PRICE SETTING AND RAPID TECHNOLOGY ADOPTION: THE CASE OF THE PC INDUSTRY

Adam Copeland and Adam Hale Shapiro*

Abstract—We examine how the confluence of competition and upstream innovation influences downstream firms' profit-maximizing strategies. We focus on personal computers and use two novel data sets to describe the dramatic fall in both price (27% at an annual rate) and sales of a computer over its product cycle. Further, we document that computers are typically sold for only four months before being replaced by a higher-quality product. To explain these facts, we develop and calibrate a vintage capital model that combines a competitive market structure with an exogenous rapid rate of innovation.

I. 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 by documenting several stylized facts about prices, sales, and consumers' income over the product cycle. To explain these time series, we develop and calibrate a vintage capital model that combines a competitive market structure with a rapid rate of innovation. We show that the model is well able to explain the observed paths of prices, sales, and consumer income over a typical PC's product cycle. Although the simplicity of the model leaves ample room for extensions to capture other important features of the PC market, we argue that 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 9%. Furthermore, we find that sales rapidly decline over the product cycle, and firms frequently introduce new, higher-quality products. Finally, we show that the 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 characterized by a narrow income distribution.

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 and the nonlinear price declines over this cycle, we argue that falling input costs or process innovations are not a main driver of the 27% (annual rate) decline in PC prices.[1] 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. As we explain in more detail in the paper, however, Stokey (1979) showed that this type of price discrimination is profit maximizing only under very strict assumptions, which we argue do not hold in the PC retail market.

In contrast, competition 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 explain the dramatic declines in PC prices, which in turn can cause consumers with lower incomes to purchase computers later in the product cycle. To see if a model combining competition and innovation can quantitatively match the set of stylized facts, we develop and calibrate a vintage capital model. 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 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

Received for publication October 8, 2013. Revision accepted for publication February 20, 2015. Editor: Philippe Aghion.

* Copeland: Federal Reserve Bank of New York; Shapiro: Federal Reserve Bank of San Francisco.

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 we both worked at the Bureau of Economic Analysis. The views expressed here are our own 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.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00539.

[1] Certainly, however, these factors may be important for explaining longer-term price trends in this industry (Aizcorbe, Flamm, & Khurshid, 2007).

The Review of Economics and Statistics, July 2016, 98(3): 601–616
© 2016 Federal Reserve Bank of New York
doi:10.1162/REST_a_00539

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 markup. 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.

With the calibrated model, we then explore two counterfactual scenarios. First, we explore how the PC retail market changes with different rates of innovation. To that end, we simulate the model with both a lower and higher rate of innovation and find a positive relationship between firms' markups and the rate of innovation. Specifically, more rapid innovation implies that different vintages of computers are further apart on the quality ladder. This greater vertical differentiation leads to a higher introductory price for a computer (and larger markups), followed by bigger price declines.

Second, we analyze the PC retail market for the monopoly case. We find that under this scenario, the firm's pricing strategy radically changes—pricing is flat over the product cycle. Furthermore, the monopolist upgrades its product less frequently compared to the competitive case, although when the monopolist upgrades its product, the jump up the quality ladder is much larger.

The paper is structured as follows: Section II reviews the related literature. Section III describes the data from the NPD Group and Metafacts and provides a description of the stylized facts for the PC market. In section IV, we present the model, and in section V, we take the model to the data. In section VI, we describe an out-of-sample exercise and then conclude in section VII.

II. Related Literature

This paper builds on the literature analyzing the effect of competition on pricing behavior,[2] as well as studies of the prices for durable technological goods.[3] 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 on 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.[4] Schumpeter conjectured that firms with larger market power would more aggressively 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. Although there is considerably more risk in undertaking R&D investment than adopting technology, there are similarities between both activities. In particular, both R&D investment and technology adoption allow the firm to improve the product's quality, and so differentiate its product from competitors. In this way, our work is related to this larger literature. Overall, our main result—that PC manufacturers seek to embed innovations into their retail products in order to leapfrog their competitors and (temporarily) grab market share—is more in line with Arrow's work.

Finally, our paper builds on 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 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 et al. 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,

[2] See, for example, Borenstein and Rose (1994) and Gerardi and Shapiro (2009).

[3] See Berndt and Rappaport (2001), Pakes (2003), Aizcorbe (2005), Erickson and Pakes (2011), Gowrisankaran and Rysman (2012), and Conlin (2010).

[4] See Dasgupta and Stiglitz (1980), Gilbert and Newbery (1982), Aghion and Howitt (1992), Greenstein and Ramey (1998), Aghion et al. (2009), Biesebroeck and Hashmi (2016), Goettler and Gordon (2009), and Nosko (2010).

which may downplay the importance of vertical differentiation and "racing to the frontier." Another major difference is that Bresnahan et al. 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 focus less on the manufacturer's brand.

III. 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-sale[5] transaction data (i.e., scanner data) sent to NPD Techworld weekly through automatic feeds from its participating outlets.[6] The data cover the course of 90 months, November 2001 to April 2009, and consist of sales occurring at outlet stores.[7] Thus, manufacturers such as Dell that primarily sell directly to the consumer are not included. Because our analysis focuses on the price and sales of a typical PC over the product cycle, not having Dell in our data could be problematic if both Dell's pricing strategy systemically different from its competitors and Dell is a big PC retailer relative to those in our sample.

Each observation consists of a model identification number (SKU), specifications for that model, the total units sold, and revenue. From units sold and revenue, we calculate a unit price of each PC sold.[8] In table 1 we report summary statistics and note that our sample covers more than 85 million sales. Table 2 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%, respectively.[9]

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, 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.[10]

TABLE 1.—SUMMARY STATISTICS

	Desktop Computers	Notebook Computers
Total number of distinct SKUs	1,824	2,993
Total number of units sold	35,759,160	50,004,428
Unit price (dollars)	728.7	985.4
	(381.3)	(482.4)
Units sold per Month	1,109.2	1,073.5
	(3,960.3)	(3,724.2)
Speed (MHz)	2,249.0	1,869.9
	(510.0)	(383.2)
Memory (MB)	1,035.8	1,455.7
	(1,208.0)	(1,147.0)
Hard drive (GB)	187.9	134.3
	(153.0)	(82.1)
Display (inches)		14.8
		(1.7)

Means reported; standard deviations in parentheses.
Source: NPD Group.

TABLE 2.—MARKET SHARES 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
E-machines	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

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

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.[11] 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 in sections IIIA and IIIB.

The exceptions 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.

[5] Point-of-sale means that any rebates or other discounts (e.g., 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.

[6] The 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.

[7] This includes sales on outlet stores' websites.

[8] We keep only SKUs with at least 1,000 units sold over the product cycle. We also drop SKUs that are left-censored—those SKUs whose first observation occurred in October 2001.

[9] HP and Compaq merged in 2003, but we chose to keep the brands separate in our analysis.

[10] All observations are reported on the user's primary computer. An observation in these 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.

[11] Prices after the cumulative density function (CDF), in terms of units sold, that reached 90% 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 fewer 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

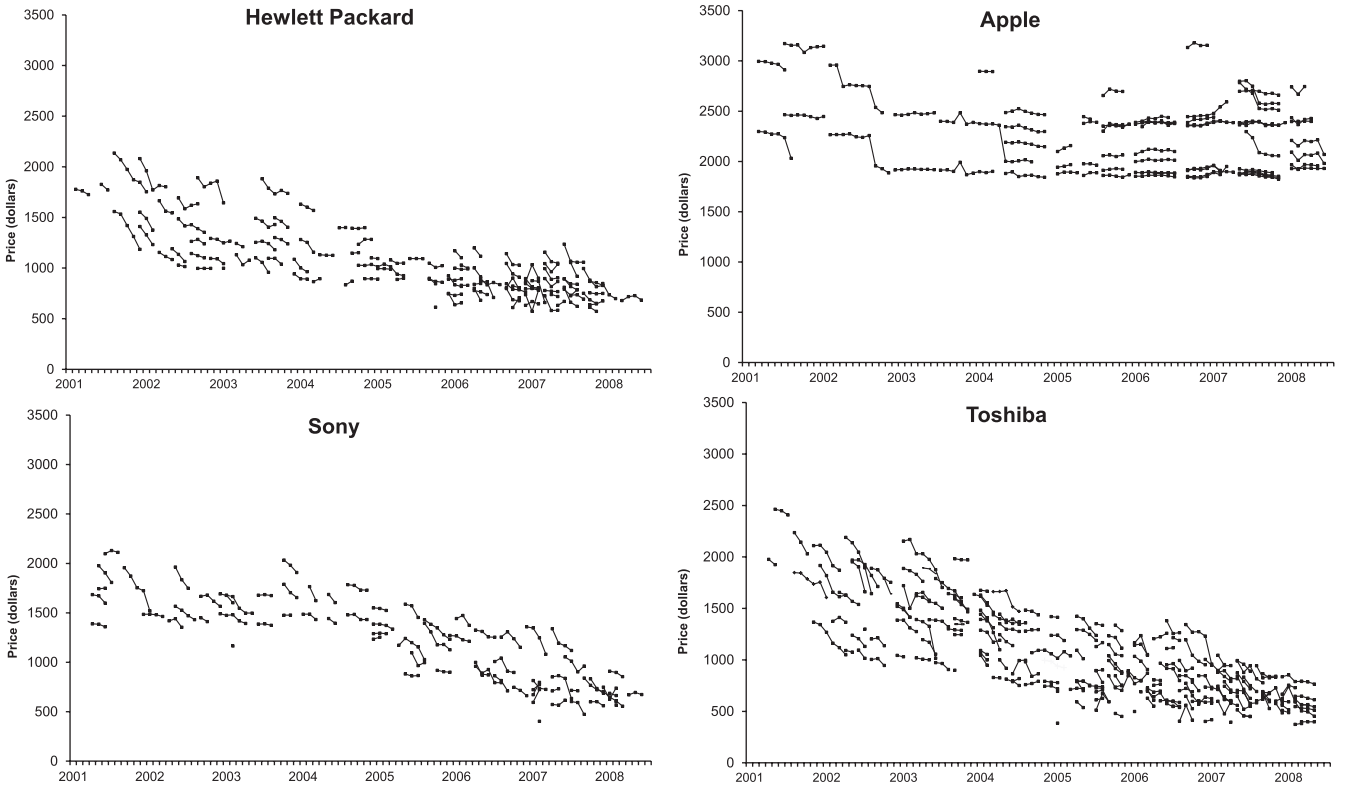

Depicted are the price contours of all 15-inch notebook computers sold by Hewlett Packard, Sony, Toshiba, and Apple computers over the course of the sample period. Prices after the the sales CDF reached 90% for each model were omitted.
Source: NPD Group.

Because of their unique operating system, Apple products are not close substitutes for other manufacturers' computers. HP, Sony, and Toshiba are more substitutable with one another because they offer the same the operating system (Microsoft Windows) and CPU (e.g., Intel chips) bundle. 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% of all sales. As we explain later, however, the pricing and technology-adopting strategy pursued by Apple helps inform our analysis.

A. *Pricing Patterns*

Figure 1 highlights some key features of price dynamics in the PC industry. Generally 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 for each of the seven major brands. The regression specification is $\ln P_{ik} = \alpha + \sum_{k=2}^{9} \beta_k D_k + \gamma_i + \varepsilon_{ik}$, 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, γ_i is a model number fixed effect, and ε_{ik} is the error term. Specifically, $k = z$ indicates that product i lies between deciles $z-1$ and z on product i's CDF.

FIGURE 2.—PRICE DECLINES OVER PRODUCT CYCLE

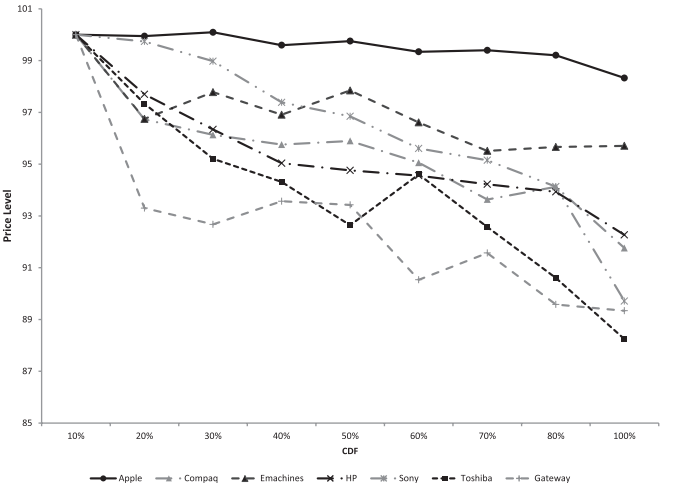

Depicted are the fitted values of a fixed effects (using model number as the fixed effect) regression of the logarithm of price on CDF decile dummy variables. See section D of the online appendix for coefficients and standard errors for a regression on all PCs, as well as Apple computers.
Source: NPD Group.

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. It is clear from figure 2 that PC

FIGURE 3.—ENTERING PC IS ALSO HIGHEST PRICED: 15-INCH 512 MB NOTEBOOKS

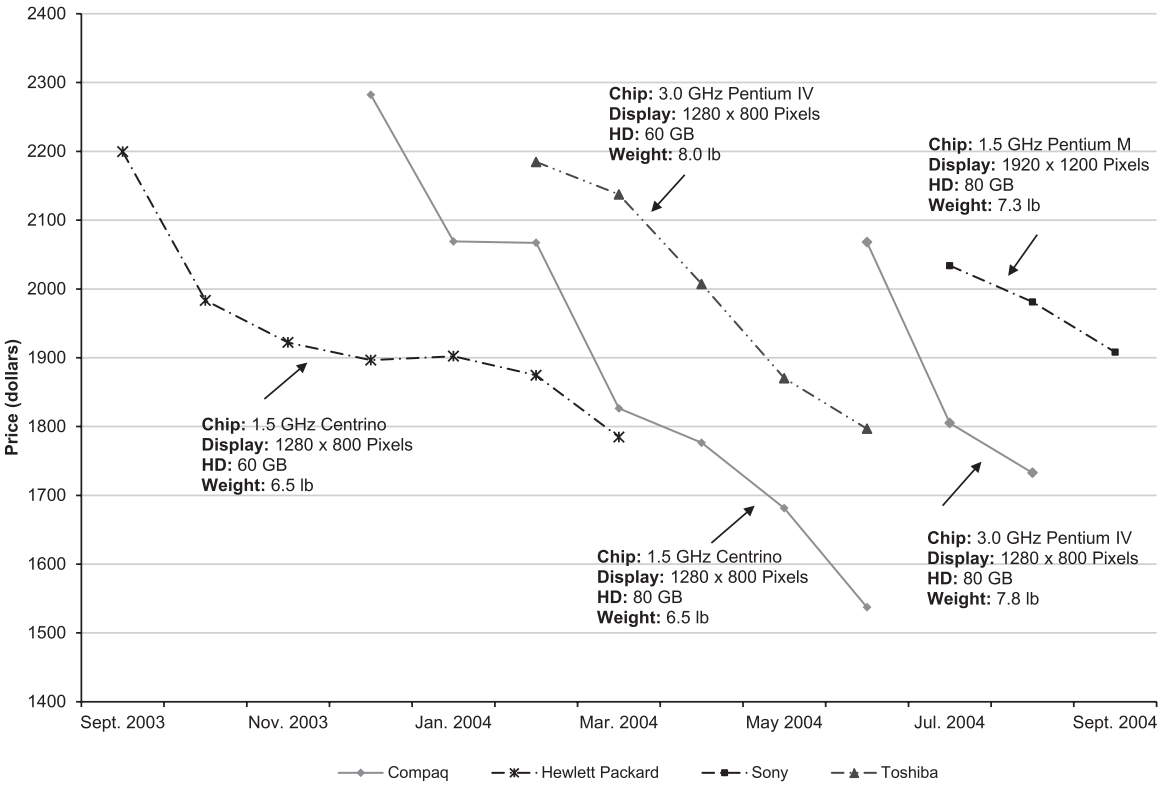

Computer models shown are models in which the entering PC happened to have the highest price in the category of 512 MB RAM 15-inch notebook computers. Prices after the unit's CDF reached 90% for each model were omitted.
Source: NPD Group.

prices fall quite rapidly over the product cycle. By the end of the cycle, they fall anywhere from 4% (Emachines) to 12% (Toshiba). The figure shows that relative to the PC manufacturers, Apple maintains a flat price profile over its product cycle. Most of the brands show a steep initial price decline. For instance, Compaq and Toshiba prices fall approximately 3% after the first 10% of computers are sold. To obtain an estimate of the average price decline, we estimate the price regression for all PCs (excluding Apple).[12] We find the typical PC declines in price by 9% over its product cycle.

Our regression specification does not control for the general rise in demand for PCs over our sample period. This could be a concern because increasing demand could cause us to underestimate the price declines over the product cycle. We find, however, that controlling for increasing demand hardly changes the estimates of price declines over the product cycle (see section D of the online appendix for details).

Figure 1 also highlights the 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. This line of computer represented 40% of all notebook units sold in our NPD sample during this time period. 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. To support our claim that newer products are of higher quality, we 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 and then fall rapidly. For instance, after its first month on the market, a Toshiba 15-inch notebook computer's price has fallen 4.2% on average from its initial price of $896. By the fourth month, its price has fallen 9.4 percent on average.

B. Sales Patterns, Product Cycle Length, and Technology Adoption

To get a better sense of the timing of sales along the product cycle, we depict PDFs of units sold in figure 4 for each

[12] Regression coefficients on this fixed effects model of all PCs are displayed in section D of the web appendix.

FIGURE 4.—PDF OF UNITS SOLD

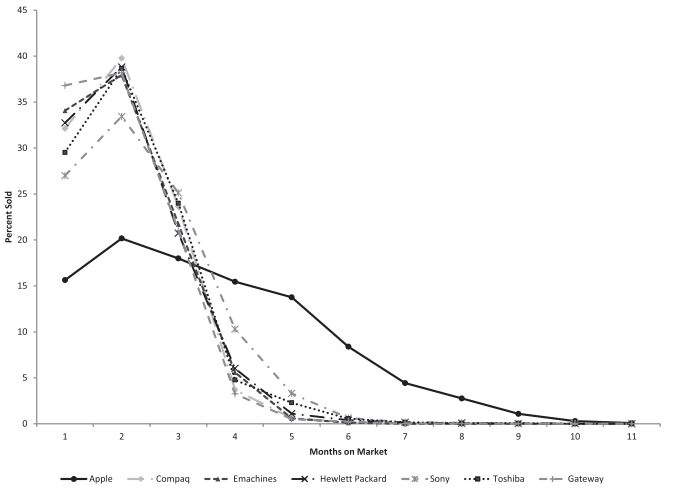

The PDF is calculated as the total units sold for a given age of an SKU (in months).
Source: NPD Group.

FIGURE 5.—ADOPTION OF INTEL CPUS

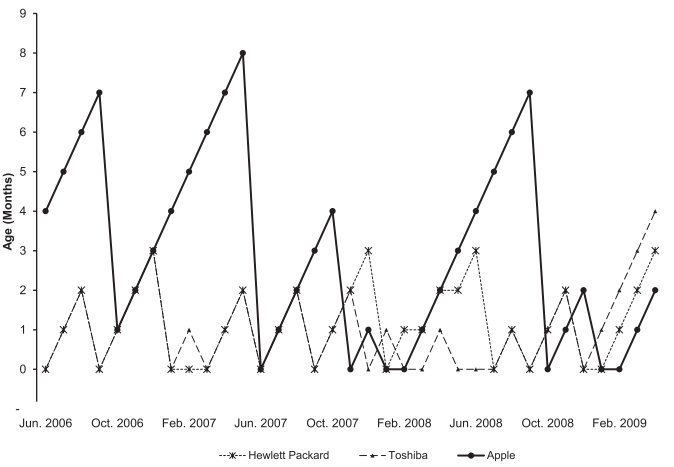

Depicted is the age of the newest Intel CPU for each month of the post-PowerPC CPU period (June 2006–April 2009) by computer manufacturer.

TABLE 3.—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
E-machines	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

This table reports the fraction of months in the sample period when no new model (columns 1–3) was introduced, as well as the maximum time period (in months) for which no new model was introduced (columns 4–6). 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.

brand. The PDF is calculated as the fraction of total units sold for a given brand and age of an SKU (in months). PC manufacturers sell the bulk of their product early in the product cycle. Generally computer manufacturers sell half of their units by the second month on the market. By month 4, they typically sell 90% of their product. Apple, however, keeps its computers on the market about twice as long as the other PC manufacturers.

Alongside the short product cycles, PC manufacturers introduce new products frequently. Indeed, manufacturers often release a new product nearly every month. Table 3 reports the fraction of months in the sample where no new computer was introduced. Both HP and Gateway almost always introduce a new computer model each month; the fraction of months where they did not introduce a new model is 1% and 2%, respectively. Apple continues to be the outlier, having the largest proportion of months in which no new models are introduced (28%).

Table 3 also depicts the maximum amount of time the manufacturer goes without introducing a new model. These statistics reinforce the finding that PC manufacturers frequently introduce new computer models, whereas Apple is relatively slow. 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 new notebook computer—by far the longest periods in the sample.[13]

Of the many components that make up a computer, CPUs are often considered one of the most important inputs. To gauge how frequently manufacturers adopt new CPUs into their products, we plot the age of the newest Intel CPU by month for the post-PowerPC period for Hewlett Packard, Toshiba, and Apple notebook computers in Figure 5.[14] 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, the age of the newest CPU available in Hewlett Packard and Toshiba notebooks almost never

[13] The 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.

[14] The 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.

TABLE 4.—CONSUMER INCOME DISPERSION AND LEVELS

	Gini Coefficient	Median Income
Apple	0.20	$65,000
Compaq	0.31	42,500
E-machines	0.30	42,500
Hewlett Packard	0.30	42,500
Sony	0.26	55,000
Gateway	0.29	47,500
Dell	0.28	55,000
IBM	0.32	55,000

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

exceeds three months. By contrast, Apple's newest CPU available was seven months old on three occasions.[15] In comparison to PC manufacturers, Apple's strategy is to adopt technology less frequently but with larger jumps in quality.

C. 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 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 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.[16] The survey data show that Apple has the highest median income, followed by Sony, Dell, and IBM. Furthermore, consumers of Apple have a very narrow distribution (0.195 Gini coefficient) relative to the other manufacturers.[17] IBM's higher Gini coefficient of 0.32 means that although it sells to a relatively high-median-income consumer, compared to Apple it sells to consumers with a wider range of incomes. This is attributable in part to the declining pricing pattern of IBM's computers, which encourages lower-income consumers to purchase IBM computers that have been on the market a few months.

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% fall in a consumer's income is correlated with a 0.9% 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 product cycle.[18] Naturally, because Apple's prices remain flat over the computer's product cycle, there is no correlation between price and income.

D. Qualitative Analysis

With this description of the set of stylized facts, we now turn to analyzing the economic determinants driving the observed behavior. Focusing on the dramatic price declines, several theories explain the falling pricing patterns for durable goods. Process innovation, falling input costs, and intertemporal price discrimination are three explanations that have received much attention for goods such as personal computers.

For three reasons, we argue that falling input costs or process innovation are not a main driver of the observed price declines for PCs over their product cycle. First, the observed pricing of Apple computers, which maintain almost flat prices over the product cycle, is not consistent with falling input costs or process innovation playing a significant role in pricing over the product cycle. Apple purchases many of the same intermediate inputs used by PCs, and so if input costs are dramatically falling over the product cycle, then Apple, similar to PC manufacturers, would be affected.[19] Standard models of a firm's pricing link pricing to costs, and so predict that falling input costs should be reflected, at least somewhat, in Apple's retail prices. Apple's strategy of almost flat pricing over the product cycle strongly suggests that input costs are not rapidly falling over the product cycle and so are not a main driver of PC price declines over the product cycle. Second, data from the U.S. Bureau of Economic Analysis's input-output tables and prices from the U.S. Bureau of Labor Statistics suggest that input costs are not driving PC prices over the product cycle. Using these data, we calculate that the declining prices of the main inputs to PCs translate into about a 2% annual price decline of PCs—far less then the

[15] We report a table in section D of the online appendix that gives CPU adoption statistics for the entire sample of notebook computers and shows that, on average, Apple offers the oldest CPU for this sample period.

[16] The 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.

[17] We 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 measure.

[18] Interactions 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.

[19] For example, both PCs and Macs use Intel CPUs, AMD graphics cards, and Samsung memory chips and are supplied by some of the same Taiwanese original design manufacturers (ODMs) such as Quanta and Asus.

27% annual rate we see in the NPD data.[20] Finally, it is hard to reconcile input cost declines with the within-product cycle pattern of PC price declines. Namely, we find that a majority of a PC's total price decline occurs in the first month after its introduction.[21] If input cost declines are important drivers of PC price declines, for some reason they would have to decline more rapidly the first month the PC is on the market. Further, discussions with market participants indicate that contracts between PC manufacturers and input suppliers are often negotiated annually, sometimes with clauses that allow cost reductions every quarter.[22] This anecdotal evidence supports the claim that input costs remain fixed over most of the PC's four-month product cycle.

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 candidate at first glance. This theory would also explain the decline in computer purchasers' income over the product cycle. Stokey's (1979) groundbreaking analysis, however, showed that such 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.[23] Thus, if the price falls in the PC market are due to intertemporal price discrimination, then it must be the case that for some particular reason, Stokey's conditions are met in the PC market but not the Apple market. We have no reason to believe that willingness to pay is more correlated with time preference for consumers in the PC market than for consumers in the Apple market. Furthermore, market power should be positively correlated with a firm's ability to commit to a price or production schedule. Unless Apple has less market power than the PC manufacturers, intertemporal price discrimination seems to be an unlikely candidate for the declining pricing patterns.

Unlike declining input costs and price discrimination, competition seems to be a plausible force behind declining PC prices over the product cycle. The frequent introduction of higher-quality PCs drives down the prices of PCs currently on the market. Furthermore, the combination of competition and rapid innovation could explain the short four-month product cycle for PCs. Finally, these price declines could lead to the sorting of consumers such that lower-income households purchase a PC later in its product cycle.

The rapid introduction of new innovations, combined with a competitive marketplace, is a plausible theory of the economic determinants driving the stylized facts documented above. To see if such a theory can match these facts quantitatively, we develop and calibrate a formal model of the PC industry in the rest of the paper.

IV. Model of a Competitive Industry

Our approach to modeling the PC industry takes a middle ground between purely theoretical models and detailed empirical models aimed toward estimating structural parameters. Our goal is to reveal how the basic forces of competition and rapid innovation can drive the main stylized facts presented above in a transparent manner. With this approach, we balance presenting a model of the PC retail market that is fairly intuitive yet still realistic enough to capture actual features of the marketplace. Consequently, we do not develop and estimate a more detailed structural model of the PC retail industry, such as Goettler and Gordon (2009) do for the PC microprocessor industry. Nor do we write a more theoretical model of competition and innovation, such as Johnson and Myatt (2003) do for studying multiproduct firms' upgrading decisions, or Aghion, Dewatripont, and Rey (1999) or Segal and Whinston (2007) do for firms undertaking risky R&D, because such models are usually not designed to match a set of stylized facts such as those presented in section III.

Our middle-ground approach employs a calibration exercise similar to numerous studies in the macroeconomics literature. Our model is based on the well-known Shaked and Sutton demand system with a vintage capital model. Computers are differentiated by their vintage ν, where ν equals the date at which the 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, implicitly assuming that the fixed costs of entering this industry dissuade more firms from entering. Further, we assume that each firm produces at most one computer and so ignore any joint maximization problem of a firm with multiple product lines. 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 high-quality 15-inch laptop computers depicted in figure 3. This last assumption is made to keep the model tractable. Nevertheless, given our focus on pricing and sales over

[20] Further, Dedrick, Kraemer, and Linden (2010) show that CPUs make up 14% of the retail price of Lenovo ThinkPad and HP nc6230 notebooks, which limits how much this input can drive PC retail prices over the product cycle. For example, if CPU prices decline at 25% per year (Byrne, Oliner, & Sichel, 2013) and manufacturers passed this gain through to consumers, this decline in input costs would translate into PC price declines of only 3.5% per year.

[21] The PDF of PC sales illustrates that at least 25% to 35% of a typical PC's total sales occur within the first month it is sold (see figure 4). Given this result, our regression results imply that about 40% of the price decline in a PC occurs between the introduction of the PC and end of the first month of its product cycle (see table 1 in section D of the online appendix).

[22] We thank employees from Hewlett Packard for discussions.

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

the product cycle, we do not anticipate that accounting for multiple product lines would add much economic insight.[24]

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 to incorporate the latest, most innovative inputs (e.g., lighter batteries, higher-resolution screens, or better chips) into their products.

A. 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 there is no accumulation of consumers across periods. We normalize the price of this alternative computer-related good (that is, the numeraire commodity) to 1. 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

$$\breve{u}_t \times y. \tag{2}$$

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

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

where $\bar{\nu}_t$ denotes the set of available computers in period t. The resulting demand function is straightforward.[25] 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:

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

[24] In contrast, accounting for multiple product lines would be more relevant for questions focused on the cross-section of PC prices or the entry and exit of product lines (e.g., see Johnson and Myatt, 2003).

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

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

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

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

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

B. Supply

A firm makes two decisions at the beginning of each period. First, it decides whether to adopt a new technology (i.e., 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 $i = 1, 2, \ldots, N$ 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 $s_t = (\nu_1, \nu_2, \ldots, \nu_N, \breve{u}_t)$, 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

$$\begin{aligned} V_i(s_t) = \max_{p_{\nu_i t}, p_{\nu_i' t}, d_{it}} & \\ \times \Big\{ (1 - d_{it}) E_{s'} \Big[(p_{\nu_i t} - c) Q_{\nu_i t}(p_{\nu_i t}, p_{\nu_{-i} t}; s') + \delta V_i(s') \Big] & \\ + d_{it} E_{s''} \Big[(p_{\nu_i' t} - c) Q_{\nu_i' t}(p_{\nu_i' t}, p_{\nu_{-i} t}; s'') - \phi + \delta V_i(s'') \Big] \Big\}, & \end{aligned} \tag{5}$$

where $\nu' = t$ is the latest (and highest-quality) vintage and $p_{\nu_{-i}t}$ 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. $Q_{\nu_i t}$ is the demand for product ν_i at time t, given prices and the outside option.

Although 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. Because 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 0 and the computer will have effectively exited the market. Of course, a firm may decide to upgrade its computer before demand reaches 0. The product cycle of a computer, then, starts with its introduction to 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.

C. *Equilibrium*

We use a Markov perfect equilibrium concept where firms are deciding at what price to sell their computer and whether 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.

Within this framework, a multitude of possible equilibria could be considered. To keep the analysis tractable, we restrict our focus by ruling out mixed strategies and considering the stationary case. An equilibrium is defined to be stationary if the upgrading decision and pricing of a computer over its product cycle are independent of time. Prices, and the resulting sales, then, are determined by a product's position within the product cycle (i.e., the first, second, third, ... month of the product cycle). To obtain this stationary equilibrium, we make an additional assumption: the ratio of the utility provided by the outside good over the utility associated with a computer embodying the frontier technology, denoted ν_t, remains constant over time. Formally,

$$\frac{\breve{u}_t}{\nu_t} = \zeta \ \forall t, \tag{6}$$

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

For a given N firm, 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. Under this strategy, one of the N firms adopts in a given period, resulting in a product cycle of N periods. 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.

A stationary equilibrium exists only if marginal cost is below the maximum price that 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. In section A of the online appendix, we prove that this equilibrium exists for the case where $N = 2$ and argue that the logic of the proof holds for larger N. Finally, the number of active firms in equilibrium will be determined by profit conditions. If N firms produce and sell computers in a stationary equilibrium, then the net present value of profits of each firm must be nonnegative. Further, the net present value of profits given $N + 1$ firms is negative.

A feature of the stationary case is that manufacturers leapfrog one another. This leapfrogging outcome is not unusual, showing up in a variety of other environments that incorporate innovation (e.g., in the study of R&D—Grossman & Helpman, 1991, and Giovannetti, 2001—and in the study of technology diffusion, Chari & Hopenhayn, 1991). In the data, we do not have definitive evidence that manufacturers always coordinate their adoption decisions by taking turns. Indeed, in some cases we see different manufacturers adopting the newest Intel CPU in the same month.

Nevertheless, we argue that this stylized feature of our model is plausible for the retail PC market. As discussed in section III, PC manufacturers incorporate quality-enhancing components into their computers through new product introductions. For an individual firm, then, the introduction of new computers moves that firm up the quality ladder. Looking across PC manufacturers, there is evidence that firms are staggering their movement up the quality ladder. Rather than releasing all their models in the same month, we see manufacturers introducing models at various times.[26] Further,

[26] In contrast, automakers coordinate and release the vast majority of their new product line in the fall of each year.

manufacturers are introducing the most innovative intermediate components at different times. As evidence, we found it very difficult to find two or more computer manufacturers offering computers with exactly the same observable specifications.[27] So while there may be instances in the data where firms introduce the same observable product at the same time or the same firm completes consecutive jumps up the quality ladder, the typical case looks like firms leapfrogging one another up a quality ladder.

V. Empirical Work

In this section we calibrate the model and then demonstrate that it can quantitatively match the stylized facts presented earlier. We then use the model to further develop the intuition behind the observed price, sales, and consumer income dynamics.

A. *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 determines the cost structure of the firm.

We use three parameters to characterize computer quality: the level of the highest-quality product; the monthly growth rate of the frontier technology, γ; and 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%, 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 (see section B in the online appendix for details on how this number is calculated). As we show later, however, our results are robust to significantly higher values of the outside good's utility (see section VC).

Turning to the consumers' budgets for computers and related products, we normalize the mass of consumers to 1 and assume the density of consumers' budgets is characterized by the beta distribution over the interval $[a, b]$. We fix a to be 1 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, whereas the parameters (κ_1, κ_2) affect the relative prices among vintages.[28]

Finally, we assume a simple cost structure for the firm. The firm pays a fixed cost, $\phi > 0$, to adopt a new technology, and 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 on the value of ϕ.

In taking 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, the computer manufacturer with the lowest-quality product upgrades its product every month.

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 (a) the number of months a computer is sold, (b) the computer's sales path over the product cycle, and (c) 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. 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 thirty days a computer is sold account for a substantial portion of a computer's product cycle, we cannot drop or otherwise ignore this measurement problem.

We address this measurement issue by constructing bounds on the sales CDF of a typical PC computer over the product cycle. We compute the sales CDF by summing sales by age for all PC manufacturers (i.e., excluding Apple), where age is the month since a computer model is first observed in the data. Because computers do not necessarily enter the market at the beginning of the month, the first month of data will include anywhere between one day and thirty days' worth of units sold. Thus, we create upper and lower bands for the sales CDF. The lower bound of the CDF is then equal to the sum of total units by age of the computer (in months). The upper bound is calculated in a similar manner; we assume that sales in month t are equal to total unit sales of computers in month t and $t + 1$.

[27] For example, less than 6% 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.

[28] In 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 significantly faster than we observe in the data.

TABLE 5.—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

This table reports the parameter values of the model. Flexible parameters are those chosen such that the moments implied by the model match those observed in the data, whereas fixed parameters reflect normalizations or are set based on prior knowledge.

TABLE 6.—TIME SERIES OF PRICES AND SALES, DATA, AND MODEL

Month	1	2	3	4
Sales CDF				
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
Sales PDF				
Data–lower bound	0.217	0.298	0.102	0
Model	0.464	0.351	0.154	0.031
Data–upper bound	0.600	0.383	0.298	0.102
Percent price decline, relative to first month				
Data	−6.9	−8.6	−9.1	−8.2
Model	−6.9	−8.6	−9.1	−8.2

This table compares the model's output of a computer's sales and prices over the product cycle to the data.

The sales CDF informs us about both the number of months a computer is sold and 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 IIIA. This regression provides an estimate of how much prices have declined from the initial price to a given point on the sales CDF. 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 are provided in section C of the online appendix.

We find a wide distribution in budgets for computers and related products (see table 5). Consumers with the largest budgets are willing to spend more than ten times as much as 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.

B. *Comparing the Model to the Data*

As displayed in table 6, our model fits the data on PCs well along both the 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 the upper panel of table 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.[29] Our model also performs well along the sales PDF dimension (see the middle panel of table 6). Specifically, it captures well the burst of sales when a computer is introduced, followed by a decline in the sales rate over time.

Turning to prices, we find the model exactly matches the price declines seen in the data (see the lower panel of table 6). 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.

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%. Once a newer vintage enters the market, this markup plummets. The second-highest-quality computer has a markup of almost 3%, and the third- and fourth-highest-quality computers have markups of 0.8% and 0.3%, 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 III, 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% fall in price is correlated with a 10% fall in income. Using the change in price observed in the data, the correlation from Aizcorbe and Shapiro (2010) suggests that incomes of consumers who buy the highest-quality computer should be almost twice as big as those who buy the lowest-quality computer.

The model predicts a larger difference in consumer budgets. As shown in table 7, the average budget of consumers who purchase the computer when it is 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

[29] We can construct bounds on the sales PDF using the sales CDF. The implied minimum sales PDF, however, is always equal to 0, 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.

TABLE 7.—AVERAGE BUDGET OF CONSUMERS OVER THE PRODUCT CYCLE

Month in the product cycle	1	2	3	4
Mean budget of purchasers	5.54	2.26	1.22	1.01

This table presents the average budget of those consumers purchasing a given computer over that computer's product cycle. Budget is in units of the outside good.

FIGURE 6.—DISTRIBUTION OF CONSUMERS OVER THE PRODUCT CYCLE

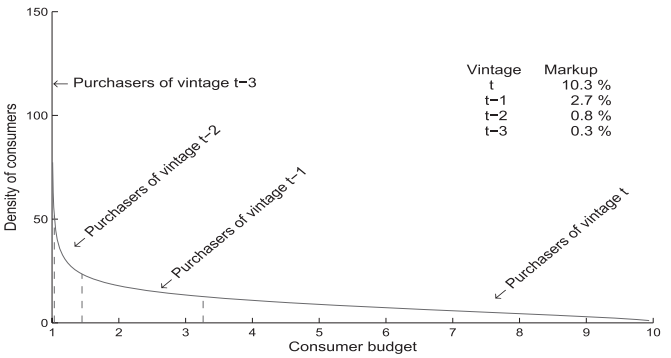

Displayed is the distribution of consumer budgets and the purchase decisions of consumers in equilibrium. Vertical dotted lines represent the marginal consumer, who is indifferent between purchasing two adjacent vintages. Because consumers with budgets greater than (less than) the marginal consumer prefer the newer (older) vintage, the marginal consumers are used to calculate the mass of consumers purchasing a computer over its product cycle. For example, the area under the curve to the right of the marginal consumer with a budget of roughly 3.2 is the mass of consumers who purchase the computer when it is introduced. Markup is the firm's markup relative to marginal cost.

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).

C. *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 competitive.

The distribution of consumers over the product cycle can be seen graphically in figure 6. The curved 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 large sales while pricing 10% above marginal cost. Referring back to the PDF of sales over the product cycle, the model predicts that 46% of a computer's total sales are completed in the first month (see the middle panel of table 6). Being replaced as the highest-quality computer reduces both a manufacturer's sales and its markup (from 10.3% to 2.7%). Market-specific obsolescence forces the manufacturer to appeal to consumers with lower budgets for computers, driving the manufacturer to slash its markup by three-fourths.

The combination of rapid obsolescence and large gains from being the highest-quality computer leads 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 preinnovation rents. In the spirit of Arrow (1962) and Aghion et al. (2001, 2005), then, computer manufacturers incorporate innovative intermediate inputs into their products to 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.

VI. Counterfactuals

Given the calibrated parameters, we conduct two counterfactuals. We first analyze how PC manufactures would react to different rates of upstream innovation. Second, we consider the case when the cost of entering the market is high enough that only one manufacturer profitably enters the market.

A. *Altering the Rate of Upstream Innovation*

A main force in the model is the exogenous rate of innovation supplied by upstream manufacturers. We assume this rate of innovation is equal to Moore's law when calibrating the model. But an interesting question is how pricing and upgrading decisions are altered with different rates of growth in innovation. Given concerns about slowing rates of innovation, it is useful to understand how such changes affect the PC retail market.[30]

With this motivation, we simulate the model with under both lower and higher rates of growth, relative to the calibrated 2.9% growth rate. We start by doubling the benchmark growth rate to 6%. The 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 7, we plot the model's predictions of markups for this high-growth-rate case against the benchmark case of a 2.9% growth rate. As highlighted in

[30] For example, see "Has the Ideas Machine Broken Down." *Economist*, January 12, 2013.

FIGURE 7.—PRICE DECLINES WITH DIFFERENT RATES OF INNOVATION

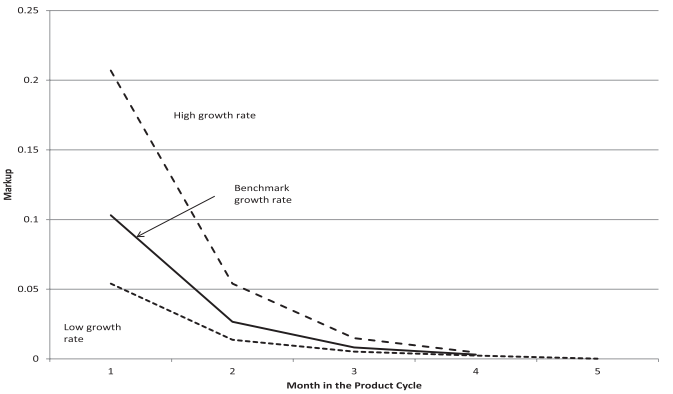

Displayed are computer manufacturers' markups over the product cycle in equilibrium, given different exogenous growth rates in innovation. The benchmark growth rate is 2.9%, whereas the high and low growth rates are 6% and 1.5%, respectively.

the figure, the initial markup is an enormous 21%, roughly double the value in the benchmark case.

In figure 7 we also plot the model's prediction of price over the product cycle given a slower growth rate of 1.5%. 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.

Changes in the rate of innovation from upstream firms, then, have a substantial impact on the markups that manufacturers charge. In particular, markups significantly increase in the rate of innovation.

B. *Altering the Competitive Environment*

In this second counterfactual, we explore the importance of competition on manufacturers' upgrading decisions. As described section II, a substantial literature looks at the effects of competition on innovation. Here we measure the impact of competition on the upgrading decision by comparing the model's results against the monopolist case. Given Apple's high degree of production differentiation through its distinct operating system, assessing the monopolist case may also give insight into to why Apple's prices and sales patterns over the product cycle differ from those of the other PC manufacturers.

Although the number of firms is endogenous in our model, by assuming that entry costs are large enough, we can impose that only one firm enters and earns positive profits. 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.

The monopolist's problem. The timing of the monopolist's 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 the outside option, $\breve{u}_t$. The monopolist's problem is

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

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.

Predictions of the monopolist model. Given the 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.[31]

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 0, for example, then the monopolist will update every period to take advantage of the high value that 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.

We set ϕ to 0.048, which is equal to 93% of the expected discounted profits of a computer manufacturer in the competitive market. This value satisfies the equilibrium constraints in the competitive setting and is consistent with the underlying idea that the PC market is quite competitive.[32] With $\phi = 0.048$, the monopolist's upgrade replacement cycle is eight months. This is the same length as that observed for the typical Apple computer, and so it is not an implausible

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

[32] Consistent with our result, Dedrick et al. (2010) find that "a large share of the PC industry profits are siphoned off to Microsoft and Intel" (p. 3).

TABLE 8.—MONOPOLIST: PRICES AND SALES OVER THE PRODUCT CYCLE

Percent Price Decline, Relative to First Month								
Month	1–2	1–3	1–4	1–5	1–6	1–7	1–8	Average
Model	−0.08	−0.17	−0.25	−0.34	−0.43	−0.53	−0.63	−0.35
Sales CDF								
Month	1	2	3	4	5	6	7	8
Model	0.125	0.250	0.375	0.500	0.625	0.750	0.875	1

This table presents computer sales and prices over the product cycle for the counterfactual where there is only one manufacturer.

prediction. Importantly, when the monopolist upgrades, it jumps to the frontier technology. Hence, the model is predicting less frequent upgrades under a monopolist, but with each upgrade being a larger leap in terms of quality.

With these parameters, the model predicts a flat price profile over the product cycle (see table 8). 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. These predictions highlight the central role of competition in driving price declines over the product cycle (alongside falling sales and purchasers' average incomes). 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 is the main driver of computer prices and sales over the product cycle.

Overall, then, the model predicts that the monopolist upgrades roughly half as frequently as a firm in a competitive environment (eight months versus four months). Competition, then, has a substantial impact on manufacturers' upgrading decisions, driving them to incorporate the latest innovative intermediate inputs into computers at twice the rate of a monopolist. There is not, however, a long-run divergence in computer quality between the competitive and monopolist case, because the monopolist jumps to the frontier when it decides to upgrade. Further, the model predicts that the monopolist essentially keeps price flat over the product cycle, a dramatically different price profile compared to the competitive case.

VII. Conclusion

In this paper we document a number of stylized facts about the PC retail industry. We show that prices and sales of computers rapidly decline over the product cycle. Further, consumers with lower incomes tend to purchase computers later in their product cycle. To explain these time series, we develop and calibrate a vintage capital model that combines rapid rates of innovation with a competitive market. We show that our model can quantitatively match the data.

We then use the model to consider two counterfactuals. We show that the slower rates of innovation will make the PC retail market more competitive, in that firms' markups will fall. We then explore how this industry would change if there was a monopolist PC manufacturer. Besides the expected results that prices increase and output falls, we show that the monopolist will update its product less frequently. Hence, compared to the competitive environment, the monopolist upgrades its product less often but with each upgrade being a larger jump up the quality ladder.

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 dynamic demand into this environment to answer questions about the evolution of the PC industry. Nevertheless, because of its simplicity and ability to match the main stylized facts of the PC retail market, this model provides a useful benchmark against which to compare more complicated models.

Finally, this paper touches on a debate about the impact of information technology on labor productivity (Byrne et al., 2013). This issue hinges on several measurement issues, including how much quality-adjusted prices in the information technology sector (e.g., semiconductors and personal computers) are falling. Our paper informs this debate in that our vintage capital model provides a link between prices and quality that supports the use of a matched model approach to measure price change.

REFERENCES

Aghion, Philippe, Nicholas Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt, "Competition and Innovation: An Inverted-U Relationship," *Quarterly Journal of Economics* 120:2 (2005), 701–728.

Aghion, Philippe, Richard Blundell, Rachel Griffith, Peter Howitt, and Susanne Prantl, "The Effects of Entry of Incumbent Innovation and Productivity," this REVIEW 91:1 (2009), 20–32.

Aghion, Philippe, Mathias Dewatripont, and Patrick Rey, "Competition, Financial Discipline and Growth," *Review of Economic Studies* 66:4 (1999), 825–852.

Aghion, Philippe, Christopher Harris, Peter Howitt, and John Vickers, "Competition, Imitation and Growth with Step-by-Step Innovation," *Review of Economic Studies* 68:3 (2001), 467–492.

Aghion, Philippe, and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica* 60:2 (1992), 323–351.

Aizcorbe, Ana, "Moore's Law, Competition, and Intel's Productivity in the Mid-1990s," *American Economic Review: Papers and Proceedings* 95:2 (2005), 305–308.

Aizcorbe, Ana, Kenneth Flamm, and Anjum Khurshid, "The Role of Semiconductor Inputs in IT Hardware Price Decline" (pp. 351–382), in E. Berndt and C. Hulten, eds., *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches* (Chicago: University of Chicago Press, 2007).

Aizcorbe, Ana, and Samuel Kortum, "Moore's Law and the Semiconductor Industry: A Vintage Model," *Scandinavian Journal of Economics* 107:4 (2005), 603–630.

Aizcorbe, Ana, and Adam Shapiro, "Implications of Consumer Heterogeneity on Price Measures for Technology Goods," Bureau of Economic Analysis working paper (2010).

Arrow, Kenneth, "Economic Welfare and the Allocation of Resources for Invention" (pp. 609–625), in R. Nelson, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors* (Princeton, NJ: Princeton University Press, 1962).

Berndt, Ernst, and Neal Rappaport, "Price and Quality of Desktop and Mobile Personal Computers: A Quarter-Century Historical Overview," *American Economic Review* 91:2 (2001), 268–273.

Biesebroeck, Johannes Van, and Aamir Hashmi "Market Structure and Innovation: A Dynamic Analysis of the Global Automotive Industry," this REVIEW 98 (2016), 192–208.

Borenstein, Severin, and Nancy Rose, "Competition and Price Dispersion in the U.S. Airline Industry," *Journal of Political Economy* 102:4 (1994), 653–683.

Bresnahan, Timothy, Scott Stern, and Manuel Trajtenberg, "Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s," *RAND Journal of Economics* 28:0 (1997), S17–S44.

Bulow, Jeremy, "Durable-Goods Monopolists," *Journal of Political Economy* 90 (1982), 314–332.

Byrne, David, Stephen Oliner, and Daniel Sichel, "Is the Information Technology Revolution Over?" *International Productivity Monitor* 25 (2013), 20–36.

Chari, V. V., and Hugo Hopenhayn, "Vintage Human Capital, Growth, and the Diffusion of New Technology," *Journal of Political Economy* 99:6 (1991), 1142–1165.

Conlin, Christopher, "A Dynamic Model of Costs and Margins in the LCD TV Industry," Columbia University working paper (2010).

Dasgupta, Partha, and Joseph Stiglitz, "Industrial Structure and the Nature of Innovative Activity," *Economic Journal* 106 (1980), 925–951.

Dedrick, Jason, Kenneth Kraemer, and Greg Linden, "Who Profits from Innovation in Global Value Chains? A Study of the iPod and Notebook PCs," *Industrial and Corporate Change* 19:1 (2010), 81–116.

Erickson, Timothy, and Ariel Pakes, "An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods," *American Economic Review* 101:5 (2011), 1707–1738.

Gerardi, Kristopher, and Adam Shapiro, "Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry," *Journal of Political Economy* 117:1 (2009), 1–37.

Gilbert, Richard, and David Newbery, David, "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review* 72 (1982), 514–527.

Giovannetti, Emanuele, "Perpetual Leapfrogging in Bertrand Duopoly," *International Economic Review* 42:3 (2001), 671–696.

Goettler, Ronald, and Brett Gordon, "Does AMD Spur Intel to Innovate More?" *Journal of Political Economy* 119:6 (2009), 1141–1200.

Gowrisankaran, Gautam, and Marc Rysman, "Dynamics of Consumer Demand for New Durable Goods," *Journal of Political Economy* 120 (2012), 1173–1219.

Greenstein, Shane, and Garey Ramey, "Market Structure, Innovation, and Vertical Product Differentiation," *International Journal of Industrial Organization* 16 (1998), 285–311.

Grossman, Gene, and Elhanan Helpman, "Quality Ladders in the Theory of Growth," *Review of Economic Studies* 58 (1991), 43–61.

Johnson, Justin, and David Myatt, "Multiproduct Quality Competition: Fighting Brands and Product Line Pruning," *American Economic Review* 93:3 (2003), 748–774.

Nosko, Chris, "Competition and Quality Choice in the CPU Market," manuscript, Harvard University (2010).

Pakes, Ariel, "A Reconsideration of Hedonic Price Indices with an Application to PCs," *American Economic Review* 93:5 (2003), 1578–1596.

Prescott, Edward, and Michael Visscher, "Sequential Location among Firms with Foresight," *Bell Journal of Economics* 8:2 (1977), 378–393.

Schumpeter, Joseph, *The Theory of Economic Development* (Cambridge, MA: Harvard University Press, 1934).

——— *Capitalism, Socialism, and Democracy* (New York: Harper and Row, 1942).

Segal, Ilya, and Michael Whinston, "Antitrust in Innovative Industries," *American Economic Review* 97:5 (2007), 1703–1730.

Shaked, Avner, and John Sutton, "Relaxing Price Competition through Product Differentiation," *Review of Economic Studies* 49:1 (1982), 3–13.

——— "Natural Oligopolies," *Econometrica* 51:5 (1983), 1469–1483.

Stokey, Nancy, "Intertemporal Price Discrimination," *Quarterly Journal of Economics* 93:3 (1979), 355–371.