often the first to adopt a new CPU (12 and 14 months out of 35, respectively) as Apple (7 out of 35 months). Second, Hewlett Packard and Toshiba almost never exceed three months to adopt a new Intel CPU. By contrast, Apple’s newest CPU available was seven months old on three occasions.¹⁹ In comparison to PC manufacturers, Apple’s strategy is to adopt technology less frequently but with larger jumps in quality.


3.3 Demographics

In addition to the firm side, there are important features of the personal computer industry on the consumer side. Using the TUP survey data we highlight some facts about the income distribution of consumers who purchase PCs. We focus on consumer income because it is typically closely linked to reservation price, and therefore product choice, in

¹⁸The age of the CPU was calculated by subtracting the current time period from the period in which the chip first appears in our sample. Apple switched from Motorola/IBM PowerPC chips to Intel chips in June 2006. We depict notebook computers in the figure because Apple’s desktops use Intel Xeon processors, which cannot be differentiated by processor name in the data.

¹⁹Table C in the Appendix, shows CPU adoption statistics for the entire sample of notebook computers and shows that, on average, Apple offers the oldest CPU for this sample period.

Figure 5: Adoption of Intel CPUs



Notes: Depicted is the age of the newest Intel CPU for each month of the post PowerPC CPU period (i.e. 2006m6 to 2009m4) by computer manufacturer.

most econometric studies and economic models. The survey data reveal that both the levels and the distributions of income differ across brands in the industry. Furthermore, we also document that income is correlated with the price paid, holding fixed the characteristics of the computer.

There are large differences in the income distribution by computer manufacturers. Table 3 highlights these differences by showing the median income and dispersion of income (represented by the Gini coefficient) for each brand in the TUP survey data.²⁰ The survey data show that Apple has the highest median income, followed by Sony, Dell, and IBM. Important to our study, consumers of Apple have a very narrow distribution (0.195 Gini coefficient) relative to the other manufacturers.²¹ This means that although Sony, for

²⁰The Gini coefficient represents twice the expected absolute difference between two individuals' income drawn randomly from the population. Thus, the larger the Gini-coefficient, the wider the degree of dispersion.

²¹We note that these dispersion statistics are somewhat prone to measurement error due to the placement of income levels into bins. Each income level represents the midpoint of the bin, except for the last bin, which is \$150,000 and greater. Therefore, if a large proportion of Apple's consumers have incomes much greater than \$150,000, the Gini coefficient on Apple could realistically be somewhat larger than what we

instance, sells to a high median consumer, it also sells to low-income consumers. This is attributable to the declining pricing pattern of Sony’s computers whereby low-income consumers purchase computers that have been on the market a few months.

Table 3: Consumer Income Dispersion and Levels

	Gini Coefficient	Median Income
Apple	0.20	65000
Compaq	0.31	42500
Emachines	0.30	42500
Hewlett Packard	0.30	42500
Sony	0.26	55000
Gateway	0.29	47500
Dell	0.28	55000
IBM	0.32	55000

Notes: This table shows the median consumer income and Gini coefficient of consumer income for each manufacturer. Source: TUP survey data.

The stylized fact that low-income consumers purchase computers later in the product cycle is formally shown in Aizcorbe and Shapiro (2010). Using the same TUP survey data that we use, Aizcorbe and Shapiro find that higher-income consumers pay a higher price for the same computer than do lower-income consumers. Specifically, a regression is run of income and other demographic variables on the logarithm of price, holding fixed the attributes of the computer purchased. The study finds that the coefficient on income is .09, indicating that a 10 percent fall in a consumer’s income is correlated with a 0.9 percent fall in the price paid for a given computer. Combining these results with the price declines observed for PCs shows that high-income consumers are presumably purchasing early in the model’s life cycle.²² Naturally, because Apple’s prices remain flat over the computer’s product cycle, there is no correlation between price and income.

It is not obvious why the income distributions between the two markets differ. It could very well be that Apple targets high-income consumers, while PC manufacturers target an array of consumer types. The model we develop in the next section, however, suggests this

measure.

²²Interactions between brand and income also verify that the correlation between income and price in the TUP data is stemming from those brands with large price declines in the NPD data.

is not the case. The model posits that competitive forces lower prices of a computer over the product cycle, drawing in lower-income consumers to purchase the product.

4 Model of a Competitive Industry

We model the computer industry using an infinite-period vintage-capital model. Computers are differentiated by their vintage ν , where ν equals the date at which a vintage is the frontier technology; at time t , the frontier technology is $\nu = t$. There is an outside option, which provides utility, \check{u}_t , to a consumer in period t . The utility of the outside option increases over time at an exogenous rate.

Because our analysis is over the short run (the lifetime of specific product), we fix the number of firms in the market to be N . Further, we simplify the problem by assuming that each firm produces at most one computer and so ignore any joint maximization problem of a multiple product-line firm. Thus, we can think of the model as characterizing firms competing with one another over vertical quality within a specific “product line,” such as the 15-inch laptop computers depicted in Figure 3.

Innovations arrive exogenously every period in the form of higher quality intermediate inputs. As described more formally below, computer manufacturers can upgrade their computers by deciding to pay a fixed cost and incorporate into their products the latest, most innovative inputs (e.g., lighter batteries, higher resolution screens, or better chips).

4.1 Demand

Each period, a mass M of consumers enters the market. Consumers have a budget to purchase a computer and related products. Consumers are differentiated by the size of their budget, denoted y , which is drawn from a distribution h . Given their budget, consumers either buy one computer and use the remainder of their income on the outside option, or just choose the outside option, where the outside option is an alternative computer-related product. In either case, consumers leave the market at the end of the period, so that there is no accumulation of consumers across periods. We normalize the price of this alternative computer-related good (that is, the numeraire commodity) to one. Let $p_{\nu t}$ denote the time t price of vintage ν .

Following Shaked and Sutton (1982, 1983) we assume that the consumer’s utility from purchasing the computer of vintage ν is

$$U(y, \nu; \bar{p}_t) = u_\nu \cdot (y - p_{\nu t}), \tag{1}$$

where \bar{p}_t is a vector of prices and u_ν represents the quality of a computer of vintage ν . We make the natural assumption that newer vintages are preferred to older ones, and thus $u_t > u_{t-1} \forall t$. The utility from just purchasing the outside good is

$$\check{u}_t \cdot y. \quad (2)$$

Given prices, the consumer's utility-maximization problem is:

$$\max \left\{ \max_{\nu \in \bar{\nu}_t} U(y, \nu; \bar{p}_t), \check{u}_t \cdot y \right\}, \quad (3)$$

where $\bar{\nu}_t$ denotes the set of available computers in period t . The resulting demand function is straightforward.²³ We order the vintages by their utility levels and consider the neighboring vintages ν_k and ν_j where $u_k < u_j$. Given that the lower-quality computer has a lower price, $p_{\nu_k} < p_{\nu_j}$, there is a marginal consumer with income \hat{y} who is indifferent between them:

$$u_{\nu_k} \cdot (\hat{y} - p_{\nu_k, t}) = u_{\nu_j} \cdot (\hat{y} - p_{\nu_j, t}). \quad (4)$$

For this marginal consumer, the utility gained from having the higher quality of computer j relative to computer k is exactly offset by the price difference.

All consumers with income less than \hat{y} prefer ν_k over ν_j and all those with income more than \hat{y} prefer ν_j over ν_k ; denote this marginal consumer y_{ν_k, ν_j} . Repeating this exercise across all pairs of neighboring vintages, we can define a set of marginal consumers from which demand for each computer vintage can be computed. Consumers between the marginal consumers $(y_{\nu_l, \nu_k}, y_{\nu_k, \nu_j})$ will purchase vintage ν_k . The demand for ν_k is then simply

$$Q_{\nu_k} = \int_{y_{\nu_l, \nu_k}}^{y_{\nu_k, \nu_j}} h(x) dx,$$

where h is the distribution of consumers' income. Given the ordering of vintages, ν_1 is the best available product. Its demand is given by

$$Q_{\nu_1} = \int_{y_{\nu_2, \nu_1}}^{\infty} h(x) dx.$$

Similarly, let ν_N be the lowest-quality product. It competes directly with the outside option, and its demand is given by

$$Q_{\nu_N} = \int_{y_{\bar{u}, \nu_N}}^{y_{\nu_N, \nu_{N-1}}} h(x) dx,$$

²³This is the demand system of Prescott and Visscher (1977), which has been well studied in the vertical product differentiation literature.

where $y_{\check{u}, \nu_N}$ solves

$$u_{\nu_N} \cdot (y - p_{\nu,t}) = \check{u}_t \cdot y.$$

4.2 Supply

A firm makes two decisions at the beginning of each period. First, the firm decides whether to adopt a new technology (that is, upgrade its product). If the firm adopts, it pays a fixed cost $\phi > 0$ and upgrades its computer so that the computer embodies the latest technology. Otherwise, the firm continues to sell its current computer. Letting $i = 1, 2, \dots, N$ denote a firm, we label the decision to adopt the latest technology as $d_{it} \in \{0, 1\}$, where $d = 1$ signifies adoption. Second, the firm sets a price for its computer. We assume firms have a constant marginal cost $c \geq 0$ and no capacity constraints. The state variables are $s_t = (\nu_1, \nu_2, \dots, \nu_N, \check{u}_t)$, which consist of all the firms' products and the outside option. Let $\delta = 0.99$ denote the discount rate, then firm i 's profit-maximizing problem is:

$$V_i(s_t) = \max_{p_{\nu_{it}}, d_{it}} \left\{ (1 - d_{it}) E_{s'} \left[(p_{\nu_{it}} - c) Q_{\nu_{it}}(p_{\nu_{it}}, p_{\nu_{-it}}; s') + \delta V_i(s') \right] + d_{it} E_{s''} \left[(p_{\nu'_{it}} - c) Q_{\nu'_{it}}(p_{\nu'_{it}}, p_{\nu_{-it}}; s'') - \phi + \delta V_i(s'') \right] \right\}, \quad (5)$$

where $\nu' = t$ is the latest (and highest-quality) vintage and $p_{\nu_{-it}}$ denotes all other firms' prices in time t , given the state variable. The expectations are taken over other firms' updating decisions, where s' denotes the case where firm i does not update its product, while s'' is the case where it does. $Q_{\nu_{it}}$ is the demand for product ν_i at time t , given prices and the outside option.

While the firm's price-setting decision is static, its adopting decision is dynamic. Because consumers value quality, updating to the latest technology generates higher revenues for the firm, holding all else constant. As the firm pays a fixed cost ϕ to acquire the latest technology, it must balance the gains to adopting in the current period against the option value of continuing to sell its computer and upgrading in the future.

Rather than physically depreciating, a computer faces two sources of obsolescence over time. First, the outside product is assumed to improve over time, while, in each successive period, a computer with vintage ν maintains the same utility value to consumers. This *general* obsolescence places downward pressure on prices of existing computers. Second, with each successive period other firms may update their computers. Newer vintages,

embodying better technologies, directly compete with a vintage ν and drive down its price. We label this second source of obsolescence *market-specific* obsolescence.

Either source of obsolescence ensures that a computer is sold for a finite number of periods. After some point, the demand for a product when priced at marginal cost will equal zero, and the computer will have effectively exited the market. Of course, a firm may decide to upgrade its computer before demand reaches zero. The life cycle of a computer, then, starts with its introduction into the market and ends when either the firm upgrades or there is no longer demand for the computer at a price weakly greater than marginal cost.

4.3 Equilibrium

We use a Markov perfect equilibrium concept where the strategy space includes setting the price and the decision to adopt the latest available technology. Firms' actions are functions of the current vintages of computers offered, along with the utility value of the outside option. As described by equation (5), firms maximize the expected discounted value of profits, conditional on their expectations of the evolution of the state variables and competing firms' strategies. Equilibrium occurs when all firms' expectations are consistent with the evolution of both the outside good's utility and the optimal pricing and adopting policies of their competitors.

To keep the analysis tractable, we consider a stationary Markov perfect equilibrium and rule out mix strategies. The model will be stationary in the sense that prices and sales of a computer over its product cycle will be independent of time. To obtain a stationary equilibrium, we make an additional assumption: the ratio of the utility associated with a computer embodying the frontier technology over the utility provided by the outside good remains constant over time. Formally,

$$\frac{\check{u}_t}{\nu_t} = \zeta \quad \forall t, \quad (6)$$

where $\zeta \in (0, 1)$.

A stationary equilibrium occurs when all firms use the following strategy: the firm with the lowest-quality computer upgrades its computer to the frontier technology and all other firms do not upgrade their products. For pricing, each firm's strategy is to use its best-response function. From Shaked and Sutton (1983), we know that, given a set of products, there exists a Nash equilibrium in prices.

Under this strategy, one of the N firms adopts in a given period. If any firm deviates from this strategy and upgrades its computer when it is not the lowest-quality vintage, then

both the deviating and lowest-quality firm will simultaneously upgrade their computers and sell the same quality product. Because firms compete in price, the Nash-equilibrium price for both firms' products will be equal to marginal cost. Thus, any deviation in this strategy will cause a firm to pay a fixed cost in exchange for earning revenues equal to marginal cost, which will result in negative profits.

An equilibrium exists only if marginal cost is below the maximum price consumers with the largest budgets will pay. Further, the fixed cost of updating must be less than the net present value of profits over a computer's product cycle. Finally, the number of active firms in equilibrium will be determined by profit conditions. If N firms manufacture computers in a stationary equilibrium, then the net present value of profits of each firm must be non-negative. Further, the net present value of profits given $N + 1$ firms is negative.

4.3.1 Discussion of the Equilibrium Leap frogging Result

An outcome of this stationary equilibrium is that manufacturers leap frog one another systematically. Although this may seem a bit stylized, the stationary equilibrium concept will be necessary for our empirical approach in Section 5. To be clear, we have no definitive evidence that manufacturers always coordinate their adoption decisions by taking turns. Indeed, in some cases we see different manufacturers adopting the same technology during the same month. For example, in many cases PC manufacturers adopt the newest Intel CPU in the same month. Nevertheless, the leap frogging assumption seems plausible for at least two reasons.

First, there is evidence that PC manufacturers vertically differentiate themselves, implying that manufacturers do in fact align their products along a quality ladder. Specifically, we found it very difficult to find two or more computer manufacturers offering computers with *exactly* the same observable specifications. For example, less than 6 percent of all notebook units sold in our sample had the same observable specifications (i.e., CPU, display size, hard drive size, pixel ratio, memory, DVD format, weight) as another manufacturer's product sold in a given month. Given that many of these components likely add dimensions of vertical quality, and that the highest-quality technology is also usually the newest, adoption will necessarily push down the relative quality of an existing product along the quality ladder.

Second, Table 2 shows that new models are being introduced very frequently. If we narrow the product category to 15-inch notebooks, something we think is closer to a product line, new models are introduced about every three to four months. This implies manufacturers are not introducing new models into the same product line every period,

but waiting and perhaps staggering as implied by Figure 3.

5 Empirical work

In this section, we describe the calibration exercise. We demonstrate that the model fits the data well, and then use the model to understand the intuition behind the price and sales dynamics.

5.1 Calibrating the Model

The parameters of the model can be categorized into three groups. The first set of parameters determines the quality level of computers relative to each other and the outside good; the second set characterizes the consumers' budget distribution; and the third set details the cost structure of the firm.

We use three parameters to characterize computer quality: (1) the level of the highest-quality product; (2) the monthly growth rate of the frontier technology, γ ; and (3) the ratio of the outside good's utility to the highest-quality product's utility, ζ . We fix the utility level of the highest-quality technology to 10. We then set the monthly growth rate of computer quality, γ , to 2.9 percent, based on the assumption that the exogenous upstream technological progress follows Moore's law. The substitutability of products across vintages is determined by this growth rate. Raising γ increases the difference in quality between a newer and an older vintage, leading to a decrease in the substitutability across vintages. Finally, using data from the TUP surveys described earlier, we calculated the ratio of the outside good's utility to the highest-quality technology to be 0.033, making the outside option a fairly poor substitute for a new computer.²⁴ As we show later, however, our results are robust to significantly higher values of the outside good's utility (see section 5.3).

Turning to the consumers' budgets for computers and related products, we assume this variable has the Beta distribution over the interval $[a, b]$. We fix a to be one and leave b free. The density of consumers over $[a, b]$ is given by the two parameters that characterize the Beta distribution, (κ_1, κ_2) . In calibrating the model, the size of b plays a significant role in the length of the product cycle, while the parameters (κ_1, κ_2) affect the relative prices among vintages.²⁵

²⁴See appendix A for details on how this number is calculated.

²⁵In Shaked and Sutton (1983) the distribution of consumer tastes, the equivalent of our distribution of consumer budgets, is assumed to be uniform. With a uniform distribution of budgets, the model much less closely matches the data. In particular, a computer is sold for five periods, and model prices decline

Finally, we assume a simple cost structure for the firm. The firm pays a fixed cost, $\phi > 0$, to adopt a new technology, while the firm’s current product can be produced at a constant marginal cost, $mc > 0$. The value of ϕ is not pinned down when we take the model to the data. Rather, equilibrium conditions impose an upper bound to the value of ϕ , whereby ϕ must be less than a product’s present discounted value of profits.

In matching the model to the data, we need to choose the length of a period in the model. Given the wide array of inputs used in computers along with the substantial innovation in computer hardware, we believe that it is reasonable to assume that computer manufacturers have the option of incorporating new technology into their products every month. Consequently, we set a period in the model to correspond to a month. In a stationary equilibrium, then, every month the computer manufacturer with the lowest-quality product upgrades its product.

Our model, then, has four free parameters, $\theta = \{b, \kappa_1, \kappa_2, mc\}$. Given these parameters and considering a stationary equilibrium, the model predicts the number of months a computer is sold; the paths of prices, sales, and markups over a computer’s product cycle; and the relationship between the size of a consumer’s computer budget and the computer vintage purchased. Our target moments are (1) the number of months a computer is sold; (2) the computer’s sales path over the product cycle; and (3) the computer’s price path over the product cycle.

Matching the model’s predictions of the prices and sales over the product cycle to the data is not straightforward. As discussed in Section 3.2, we are not certain how many days are included in the first “month” of data because computers could be introduced any day of a month. Because the first 30 days a computer is sold accounts for a substantial portion of a computer’s product cycle, we cannot drop or otherwise ignore this measurement problem.

We address this data issue for sales by constructing bounds on the sales CDF of a typical PC computer over the product cycle (see Section 3.2 and Figure 4). The sales CDF informs us both about the number of months a computer is sold as well as the path of sales over the product cycle (the first two target moments mentioned earlier). We look at whether the model’s predicted sales CDF is within the bounds inferred from the data, and, if not, by how much it deviates, to guide our choice of θ .

For price declines, we address the data issue by relying on the predictions of a fixed-effects regression of price on dummy variables representing deciles along the sales CDF, as described in Section 3.1. This regression provides an estimate of how much prices have declined from the initial price to a given point on the sales CDF. For example, suppose the

significantly faster than we observe in the data.

model’s sales CDF is equal to 0.6 in the second month. From the fixed-effects regression, we know the relevant price decline is -6.1 percent (see Figure 2). We use the distance between the price declines given by the model and those inferred from the data to guide our choice of θ . Note that we do not attempt to match sales or price levels, but focus instead on the dynamics of sales and prices over the product cycle. Details on exactly how we minimize the difference between the model’s predictions and the data are provided in Appendix B.

Table 4: Model Parameters

Product quality			
Frontier value	\bar{v}	fixed	10
Quality growth rate	γ	fixed	1.029
Utility ratio	ζ	fixed	0.033
Income (Beta distribution)			
Lower bound	a	fixed	1
Upper bound	b	flexible	9.949
Density	κ_1	flexible	0.697
Density	κ_2	flexible	1.718
Cost			
marginal cost	mc	flexible	0.873

We find there is a wide distribution in budgets for computers and related products (see Table 4). Those consumers with the largest budgets are willing to spend more than 10 times as much as those consumers with the smallest budgets. Further, the density of consumers is characterized by the Beta distribution with parameters (0.697,1.718). This particular distribution is decreasing in consumers’ budgets and looks like a Pareto distribution. Finally, the marginal cost parameter is 0.87, below the lower bound on a consumer’s budget.

5.2 Comparing the Model to the Data

As displayed in Table 5, our model fits the data on PCs well along both price and sales dimensions. Looking first at sales, the model generates a four month product cycle, in line with the data. The model’s sales CDF falls within the CDF bounds estimated from the

data for all four periods, although admittedly the sales CDF bounds are fairly far apart (see line labeled “competition model” in Figure 6). We can further test the fit of the model by looking at the sales probability density function (PDF). Under the reasonable assumption that true PC sales decrease over time, we construct upper and lower bounds on the sales PDF, which are tighter than those implied by our CDF bounds.²⁶ Our model also performs well along the sales PDF dimension (see Figure 7). Specifically, it captures well the burst of sales when a computer is first introduced, followed by a decline in the sales rate over time. Overall, then, the model does a nice job of matching the sales data along both the CDF and the PDF dimensions.

Table 5: Time-series of price and sales, data and models
Percent price decline, relative to first month

<i>Month</i>	1-2	1-3	1-4	average
Data	-6.9	-8.6	-9.1	-8.2
Model	-6.9	-8.6	-9.1	-8.2

Sales CDF


<i>Month</i>	1	2	3	4
Data–lower bound	0.217	0.600	0.898	1
Model	0.464	0.815	0.969	1
Data–upper bound	0.600	0.898	1	1

Turning to prices, we find the model exactly matches the price declines seen in the data (see Table 5). In particular, the model is able to capture the large initial price decline between months 1 and 2, followed by smaller price declines in months 3 and 4. We provide a visual display of the model’s fit to the data in Figure 8, where we plot the price declines for the competitive model (dashed line) and data (circles) given the associated sales CDF value.

To check the model’s fit to the data, we consider the markups and the timing of consumer purchases implied by the model. The markups implied by the model are reasonable. When a computer is the newest vintage available, manufacturers charge a markup of 10

²⁶We can construct bounds on the sales PDF using the sales CDF. The implied minimum sales PDF, however, is always equal to zero and the maximum sales PDF is quite large, making these PDF bounds uninformative. Assuming that true PC sales are decreasing over time allows us to use sales in months t and $t + 1$ as upper and lower bounds, respectively, for sales in month t . For the first month, we use the sum of sales in months 1 and 2 as an upper bound and the observed level of sales as a lower bound.

Figure 6: Model versus Data: CDF Sales over the Product Cycle



Notes: Monopolist model is discussed in Section 6

percent. Once a newer vintage enters the market, however, this markup plummets. The second-highest quality computer has a markup of almost 3 percent, while the third- and fourth-highest quality computers have markups of 0.8 and 0.3 percent respectively.

We next compare the model's predictions on the timing of a household's purchase decision, conditional on its budget. As described in Section 3, the TUP survey data imply that households with higher incomes purchase computers earlier in the product cycle. Aizcorbe and Shapiro (2010) find that a 0.9 percent fall in price is correlated with a 10 percent fall in income. Using the change in price observed in the data, the correlation from Aizcorbe and Shapiro (2010) suggests that incomes of those consumers who buy the highest-quality computer should be almost twice as big as those who buy the lowest-quality computer.

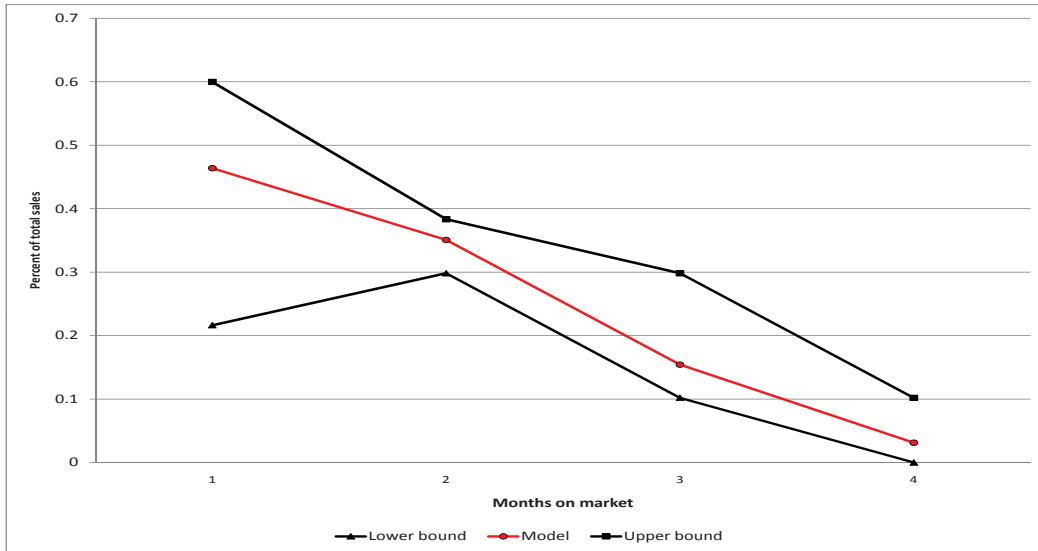
Table 6: Average Budget of Consumers over the Product Cycle

Month	1	2	3	4
Budget	5.54	2.26	1.22	1.01

Notes: Budget is in units of the outside good

The model predicts a larger difference in consumer budgets. As shown in Table 6, the

Figure 7: Model versus Data: PDF Sales over the Product Cycle




Notes: The upper and lower bounds of sales have been constructed under the assumption that true sales are declining over time.

average budget of consumers who purchase the computer when it is first introduced is predicted to be more than five times as large as the average budget of those who purchase the computer in the last month of its product cycle. Household income, however, is obviously different from a consumer's budget for computers and related products. It seems plausible that the relative range in budgets is greater than the range of household incomes. Hence, we argue that the model's predictions of the correlation between price and budget are consistent with the results in Aizcorbe and Shapiro (2010).

5.3 Analysis

As calibrated, the model rationalizes the declining price and sales paths over the product cycle through the mix of competitive pressures and rapid innovation. Because consumers value quality, there are large returns to incorporating the latest innovation into a computer and selling the highest-quality product on the market. Due to the competitive nature of the market, however, a computer can only temporarily maintain highest-quality status. With newer and better competing computers being introduced every month, a computer manufacturer has to quickly drop the price of its products to remain competi-

Figure 8: Model versus Data: Price Declines over the Product Cycle



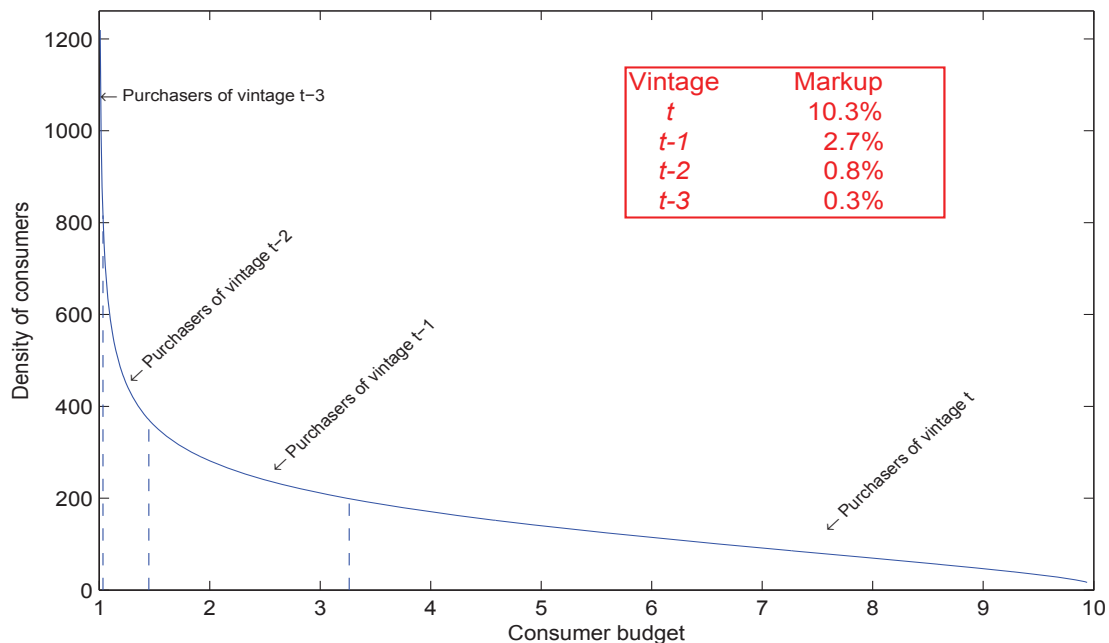
Note: For the data, we plot polynomials that were fitted to the estimated coefficients from the fixed-effects price regression. Monopolist model is discussed in Section 6.

tive.

The distribution of consumers over the product cycle can be seen graphically in Figure 9. The curved solid line represents the mass of consumers over the range of possible budgets as given by $\beta(0.697, 1.718)$. The vertical lines denote the marginal consumer between two different vintages. When offering the highest-quality computer, a manufacturer can generate a large amount of sales while pricing 10 percent above marginal cost. Referring back to the PDF of sales over the product cycle, the model predicts that 46 percent of a computer's total sales are completed in the first month (see Figure 7). Being replaced as the highest-quality computer reduces both a manufacturer's sales and its markup (from 10.3 to 2.7 percent). Market-specific obsolescence forces the manufacturer to appeal to consumers with lower budgets for computers, driving the manufacturer to slash its markup by three quarters.

The combination of rapid obsolescence and large gains from being the highest-quality computer lead the manufacturer to upgrade its computer on a fairly frequent basis. A main driver of a firm's adoption decision is the size of the difference between post- and pre-innovation rents. In the spirit of Arrow (1962) and Aghion et al. (2001, 2005), then, computer manufacturers incorporate innovative intermediate inputs into their products to

Figure 9: Distribution of Consumers over the Product Cycle



Notes: Displayed is the market equilibrium given the parameters reported in Table (4).

distinguish themselves from the competition.

Finally, we consider the role of the outside option. Given the intense competition among manufacturers, computers face a large amount of market-specific obsolescence. Consequently, the impact of the outside option as the driver of general obsolescence is minimal. Indeed, it is not until the utility value of the outside option exceeds 2.0 (from its current value of 0.33), that the outside option affects the model's results. In this case, the relative attractiveness of the outside option draws consumers with the lowest budgets. Consequently, the length of a computer's product cycle is shortened.

6 Alternative Scenarios

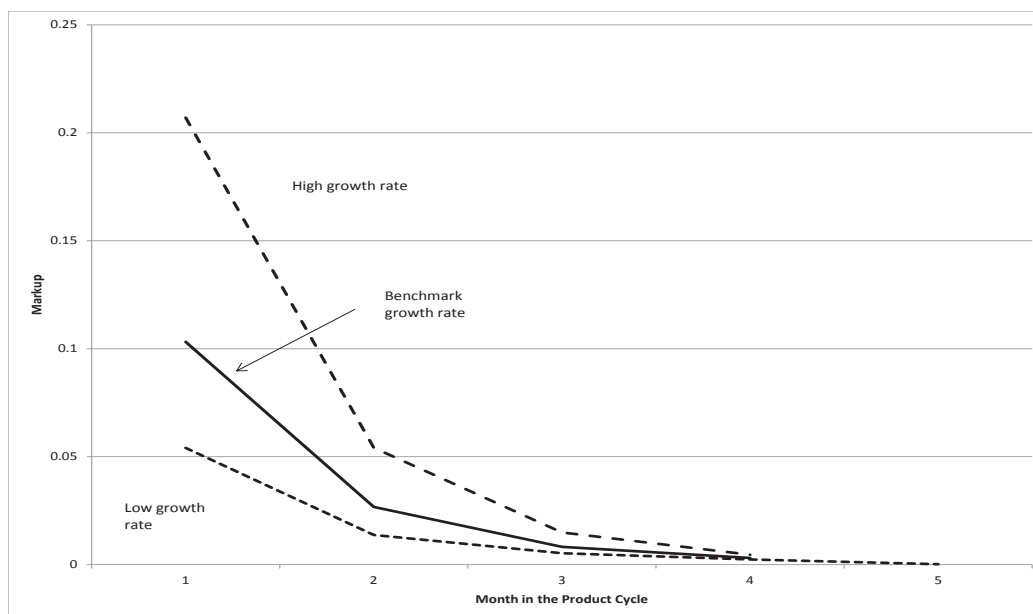
To assess the roles of innovation and competition in generating the model's fit, we assess the price and sales dynamics of the model under alternative scenarios. First, we perform a counterfactual analysis by altering the rate of innovation growth in the competitive setting. Second, we consider the case when the cost of entering the market is high enough that only one manufacturer profitably enters the market. This case serves as a counterfactual highlighting competition's effect on pricing. In addition, we use this case as an out-of-

sample test, by comparing the model's predictions to what we observe for Apple computers.

6.1 Altering the rate of upstream innovation

We explore how different growth rates in product quality affect the model's predictions. We start by doubling the growth rate to 6 percent. This higher rate increases the degree of differentiation between vintages of a computer. This implies an even greater return to being the highest-quality product and increases the role of market-specific obsolescence, leading to a faster decline in price. In Figure 10 we plot the model's predictions of markups for this high growth rate case against the benchmark case of a 2.9 percent growth rate. As highlighted in the figure, the initial markup is an enormous 21 percent, roughly double the value in the benchmark case.

Figure 10: Price Declines with Different Rates of Innovation



In Figure 10 we also plot the model's prediction of price over the product cycle given a slower growth rate of 1.5 percent. Computers are less differentiated in this case, leading to lower markups throughout the product cycle. Furthermore, given the weaker force of market-specific obsolescence, the product cycle is lengthened to five months. This is because, unlike in the benchmark case, five month-old computers can generate sales at a

price above marginal cost.²⁷

6.2 Altering the competitive environment

To explore how the competitive environment affects the model’s prediction, we consider altering the competitive environment to the case of monopoly. Recall that the number of firms is endogenous in our model. However, by assuming that entry costs are large enough, we can impose that only one firm can enter and earn positive profits. We use this monopoly case to highlight the role of competition in driving price declines over the product cycle.

We then use this monopoly case as an out-of-sample test of the model’s validity. As mentioned previously, we excluded Apple from our calibration exercise because its products, with their unique operating system, substantially differ from other PCs. From the perspective of our model, Apple is a monopolist—a manufacturer that does not face product-specific obsolescence. As an out-of-sample exercise, then, we evaluate how well the model’s predictions of prices and sales over the product cycle under a monopolist match the observed time series for Apple computers.

We first describe formally the monopolist’s problem and then present the results. The monopolist’s problem is equivalent to the competitive case outlined above; however, it does not face competing product introductions by other firms. Hence, the monopolist is concerned only about general obsolescence stemming from growth in quality of the outside good. Consistent with the competitive case, we assume that the monopolist sells only one product at a time.

6.2.1 The Monopolist’s Problem

The timing of the monopolist problem is the same as the competitive case. At the beginning of the period, the monopolist chooses whether to upgrade its product and then sets the price. The state space of the monopolist in period t is its existing product ν_s and

²⁷We also examined how expanding or contracting the distribution of consumer income affects pricing. We increase b , the highest level of income, while keeping the shape of the income density distribution fixed. Changing b effectively makes consumers more heterogeneous. Increasing b to $b = 20$, the producer of the highest-quality computer significantly raises its price. This, in turn allows all other producers to raise their prices. Furthermore, this general increase in price leads to a lengthening of the product cycle. These results are in line with those in Shaked and Sutton (1983), which linked increases in the heterogeneity among consumers to increases in the number of product qualities supported in equilibrium. Setting $b = 5$ has the opposite effect, causing prices (and so markups) to fall dramatically relative to the benchmark case.

the outside option, \check{u}_t . The monopolist's problem is

$$V(\nu_s, \check{u}_t) = \max_{p_{\nu_s}, d_t} \left\{ (1 - d_t) \left((p_{\nu_s} - c) Q_t(p_{\nu_s}; \nu_s, \check{u}_t) + \delta V(\nu_s, \check{u}_{t+1}) \right) + d_t \left((p_{\nu'} - c) Q_t(p_{\nu'}; \nu', \check{u}_t) - \phi + \delta V(\nu', \check{u}_{t+1}) \right) \right\}, \quad (7)$$

where $\nu' = t$ is the latest vintage. Like the competitive firm, the monopolist's adoption problem balances the gains from introducing a computer in the current period against waiting a period (or more) to do so. Bringing out a new vintage increases profits because consumers are willing to pay more for a superior product. However, the introduction of a new computer entails paying a fixed cost, ϕ .

6.2.2 Predictions of the Monopolist Model

Given these parameters calibrated in the competitive case, we solve for the monopolist's optimal pricing and updating strategy. Formally, we keep the same innovation rate and consumer preferences as in the competitive case, but simply change the market structure by restricting the number of firms to one.²⁸

In taking the monopolist's problem to the data, we need to choose a value for the fixed cost of upgrading, ϕ . This cost parameter is a main driver for the length of the product cycle. If this cost is near zero, for example, then the monopolist will update every period to take advantage of the high value consumers place on quality. We expect the cost of updating to be similar for all computer manufacturers, especially because the manufacturers are pulling in similar innovative intermediate inputs (e.g., lighter batteries or higher resolution laptop screens). From the competitive case, then, we have an upper bound on ϕ —the discounted flow of profits from selling a PC over its four-month product cycle.

We set ϕ to 0.75 so that the model's prediction on the monopolist's timing of replacement matches the eight-month long product cycle we observe in the data. This value of ϕ is 93 percent of the expected discounted profits of a computer manufacturer in the competitive market and so is consistent with the calibration done for the competitive model. Further, this parameter value, and its implication that PC manufacturers do not reap high profits, is consistent with the underlying assumption about the competitiveness of the PC market. Intel dominates the supply of CPU chips, and Microsoft for operating system

²⁸This restriction is motivated by assuming that the costs of entering the PC manufacturing industry are large enough that only one firm can enter profitably.

software for PC manufacturers. Given that both these inputs are crucial components of any PC, it is not surprising that a substantial portion of a PC manufacturer’s operating profits would be transferred up the supply chain.²⁹

With these parameters, the model predicts a flat price profile over the product cycle (see Table 7). Given the static nature of the consumer’s problem, this price profile generates a near-constant flow of sales and little variation in the average income of purchasers over the product cycle. Specifically, consumers purchasing the product have a budget of six units of the outside good each period. This monopoly case highlights the central role of competition in driving price declines over the product cycle (alongside falling sales and purchasers’ average incomes). Indeed, given the endogeneity of adopting new technology, rapid innovation alone is not sufficient to generate the rapid declines in prices over the product cycle. Rather, the combination of a competitive market structure and a rapid rate of innovation, together, are the main drivers of computer prices and sales over the product cycle.

Table 7: Monopolist: Price Declines over the Life Cycle
Percent Price decline, relative to first month

<i>Month</i>	1-2	1-3	1-4	1-5	1-6	1-7	1-8	average
Data	-0.04	-0.23	-0.34	-0.46	-0.69	-1.13	-1.86	-0.68
Model	-0.08	-0.17	-0.25	-0.34	-0.43	-0.53	-0.63	-0.35

Sales CDF

<i>Month</i>	1	2	3	4	5	6	7	8
Data: lower bound	0.112	0.304	0.504	0.661	0.801	0.897	0.973	1
Model	0.125	0.250	0.375	0.500	0.625	0.750	0.875	1
Data: upper bound	0.304	0.504	0.661	0.801	0.897	0.973	1	1

We also use the monopoly case to conduct an out-of-sample exercise to test the validity of the model. In particular, we compare the model’s prediction’s to the average Apple computer. We find that the model’s prediction of a flat price profile over the product cycle (see Table 7) is a close match to the prices of Apple computers seen in the data.³⁰ In the data we observe an average monthly price decline of -0.68 percent, while the model predicts a decline of -0.35 (a difference of only one-third of a percentage point). To better

²⁹Consistent with our result, Dedrick, Kraemer, and Linden (2010) find that “a large share of the PC industry profits are siphoned off to Microsoft and Intel” (page 3).

³⁰We use the same technique described in Section 5 to match the model and data price declines.

illustrate how the model captures the shape of the price declines observed in the data, in Figure 8 we plot the price declines predicted by the model against those observed in the data for a typical Apple computer over the product cycle. Given that this is an out-of-sample exercise, the model does a good job of matching the lack of price declines over the product cycle. The largest miss, in month 8, is by less than 1.3 percentage points (-1.86 percent in the data versus -0.63 percent predicted by the model).³¹

The model's predictions of uniform sales over the product cycle, shown in Figure 6, are only somewhat close to the data. In the data, we see a slight clustering of sales at the beginning of the product cycle. But these differences do not seem large, especially given that this is an out-of-sample exercise.

Reinforcing the validity of the model, its predictions about the timing of consumer's purchases accord with the data. The income distribution of Apple purchasers is higher and less variable than that of PC purchasers. Under the monopoly case, the model generates this result by predicting that Apple sells only to the high-income consumers each period, with a budget of six units of the outside good each period. This result highlights how the set of actual consumers can differ from the set of potential ones. In contrast to the monopoly case, PC manufacturers end up selling to all potential consumers because of the competitive pricing pressure.³²

Overall, with only a change in market structure, we find that the parsimonious vintage-capital model is able to roughly match the stylized facts for the Apple case. This is a strong test of the model's validity because this exercise forces the model to predict prices and sales under a monopoly, a starkly different market structure compared to the PC case.

7 Conclusion

The goal of our analysis is to determine if upstream innovation and a competitive market can explain the observed time series of prices, as well as sales and consumers' incomes, over the product cycle for personal computers. Our strategy is to develop and

³¹One way to generate steeper price declines in the model is to increase the utility from the outside option. The steeper price declines, however, are associated with lower per period profits for the monopolists. Consequently, increasing the outside good's utility induces the monopolist to upgrade its product more frequently, all else equal.

³²The model predicts that the monopolist sets a high price for its product. This price level and the resulting markup are not realistic. But this is mainly driven by the assumption that the monopolist faces the same set of potential consumers as PC manufacturers. Reducing the potential demand for the monopolist's computer naturally results in a lower price level.

calibrate a vintage-capital model to match stylized facts on prices and sales from the PC retail market.

As calibrated, the model matches the observed stylized facts and performs well in an out-of-sample test. The main result is that the competitive structure of the PC market along with the rapid rate of innovation are important elements in this market. These two forces, as formalized in a simple vintage-capital model, can explain in large part the strategies computer manufacturers' use for pricing over the product cycle and for adopting technology.

Because our model is fairly stylized, there is ample room for extending it to account for other important features of the PC or other innovative markets. For example, it may be fruitful to incorporate a dynamic demand into this environment. Nevertheless, because of its simplicity, this model provides a useful benchmark against which to compare more complicated models.

References

- [1] Aghion, P., Bloom, N., Blundell R., Griffith R., and Howitt, P. 2005. "Competition and Innovation: An Inverted-U Relationship," *The Quarterly Journal of Economics*, MIT Press, vol. 120(2), pages 701-728, May.
- [2] Aghion, P., Blundell R., Griffith R., Howitt, P., and Prantl, S. 2009. "The Effects of Entry of Incumbent Innovation and Productivity," *The Review of Economics and Statistics*. vol. 91(1), pages 20-32.
- [3] Aghion, P., Harris, C., Howitt, C., and Vickers, J. 2001. "Competition, Imitation and Growth with Step-by-Step Innovation," *Review of Economic Studies*, Blackwell Publishing, vol. 68(3), pages 467-92, July.
- [4] Aghion, P. and Howitt, P., 1992. "A Model of Growth through Creative Destruction," *Econometrica*, Econometric Society, vol. 60(2), pages 323-51, March.
- [5] Aizcorbe, A. 2005. "Moore's Law, Competition, and Intel's Productivity in the Mid-1990s," *American Economic Review* vol 95(2), pages 305-308, Papers and Proceedings of the 117th Annual Meeting of the American Economic Association.
- [6] Aizcorbe, A. and Kortum, S. 2005. "Moore's Law and the Semiconductor Industry: A Vintage Model," *Scandinavian Journal of Economics*, Blackwell Publishing, vol 107(4), pages 603-630.

- [7] Aizcorbe, A., and Shapiro, A. 2010. "Implications of Consumer Heterogeneity on Price Measures for Technology Goods" Bureau of Economic Analysis Working Paper.
- [8] Arrow, K., 1962. "Economic Welfare and the Allocation of Resources for Invention, in R. Nelson, ed., *The Rate and Direction of Inventive Activity* (Princeton, NJ: Princeton University Press), pp. 609-625.
- [9] Berndt, E. and Rappaport, N. 2001. "Price and Quality of Desktop and Mobile Personal Computers: A Quarter-Century Historical Overview," *American Economic Review*, American Economic Association, vol. 91(2), pages 268-273, May.
- [10] Biesebroeck, J.V. and Hashmi A. 2009 "Market Structure and Innovation: A Dynamic Analysis of the Global Automotive Industry," Katholieke Universiteit manuscript.
- [11] Borenstein, S., and Rose, N., 1994. Competition and Price Dispersion in the U.S. Airline Industry. *The Journal of Political Economy* Vol. 102(4), 653-683.
- [12] Bresnahan, T., Stern, S., and Trajtenberg, M. 1997. "Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s" *RAND Journal of Economics*, The RAND Corporation, vol. 28(0), pages S17-S44.
- [13] Bulow, J. 1982. "Durable-Goods Monopolists" *Journal of Political Economy*, vol 90, pages 314-332.
- [14] Conlin, C. 2010. "A Dynamic Model of Costs and Margins in the LCD TV Industry" Columbia University working paper.
- [15] Dasgupta, P. and Stiglitz, J. 1980. "Industrial Structure and the Nature of Innovative Activity." *Economic Journal* 106 (July): 925-51.
- [16] Dedrick, J., Kraemer, K., and Linden, G. 2010. "Who profits from innovation in global value chains? A study of the iPod and notebook PCs". *Industrial and Corporate Change* 19(1): 81-116.
- [17] Erickson, T. and Pakes, A. 2008. "An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods", NBER Working Paper. September.
- [18] Feenstra, R. and Knittel, C. 2009. "Re-Assessing the U.S. Quality Adjustment to Computer Prices: The Role of Durability and Changing Software," NBER Chapters, in: Price Index Concepts and Measurement, pages 129-160.

- [19] Gerardi K. and Shapiro, A. 2009. "Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry," *Journal of Political Economy*, University of Chicago Press, vol. 117(1), pages 1-37, 02.
- [20] Gilbert, R.J. and Newbery, D.M.G. 1982. "Preemptive Patenting and the Persistence of Monopoly." *American Economic Review* 72 (June): 514-27.
- [21] Goettler, R. and Gordon, B. 2009 "Does AMD Spur Intel to Innovate More?" *Journal of Political Economy*.
- [22] Gordon, B. 2009 "A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry" *Marketing Science* 28(5), pages 846-867.
- [23] Gowrisankaran, G. and Rysman, E. 2009. "Dynamics of Consumer Demand for New Durable Goods", NBER Working Paper w14737. February.
- [24] Greenstein, S. and Ramey, G. 1998. "Market Structure, innovation, and vertical product differentiation." *International Journal of Industrial Organization* 16, pages 285-311.
- [25] Nosko, C. 2010. "Competition and Quality Choice in the CPU Market", Harvard University manuscript.
- [26] Pakes, A. 2003. "A Reconsideration of Hedonic Price Indices with an Application to PCs" *American Economic Review*, American Economic Association, vol.93(5) pages 1578-1596.
- [27] Prescott, E. and Visscher, M. 1977. "Sequential Location among Firms with Foresight," *Bell Journal of Economics*, The RAND Corporation, vol. 8(2), pages 378-393, Autumn.
- [28] Schumpeter, J. 1934. *The Theory of Economic Development*, Cambridge: Harvard University Press.
- [29] Schumpeter, J. 1942. *Capitalism, Socialism, and Democracy*. New York: Harper and Row.
- [30] Shaked, A. and Sutton, J. 1982. "Relaxing Price Competition Through Product Differentiation," *Review of Economic Studies*, The Review of Economic Studies Ltd, vol. 49(1), pages 3-13, January.

- [31] Shaked, A. and Sutton, J. 1983. "Natural Oligopolies," *Econometrica*, Econometric Society, vol. 51(5), pages 1469-83, September.
- [32] Stokey, N., 1979. "Intertemporal Price Discrimination," *The Quarterly Journal of Economics*, MIT Press, vol. 93(3), pages 355-71, August.

Appendix

A Calculating ζ

In all four surveys, households, that had purchased a new computer were asked about their old computer. Of the 36,137 responses, 25 percent reported this was the first computer ever owned, 26 percent stated they still used their old computer, and 48 percent answered their old computer was no longer owned or used. For households that did not have an old computer or that no longer used their old computer, we assign a zero utility to their outside option. Given the lack of a secondary market for personal computers, we assume these households' outside option is to use a computer set up for the general public (e.g., at a library). Because users of public computers cannot install software or save their work for extended periods, and often face time restrictions, we believe this alternative to be an extremely poor substitute for a new computer, hence the zero utility. For those households that continue to use their old computer, we assume this is their outside option.

Gordon (2009) reports that households replace their computers every three to four years, and so we assume the old computer is four years old. Using Moore's law and the normalization that a new computer provides 10 utils, the utils from an old computer would be 2.5. However, we view 2.5 as an upper bound since older computers face software compatibility constraints.³³ Hence, for this group of households, the utility value of an computer lies between 0 and 2.5. We assume that these households' average outside value is equal to the midpoint of this range, or 1.25. Pulling everything together, 26 percent of households continue to use an old computer with an outside option averaging 1.25. The remaining households that use library computers have an outside option of 0. Taking a weighted average, we calculate that the outside option of an average household is $(0.26 * 1.25 + 0.74 * 0) = 0.33$. The ratio of the outside option to the frontier technology is then $\frac{0.33}{10} = 0.033$.

³³Feenstra and Knittel (2009) show that consumer surplus gains are lower with a price index accounts for compatibility issues between hardware and software.

B Algorithm for Solving the Vintage Capital Model

B.1 Demand side

To solve the demand side of the model in a period t , we need to know the price and associated utility of each available vintage, along with the utility of consuming the outside good. Given these two vectors, we implement the algorithm described in Shaked and Sutton (1982, 1983) to compute each vintage's demand. Essentially, we order the vintages from newest to oldest and use the indifference condition, equation (4), to find the marginal consumer between each pair of vintages. Prices can result in vintages not having any demand. Given the set of marginal consumers between all the neighboring vintages, we use the Beta density of consumers over budgets to compute demand for each product.

B.2 Supply side: competitive case

On the supply side for the competitive case, we solve for the stationary equilibrium where the lowest-quality firm updates in every period. To find the price and sales over the product cycle for the case where there is entry every period, we only need to solve the model once. To see this, suppose the parameters are such that three vintages are sold (that is, there is no demand for the fourth vintage when its price is equal to the marginal cost). At time t the utility levels associated with the three computers and the outside good are:

$$\bar{u}_t = \begin{vmatrix} u * 1.1029^2 \\ u * 1.029 \\ u \\ \tilde{u} \end{vmatrix}, \quad (8)$$

where $u > \tilde{u} > 0$. The second-highest quality computer is 1.029 better than the lowest-quality computer because we have assumed that the monthly growth in quality is 1.029, based on Moore's law. In the following period, the lowest quality computer is upgraded and the outside good improves in quality. The resulting utility vector is

$$\bar{u}_{t+1} = \begin{vmatrix} u * 1.1029^3 \\ u * 1.029^2 \\ u * 1.029 \\ \tilde{u} * 1.029 \end{vmatrix}. \quad (9)$$

Comparing these two vectors, we see that $\bar{u}_{t+1} = 1.029 * \bar{u}_t$. Since only relative utilities matter in the model, the Nash equilibrium prices and sales are the same across both

examples. Consequently, prices over a computer's product cycle are equal to prices in the cross-section.

Given static demand, the firm's pricing problem is also static. From Shaked and Sutton (1983), we know there exists a Nash equilibrium in prices. We find this equilibrium through an iterative technique. We begin with each firm pricing at marginal cost. We then use the manufacturer offering the newest vintage to find the price that maximizes its profit holding all other manufacturers' prices fixed. We update the price vector to reflect this new price and repeat the exercise with the next newest vintage. Once we reach the oldest available vintage, we repeat the loop, cycling over firms until a Nash equilibrium is reached. For the calibrated parameters, we did not find multiple equilibria. When using different starting price vectors and different looping techniques, the program converged to the same price equilibrium.

B.3 Supply side: monopoly case

The monopolist makes two decisions to maximize profits: (1) it sets price and (2) it decides on the profit-maximizing product replacement strategy.³⁴ The pricing decision is static and straightforward to solve since there is no strategic behavior. The pricing tradeoff confronting the monopolist is that higher prices push marginal consumers to purchase the outside option. The product replacement strategy is more complicated since it is a dynamic problem.

We find the optimal replacement strategy through simulation. Given a replacement strategy, we calculate the monopolist's per period profits for 200 periods. We then compute the present discounted profits associated with the chosen replacement strategy. We compute this present discounted profit result when the monopolist replaces the product every period, every second period, every third period, etc. We evaluate longer and longer replacement cycles until the present discounted profits start to decline. We then return the product replacement strategy with the highest present discounted profits as the optimal strategy.

The sunk cost of entry is a main determinant of the length of the optimal replacement strategy. From the competitive case, we only know the upper bound on this value. Hence, we pin down this parameter by choosing the sunk cost such that the monopolist's optimal replacement strategy matches the data.

³⁴Note, we do not allow the single-product monopolist to sell two vintages at the same time.

B.4 Least distance criterion

For the competitive case, we use a least distance criterion to find the value of the free parameters to most closely align the model with the data. Our criterion has three parts. The first portion is the squared difference between the length of the product cycle in the model and the data. The second portion is the sum of squared differences between the model's sales CDF and that in the data. The third portion of the criterion is the sum of squared differences between the model's price declines and those in the data.

Given a vector of parameters θ , we compute the criterion by following these steps:

1. Solve the competitive case of the model and record the price and sales of a computer over its product cycle along with the number of periods in a product cycle.
2. Use the price series to compute the decline in price over the product cycle, and the sale series to construct the sales CDF.
3. Compute the first portion of the criterion, which is the squared difference between the model's product cycle length and 4, the number of months a typical PC is sold in the data.
4. Compute the second portion of the criterion. Let ω denote a four element vector, where each element corresponds to the first four months of the product cycle. For a given month, if the model's sales CDF is within the minimum and maximum sales CDF bounds implied by the data, then the corresponding element of ω is set to zero. If the model's sales CDF is outside these bounds, then the element of ω is equal to the squared difference between the model's sales CDF and the closest bound implied by the data. The second portion of the criterion is then equal to the sum of the elements of ω .
5. Compute the third portion of the criterion. For months two through four of the product cycle, we want to compute the squared difference between the model's price declines and those observed in the data. Recall the price declines are relative to the price in the first month of the product cycle. From a fixed-effects regression (see Section 3.1), we know the price declines associated with different deciles of a computer's sales CDF. These inferred price declines by sales CDF are listed in Table 8.

To match the model and the data, we take months two through four of the *model's* sales CDF and find the associated inferred price declines from the data. Hence, if

Table 8: Estimated Price Declines for PCs over the Life Cycle

Sales CDF	Price Decline (percent)
20	-0.026579
30	-0.037559
40	-0.046736
50	-0.05204
60	-0.055562
70	-0.062674
80	-0.06973
90	-0.074973
100	-0.091519

Notes: These estimated price declines are plotted in Figure 2

the model generates a sales CDF of 0.6 in the second month of the product, we would compare the model's price decline to -0.055562. In practice, we estimate a third order polynomial over the eight data points relating price declines and the sales CDF. Using the model's sales CDF as an input, this function returns the price declines inferred from the data. The third portion of the criterion is equal to the sum of squared differences between the model's price declines and those in the data.

C Age of Newest CPUs (Notebook Computers)

	All CPUs		Intel CPUs	
	Avg.	Max	Avg.	Max
Hewlett Packard	0.8	3	1.1	3
Compaq	1.1	4	1.4	5
Toshiba	0.9	4	1.0	4
Apple	2.9	8	2.9	8
Gateway	1.5	5	1.5	5
Sony	1.5	7	1.5	7

Notes: Values in the table represent the average and maximum age (in months) of the newest chip for each month of the post-PowerPC period (i.e. 2006m6- 2009m4). The sample is taken over notebook computers.

Source: NPD.