

The Impact of Competition on Technology Adoption: An Apples-to-PCs Analysis*

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

January 29, 2010

Abstract

We study the effect of market structure on a firm's decision to adopt a new technology in the personal computer industry. This industry is unusual because there exists two horizontally segmented retail markets with different degrees of competition: the IBM compatible (or "PC") platform and the Apple platform. We first document that relative to Apple, producers of PCs have more frequent technology adoption, shorter product cycles, and steeper price declines over the product cycle. We then develop a parsimonious vintage-capital model which is able to match the different stylized facts of PC and Apple producers. The model predicts that competition is the key driver of the rate at which technology is adopted.

Key words: innovation, market structure, computers

JEL classification: D40, L10, L63, O30

*We thank Ana Aizcorbe, Olivier Armantier, Ben Bridgman, Kyle Hood, Matt Osborne, and Jeff Prince for their comments and suggestions. 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, or the Federal Reserve System.

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

[‡]Bureau of Economic Analysis; e-mail: adam.shapiro@bea.gov

1 Introduction

There has been an important historical debate about the role of competition on innovative activity in the marketplace. Schumpeter argued that competition impedes innovation in the sense that rents accrued to the innovative firm cannot be sustained in a competitive environment. The alternative view, associated with Arrow, argues that competition acts to reduce a firm's pre-innovation rent relative to its post-innovation rents, thereby enticing firms to innovate. The debate has important implications for many economic questions, including the sources of economic growth as well as antitrust and other competition policy issues. There has thus been an ongoing line of research dedicated to the topic.¹

In the environment to which Schumpeter and Arrow were referring, research and development (R&D) is undertaken directly by the innovating firm. A large number of industries, however, consist of firms of final-consumer goods that *adopt* and assemble components which embed innovations generated by intermediate manufacturers. Producers of digital cameras, cell phones, stereos, and personal computers (PCs) are examples of such goods which can be thought of as an assembly of individual innovative components.² Retail goods manufactures of these types of goods can be thought of as "technology adopters" in the sense that their technological limit, or frontier technology, is dictated by the underlying component manufacturers. In this setting, although much of the large R&D costs are not directly taken by consumer-goods producing firms, they still incur similar costs of entry stemming from marketing, advertising, assembly, and establishing retail channels. Thus, for these consumer-goods firms, the classic question of market structure's effect on innovative activity applies since the firm must weigh the cost against the marginal profit of adopting the new technology.

This paper focuses on an important consumer-goods industry, personal computers, and analyzes the relationship between market structure and innovation. Our work makes

¹See Aghion, Harris, Howitt, Vickers (2001), Aghion, Bloom, Blundell, Griffith, Howitt (2005), Aghion, Blundell, Griffith, Howitt, Prantle (2009), Aghion and Howitt (1992), Biesebroeck and Hashmi (2009), Dasgupta and Stiglitz (1980), Gilbert and Newbery (1982), and Goettler and Gordon (2009).

²The central processing unit (CPU), for instance, which is a key component of many of these products, is in many cases produced by chip manufacturers.

two contributions. First, we provide a descriptive analysis of the industry, measuring rates of technology adoption under different market structures as well as documenting prices and sales over a typical computer's life cycle. Second, we use a vintage-capital model to quantify the importance of market structure on the rate of technology adoption.

Similar to many other technological industries, the personal computer market is one in which innovations are extraordinarily frequent and new products almost always incorporate new technology advances. Unlike most industries, however, an important feature of the personal computer industry is the existence of two segmented retail markets. The industry is divided between two platforms: the IBM compatible platform (or simply the "PC") and the Apple platform.³ While there are many differences between these platforms, arguable the largest difference is their operating systems; PCs run on Microsoft Windows while Apples run on Mac OS.⁴ The horizontal differentiation between PCs and Apple therefore segments the retail market, in the sense that consumers of PCs do not consider Apple products as close substitutes and vice-versa.⁵ This segmentation is crucial for our analysis, because there are many firms that manufacture PCs, while only one firm that manufactures Apple computers. Thus, one can think of Apple as having a great deal of market power along this horizontal dimension of operating system such that it has no close vertical substitutes by other computer manufacturers for any specific product line (e.g. 15-inch notebook computers). Such segmentation allows us to compare and contrast the technology adoption decisions for a computer manufacturer in a competitive and monopolistically structured market.

Using NPD Group scanner data from 2001 to 2004, we compute various statistics

³For a general overview of the history of competing platforms see Bresnahan and Greenstein (1999). Given the evolution of this industry, the IBM compatible platform is also referred to as the Wintel platform, a label that alludes to the Windows operating system and Intel processor combination used by the vast majority of these computers.

⁴Stavins (1995) refers to this type of platform differentiation as horizontal differences in the order of quality. This is in contrast to vertical differences (e.g. speed and memory) along a certain platform line.

⁵Bresnahan, Stern and Trajtenberg (1997) differentiate the source of transitory market power in the personal computer industry between frontier (vertical quality) and brand (horizontal quality). The authors find that the effects the competition on market power were confined to substitution clusters, and brand differentiation provides protection from competition along similar vertical dimensions.

on the rate of product entry across firms. Further, we present evidence that new computers embody new technologies, making them higher-quality products relative to the existing set of comparable computers. An interesting feature of the data is that they unequivocally show that Apple is slower to adopt a new technology than the other PC manufacturers. The rate of introduction of new computer models is much higher for PCs than for Apple, and in addition, Apple keeps its product on the market for roughly twice the length of time, on average. Thus, for a given period of time, the PC manufacturers are introducing significantly more products with shorter life spans relative to Apple. Thus, the data suggest that a more competitive market structure acts to increase the rate at which firms adopt new technological innovations.

The difference in adoption decisions between PC manufacturers and Apple is likely to be function of many factors. To better understand the importance of market structure in explaining the differences in the rates of adoption across Apples and PCs, we use a parsimonious vintage-capital model. Our strategy is to parameterize the model to match the stylized facts on adoption, pricing and sales for PCs, where many firms compete. We then use the model to make an out-of-sample prediction on adoption, prices and sales given a monopolistic market structure. Surprisingly, the model's predictions closely match the stylized facts for Apple computers. Essentially, keeping preferences and technology fixed, and only changing market structure, the model is able to match the different stylized facts for both PC and Apple computers. Consequently, a main result of the paper is that market structure plays a major influence on the rate of technological adoption in the market for personal computers.

Importantly, the model provides insight as to how and why the competitive market structure generates higher rates of technology adoption. On the demand side, we use the quality-ladder framework of Shaked and Sutton (1982, 1983) (henceforth, SS). On the supply side, firms offer computers of different vintages and compete in price.⁶ Computer manufacturers face a constant marginal cost and pay a fixed cost to introduce a new product. Within this environment, the introduction of a new, higher-quality computer places enormous pressure on all existing computers. Lower-quality computers capture

⁶Our supply-side is similar to Aizcorbe and Kortum (2005), who use a vintage-capital model to analyze pricing and production in the semiconductor industry. Besides focus, our papers differ in our incorporation of consumer heterogeneity.

a relatively small portion of the market and charge minimal markups. Hence, once a firm’s computer has been displaced as the highest-quality product, its per-period profits are fairly small. Given the rapid entry of new computers, the model predicts steep falls in price, sales, and consequently profits, over a computer’s product cycle. This “market-specific” obsolescence generates a significant premium on having the highest-quality product in the market. Following the reasoning originally posed by Arrow and more currently by Aghion et al. (2001, 2005), competitive forces encourage rapid rates of innovation because having the highest-quality computer, and so differentiating one’s product from the competition, generates large profits.

In contrast, in the monopoly case there is no market-specific obsolescence. While all computers faces general obsolescence (e.g. new software is often not compatible with older computers), this force places much less pressure on prices and sales. Because the monopolist is able to maintain a high level of revenue for its computer over time and given the fixed cost of introducing a newer higher-quality computer, the monopolist introduces newer computers less frequently than its counterparts in the competitive industry. Consequently, the monopolist adopts newer technologies at a slower pace.

In addition to the literature on innovation and market structure, this paper is related to research on pricing dynamics. In this respect, our study is closely tied with research analyzing the effect of competition on pricing dynamics.⁷ Prices for technological goods are well known to fall over the life of the product cycle.⁸ There are a host of explanations for these price declines, including inter-temporal price discrimination (Stokey (1979)) and input-cost declines. Our work focusing on the role of competition, demonstrating that the introduction of close substitutes by competing firms can be the main driver of declining prices and sales over the product cycle.⁹

The study is structured as follows: Section 2 contains a detailed overview of the NPD and TUP survey data and a description of the pricing and sales dynamics of PCs and Apple computers. In Section 3 we describe the competitive vintage-capital model,

⁷See, for example, Borenstein and Rose (1994), Aizcorbe (2005), Gowrisankaran and Rysman (2009), and Gerardi and Shapiro (2009)

⁸See Erickson and Pakes (2008), Berndt and Rappaport (2001), and Pakes (2003)

⁹This result is in-line with Aizcorbe and Kortum (2005), who show that the introduction of newly introduced semiconductor chips are a main force behind the rapid fall in semiconductor prices.

analyze its fit with the PC data under the stationary case, and discuss its implication for dictating product cycle length. In Section 4 we examine the monopolistic industry and measure its fit with the data on Apple computers. We conclude in Section 5.

2 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-sale¹⁰ transaction data (i.e., scanner data) sent to NPD Techworld via automatic feeds from their participating outlets on a weekly basis.¹¹ The data cover the course of 34 months, November 2001 to August 2004 and consist of sales occurring at outlet stores. Thus, manufacturers such as Dell and Gateway that primarily sell directly to the consumer are not included. Each observation consists of a model identification number, specifications for that model, and the total units sold and revenue. From units sold and revenue we calculate a unit price of each PC sold.¹² Table 1 displays the share of units sold in the data for the entire sample and as well as for the notebook and desktop subsamples. Compaq and HP make up the bulk of computers sold in the data, at 23 and 25 percent respectively, while Apple’s market share is smallest at 4 percent.

In the TUP survey data, we have access to two annual surveys that were conducted in 2003 and 2004. TUP is a detailed two-stage survey of American’s 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

¹⁰“Point-of-sale means” that any rebates or other discounts (coupons, for example) 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.

¹¹The weekly data are organized into monthly data using the Atkins Month Definition, where the first, second and third weeks of the quarter include four, four and five weeks, respectively.

¹²We omitted observations with a high probability of measurement error. Specifically, these were computer models in which the total number of units sold over its life cycle was less than 1000 as well as individual observations where the units CDF was less than 0.01 (that is, where less than 1 percent of the units sold were in the first month on the market.) In the latter, we are trying to isolate instances where manufacturers often test the product in only few markets.

(such as income, education level, marital status, 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, to the head of household’s income. All 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.¹³ We isolate observations where PC is used at home, and we drop observations where the specification of PC is not reported.

2.1 NPD Data

Since the NPD data is monthly, it allows us to take a close look at the dynamics of sales and prices within the PC industry at a relatively high frequency. We generate statistics from unit sales and unit prices over the life of all computers in the sample period and then compare these statistics between different brands in the industry. Overall, the data show that PC manufactures all display similar patterns of sales and price dynamics while Apple’s dynamics are quite distinct. We expound on these findings below.

2.1.1 Sales and Adoption Dynamics

We first highlight a few facts about the dynamics of unit sales where we focus on the length of the computer’s lifespan on the market as well as the frequency of turnover. The length of product cycle is measured by constructing average CDFs by brand as depicted in Figure 1.¹⁴ Since computers do not necessarily enter the market at the beginning of the month, the first month of data will include less than 31 days worth of units sold. We can therefore interpret the CDF function in terms of bands as opposed to explicit months on the market. The figure shows that on average, most brands of PC platforms sell well over half (62 percent for all PCs) of their units by the second month on the market. By the third month, over 80 percent of units have been sold on average for most

¹³We drop observations where we believe measurement error is present. Specifically, we drop observations where the price paid is reported less than \$100. We also drop observations where the PC was purchased prior to 1997.

¹⁴We measure the CDF by running a regression of monthly unit sales on “months on market” dummies creating predicted values of the number of units sold given months on the market.

manufacturers. Apple on the other hand keeps its products on the market for about twice the amount of time as the other PC manufacturers. By the third month, Apple has sold less than half of its units (46 percent) over any given product's lifespan. Apple's computers do not reach the reach the 80 percent point until after 6 months of the product being on the market.

Table 2 displays statistics that describe the frequency of product turnover for the entire sample: the average number of new models per month and the number of months without a new model introduced. Both of these statistics indicate that Apple has the lowest product turnover frequency of the brands in the sample, while Hewlett Packard has the highest. Specifically, Apple introduces the lowest number of new models per month over the sample, (1.94 models per month) and has the largest fraction of months without a new model introduced (34 percent or 12 out 34 months without a new model introduced). In contrast, Hewlett Packard, has averaged 9 new models per month over the sample period and experienced only one month without introducing a new computer (6 percent of months without a new model adoption). The other PC manufacturers offer between 3 and 6.5 new models per month and range between 2 and 6 months out of the total 34 without offering a new model. Splitting the sample between desktop and notebook computers reveals the same adoption dynamics. Apple adopts one desktop and one notebook per month on average, while most of the other PC platform firms adopt at least twice as many of each relative to Apple.¹⁵

The adoption statistics described above show that the "PC" manufacturers adopt at a more rapid pace than Apple. This suggests that the PC manufacturers are speedily trying to adopt better quality products to the market. One way to test this idea is to analyze which manufacturers are adopting to the frontier quality and when they are choosing to do so. Although the quality of a computer is a multi-dimensional vector of characteristics, for expositional purposes in Figure 2 we proxy the technological frontier with CPU clock speed.¹⁶ Specifically, this figure shows the frontier CPU clock speed for both notebook and desktop computer samples for each manufacturer in the NPD sample. The main

¹⁵Toshiba and Emachines are excluded from the desktop and notebook sample respectively because Toshiba does not produce desktop computers and only 4.3 percent of Emachines computers sold are notebooks.

¹⁶Clock speed is the rate at which a processor can complete a processing cycle.

result from the figure we wish to highlight is that for both notebooks and desktops, the PC manufacturers leap frog one another quite often, with each manufacturer taking over the leadership position on this speed quality dimension. Thus, even looking at only one dimension of quality, there seems to be fierce competition over quality leadership between the PC manufacturers.

In contrast, Apple does not participate in the leap-frog environment with the other manufacturers and is slower to adopt any given CPU clock speed. While clock speed does not necessarily dictate a given CPU's entire performance, it can be very important from a marketing standpoint where simplification of information about quality is relevant to attracting consumers.¹⁷ As we show in the next section, Apple charges high prices relative to the other computer manufacturers. Thus, at the very least, this figure demonstrates that Apple has sufficient product differentiation from the other PC manufacturers that it can charge high prices even though it is disadvantaged from a marketing standpoint. At the very most, it demonstrates that Apple has enough product differentiation that it can charge high prices with lower quality products altogether.

Table 4 highlights many of the features seen in Figure 2. For instance, in the desktop market, Hewlett Packard held the sole leadership position 46 percent of the months in the sample, while Sony held the position 17 percent of the time. In the notebook market, Toshiba held a leadership position 46 percent of the months, while Compaq and Sony held it for 6 and 3 percent, respectively. The table also shows (column 1) that Compaq, Hewlett Packard, Sony, and Toshiba all spent a large portion of the sample period at the frontier speed.

2.1.2 Pricing Dynamics

In addition to sales, there are also stark differences in the dynamics of prices between manufacturers in the computer industry. Specifically we look at the extent as to which

¹⁷One apparent reason for the discrepancy in CPU speed is that, during this sample period, Apple was manufacturing computers with IBM PowerPC processors, while the other manufacturers were primarily using Intel chips. In his MacWorld 2001 keynote address, Apple's CEO Steve Jobs labeled the fact that PowerPC processors operate at lower clock speeds the "Megahertz myth," explaining that clock speed is only one of four key contributing factors that vary a CPU's performance. In January 2006, Apple switched from IBM to Intel CPUs.

prices fall over the life of computer by estimating a fixed effects regression of the logarithm of price on “months on market” variables for each brand in the sample. Figure 3 depicts the results of these regressions and highlights the extent to which each brand reduces the price of the computer over the life cycle. Depicted are the estimated coefficients for the first six months of the life of the good. The omitted dummy variable is the first month of entry indicating that the subsequent coefficients represent the percentage change between the given month and the first month.¹⁸

It is clear from Figure 3 that Apple’s prices fall relatively slow and less extensively than do the prices of the PC manufacturers. Hewlett Packard, Compaq, Sony, Emachines and Toshiba’s prices fall by roughly 8 to 11 percent between the fourth and first month. In contrast, Apple’s prices show hardly any fall at all over this time horizon (measured as a fall of 0.4 percent). In fact, Apple’s prices remain relatively flat for the life of the good, falling by 3.3 percent by the end of the sixth month relative to the first month.

To get a better idea of the dynamics of prices in the industry, Figure 4 depicts 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.¹⁹ The statistics described above become quite transparent in this graph. The PC manufacturers have numerous models within the 15-inch notebooks category, frequent entry, and declining prices over the life of the good. Apple, in contrast, sells only two 15-inch notebook models simultaneously (a 256MB RAM and 512MB RAM) in any given month. Furthermore, Apple’s prices are generally flat over the life of the good, and there is little, if any, overlap of new and incumbent models. The figure also makes clear that Apple terminates an entire line of products, at the same time it introduces another, while the other PC manufactures introduce models in a more staggered fashion.

We narrow down the product space of Figure 4 in order to demonstrate some of the competitive dynamics involved between manufacturers. Figure 5 isolates 512 MB

¹⁸We describe the fixed-effect regression analysis in Appendix A where we provide a table of estimates and their standard errors for months 2 through 12 of the product cycle. Because many of the price declines in the later part of a product’s life cycle are due to stock-out sales, we also report estimates where we omitted observations beyond the 90th percentile of the units CDF function.

¹⁹For ease of view, prices after the the units CDF reached 90 percent for each model were omitted. In this figure, each dotted line represents an individual computer model.

RAM 15-inch notebooks where the entering PC happened to be the highest price in this product line. This exercise is attempting to isolate the computer models with both the newest (entering goods) and highest quality technology under the assumption that the computer with the highest quality is also the highest priced.²⁰

The price contours generated from this exercise are quite interesting. First off, as expected, the figure shows that quality is generally increasing in price across at least one dimension. Second, in all cases, the prices of incumbent computer models dip below those of the entering model. Generally, the new model is priced at similar price to introductory price of the model that preceded it. Finally, in some cases, the introduction of the new model seems to put downward pressure on the preceding model as is the case with the introduction of Sony's product. In contrast, as seen in Figure 4 Apple's prices of this particular product decline only when it cannibalizes itself.²¹ The figure makes clear that there are strategic interactions between firms in how the products are priced with one another. Specifically, the prices of incumbent computers are generally lowered to make room for the relatively higher prices of the newer higher quality counterparts.

2.2 TUP Survey Data

As well as the firm side, there are important features of the personal computer industry on the consumer side. In this section we highlight some facts about the underlying distribution of consumers purchasing PCs by looking at the TUP survey data. We focus on the consumer income as this is the demographic variable most tied to reservation price, and hence product choice in most econometric studies and economic models.²² The survey data reveal that both the levels and 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.

²⁰Since there are so many attributes to an underlying computer model is it not straightforward in choosing the computer with the highest quality.

²¹Although this is only apparent in a few of the contours in Figure 4, relaxing the 90 percent cutoff of the CDF makes this more clear as can be seen in Appendix B.

²²Sources

2.2.1 Income Distributions and Levels

There are large differences in the income distribution of consumers between Apple and PCs. Table 4 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.²³ The survey data show that consumers of Apple have narrow income dispersion (0.195 gini coefficient) around a relatively high median level (\$65,000).²⁴ In contrast, the “PC” brands have consumers with much wider income dispersions around lower median levels. These data make it apparent that PC firms are competing over a wider distribution of consumers relative to Apple, which sells mainly to high-income consumers.

2.2.2 Income and Price Dynamics

Using the same TUP survey data that we use in this study, Aizcorbe and Shapiro (2009) find that higher income consumers pay a higher price for the same computer than do lower income consumers. Specifically, a fixed effects regression is run— holding fixed the attributes of the computer purchased— of income and other demographic variables on the logarithm of price. The study finds that the coefficient on income is .125 indicating that a 10 percent fall in a consumer’s income is correlated with a 1.25 percent fall in the price paid for a given computer. These results, combined with the declining nature of prices over the cycle of the good, show that high income consumers are presumably purchasing early in the model’s life cycle.²⁵

The particular reason for this correlation is not immediately transparent. One explanation is that low-income consumers are waiting for the price of the good to fall as

²³The 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.

²⁴We 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 on Apple could realistically be somewhat larger than what we measure.

²⁵Interactions 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.

the firm implements intertemporal price discrimination. Such an explanation, however, seems unlikely given that there is no reason why Apple would chose not to implement such a strategy if it is in fact profit maximizing for all of the other manufacturers.

The model we develop below attributes this correlation between income and price to a byproduct of the manner in which firms compete over a distribution of heterogeneous consumers through time. In particular, the model shows that one can generate the dynamics of prices and income of both the monopoly (Apple) and the competitive side (PCs) when competition over the product space is the root cause of such price declines. The result of the model is that in a highly competitive environment, high income consumers pay a large markup for the highest quality goods early in its product cycle.

3 The Competitive Industry

In this section, we consider the competitive industry where firms compete with one another over products that are perfect horizontal substitutes (e.g. the same product line such as 15 inch notebook computers running on the IBM platform) but imperfect vertical substitutes (e.g. differing degrees of quality such as CPU clock speed).²⁶ The competitive industry is therefore oligopolistic in the sense that, for a specific product line, each firm offers a product that differs in its vertical attributes which allows for differing degrees of substitution between products. As in SS, single product firms compete with one another in price while taking into account the distribution of consumer taste over quality. The novelty of our approach is the interpretation of this static model into that of a dynamic one by imposing that products of specific quality can be identified by their vintage. The result is that in the steady state, with constant growth in quality, a product of a specific product line will follow a specified price and sales path over its life cycle.

After introducing the model, we measure the model's fit with the data by comparing the price and sales paths generated by the model with that actually seen in data. Our results indicate that the model does a nice job of explaining the data for IBM platform PCs. We then discuss the model's implications about the effects of product entry on

²⁶In the following section, we look at the monopolistic industry where the firm is the sole provider of the horizontally differentiated product.

pricing and sales dynamics.

3.1 The Competitive Model

3.1.1 Demand

We model the competitive industry using an infinite-period vintage-capital model, where the discount rate is $\delta = 0.99$ and there is a fixed cost of entry, $\phi > 0$. Products are differentiated by their vintage ν , where ν equals the date at which a vintage is the frontier technology. Hence, 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 .

Each period, a mass of consumers enter the market. Consumers either buy 1 unit of a product, or choose the outside option. In both cases, consumers leave the market at the end of the period, and so there is no accumulation of consumers across periods. Consumers are differentiated by their budget for computer and related products, denoted y .

Consumers gain utility from purchasing a computer and from using the remainder of their income to purchase some alternative computer-related good such as software. Let \mathcal{V}_t denote the set of available computers in period t and let p_ν^t denote the time t price of vintage ν . We normalize the price of the outside good to zero and denote the full set of choices available to consumers as $\Gamma_t = \{\mathcal{V}_t, \check{u}_t\}$.

Following SS we assume the consumer's utility from purchasing the product ν is

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

where p^t is a vector of prices and u_ν represents the utility of purchasing vintage ν . The utility from purchasing the outside good is

$$\check{u}_t \cdot y. \tag{2}$$

The consumer's utility-maximization problem is:

$$\max \left\{ \max_{\nu \in \mathcal{V}_t} U(y, \nu; p^t), \check{u}_t \cdot y \right\}. \tag{3}$$

We make the natural assumption that newer vintages are preferred to older ones. The resulting demand function is straightforward.²⁷ We order the vintages by their utility levels and consider the neighboring vintages ν and $\nu + 1$. There is a consumer, with income \hat{y} , who is just indifferent between them,

$$u_\nu \cdot (\hat{y} - p_\nu^t) = u_{\nu+1} \cdot (\hat{y} - p_{\nu+1}^t). \quad (4)$$

All consumers with income less than \hat{y} prefer ν over $\nu + 1$ and all those with income more than \hat{y} prefer $\nu + 1$ over ν . Denote this marginal consumer $y_{\nu, \nu+1}$. 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 $\{y_{\nu-1, \nu}, y_{\nu, \nu+1}\}$ will purchase vintage ν . The demand for ν is then simply

$$Q_\nu = \int_{y_{\nu-1, \nu}}^{y_{\nu, \nu+1}} h(x) dx,$$

where h is the distribution of consumer's income. If ν is the best available product, demand is given by

$$Q_\nu = \int_{y_{\nu-1, \nu}}^{\infty} h(x) dx.$$

If ν is the lowest quality product, it competes directly with the outside option, and its demand is given by

$$Q_\nu = \int_{y_{\check{u}, \nu}}^{y_{\nu, \nu+1}} h(x) dx,$$

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

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

²⁷This is the demand system of Prescott and Visscher (1977), which has been well-studied in the vertical product differentiation literature.

3.1.2 Supply

To model supply, we first describe how firms choose their product's price given a set of competing computers. We then turn to the dynamic problem which entails specifying when a firm should introduce a new vintage into the market. We assume that each product is supplied by a distinct firm, and so ignore any joint maximization problem of a multi-product line firm. Thus, we can think of the model as characterizing firms competing over a specific product line, such as the 15-inch 512MB laptop computers depicted in Figure 5.

The firm's price-setting problem is to maximize profits in period t . We assume firms have a constant marginal cost $c \geq 0$ and no capacity constraints. Firms compete in prices, a la Bertrand. Given competing vintages' prices, denoted as $p_{-\nu}^t$, the firm's problem is

$$\pi_C(\nu; \Gamma_t, p_{-\nu}^t) = \max_{p_\nu^t} \left\{ (p_\nu^t - c) Q_\nu^t(p_\nu^t, p_{-\nu}^t; \Gamma_t) \right\}, \quad (5)$$

where Q_ν^t is the demand for product ν at time t .

While the firm's price-setting decision is static, its entry decision is necessarily dynamic. Within our vintage-capital framework, an exogenous growth rate in the technological frontier implies that the life-cycle of the firm's product evolves deterministically over time. When a firm enters the market, we assume it chooses the frontier technology and so offers the most preferred product that period. We assume that only one firm can enter the market in a period, and that the firm enters only if the expected profits of doing so exceeds the cost of entry.²⁸

Rather than physically depreciate, a computer faces two sources of obsolescence over time. First, the outside product is assumed to improve over time. With each successive period, then, a computer ν maintains its utility value to consumers while the outside option becomes more attractive. This *general* obsolescence places downward pressure on the price a firm can charge for ν . Second, with each successive period new computer vintages may enter. Newer vintages, embodying better technologies, compete with a

²⁸If two firms enter and offer the newest vintage, Bertrand competition drives the price to marginal cost. Without a markup over marginal cost, a firm cannot recover the fixed cost of adoption.

vintage ν and drive down its price. We label this second source of obsolescence *market-specific* obsolescence.

Both sources of obsolescence ensure 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 effectively, the product will have exited the market. The life-cycle of a computer, then, starts with its adoption into the market and ends when there is no longer demand for the computer at a price greater than marginal cost.

We assume that firms are rational and so can predict the deterministic evolution of Γ_t over time. Furthermore, firms are able to compute the vector of Nash equilibrium prices given the set of available products. The lifetime profits of entering the market at time t are then

$$\Pi_C(\nu = t) = \sum_{s=t}^{\infty} \delta^{s-t} \pi_C(\nu = t | \Gamma_s, p_{-\nu=t}^s). \quad (6)$$

We arrive at the usual condition: New products are only introduced if future discounted profits are greater than the fixed cost of entry, or

$$\Pi_C(t) > \phi. \quad (7)$$

3.1.3 Competitive equilibrium

A competitive equilibrium is a collection of functions over all dates $t \geq t_0$, of prices, \bar{p}^t , and products for sale, Γ_t , such that

1. the initial vector of available products is given by Γ_{t_0}
2. for all $t \geq t_0$, \bar{p}^t describes a Nash equilibrium in prices for the set of products contained in Γ_t .
3. for all $t \geq t_0$, firms enter the market only if the entry condition, equation 7, holds. Consequently, the evolution of Γ_t is given by

$$\Gamma_t = \begin{cases} \{\mathcal{V}_{t-1}, \check{u}_t\} & \text{if } \Pi_C(t) \leq \phi \\ \{\nu = t, \mathcal{V}_{t-1}, \check{u}_t\} & \text{if } \Pi_C(t) > \phi. \end{cases} \quad (8)$$

From SS, we know that given a set of products, there exists a Nash equilibrium in prices that is not, in general, unique. Consequently, a competitive equilibrium exists for our extension of their model, and it is not necessarily unique.

3.2 Empirical Fit

To fit the model to the data, we consider the stationary equilibrium. Given a vector of parameters, we compare the model's predictions of prices and sales of a computer over its product cycle to prices and sales observed in our data. We parameterize the competitive model to match the sales and prices of the average PC.

3.2.1 Description of Parameters

To obtain a stationary equilibrium, we assume the ratio of the utility associated with the product embodying the frontier technology over the utility provided by the outside good remains constant over time. Furthermore, in a stationary equilibrium, the entry pattern of new computers needs to follow a regular cycle. With these assumptions and given a fixed distribution of consumer income, the equilibrium number of products sold and the vector of Nash-equilibrium prices are time-invariant across entry pattern cycles.

The parameters of the model can be categorized into three groups. The first set of parameters determine the quality level of products relative to each other and to the outside good, the second set characterize the consumers' income distribution parameters, and the third set detail the cost structure of the firm. We use three parameters to characterize the products quality levels: (i) The level of the highest quality product, \bar{v} , which we normalize to 10, (ii) the growth of the frontier technology, γ , which, based on data measuring the growth rate of CPU, we set to 1.7 percent, and (iii) ratio of the outside good's to highest quality product's utility, ζ which we set to 0.1.

The substitutability of products across vintages is determined by γ . Raising γ enlarges the difference in utility associated with the newer vintage relative to the older vintage increases, decreasing substitutability across vintages. The attractiveness of the outside good relative to products is determined by ζ . As ζ approaches one from below, the outside good becomes more attractive relative to the frontier technology.

We assume consumers' budget for computers and related products is drawn from the Uniform distribution over the interval $[a, b]$. We normalize a to be one, but leave b free. The upper bound of income plays a large role in determining the total number of products the market is able to support (see SS). The density of consumers over $[a, b]$ is given by $\kappa > 0$.

Finally, we assume a simple cost structure for the firm. The firm pays a fixed cost, $\phi > 0$ to enter the market. Upon entry, the firm's production technology is constant marginal cost, $mc > 0$. The model provides an upper bound to the value of ϕ with its prediction of the product's present discount value of lifetime profits,. The firm's monthly discount rate is set to 0.99.

In summary, there are 3 free parameters in our model, $\theta = \{b, \kappa, mc\}$. Given these parameters and a fixed pattern of entry, the model yields predictions about price and sales for a product over its life-cycle in the stationary case. In the data, it is difficult to estimate the entry rate of new computers because it is not obvious what is the relevant market. For example, while we are confident that the introduction of a desktop would have little impact on a consumer looking at laptops, it is much harder to define an appropriate market within the set of laptops. As such, we consider four different entry patterns for our model: entry every period, every second period, every third period, and every fourth period. Given θ , we solve the model for each entry pattern and record the price and sales path of a product over its life cycle. We then take a weighted average of the four price and sales paths, using the weights $\omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$. Interpreting the weights as probabilities, this average is the expected price and sales paths of a computer over its life cycle.

3.2.2 Results

To compare the model to the data, we compute price declines and the sales cdf from the expected price and sales paths. In the data, we estimate these average price declines and sales cdf using fixed-effect regressions (see Section 2.1). To most closely align the model with the data, we find the pair $\{\theta, \omega\}$ that minimizes a least-distance criterion.²⁹

The parameters that best match the model to the data are displayed in table (5).

²⁹Details are provided in Appendix C

There is a wide distribution in budgets for computers and related products. Those consumers with the biggest budgets are willing to spend more than 63 times the amount of those consumers with the smallest budgets. Further, marginal cost is low enough that all consumers purchase a computer. The estimated weights across the four entry patterns place about equal weight across the 4 entry patterns.

As displayed in table (6), the model fits the data well along both the price and sales dimensions. The model matches the price declines in the data at the beginning of the product cycle, although it underestimates the decline in months 4 and 5. For months 6 through 8, the model predicts a tiny amount of sales. Overall, only about 5 percent of a computer's total expected sales will occur after month 5. Hence, while we report the price declines for these sales late in the computer's product cycle we do not attach much importance them. To provide a visual display of the model's fit to the data, in figure (6) we plot the price levels for the model and data fixing the price in period one to 100 for both cases.

Along the sales dimension, the model tracks the data well overall. The model matches the data by generating most of a computer's sales in the first two months of its product cycle. Further, as mentioned above, the model predicts that 95 percent of a computer's sales occur by month 5 of the product cycle, closely matching what we observe in the data. To better see how the model matches the shape of sales in the data, in figure (7) we plot the sales pdf for the model and the data. The figure shows how the model matches the decline in sales over the product cycle that we observe in the data.

These results demonstrate that the competitive model is able to closely match the price and sales patterns observed in the IBM platform PC data. In a goodness-of-fit measure, we compare the model's predictions on the timing of household's purchase decision. As described in section (2), the TUP survey data imply that households with higher incomes purchase computers earlier in the product cycle. Aizcorbe and Shapiro (2009) find that a 1.25 percent fall in price is correlated with a 10 percent fall in income. In line with these results, our model predicts that households with higher budgets for computer and related software purchase computers earlier in the product cycle. In table (7) we compute the average budget of consumers for a computer over the product cycle for each entry pattern. The average budget of consumers who purchase the highest-quality computer

is 34.3, while those consumers who purchase the second-highest quality computer have a budget of 2.7, a fall of over 90 percent. While the household income in the data is not equivalent to consumer's budget in the model, the model seems to over-predict the fall in consumers' budgets the product cycle. We believe this is mainly due to our assumption that consumers' budgets are uniformly distributed.³⁰ A more flexible distribution may preserve the model's fit to the price and sales moments, while generating a more gentle decline in consumers' budgets over the product cycle.

Finally, the model implies markups that are quite reasonable. Specifically, when a computer is the newest vintage available, manufacturers charge a markup of almost 20 percent. Once a newer vintage enters the market, however, this markup falls dramatically to about 1 percent.

3.3 The Effect of Product Entry on Price and Sales Dynamics

The model gives clear insight as to how and why product entry acts to shorten the product's lifespan. Interestingly, the underlying intuition is very similar to the reasoning behind the role of competition and innovation originally posed by Arrow (1962) and more currently in Aghion et. al (2001, 2005). Aghion et. al show that what is important for innovation incentives for incumbent firms is the size of the difference between post- and pre-innovation rents—the incremental profit from innovating relative to not innovating.³¹ A larger difference between post- and pre-innovation rents implies a larger incentive to undergo the cost of innovation.

It is quite straightforward to make the analogy between the concepts of pre- and post-innovation rents with the model described above. Post-innovation rent is defined as the profit accrued between the introduction of an innovation and its subsequent diffusion—in our model we can think of this as the profit made when the firm has sole possession of the highest quality product. Pre-innovation rent is the profit earned before innovation—in our model, we can think of this as the profits earned in periods subsequent to its

³⁰In the model, the distribution of consumers' budgets is uniform over $[1, 63.2]$ with a density of 0.065.

³¹This is a different approach to the original Schumpeter idea which was looking at the incentives for *outsider* firms—firms just entering the market—where only the level of the post-innovation is important, since the pre-innovation rent is always zero.

quality leader position.

To show the effect of product entry on price and sales dynamics, we compare the model’s predictions across the four entry patterns used to parameterize the model above. Specifically, in figures (8) and (9) we plot prices and sales over the product cycle for each case, alongside the average across all cases.³² We can think of more rapid entry as either being due to faster introduction of the underlying components by the intermediate-goods manufacturer or being due to perhaps a specific product line being more conducive to quick adoption given a new component.

Turning first to prices, we see that the introduction price is rising across the 4 cases. This pattern reflects the level of substitutability across vintages. When entry occurs less frequently, there are large differences in quality among computers on the market, decreasing their substitutability. A lesser degree of competition therefore allows the manufacturer to charge a higher initial (i.e. post-innovation) price for its new product. Comparing the four entry cases reveals two basic levels of sales. When a product is the newest vintage available (i.e. accruing its post-innovation rent), it can commandeer sales around 3.8, which is over 90 percent of the market.

This high market share as well as the high price that the firm can charge as the quality leader is the Schumpeterian effect. As entry occurs more frequently, the Schumpeterian effect diminishes and the “Arrow effect” becomes important. With entry every period, for instance, the firm sets a relatively low initial price but in the subsequent period, specific obsolescence induced by product entry pushes both prices and sales down, and eventually profits reach zero by the third period. The figure makes clear that, in all cases, product entry causes an immediate increase in the difference between post- and pre-innovation rents. Eventually, this forces the firm to take the product off the market when pre-innovation rents are low enough that profits reach zero. In contrast to the active role of specific obsolescence, general obsolescence, captured by the growth rate of the outside good, plays at most a small role in the model’s predictions of prices and sales.

³²The “Average” line in figures (8) and (9) is the average across all four cases, using our estimated weights. These averages are the expected price and sales paths over the product cycle. Using this expected price and sales figures, we compute the price declines and sales cdf figures used to fit the model to the data.

4 The Monopolistic Industry

The monopolist's problem is distinct from the competitive case. Whereas the competitive firm faces product entry by other firms the monopolist need only worry about general obsolescence stemming from growth in quality of the outside good. In the competitive firm's pricing and adoption decision described above we considered the case in which each firm offered one product within a product line. The analogous case in the monopoly market is where the monopolist sells only one product at a time. Furthermore, we add an endogenous adoption component to the model by allowing the monopolist to decide when to introduce a new vintage. To simplify the model, we assume the monopolist phases out production of the existing vintage when it introduces a new vintage into the market. Introducing a new product, then, is equivalent to updating the existing product with the latest technology.

4.1 The Monopolist's Problem

Given a vintage ν_t , a monopolist solves a static problem to maximize profits,

$$\pi_M(\nu_t; \check{u}_t) = \max_{p_\nu^t} \{(p_\nu^t - c)Q^t(p^t; \nu_t, \check{u}_t)\} \quad (9)$$

Demand for vintage ν_t is given by Q^t , which accounts for the price of the computer as well as the outside option.

The monopolist's entry problem balances the gains from introducing a computer in the current period, versus waiting a period (or more) before introducing a new product. 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 . The state space of the monopolist in period t is its existing product ν_s and outside option, \check{u}_t . The monopolist's entry problem is then

$$V(\nu_s, \check{u}_t) = \max \{\pi_M(\nu_s, \check{u}_t) + \delta V(\nu_s, \check{u}_{t+1}), \pi_M(\nu_t, \check{u}_t) - \phi + \delta V(\nu_t, \check{u}_{t+1})\}. \quad (10)$$

The first term above are the profits from delaying introduction. Here the monopolist continues to sell its existing computer ν_s and enters the next period with that same vintage, competing against an improved outside option. The second term considers the case

where the monopolist introduces a new computer vintage. Consumers are willing to pay more for the newer vintage, $\pi_M(\nu_t, \check{u}_t) > \pi_M(\nu_s, \check{u}_t)$, and the monopolist's continuation value is higher, $V(\nu_t, \check{u}_{t+1}) > V(\nu_s, \check{u}_{t+1})$. But entry is expensive because it requires the absorption of a fixed cost ϕ .

4.2 Out-of-Sample Prediction

We now consider the monopolist's stationary case. Using the same θ as in the competitive case, we can solve for the monopolist's optimal pricing strategies. Formally, by keeping preferences and technology fixed and changing market structure, we are measuring how much changes in the competitive structure of the market can explain differences between the PC and Apple retail markets. As in the competitive setting, choosing price is a straightforward static problem. For the monopolist, the tradeoff to raising its price is losing customers to the outside option. The frequency of replacement in the monopoly setting, however, is a dynamic problem directly linked to the fixed cost of entry, ϕ . *Ceteris paribus*, the larger the adoption cost, the longer the monopolist will keep a given product on the market. From the competitive case, we only have an upper bound on ϕ .³³ Given this constraint, we choose a fixed cost of entry such that the model's prediction on the monopolist's timing of replacement matches the data.

We illustrate the monopolist's replacement problem in figures (12) and (13). Because the outside option grows over time, the attractiveness of the monopolist's product degrades over time. As a result, the price charged and per-period profits earned by the monopolist decline with the vintage of the computer (see figure (12)). Because per-period profits fall as the product ages, the monopolist wants to update its computer sooner rather than later. Updating, or replacing, its computer, however, is costly because of the fixed cost of replacement. The larger the fixed cost, the less frequently the monopolist updates its computer. In figure 13, we plot the optimal replacement cycle length against the fixed cost of replacement. With a fixed cost of replacement of 0.03,

³³Using a discount rate equal to 0.99, we use our model's results to compute a product's lifetime profits for each entry pattern in the competitive case. We then use the estimated weights to obtain expected lifetime products from entry, $\hat{\pi}$. To ensure that a manufacturer is willing to introduce a new product, it must be that the fixed cost of entry is less than expected profits, or $\phi < \hat{\pi}$.

the monopolist chooses to update its computer every 2 periods. If the fixed cost of replacement is 0.114, the monopolist chooses to update its computer every 13 periods.

In the data, we infer that Apple’s computers are sold for 8 months on average. For the model to generate the 8 month replacement cycle in a monopolist setting, we set ϕ to be 0.35, or 18 percent of the expected discounted profits of a firm in the competitive market. Thus, the parameter value chosen for the cost of adoption is consistent with the results from the competitive case presented earlier.

Given θ and ϕ , the model predicts a flat price profile over the product cycle (see table 8), a close match to the prices of Apple computers seen in the data. Specifically, Apple’s price declines are negligible for three-quarters of the product cycle, before being heavily discounted. The reason the data show that Apple’s average prices fall so dramatically after month six is due to the occasional product replacement. When Apple introduces a new vintage, the stock of existing computer is heavily discounted.³⁴ Since our model does not account for such stock-out sales, it is unable to match the sharp falls seen later in the product’s life.³⁵ In figure (10) we plot the price levels predicted by the model and observed in the data for a typical computer over the product cycle.

The model also matches well the steady flow of sales throughout the product cycle. As seen in table (8), the sales cdf generated by the model closely follows what we observe in the data, although the model does under predict the amount of sales at the beginning of the product cycle. This can be more clearly seen in figure (11), where the sales pdf for the model and data are plotted next to one another. In the data, we see a decline in sales over the product cycle, unlike in the model. The decline in monthly sales volume, however, is actually quite small. In its first month, sales are almost 17 percent of total sales over the product cycle. In the sixth month of the product cycle, sales are still

³⁴These pricing dynamics can be seen in Appendix C, which depicts the price contours of four of Apple’s notebook product lines: 12-inch PowerBooks, 14-inch iBooks, 15-inch PowerBooks, and 17-inch PowerBooks. Prior to replacement the average price fall (where we average over number of periods) of these four product lines is 0.1, 0.02, 0.7, and 0.4 percent, respectively. However, the average price fall in the period directly after replacement (where we average over number of models) are 5.4, 6.5, 9.6, and 5.2 percent, respectively.

³⁵Because the arrival of a replacement product can only be found by visually looking at distinct product lines and using detailed knowledge of the industry, we found no systematic way to correct for stock-out sales through-out the entire dataset.

13 percent of total sales, only a small decline. Admittedly, in the last two months of the product cycle, sales decline in the data but not in our model. As mentioned above, however, we believe this behavior reflects fire-sales by Apple because of the introduction of a replacement computer; we do not account for this behavior in our vintage capital model.

In stark contrast to the competitive case, in the monopoly setting the average budget of consumers purchasing a computer at different vintages hardly changes (see table 9). Further, the average budget of a household which purchases a computer in the monopolistic setting is higher than in the competitive setting. Under the assumption that consumers' budgets for computer and related products is positively correlated with household income, our model matches the income data from TUP survey data: the distribution of consumers' income in the monopoly setting has a higher mean and a smaller variance relative to the distribution of consumers' income in the competitive setting.

5 Concluding Remarks

Our key contribution in this study is assessing the role that market structure plays in the length of the product cycle, and hence the frequency of new product adoption. The model's close fit with the PC and Apple data indicate that the slower technological adoption performed by the Apple is in fact due to the competitive structure of the market. Because the monopolist has no close substitutes, it only worries about general obsolescence and decides to keep its product on the market for a long period of time to avoid paying the fixed entry cost ϕ . Importantly, when the update occurs the monopolist adopts the latest technology and so there is a large leap forward in technology between the new computer and the one it is replacing.

In contrast, in the competitive industry, manufacturers are forced to use more frequent adoption. The reason for such a discrepancy between the two market structures is the role played by specific obsolescence which is absent in the monopolistic industry. As products from competitors are close substitutes to its own, the incumbent firm loses a great deal of markup as well as market share upon competitive product entry. Specific obsolescence induced by entry therefore pushes per period profits to zero, forcing the firm

to take the product off of the market. Thus, even though the new vintage is only slightly better than the existing product, the advantage of being in the leadership position is enough to capture large market share. It is therefore profitable for new products to be introduced frequently, even though each introduction requires paying a large fixed entry cost.

An interesting extension of our analysis is to link it with Coase's conjecture which explains that the monopoly cannot sustain its markup when consumers are forward looking. As this would imply that the monopolist must keep its price at marginal cost indefinitely, one might argue that Apple is falling into this trap. Our data and model provide a great deal of evidence to suggest that intertemporal price discrimination is not the main driver of the difference between Apple's and PCs' adoption and pricing dynamics. First off, that Apple charges higher prices for similar underlying components suggests that it is charging a higher markup than the PC manufacturers—that is, it is charging a price higher than marginal cost. Second, the ability to intertemporally price discriminate requires market power, which means that if intertemporal price discrimination was a main driver for the large price falls in the PC industry, we would more likely see them in the Apple market rather than the PC market. Finally, even with market power, Stokey (1979) shows that intertemporal price discrimination may not be profitable unless certain conditions are met—reservation prices must be correlated with time preference. Thus, if the price falls in PC market are due to intertemporal price discrimination, then it must be the case that for some particular reason Stokey's condition is met in the PC market but not the Apple market.

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., 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.

- [3] Aghion, P. and Howitt, P., 1992. "A Model of Growth through Creative Destruction," *Econometrica*, Econometric Society, vol. 60(2), pages 323-51, March.
- [4] Aghion, P., Blundell R., Griffith R., and 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.
- [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. 2009. "Are Technology-Goods Price Measures Falling Too Fast?" BEA Working Paper.
- [8] 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.
- [9] Biesebroeck, J.V. and Hasmi Aamir. 2009 "Market Structure and Innovation: A Dynamic Analysis of the Global Automotive Industry," Katholieke Universiteit manuscript.
- [10] Borenstein, S. and Rose, N., 1994. "Competition and Price Dispersion in the U.S. Airline Industry," *Journal of Political Economy*, University of Chicago Press, vol. 102(4), pages 653-83, August.
- [11] Bresnahan, T., 1987. "Competition and Collusion in the American Automobile Industry: The 1955 Price War," *Journal of Industrial Economics*, Blackwell Publishing, vol. 35(4), pages 457-82, June.

- [12] Bresnahan, T. and Greenstein, S., 1999. "Technological Competition and the Structure of the Computer Industry," *Journal of Industrial Economics*, Blackwell Publishing, vol. 47(1), pages 1-40, March.
- [13] 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.
- [14] Conlisk, J., Gerstner, E., Sobel, J. 1984. "Cyclic Pricing by a Durable Goods Monopolist" *The Quarterly Journal of Economics*, The MIT Press, vol. 99(3), pages 489-505, August.
- [15] Dasgupta, P. and J Stiglitz. 1980. "Industrial Structure and the Nature of Innovative Activity." *Economic Journal* 106 (July): 925-51.
- [16] 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.
- [17] 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.
- [18] Gilbert, R.J. and D.M.G Newbery. 1982. "Preemptive Patenting and the Persistence of Monopoly." *American Economic Review* 72 (June): 514-27.
- [19] Goettler, R. and Gordon, B., 2009 "Does AMD spur Intel to Innovate More?" University of Chicago Booth School of Business manuscript.
- [20] Gowrisankaran, G. and Rysman, E. 2009. "Dynamics of Consumer Demand for New Durable Goods", NBER Working Paper w14737. February.
- [21] Pakes, A. 2003. "A Reconsideration of Hedonic Price Indices with an Application to PC's" *American Economic Review*, American Economic Association, vol.93(5) pages 1578-1596.

- [22] 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.
- [23] Shaked, A. and Sutton, J., 1983. "Natural Oligopolies," *Econometrica*, Econometric Society, vol. 51(5), pages 1469-83, September.
- [24] 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.
- [25] Stavins, J., 1995. "Model Entry and Exit in a Differentiated-Product Industry: The Personal Computer Market," *The Review of Economics and Statistics*, MIT Press, vol. 77(4), pages 571-84, November.
- [26] Stokey, N., 1979. "Intertemporal Price Discrimination," *The Quarterly Journal of Economics*, MIT Press, vol. 93(3), pages 355-71, August.

Appendix

A Regression Analysis

A.1 Price Regression

To calculate price changes over the product cycle seen in Figure 3 we run the following regression

$$\ln Price_{iv} = \alpha + \sum_{v=2}^{31} \beta_v D_v + \gamma_i + \varepsilon_{iv} \quad (11)$$

for each manufacturer in the sample, where $\ln Price_{iv}$ is the logarithm of the price of model number i with time on the market v , D_v is a dummy variable indicating the amount of months the product has been on the market and γ_i are model-number fixed effects. We set the maximum time on market equal to 31 months, the longest any model is in the market in the sample. While for expositional purposes Figure 3 reports estimates of β_v for $v \in [2, 6]$, we provide a more detailed report of the estimates for $v \in [2, 12]$, along with standard errors in the table below. Because there is heterogeneity in the length of the life of any given computer model, stock-outs may be influencing prices at different times on the market. For robustness purposes, we therefore report estimates using two different units-CDF cutoff rules: the upper panel uses no right cutoff rule while the lower panel uses a 90 percent CDF cutoff rule (i.e. observations after the units CDF reached 90 percent were omitted from the regression). The latter is attempting to remove observations where stock-outs are likely to occur.

No Units CDF Cutoff

	PCs	Apple	Compaq	HP	Sony	Toshiba	Emachines
β_2	-0.016 (0.002)**	0.002 -0.002	-0.024 (0.003)**	-0.02 (0.005)**	-0.012 (0.003)**	-0.015 (0.003)**	0.007 -0.011
β_3	-0.048 (0.002)**	-0.003 -0.002	-0.058 (0.005)**	-0.05 (0.004)**	-0.048 (0.005)**	-0.056 (0.005)**	-0.017 -0.01
β_4	-0.1 (0.004)**	-0.004 -0.005	-0.112 (0.007)**	-0.098 (0.007)**	-0.096 (0.007)**	-0.111 (0.009)**	-0.08 (0.015)**
β_5	-0.181 (0.007)**	-0.02 (0.007)**	-0.179 (0.013)**	-0.172 (0.009)**	-0.178 (0.011)**	-0.186 (0.013)**	-0.207 (0.037)**
β_6	-0.245 (0.006)**	-0.039 (0.008)**	-0.244 (0.011)**	-0.239 (0.012)**	-0.246 (0.012)**	-0.285 (0.019)**	-0.224 (0.023)**
β_7	-0.316 (0.007)**	-0.102 (0.015)**	-0.322 (0.013)**	-0.289 (0.013)**	-0.315 (0.013)**	-0.393 (0.024)**	-0.3 (0.024)**
β_8	-0.363 (0.008)**	-0.2 (0.019)**	-0.386 (0.015)**	-0.351 (0.016)**	-0.341 (0.012)**	-0.427 (0.030)**	-0.306 (0.031)**
β_9	-0.409 (0.009)**	-0.263 (0.021)**	-0.427 (0.015)**	-0.388 (0.018)**	-0.398 (0.017)**	-0.481 (0.035)**	-0.363 (0.033)**
β_{10}	-0.478 (0.012)**	-0.291 (0.024)**	-0.522 (0.022)**	-0.454 (0.021)**	-0.435 (0.019)**	-0.569 (0.042)**	-0.416 (0.042)**
β_{11}	-0.539 (0.014)**	-0.32 (0.025)**	-0.553 (0.027)**	-0.554 (0.027)**	-0.486 (0.021)**	-0.538 (0.036)**	-0.572 (0.070)**
β_{12}	-0.602 (0.018)**	-0.397 (0.028)**	-0.627 (0.032)**	-0.549 (0.020)**	-0.576 (0.026)**	-0.723 (0.089)**	-0.604 (0.078)**
Observations	11438	1121	2986	3683	2207	1349	1213
R-squared	0.44	0.71	0.43	0.48	0.67	0.33	0.42

90 Percent CDF Cutoff

	PC	Apple	Compaq	HP	Sony	Toshiba	Emachines
β_2	-0.018 (0.002)**	0.002 -0.002	-0.024 (0.003)**	-0.023 (0.005)**	-0.013 (0.003)**	-0.014 (0.003)**	0 -0.005
β_3	-0.046 (0.002)**	-0.003 -0.002	-0.055 (0.005)**	-0.048 (0.005)**	-0.044 (0.004)**	-0.053 (0.005)**	-0.008 -0.007
β_4	-0.086 (0.004)**	-0.001 -0.004	-0.107 (0.009)**	-0.075 (0.009)**	-0.08 (0.006)**	-0.092 (0.010)**	-0.057 (0.017)**
β_5	-0.125 (0.006)**	-0.015 (0.004)**	-0.144 (0.014)**	-0.107 (0.011)**	-0.13 (0.011)**	-0.127 (0.015)**	-0.123 (0.033)**
β_6	-0.174 (0.011)**	-0.034 (0.008)**	-0.166 (0.019)**	-0.149 (0.016)**	-0.227 (0.026)**	-0.168 (0.024)**	-0.166 (0.010)**
β_7	-0.233 (0.017)**	-0.06 (0.012)**	-0.236 (0.026)**	-0.213 (0.028)**	-0.274 (0.037)**	-0.215 (0.053)**	
β_8	-0.333 (0.035)**	-0.107 (0.021)**	-0.257 (0.035)**	-0.39 (0.093)**	-0.402 (0.035)**	-0.248 (0.106)*	
β_9	-0.408 (0.051)**	-0.102 (0.028)**	-0.332 (0.048)**	-0.501 (0.116)**	-0.501 (0.091)**	-0.27 (0.082)**	
β_{10}	-0.439 (0.074)**	-0.124 (0.029)**	-0.185 (0.017)**	-0.674 (0.124)**	-0.565 (0.058)**	-0.282 (0.033)**	
β_{11}	-0.434 (0.084)**	-0.046 (0.007)**	-0.234 (0.046)**	-0.81 (0.033)**	-0.568 (0.078)**	-0.29 (0.036)**	
β_{12}	-0.495 (0.076)**	-0.054 (0.007)**	-0.361 (0.015)**	-0.809 (0.033)**	-0.58 (0.041)**	-0.323 (0.043)**	
Observations	3036	436	799	1037	556	374	270
R-squared	0.57	0.43	0.6	0.49	0.82	0.72	0.35

A.2 Unit Sales Regression

To calculate the cumulative density function (CDF) of units sold over the product cycle seen in Figure 1 we run the following regression for each manufacturer in the sample:

$$q_{iv} = \alpha + \sum_{v=2}^{31} \beta_v D_v + \varepsilon_{iv} \quad (12)$$

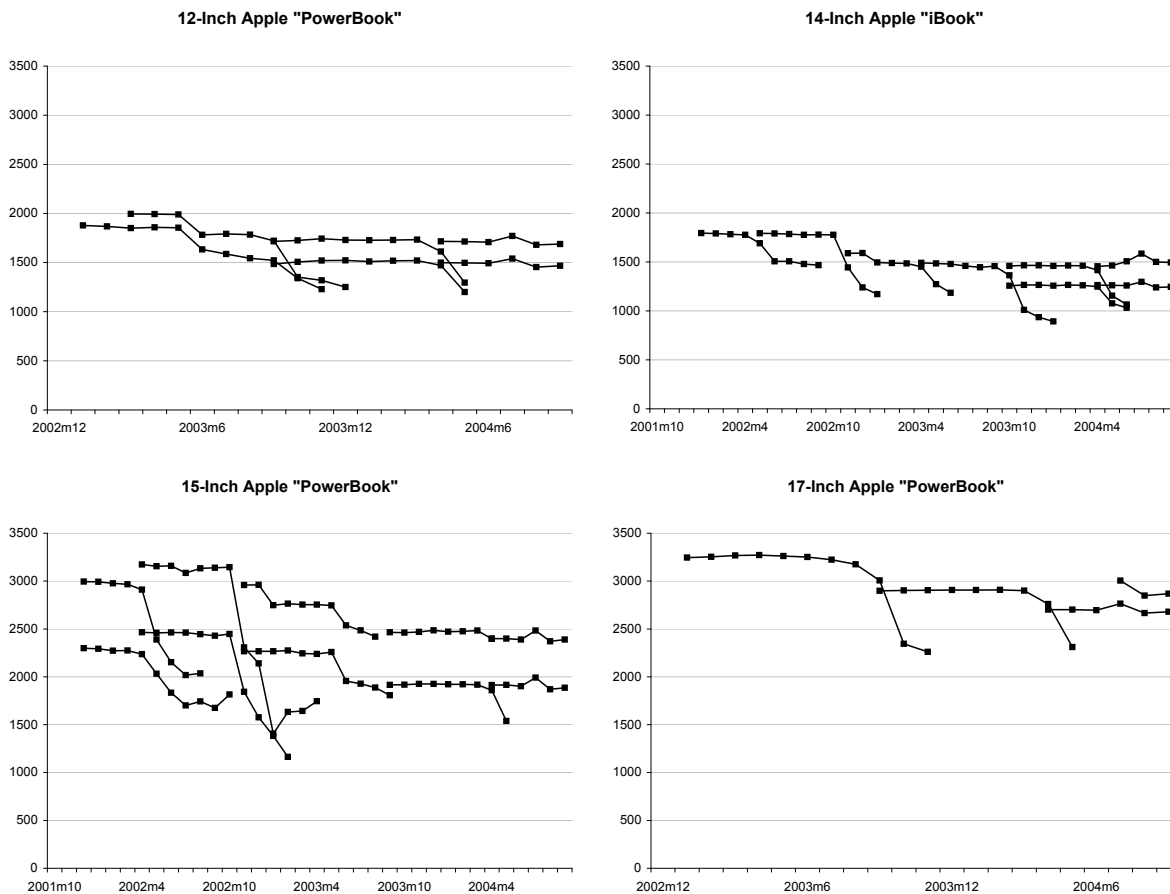
where q_{iv} is the number of units sold of model-number i with time on the market v . D_v is a dummy variable indicating the amount of months the product has been on the market. We then generate the predicted CDF for each manufacturer as follows. First, we calculate predicted unit sales as $\hat{q}_1 = \hat{\alpha}$, and $\hat{q}_v = \hat{\alpha} + \hat{\beta}_v$ for $v \in [2, 31]$, where 31 is the longest age of any computer in the sample. We then calculate the probability density function (PDF) at each time on the market point as

$$PDF_v = \frac{\hat{q}_v}{\sum_{v=1}^{31} \hat{q}_v} \quad (13)$$

The CDF can then be trivially calculated from the PDF as $CDF_v = \sum_{v=1}^v PDF_v$. We repeat this exercise for each manufacturer in the sample. We also ran this exercise for the sample of all “PCs” (all manufacturers except Apple) which was used to match the moments for sales in the model for the competitive industry. Figure 1 depicts the CDF for $v \in [1, 11]$.

B Apple Notebook Computer Prices

Apple Notebook Computers



The above figure shows prices of four different product lines of Apple notebook computers under a 98 percent CDF cutoff rule. Note that the bottom left panel depicts the same computer models as the upper right panel of Figure 4, however, with a larger units CDF cutoff rule (98 percent instead of 90 percent). The above figure shows that Apple generally offers two versions of each product line; for example, Apple offers the 15-inch PowerBook with either 512MB (higher priced) or 256MB RAM (lower priced). The figure most importantly demonstrates that when Apple introduces a new notebook computer of either RAM size, the model it replaces sharply falls in price. Specifically,

prior to replacement the average price fall (where we average over number of periods) of these four product lines is 0.1, 0.02, 0.7, and 0.4 percent for 12-, 14-, 15- and 17-inch notebooks, respectively. However, the average price fall in the period directly after replacement (where we average over number of models) are 5.4, 6.5, 9.6, and 5.2 percent, respectively.

C Algorithm for Solving the Vintage Capital Model

C.1 Demand side

To solve the demand side of the model in a period t , we need to know the price, p_ν^t , and associated utility, u_ν of each available vintage ν . 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 lowest 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 uniform density of consumers over budgets to compute demand for each product.

C.2 Supply side: competitive case

On the supply side for the competitive case, we solve four different versions of the model, reflecting the four different entry patterns. Different entry patterns influence the degree of substitutability between competing products. Suppose the model's parameters result in a market that can support 3 vintages being sold. In a stationary equilibrium where there is entry every period, we see three vintages being sold where the utility levels of the three products are $\{u, u * 1.017, u * 1.017^2\}$ for some constant u .³⁶ When there is entry every other period, the utility levels of the three products are $\{u, u * 1.017^2, u * 1.017^4\}$. In this second case, the utility levels are farther apart, implying a less substitutability between vintages.

To find the price and sales over the product cycle for the case where there entry every period, we only need to solve the model once. To see this, consider the vector of utilities for each vintage and the outside good in period t ,

$$\bar{u}_t = \begin{vmatrix} u * 1.1017^2 \\ u * 1.017 \\ u \\ \hat{u} \end{vmatrix}. \tag{14}$$

³⁶The assumed monthly growth in quality is 1.017.

In the following period, a new vintage is introduced, the outside good improves in quality and the oldest vintage exits. The resulting utility vector is

$$\bar{u}_{t+1} = \begin{vmatrix} u * 1.1017^3 \\ u * 1.017^2 \\ u * 1.017 \\ \hat{u} * 1.017 \end{vmatrix}. \quad (15)$$

Comparing these two vectors, we see that $\bar{u}_{t+1} = 1.017 * \bar{u}_t$. Since only relative utilities matter in the model, the Nash equilibrium prices and sales are the same across both examples. Consequently, a prices over a product's life cycle are equal to prices in the cross-section.

Given entry every other period, we need to solve the model twice. First we solve the model for

$$\bar{u}_t = \begin{vmatrix} u * 1.1017^4 \\ u * 1.017^2 \\ u \\ \hat{u} \end{vmatrix}. \quad (16)$$

We then solve for the case where there is no entry, but the utility of the outside good improves,

$$\bar{u}_{t+1} = \begin{vmatrix} u * 1.1017^4 \\ u * 1.017^2 \\ u \\ \hat{u} * 1.017 \end{vmatrix}. \quad (17)$$

Then in the following period, $t + 2$, when a new vintage is introduced, the resulting utility vector is a multiple of \bar{u}_t . Consequently, prices and sales in $t + 2$ are the same as those in t although the identities of the best, second-best, and third-best vintages have changed.

Following this logic, we need to solve the model 3 times when new computers enter every third period, and 4 times when vintages are introduced every fourth period.

Whichever entry case we are considering, we solve for prices in the same way. Given static demand, the firm's pricing problem is also static. From Shaker and Sutton (199?),

we know there exists a unique Nash-equilibrium in prices. We find this equilibrium through an iterative technique. We begin with each firm pricing at marginal cost. We then start with the manufacturer offering the newest vintage, and find the price which 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.

C.3 Supply side: monopoly case

The monopolist makes two decisions to maximize profits. It sets price and decides on the profit-maximizing replacement strategy.³⁷ The pricing decision is static and straightforward to solve since there is no strategic behavior. The tradeoff confronting the monopolist is that higher prices push marginal consumers to purchase the outside option. The 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 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 replacement strategy with the highest present discounted profits as the optimal strategy.

The fixed cost of entry is the 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 fixed cost such that the monopolist's optimal replacement strategy matches the data.

³⁷Note, we do not allow the single-product monopolist to sell two vintages at the same time.

C.4 Least distance criterion

For the competitive case, we use a least distance criterion to find the value of the free parameters and entry pattern weights to most closely align the model with the data. The free and fixed parameters are listed in table 5. Computing this criterion involves several steps:

1. Choose a vector of parameters θ and weights ω .
2. Solve the competitive case of the model for each the 4 entry patterns, recording the price and sales of a computer over its product cycle.
3. Use ω to compute weighted averages of price and sales over the product cycle.
4. Use the average price series to compute the decline in price over the product cycle, and the average sale series to compute the sales cdf. In the data, we construct similar measures of the price declines and sale cdf for the average computer using a fixed effect regression. To get the price declines and sale cdf on the same scale, we divide through both the model and data price decline numbers by the largest price decline seen in the data.
5. Let $g(\bar{p}), g(\bar{s})$ denote the difference between the model and data price declines and sales cdf, respectively. Denoting $G = \begin{vmatrix} g(\bar{p}) \\ g(\bar{s}) \end{vmatrix}$, our criterion is:

$$G' W G \tag{18}$$

where W is a matrix with zeros everywhere except the diagonal. The criterion weights interacting with $g(\bar{p})$ are the appropriate elements from the sales pdf. Hence price declines for periods with more sales get a higher weight (i.e. price declines at the beginning of the product cycle are weighted more than price declines at the end of the price decline). The criterion weights interacting with $g(\bar{s})$ are equal to 1.

Table 1: Market Share in NPD Sample

	All	Desktops	Notebooks
Apple	0.04	0.03	0.06
Compaq	0.23	0.24	0.21
Emachines	0.17	0.25	0.02
HP	0.35	0.40	0.25
Sony	0.11	0.08	0.15
Toshiba	0.11	0	0.32

Notes: Marketshares are based on units sold in each of three samples: all computers, desktop computers, and notebook computers.

Table 2: Adoption of New Models

	New Models per Month (Avg)	Fraction of Months with no New Models
All		
Apple	1.94	0.35
Compaq	6.50	0.06
Emachines	2.97	0.12
Hewlett Packard	9.15	0.03
Sony	4.09	0.15
Toshiba	2.82	0.18
Desktops		
Apple	0.97	0.59
Compaq	3.65	0.24
Emachines	2.56	0.15
Hewlett Packard	5.76	0.03
Sony	2.09	0.24
Notebooks		
Apple	0.97	0.62
Compaq	2.85	0.12
Hewlett Packard	3.38	0.21
Sony	2.00	0.24
Toshiba	2.82	0.18

[h] Notes: This table reports the average number of new models per month and the fraction of months with no new model for each manufacturer in the NPD sample. The exercise is repeated for the sample of desktop computers and the sample of notebook computers.

Table 3: Frontier CPU Clock Speed

	Fraction of Months at Frontier	Fraction of Months as Sole Frontier Leader
All		
Apple	0	0
Compaq	0.11	0
Emachines	0	0
Hewlett Packard	0.83	0.46
Sony	0.54	0.17
Toshiba	0.03	0
Desktops		
Apple	0	0
Compaq	0.11	0
Emachines	0	0
Hewlett Packard	0.83	0.46
Sony	0.54	0.17
Notebooks		
Apple	0	0
Compaq	0.23	0.06
Hewlett Packard	0.34	0
Sony	0.17	0.03
Toshiba	0.77	0.46

Notes: The first column of this table reports the fraction of months at the manufacturer was at the frontier CPU clock speed in the NPD sample. The second column reports the fraction of months the manufacturer held the frontier CPU clock speed by itself. The exercise is repeated for the sample of desktop computers and the sample of notebook computers.

Table 4: 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 income and gini-coefficient of income for each manufacturer in the TUP survey data.

Table 5: Model Parameters

Product quality			
Frontier value	$\bar{\nu}$	fixed	10
Quality growth rate	γ	fixed	1.017
Utility ratio	ζ	fixed	0.1
Income distribution			
Lower bound	a	fixed	1
Upper bound	b	flexible	63.160
Density of consumers	κ	flexible	0.065
Cost			
marginal cost	mc	flexible	0.827
Entry pattern weights			
entry every period	ω_1	flexible	1.010
entry every 2 periods	ω_2	flexible	1.012
entry every 3 periods	ω_3	flexible	1.020
entry every 4 periods	ω_4	fixed	1

Table 6: Time-series of price and sales, data and models
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	-1.6	-4.8	-10.1	-18.1	—	—	—	-7.2
Model	-1.6	-4.7	-7.1	-15.3	15.3	14.8	14.8	-6.0

Note: "Average" is the average price decline computed using sales weights.

Sales cdf

<i>Month</i>	1	2	3	4	5	6	7	8
Data	0.314	0.559	0.793	0.950	1	—	—	—
Model	0.352	0.621	0.802	0.931	0.945	0.959	0.979	1

Table 7: Average Budget of Consumers Over the Product Cycle

Product Cycle	1	2	3	4	5	6	7	8
Entry every period	34.3	2.7	1	-	-	-	-	-
Entry every 2 periods	34.3	34.3	2.7	2.7	1	1	-	-
Entry every 3 periods	34.3	34.3	34.3	2.7	2.7	2.7	-	-
Entry every 4 periods	34.3	34.3	34.3	34.3	2.7	2.7	2.7	2.7

Table 8: Monopolist: Price Declines over the Life-cycle
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	-0.3	-0.6	-1.8	-4.0	-10.7	-20.3	-4.0
Model	-0.0002	-0.0003	-0.0005	-0.0007	-0.0009	-0.0011	-0.0012	-0.0061

Note: "Average" is the average price decline computed using sales weights


Sales cdf

<i>Month</i>	1	2	3	4	5	6	7	8
Data	0.165	0.309	0.446	0.584	0.730	0.856	0.940	1
Model	0.125	0.250	0.375	0.500	0.625	0.750	0.875	1

Table 9: Monopolist: Average Budget of Consumers Over the Product Cycle

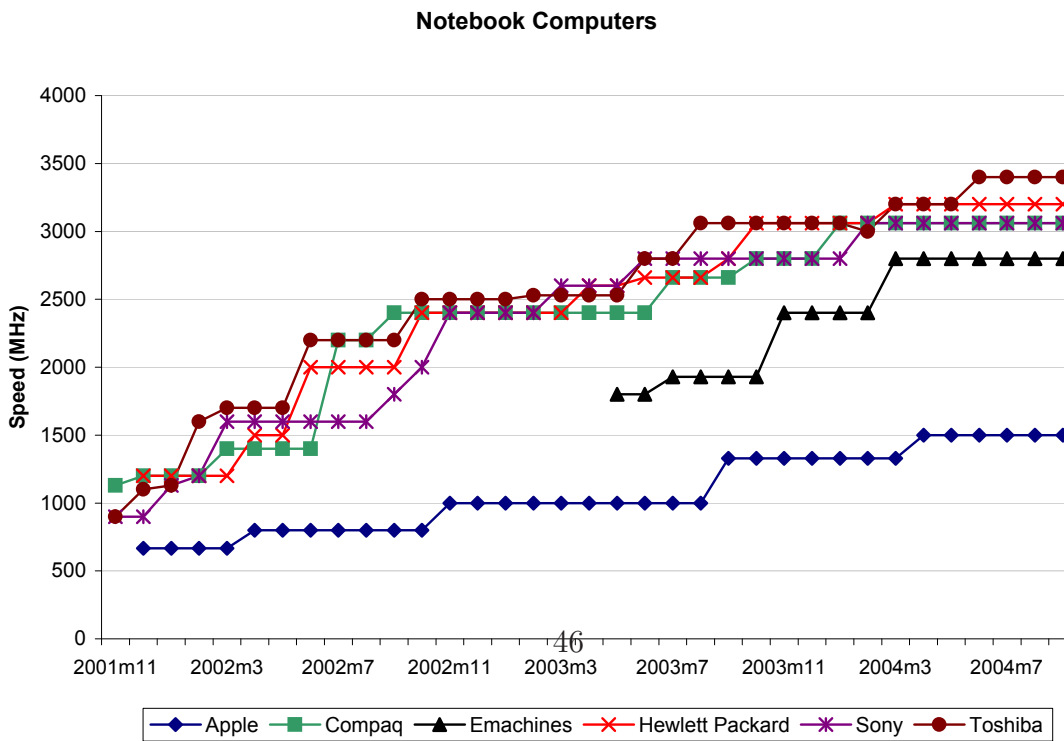

Product Cycle	1	2	3	4	5	6	7	8
	32.5	32.5	32.5	32.5	32.5	32.5	32.5	32.5

Figure 1: CDF of Units Sold



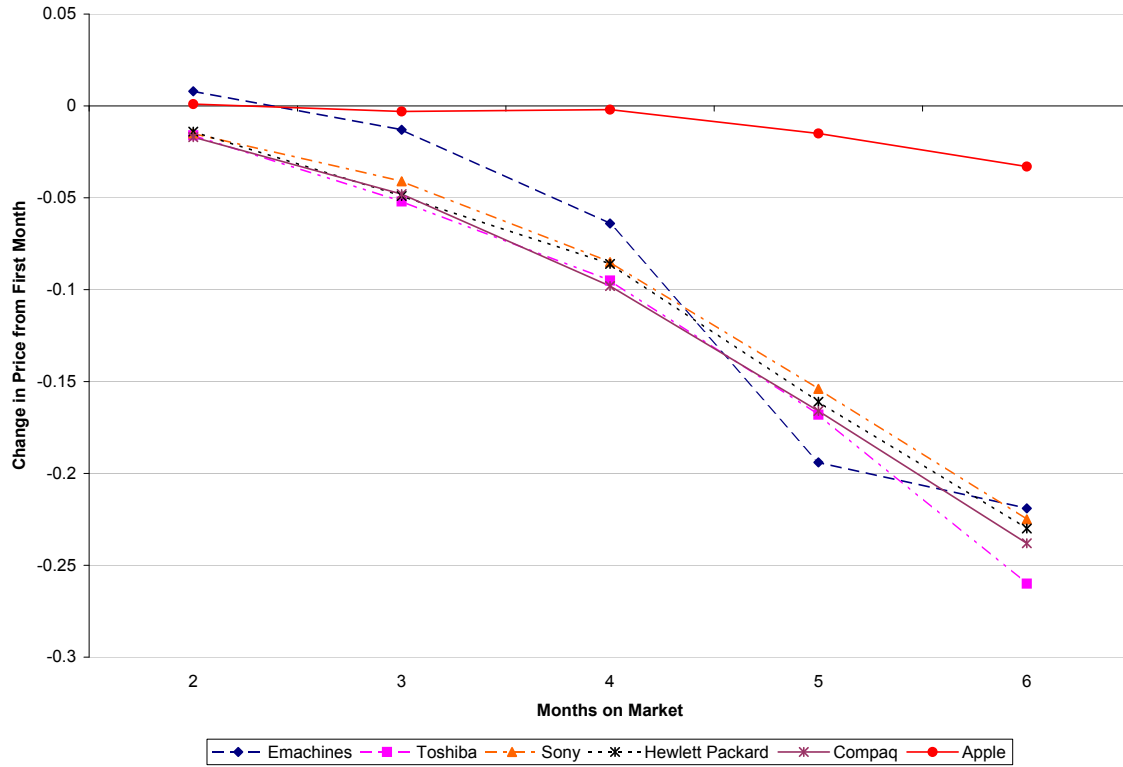
Notes: Depicted are the cumulative density functions (CDF) of the each computer manufacturer in the NPD data. We measure the CDF by running a regression of monthly unit sales on “months on market” dummies creating predicted values of the number of units sold given months on the market. Since computers do not necessarily enter the market at the beginning of the month, the first month of data will include less than 31 days worth of units sold. We can therefore interpret the CDF function in terms of bands as opposed to explicit months on the market.

Figure 2: Frontier CPU Clock Speed by Computer Manufacturer



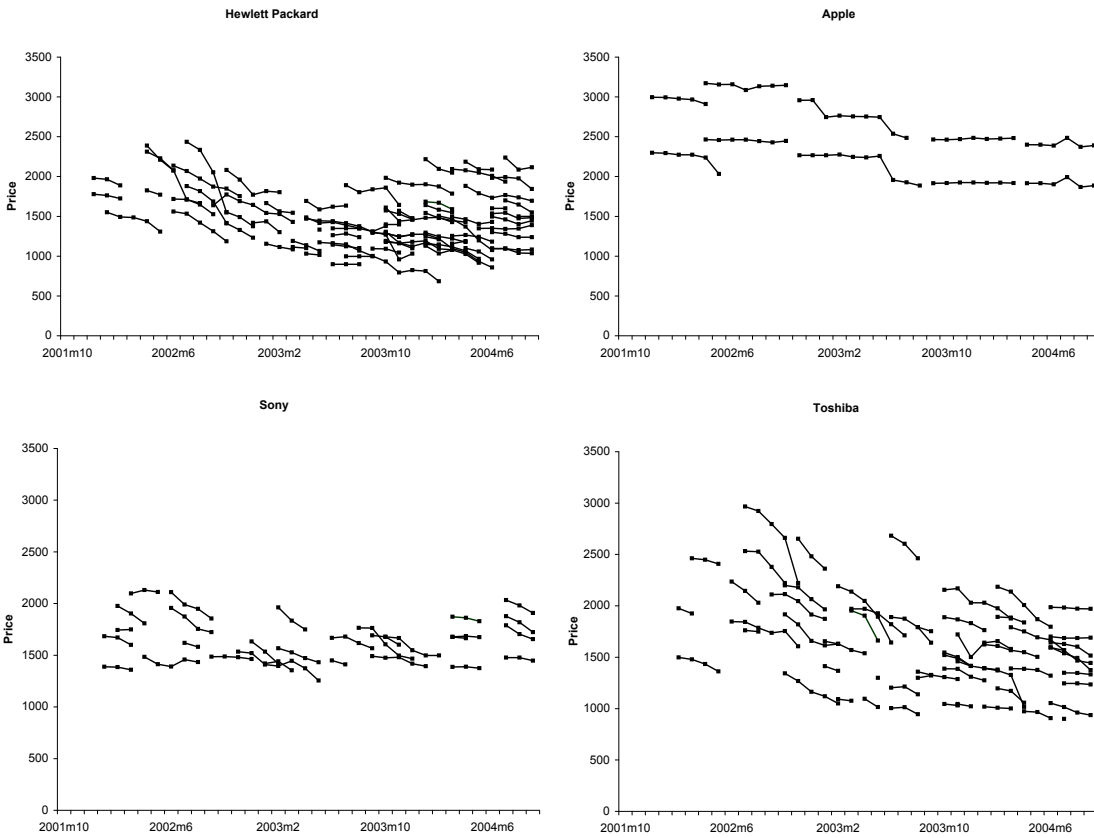
Notes:

Figure 3: Price Declines over Product Cycle



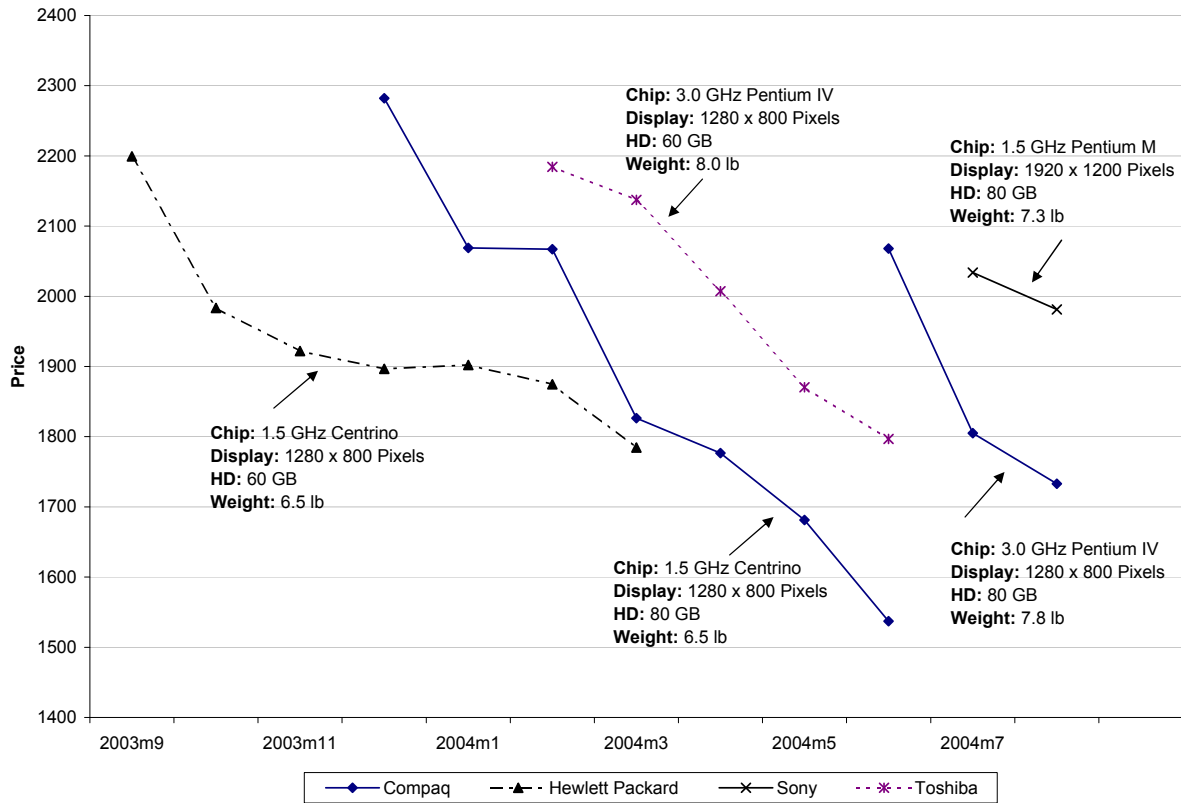
Notes: The figure shows the point coefficients of a fixed-effect (using model number as the fixed effect) regression of the logarithm of price on months-on-market dummy variables. The omitted dummy variable is the first month of entry indicating that the subsequent coefficients represent the percentage change between the given month and the first month. For ease of display we show only the first six months of life on the market. At this point, 99% of a given PC model and 80% of a given Apple model's total units have been sold.

Figure 4: 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. For ease of view, prices after the the units CDF reached 90 percent for each model were omitted. Each dotted line represents an individual computer model.

Figure 5: Entering PC is also Highest Priced: 15-Inch 512MB Notebooks



Notes: The figure shows price contours where the entering PC happened to be the highest price in the category of 512 MB RAM 15-inch notebook computers. For ease of view, prices after the the units CDF reached 90 percent for each model were omitted.

Figure 6: Model versus Data: Prices over the Product Cycle



Notes: Price of the first period has been normalized to 100.

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




Figure 8: Prices over the Product Cycle, by entry case




Figure 9: Sales over the Product Cycle, by entry case


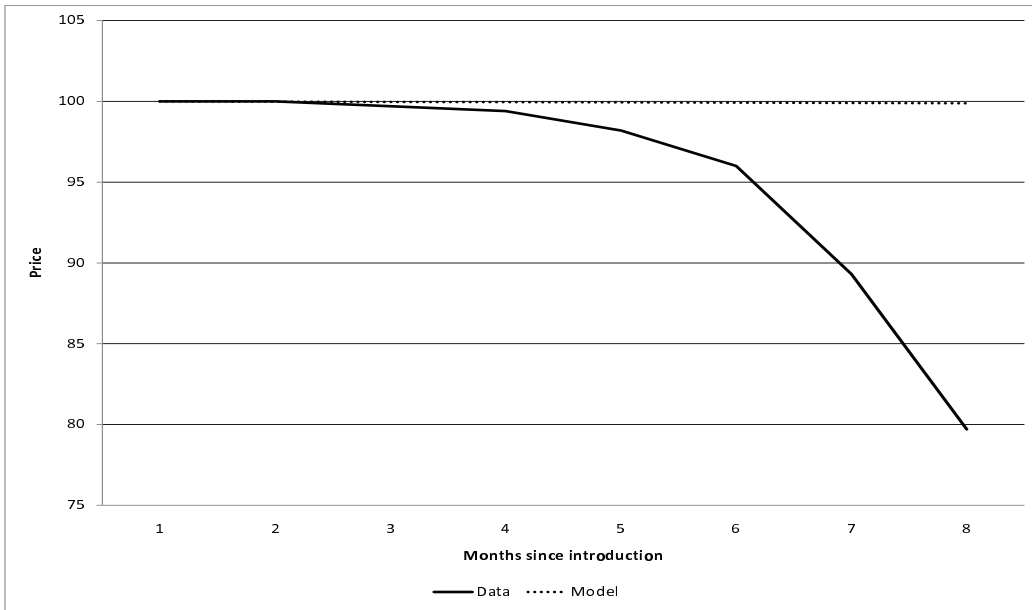


Figure 10: Monopoly Prices over the Product Cycle



Notes: Price of the first period has been normalized to 100.

Figure 11: Monopoly Sales over the Product Cycle




Figure 12: Monopoly Price and Period Profit by Vintage




Figure 13: Fixed Costs and Optimal Replacement Cycle Length

