

Durable Goods Oligopoly with Innovation: Theory and Empirics*

Ronald Goettler[†]

University of Chicago

Brett Gordon[‡]

Columbia University

September 25, 2008

Abstract

We propose and estimate a model of dynamic oligopoly with durable goods and endogenous innovation. Firms make dynamic pricing and investment decisions while taking into account the dynamic behavior of consumers who anticipate the product improvements and price declines. The distribution of currently owned products is a state variable that affects current demand and evolves endogenously as consumers make replacement purchases. Our work extends the dynamic oligopoly framework of Ericson and Pakes (1995) to incorporate durable goods. We propose an alternative approach to bounding the state space that is less restrictive of frontier firms and yields an endogenous steady-state rate of innovation. Using a simulated minimum distance estimator, we estimate the model for the PC microprocessor industry and perform counterfactuals to measure the benefits of competition. Consumer surplus is 4.1 percent higher (\$17 billion per year) with AMD than if Intel were a monopolist. Innovation, however, would be higher without AMD. We also show that prices and profits are substantially higher when firms correctly account for the dynamic nature of demand, compared to an alternative scenario in which they mistakenly ignore the effect of current prices on future demand. Finally, equilibrium prices, profits, innovation, and consumer surplus are all increasing in the consumer's discount factor.

Keywords: dynamic oligopoly, durable goods, simulation estimation, microprocessors

JEL Classification: C73, L11, L13, L40, L63

*We would like to thank Ana Aizcorbe, Stephen Ryan, Holger Sieg, and Gabriel Weintraub for helpful comments. All remaining errors are our own. Gordon appreciates financial support from the William Larimer Mellon Fund and the Center for Analytical Research in Technology at the Tepper School of Business at Carnegie Mellon University.

[†]E-mail: goettler@uchicago.edu

[‡]E-mail: brg2114@columbia.edu

1 Introduction

Competition benefits consumers for two reasons—lower prices and higher quality. High-tech aficionados often espouse the benefit of having Advanced Micro Devices (AMD) as a competitor to the giant chip-maker Intel. Some argue that without AMD, consumers would be paying twice as much for computer processors half the speed, even though AMD's market share is typically less than twenty percent. In this paper we assess the importance of competition in high-tech industries in terms of its effect on prices, innovation, profits, and consumer surplus.

The analysis is complicated by the fact that most high-tech goods are durable. In durable goods markets sellers face a dynamic trade-off: selling more today reduces demand tomorrow. One strategy used by firms to mitigate this aspect of durable good demand is to continually improve product quality to give consumers an incentive to upgrade. Despite the importance of dynamic demand and product innovation in oligopolistic, durable goods markets, the equilibrium implications of firms' and consumers' strategies in such markets remain unclear. We therefore construct a model of dynamic oligopoly with durable goods and endogenous innovation.

In our model firms make dynamic pricing and investment decisions while taking into account the dynamic behavior of consumers. In turn, consumers account for the fact that firms' strategies lead to higher quality products and lower prices when considering whether to buy now or to delay their purchase. Since each consumer's demand depends on which product they currently own (if any), the distribution of currently owned products affects aggregate demand in each period. We explicitly model the endogenous evolution of this distribution and its effect on equilibrium behavior. In particular, we show that accounting for product durability and the distribution of consumer ownership has significant implications for prices, innovation, profits, and consumer surplus.

We use a minimum simulated distance estimator to estimate the model with data from the PC microprocessor industry. This industry is well-suited for the analysis because it is a duopoly, with Intel and AMD controlling about 95 percent of the market, and

sales have been driven by rapid technological innovation and intense price competition. While our quantitative results are specific to this industry, we believe the qualitative insights are relevant for any durable goods market where innovation and obsolescence drive product replacement.

To assess the importance of competition in durable good markets, we compare social surplus across three exogenous market structures: monopoly, duopoly, and a planner who maximizes social surplus. For our estimated model, the duopoly yields 84.7 percent of the planner's social surplus whereas the monopoly yields 83.5 percent. We also compare consumer surplus, which is \$435 billion per year with AMD and \$419 billion per year without AMD. Social surplus under monopoly and duopoly is lower than the planner's level for two reasons: the planner innovates more and charges less. Innovation in the duopoly is about 76 percent as frequent as under the planner.

Gilbert (2005) reports that the potential effect of competition on innovation is increasingly cited as a concern by the Federal Trade Commission, despite the absence of conclusive theoretical or empirical evidence. Importantly, our model can generate either a positive or a negative relationship between competition and innovation. Since this relationship surely varies across industries, researchers and policymakers must be careful to let the data speak. We take a decidedly different approach from the earlier empirical literature, such as Aghion, et al. (2005), which uses variation in competition and innovation across time and/or location to identify the effect of competition on innovation. The difficulty of measuring both innovation and competition in a manner which yields clear patterns in cross-sectional analyses is indicated by Figure 1 of Aghion, et al. (2005) which we replicate below. Our structural approach instead uses data on firm's pricing and investment decisions and consumer upgrade behavior to estimate consumer preferences for quality and their willingness to pay for this quality. Combined with estimates for the degree of competition and their investment processes, we are able to compute the industry equilibrium and then evaluate how this equilibrium changes in response to changes to market structure.

Figure from Aghion, Bloom, Blundell, Griffith, and Howitt (2005)

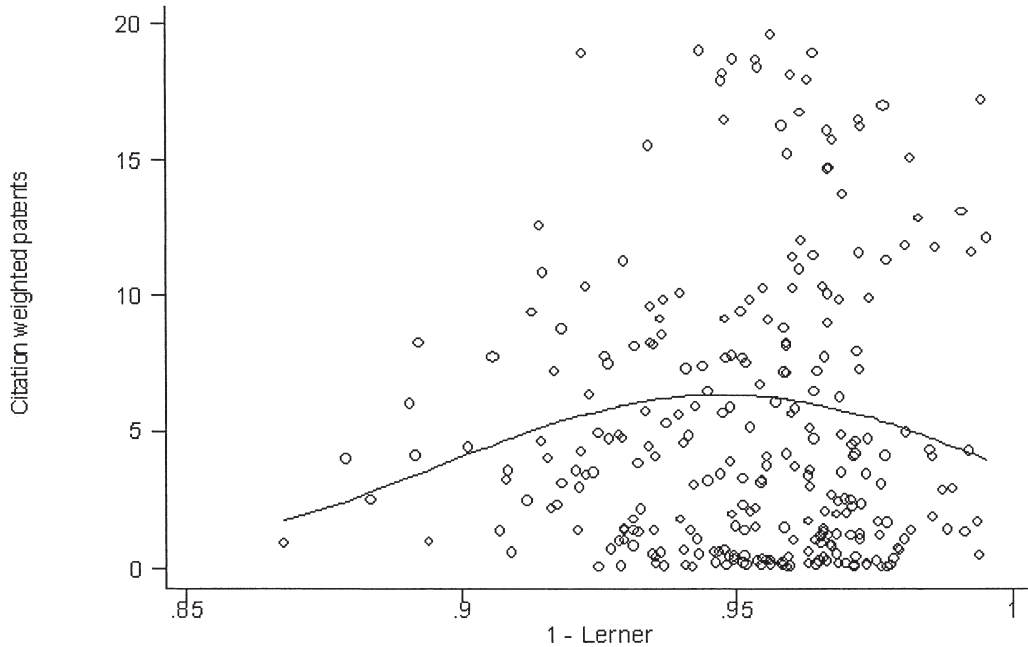


FIGURE I

Scatter Plot of Innovation on Competition

The figure plots a measure of competition on the x -axis against citation-weighted patents on the y -axis. Each point represents an industry-year. The scatter shows all data points that lie in between the tenth and ninetieth deciles in the citation-weighted patents distribution. The exponential quadratic curve that is overlaid is reported in column (2) of Table I.

We find that competition does **not** lead to higher innovation in the PC microprocessor industry. This finding highlights the “competing-with-itself” aspect of being a monopolist of a durable good: the monopolist must innovate to stimulate demand through upgrades. A second reason the monopolist has a higher rate of return on investment (and hence invests more) is that its pricing power enables it to extract much of the surplus generated by innovations.¹

To assess the importance—for firms and researchers—of accounting for a product’s durability when modeling demand, we compute equilibrium prices and investment under the hypothetical scenario in which firms ignore the effect of current prices on future demand. We find that prices and profits are substantially higher when firms correctly ac-

¹In a nondurable goods setting, Macieira (2007) finds competition increases innovation in the super-computer industry.

count for the dynamic nature of demand. One implication of this result is that marginal costs of durable goods derived from static first-order conditions will be too high: the prices are high to preserve future demand, not because costs are high. The dynamic nature of demand for durable goods also reverses the well-known “inverse elasticity rule” of optimal pricing, which establishes that firms maximize static profits by setting price such that its markup equals the inverse of the demand elasticity.² In our model demand is more elastic the higher is the level of demand. When demand is high, firms price much higher than if they were static optimizers since consumers who do not buy today will likely be buyers next period. Hence, optimal pricing of durable goods leads to a positive relationship between markups and elasticity.

A natural application of our model is the evaluation of mergers and antitrust policy for industries with durable goods. We consider a series of counterfactuals in which we vary the portion of the market in which one firm has monopoly status. This policy simulation is motivated by the recent antitrust lawsuit filed by AMD alleging that Intel has engaged in anti-competitive practices that effectively excluded AMD from part of the market.³ We find that excluding AMD from 30 percent of the market decreases consumer surplus by 1.9 percent, social surplus by 1.6 percent, and innovation by 1.4 percent. Though these percentage changes are small, the lost consumer surplus is still \$8 billion per year.

Our work incorporates durable goods into the dynamic oligopoly framework developed by Ericson and Pakes (1995) and applied to non-durable differentiated products by Pakes and McGuire (1994). To use numerical solution techniques the state space must be finite, which requires defining product qualities relative to some base product since quality improves over time, and then ensuring that relative qualities are bounded. In the Ericson-Pakes framework (hereafter, EP) consumer preferences are concave in quality measured relative to the outside good, so that a firm’s benefit to innovation goes to zero regardless of its competitors’ qualities. This specification bounds relative qualities, but

²Markup is defined here as $(p - mc)/p$, where p is price and mc is marginal cost.

³See Singer, M. and D. Kawamoto, “AMD Files Antitrust Suite Against Intel,” CNet News.com, June 28, 2005.

it also implies the industry's steady-state innovation rate equals the exogenous rate at which the outside alternative's quality improves. Since we are interested in assessing the effect of competition on both pricing and innovation, we propose an alternative approach to bounding the state space in which the industry's innovation rate is endogenous.

Related Literature: The EP framework has been applied to a variety of durable goods markets. In his study of learning-by-doing, Benkard (2004) allows the market size for airplanes to change stochastically based on current sales to mimic the dynamic implications of forward-looking consumers. Markovich (forthcoming) and Markovich and Moenius (2005) have consumers that look two periods into the future in an oligopoly model of network effects. Dubé, Hitsch, and Chintagunta (2007) investigate indirect network effects in the video game console market. Esteban and Shum (2007) consider the effects of durability and secondary markets on equilibrium prices in the automobile industry, taking durability as fixed. Chen, Esteban, and Shum (2008) assess the competitive effect of secondary markets for automobiles. Finally, Song (2006) studies investment behavior in a model of the PC industry with static consumers of non-durable goods. Our focus on innovation when consumers are forward-looking is unique in this dynamic oligopoly literature.⁴

Our work also connects to the large theoretical literature on durable goods which is reviewed nicely by Waldman (2003). Early work focused on obtaining analytical results that often required strict assumptions. The most prominent of these assumptions is that old and new goods are perfect substitutes (in some proportion), that the infinite durability of the goods eliminates the need for replacement, and that the markets are either monopolies or perfectly competitive. Later work typically investigated the robustness of

⁴Our work also relates to recent empirical models of dynamic demand that take firm behavior as exogenous. Most of these papers consider high-tech durables, but restrict their attention to the initial product adoption decision. Melnikov (2001) develops a model of demand for differentiated durable goods that he applies to the adoption of computer printers. Carranza (2005) and Song and Chintagunta (2003) apply similar models to examine the introduction of digital cameras. More recently, Gordon (2008) and Gowrisankaran and Rysman (2007) have developed models that allow consumers to replace their products over time. By estimating supply as well as demand, we are able to compute equilibrium outcomes under counterfactual scenarios of interest.

the original conclusions to the relaxation of some of these assumptions.

Two strands of this literature are most relevant. The first area, starting with the works of Kleiman and Ophir (1966), Swan (1970), and Sieper and Swan (1973), asks whether a durable goods monopolist would provide the same level of durability as competitive firms and whether such a firm would choose the socially optimal level of durability. The so-called “Swan Independence Result” states that a monopolist indeed provides the socially optimal level of durability, though under strict assumptions. Rust (1986) shows that the monopolist provides less than the socially optimal level of durability when consumers’ scrappage rates are endogenous. Waldman (1996) and Hendel and Lizzeri (1999) show that Swan’s independence result fails to hold when new and used units are imperfect substitutes that differ in quality.

In our model the good’s depreciation rate (i.e., durability) is exogenous (and set to zero for our application to computer processors). Firms do, however, choose innovation rates which in turn determine the rate at which goods become obsolete. Though durability and obsolescence are similar, they have a potentially important difference: durability entails commitment since the good is produced and sold with a given durability, while a product’s obsolescence depends on the firm’s future innovations.

The second area, beginning with Coase (1972) and followed by Stokey (1981) and Bulow (1982), among others, considers the time inconsistency problem faced by a durable goods monopolist—it would like to commit to a fixed price over time, but after selling to today’s buyers it will then want to lower the price to sell to those consumers who were unwilling to buy at the supposedly fixed price.⁵ Bond and Samuelson (1984) show that depreciation and replacement sales reduce the monopolist’s price cutting over time. Since our model is in discrete time, the degree to which a firm can commit to a given price is exogenously specified by the length of our period. That is, firms are assumed to commit to fixed prices within periods, but not across periods.⁶

⁵A related area studies the problem of a monopolist pricing a new product, such as the next generation of a durable good. See, for example, Levinthal and Purohit (1989), Fudenberg and Tirole (1998), and Lee and Lee (1998). More recently, Nair (2007) estimates consumers’ initial adoption strategies and then numerically solves for the monopolist’s optimal intertemporal pricing schedule.

⁶An interesting comparative static may be to see how industry outcomes vary with period length.

The main drawback of the literature thus far is its focus on monopoly and perfect competition, whereas most durable goods (e.g., automobiles and appliances) are provided by oligopolies. By turning to numerical methods, we are able to study the interaction of innovation and pricing behavior in a dynamic oligopoly with forward-looking consumers.

2 Model

In this section we present a dynamic model of differentiated products oligopoly for a durable good. Time, indexed by t , is discrete with an infinite horizon. Each firm $j \in \{1, \dots, J\}$ sells a single product with time-varying log-quality denoted $q_{jt} \in \{0, \delta, 2\delta, \dots\}$.⁷ In each period, firms simultaneously choose their prices p_{jt} and investment x_{jt} .⁸ Price is a dynamic control since lowering price in period t increases current sales, but reduces future demand. Investment is a dynamic control since future quality is stochastically increasing in investment. Consumers decide each period whether to buy a new product or to continue using their currently owned product (if any).⁹ Hence, the distribution of currently owned products affects current demand. We denote this endogenous distribution Δ_t .

Firms and consumers are forward-looking and take into account the optimal dynamic behavior of the other agents (firms and consumers) when choosing their respective actions. We assume the vector of firms' qualities $q_t = (q_{1t}, \dots, q_{Jt})$ and the ownership distribution Δ_t is observed by all agents. These two state variables comprise the state space of payoff relevant variables for firms. The consumer's state space consists of the

Such a comparative static, however, is not trivial to construct since it involves changing the scale of several parameters simultaneously.

⁷We could normalize q_{jt} to be positive integers, but the estimated model is easier to interpret if the quality grid (and the implied innovation process) matches the data. We restrict firms to sell only one product because the computational burden of allowing multiproduct firms is prohibitive—the state space grows significantly and the optimization within each state becomes substantially more complex.

⁸The model does not allow entry or exit, primarily because of the lack of significant entry in the CPU industry. However, one could include entry and exit in the same manner as Ericson and Pakes (1995), if desired. Our focus on the CPU industry is also why we consider price, instead of quantity, as the choice variable: Intel and AMD both publish price lists and announce revisions to these lists.

⁹Some durable good markets, such as automobiles, have established used good markets. Only a small fraction of purchases of durable goods with rapid innovation, such as CPUs and consumer electronics, transact in used markets. As such, our model does not allow for resale of used goods.

quality of her currently owned product \tilde{q}_t , the firms' current offerings q_t , and the ownership distribution Δ_t . This latter state variable is relevant to the consumer since it affects firms' current and future prices and investment levels (i.e., innovation rates).

2.1 Consumers

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

$$u_{jt} = \gamma q_{jt} - \alpha p_{jt} + \xi_j + \varepsilon_{jt} \tag{1}$$

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

XXX: add footnote explaining brand fixed effect as single shot when purchase. conceptually easy to make part of flow utility, but increases state space, and no gain. value of single shot fixed effect will be expected discounted value of flows

Utility for a consumer from the outside alternative (i.e., no-purchase option) is

$$u_{0t} = \gamma \tilde{q}_t + \varepsilon_{0t} \tag{2}$$

where \tilde{q} denotes the quality of the product the consumer will use if no purchase is made this period.

One can think of our model as having two outside alternatives—one for consumers who have purchased at least once in the past, and one for “non-owners” who have never purchased the good. For consumers with previous purchases, \tilde{q}_t is the quality of their most recent purchase. For consumers who have yet to make an initial purchase, the utility from the no-purchase option could be determined by a variety of factors. In the context of CPUs, the outside good for “non-owners” may consist of using computers at schools and libraries or using old computers given to them by family or friends who have upgraded. To capture the notion that this outside alternative for non-owners improves as the frontier's quality improves, we specify that \tilde{q}_t is $(\max(q_t) - \bar{\delta}_c)$ for non-owners.

That is, quality of the outside alternative for non-owners is always $\bar{\delta}_c$ below the frontier quality. Furthermore, since everyone has access to this non-owner outside good, its quality serves as a lower bound to \tilde{q}_t for all consumers. We denote and define this lower bound as $\underline{q}_t = \bar{q}_t - \bar{\delta}_c$, where \bar{q}_t is defined to be $\max(q_t)$. To ensure that the model's behavior is not driven by our choice of $\bar{\delta}_c$, we check that consumers upgrade frequently enough that the quality of their most recent purchase is rarely below \underline{q}_t .

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

Given the lower bound \underline{q}_t , the ownership distribution can treat all consumers with $\tilde{q} \leq \underline{q}_t$ as owning the lower bound itself. Hence, $\Delta_t = (\Delta_{\underline{q}_t,t}, \dots, \Delta_{k,t}, \dots, \Delta_{\bar{q}_t,t})$, where $\Delta_{k,t}$ is the fraction of consumers in the population whose outside option (i.e., current product) has quality q_{kt} .¹⁰

Each consumer maximizes her expected discounted utility, which can be formulated using Bellman's equation as the following recursive decision problem:

$$V(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) = \max_{y_t \in (0,1,\dots,J)} u_{y_t,t} + \beta \sum_{q_{t+1}, \Delta_{t+1}} \int V(q_{t+1}, \Delta_{t+1}, \tilde{q}_{t+1}, \varepsilon_{t+1}) f_\varepsilon(\varepsilon_{t+1}) d\varepsilon_{t+1} \\ h_c(q_{t+1}|q_t, \Delta_t, \varepsilon_t) g_c(\Delta_{t+1}|\Delta_t, q_t, q_{t+1}, \varepsilon_t) \quad (3)$$

where y_t denotes the optimal choice in period t , $h_c(\cdot|\cdot)$ is the consumer's beliefs about future product qualities, $g_c(\cdot|\cdot)$ is the consumer's beliefs about the transition kernel for

¹⁰Here, we use the subscript k instead of j because these subscripts do not necessarily refer to products currently offered by any of the J firms. Furthermore, the dimension of Δ_t is $\bar{\delta}_c/\delta + 1$ which has no relation to J .

Δ_t , and f_ε is the density of ε . The expected continuation value depends on consumer's expectations about future products' qualities and future ownership distributions because these are the state variables that determine firms' future prices and investment levels. With an appropriate distributional assumption on $\{\varepsilon_{jt}\}$, we can derive an expression for the demand for each product based on the value function governing consumer behavior. The resulting demand system implies a law of motion for Δ_t and is used below in the model of firm behavior.

If $y_t = 0$ then $\tilde{q}_{t+1} = \max(\tilde{q}_t, \underline{q}_{t+1})$, else $\tilde{q}_{t+1} = q_{y_t, t}$ (i.e., the quality of the product just purchased). Note that once a consumer purchases a product at some quality level, the brand of the product no longer matters. That is, the consumer receives a one-time utility payoff of ξ_j from purchasing a product from firm j . This payoff does not occur in future periods since the outside option depends only on \tilde{q}_t . Relaxing this assumption would require Δ_t to be brand specific, which would substantially increase the state space.

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

2.2 Firms

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

$f(\cdot|x, q_t)$.¹¹ The dependence of τ_{jt} on q_t permits spillover effects in investment, which we model by specifying τ_{jt} to be stochastically increasing in $(\bar{q}_t - q_{jt})$ —the degree to which the firm is behind the frontier. As such, innovations are easier when catching up than when advancing the frontier.

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

$$\pi_j(p_t, q_t, \Delta_t) = Ms_{jt}(p_t, q_t, \Delta_t)(p_{jt} - mc_j) \quad (4)$$

where M is the (fixed) market size, $s_{jt}(\cdot)$ is the market share for firm j , p_t is the vector of J prices, and mc_j is firm j 's constant marginal cost of production. Each firm maximizes its expected discounted profits, which for firm j yields the Bellman equation

$$\begin{aligned} W_j(q_{jt}, q_{-j,t}, \Delta_t) &= \max_{p_{jt}, x_{jt}} \pi_j(p_t, q_t, \Delta_t) - cx_{jt} + \beta \sum_{\tau_{jt}, q_{-j,t+1}, \Delta_{t+1}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}) \\ &\quad f(\tau_{jt}|x_{jt}) h_f(q_{-j,t+1}|q_t, \Delta_t) g_f(\Delta_{t+1}|\Delta_t, q_t, p_t, q_{t+1}) \end{aligned} \quad (5)$$

where c is the unit cost of investment, $h_f(\cdot|\cdot)$ is the firm's beliefs about its competitors future product quality levels, and $g_f(\cdot|\cdot)$ is the firm's beliefs about the transition kernel for Δ_t which is based on beliefs about consumers choices given prices and qualities. The dependence of g_f on q_{t+1} reflects the shifting of Δ_t when the frontier quality increases, as described below, since Δ_t is defined only for quality levels within $\bar{\delta}_c$ of the frontier.¹²

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

$$\begin{aligned} \hat{V}_j(q_t, \Delta_t, \tilde{q}_t) &= u_{jt} - \varepsilon_{jt} + \beta \sum_{q_{t+1}, \Delta_{t+1}} \log \left(\sum_{j' \in \{0, \dots, J\}} \exp \left\{ \hat{V}_{j'}(q_{t+1}, \Delta_{t+1}, \tilde{q}_{t+1}) \right\} \right) \\ &\quad h_c(q_{t+1}|q_t, \Delta_t) g_c(\Delta_{t+1}|\Delta_t, q_t, p_t, q_{t+1}) . \end{aligned} \quad (6)$$

¹¹We specify quality as being in logs, so that improvements are proportional increases in quality. Using a log scale makes more sense than a linear scale when calibrating the model to the CPU industry.

¹²Expressing the transition of Δ_t conditional on the realized q_{t+1} (which combines the realizations of τ_{jt} and $q_{-j,t+1}$) simplifies the derivation of optimal investment in Section 2.3.

The conditional choice probabilities for a consumer owning product \tilde{q} are therefore

$$s_{jt|\tilde{q}} = \frac{\exp\{\hat{V}_j(q_t, \Delta_t, \tilde{q}_t)\}}{\sum_{k \in \{0, \dots, J\}} \exp\{\hat{V}_k(q_t, \Delta_t, \tilde{q}_t)\}}. \quad (7)$$

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

$$s_{jt} = \sum_{\tilde{q} \in \{q_t, \dots, \bar{q}_t\}} s_{jt|\tilde{q}} \Delta_{\tilde{q}, t}. \quad (8)$$

These market shares translate directly into the law of motion for the distribution of ownership.¹³ Recall that Δ_t only tracks ownership of products within $\bar{\delta}_c$ quality units of the highest quality offering. Assuming this highest quality is unchanged between t and $t + 1$, the share of consumers owning a product of quality k at the start of period $t + 1$ is

$$\Delta_{k, t+1}(\cdot) = s_{0t|k} \Delta_{kt} + \sum_{j=1, \dots, J} s_{jt} I(q_{jt} = k) \quad (9)$$

where the summation accounts for the possibility that multiple firms may have quality k . For quality levels not offered in period t , this summation is simply zero. If a firm advances the quality frontier with a successful R&D outcome pushing its $q_{j, t+1}$ beyond $\max(q_t)$ then Δ_{t+1} shifts: the second element of Δ_{t+1} is added to its first element, the third element becomes the new second element (and so on), and the new last element is initialized to zero. Formally, define the shift operator Γ on a generic vector $y = (y_1, y_2, \dots, y_L)$ as $\Gamma(y) = (y_1 + y_2, y_3, \dots, y_L, 0)$. If the quality frontier advances at the end of period $t + 1$, then we shift the interim Δ_{t+1} that results from equation (9) via

$$\Delta_{t+1} = \Gamma(\Delta_{t+1}). \quad (10)$$

The continuation ownership distribution is therefore a deterministic function of prices, except for the potential shift due to the stochastic innovation of frontier products.

Finally, we note that physical depreciation of goods could be easily added to the

¹³For conciseness our notation suppresses the dependence of market shares on prices.

model by supposing that each currently owned product declines by δ (i.e., one quality grid-step) with some fixed probability. In our application of the model to the CPU industry, however, physical depreciation is zero.

2.3 Optimal Prices and Investments

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

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

$$\begin{aligned}
 -c + \beta \sum_{\tau_{jt}, q_{-j,t+1}, \Delta_{t+1}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}) \\
 h_f(q_{-j,t+1}|q_t, \Delta_t) g_f(\Delta_{t+1}|\Delta_t, q_t, p_t, q_{t+1}) f(\tau_{jt}|x_{jt}) \frac{\partial f(\tau_{jt}|x_{jt})}{\partial x_{jt}} = 0.
 \end{aligned} \tag{11}$$

Given q_t and outcomes for $(\tau_{jt}, q_{-j,t+1})$ the transition for Δ_{t+1} depends only on prices. Thus, with a suitable choice for f we can analytically compute the optimal investment as a function of price, $x_{jt}^*(p_{jt})$. We provide details in Appendix A.

2.4 Equilibrium

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

The equilibrium specifies that (1) firms' and consumers' equilibrium strategies must only depend on the current state variables (which comprise all payoff relevant variables), (2) consumers possess rational expectations about firms' policy functions (which deter-

mine future qualities and prices) and the evolution of the ownership distribution, and (3) each firm possesses rational expectations about its competitors' policy functions for price and investment and about the evolution of the ownership distribution.

Formally, a MPNE in this model is the set $\left\{ V^*, h_c^*, g_c^*, \left\{ W_j^*, x_j^*, p_j^*, h_{f_j}^*, g_{f_j}^* \right\}_{j=1}^J \right\}$, which contains the equilibrium value functions for the consumers and their beliefs h_c^* about future product qualities, beliefs g_c^* about future ownership distributions, and the firms' value functions, policy functions, beliefs $h_{f_j}^*$ over their $J - 1$ rivals' future qualities, and beliefs $g_{f_j}^*$ about the future ownership distribution. The expectations are rational in that the expected distributions match the distributions from which realizations are drawn when consumers and firms behave according to their policy functions. In particular, $h_c^*(q_{t+1}|q_t, \Delta_t, \tilde{q}) = \prod_{j=1}^J f(\tau = q_{j,t+1} - q_{jt} | q_{jt}, x_{jt}^*)$, $h_{f_j}^*(q_{-j,t+1}|q_t, \Delta_t) = \prod_{j' \neq j}^J f(\tau = q_{j',t+1} - q_{j't} | q_{j't}, x_{j't}^*)$, and g_c^* and $g_{f_j}^*$ are derived from the law of motion for Δ_t as described by equations (9) and (10).¹⁴

The functional form of the investment transition function satisfies the UIC admissibility criterion in Doraszelski and Satterthwaite (2007). To guarantee existence of equilibrium requires us to show that there is a pure-strategy equilibrium in both investment choices and prices. Ericson and Pakes (1995) and the various extensions found in Doraszelski and Satterthwaite (2007) do not consider dynamic demand. As such, they are able to construct a unique equilibrium in the product market in terms of prices or quantities (depending on the specific model of product market competition).

3 Computation

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

¹⁴Symmetry corresponds to $W_j^* = W^*$, $x_j^* = x^*$, $p_j^* = p^*$, $h_{f_j}^* = h_f^*$, and $g_{f_j}^* = g_f^*$ for all j . Symmetry obviously requires that firm specific parameters, such as brand intercepts ξ_j , are the same across firms.

3.1 Bounding the State Space

The state space in the model presented in Section 2 is unbounded since product qualities increase without bound. To solve for equilibrium, we transform the state space to one that is finite. Rather than measuring qualities on an absolute scale, we measure all qualities relative to the current period's maximum quality $\bar{q}_t = \max(q_t)$. Our ability to implement this transformation without altering the dynamic game itself hinges on the following proposition.

Proposition 1. *Shifting q_t and \tilde{q} by \bar{q}_t affects firms and consumers as follows:*

$$\begin{aligned} \text{Firms: } W_j(q_{jt}, q_{-j,t}, \Delta_t) &= W_j(q_{jt} - \bar{q}_t, q_{-j,t} - \bar{q}_t, \Delta_t) \\ \text{Consumers: } V(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) &= \frac{\gamma \bar{q}_t}{1-\beta} + V(q_t - \bar{q}_t, \Delta_t, \tilde{q}_t - \bar{q}_t, \varepsilon_t). \end{aligned} \tag{12}$$

The proof, which appears in Appendix B, rests on the following properties of the model:

1. Quality (actually log-quality) enters linearly in the utility function, so that adding any constant to the utility of each alternative has no effect on consumers' choices.
2. Innovations are governed by $f_\tau(\cdot)$ which is independent of quality levels (though $f_\tau(\cdot)$ does depend on differences in qualities).
3. Δ_t is unaffected by the shift since it is defined as the ownership shares of only those products within $\bar{\delta}_c$ of the frontier. That is, Δ is already in relative terms.

The proposition also claims that the change in the consumer's value, when her \tilde{q} and the industry's offered qualities q_t are all shifted down by \bar{q} , can be decomposed into a component driven by relative values and a component driven by absolute levels. The shift by $(-\bar{q}_t)$ to the relative qualities in the arguments of V on the right-hand side subtracts $\gamma \bar{q}_t$ from utility in each period (for all realizations of future states). To restore equality the present value of this lost utility in every period, $\frac{\gamma \bar{q}_t}{1-\beta}$, must be added back, which is accomplished by the first-term on the right-hand side.

The only subtlety in implementing the transformation is in computing the continuation values in equations (5) and (6) (or equivalently (3)). When integrating over future states, some of the possible states involve an improvement in the frontier's quality (always by δ units). In this event, the consumer's continuation value for that particular outcome is $\gamma\delta/(1-\beta) + V(q_{t+1} - \delta, \Delta_{t+1}, \tilde{q}_{t+1} - \delta, \varepsilon_{t+1})$ instead of $V(q_{t+1}, \Delta_{t+1}, \tilde{q}_{t+1}, \varepsilon_{t+1})$.

To facilitate writing the value functions in terms of a relative state space, we define $\omega_t = q_t - \bar{q}_t$ and $\tilde{\omega}_t = \tilde{q}_t - \bar{q}_t$ as analogs to the original state variables. We also define the indicator variable $I_{\bar{q}_t} = 1$ if $\bar{q}_{t+1} > \bar{q}_t$ to indicate whether the frontier product improved in quality from period t to $t + 1$. The consumer's product-specific value function in equation (6) can then be expressed using the relative state space as

$$\begin{aligned} \hat{V}_j(\omega_t, \Delta_t, \tilde{\omega}_t) &= \gamma\omega_{jt} - \alpha p_{jt} + \beta \sum_{I_{\bar{q}_t}, \omega_{t+1}, \Delta_{t+1}} \log \left(\sum_{j' \in \{0, \dots, J\}} \exp \left\{ \frac{\gamma\delta I_{\bar{q}_t}}{1-\beta} + \hat{V}_{j'}(\omega_{t+1}, \Delta_{t+1}, \tilde{\omega}_{t+1}) \right\} \right) \\ &\quad h_c(I_{\bar{q}_t}, \omega_{t+1} | \omega_t, \Delta_t) g_c(\Delta_{t+1} | \Delta_t, \omega_t, p_t, I_{\bar{q}_t}) \end{aligned} \quad (13)$$

where the outside alternative's p_{0t} is zero and, in a slight abuse of notation, $h_c(I_{\bar{q}_t}, \omega_{t+1} | \omega_t, \Delta_t)$ and $g_c(\Delta_{t+1} | \Delta_t, \omega_t, p_t, I_{\bar{q}_t})$ are the analogs of the consumer's transition kernels for q_{t+1} and Δ_{t+1} in the original state space.

Firm j 's value function in equation (5) using the relative state space becomes

$$\begin{aligned} W_j(\omega_{jt}, \omega_{-j,t}, \Delta_t) &= \max_{p_{jt}, x_{jt}} \pi_j(p_t, \omega_t, \Delta_t) - cx_{jt} + \\ &\quad \beta \sum_{\tau_{jt}, \omega_{-j,t+1}, I_{\bar{q}_t}, \Delta_{t+1}} W_j(\omega_{jt} + \tau_{jt} - I_{\bar{q}_t}, \omega_{-j,t+1} - I_{\bar{q}_t}, \Delta_{t+1}) \quad (14) \\ &\quad h_f(I_{\bar{q}_t}, \omega_{-j,t+1} | \omega_t, \Delta_t) g_f(\Delta_{t+1} | \Delta_t, \omega_t, p_t, I_{\bar{q}_t}) f(\tau_{jt} | x_{jt}) \end{aligned}$$

where $\omega_{-j,t+1}$ refers to competitors' continuation qualities prior to shifting down by δ in the event that the frontier's quality improved. Again, we slightly abuse notation by using $h_f(I_{\bar{q}_t}, \omega_{-j,t+1} | \omega_t, \Delta_t) f(\tau_{jt} | x_{jt})$ and $g_f(\Delta_{t+1} | \Delta_t, \omega_t, p_t, I_{\bar{q}_t})$ as the analogs of the firm's transition kernels for competitors' qualities and Δ_{t+1} .

Finally, we invoke a knowledge spillover argument to bound the difference between each firm's own quality and the frontier quality. We denote the maximal difference in

firms' qualities $\bar{\delta}_f$ and impose this maximal difference directly in the transition kernels $f(\cdot)$ and $h_f(\cdot)$. We choose $\bar{\delta}_f < \bar{\delta}_c$ to capture the fact that quality differences among new products are typically less than the quality difference between the frontier and the low end of products from which consumers have yet to upgrade. We also choose $\bar{\delta}_f$ to be sufficiently large that it has minimal effect on equilibrium strategies.¹⁵ In particular, we verify that investment behavior when a laggard is maximally inferior does not suggest that laggards simply “free ride” off leader’s innovations. We also check that the laggard rarely reaches this maximal inferiority state.

Comparison with Ericson-Pakes: Researchers have applied and extended the EP framework to study a variety of dynamic differentiated products industries, as detailed by Doraszelski and Pakes (2006). In each case, the state space is bounded by defining firms' qualities relative to an outside good and assuming consumers have concave preferences for this relative quality. Standard discrete choice models often specify diminishing marginal utility for absolute levels of quality, but not for quality measured relative to an outside alternative.¹⁶ This concavity implies the derivative of market share with respect to a firm’s own quality goes to zero regardless of its competitors’ qualities. Since investment is costly, there will exist a relative quality above which investment is zero, thereby establishing an upper bound. The lower bound is established by firms exiting when their relative quality gets low.

The EP approach to bounding the state space has a few potential drawbacks. Perhaps the most significant is that the industry’s innovation rate is exogenously specified by the innovation rate of the outside good. Improvements in the outside good provide a continual need for inside firms to invest to remain competitive. If the outside good never improves, the steady-state equilibrium has no investment and no innovations.

¹⁵Note that if firms were permitted to exit, then firms' relative differences would be bounded automatically by the exiting of firms with sufficiently low relative quality.

¹⁶The standard normalization in discrete choice models is to subtract the mean utility of the outside good from all options. The EP approach, however, fixes the mean utility of the outside good to zero and subtracts its (absolute) quality from firms' qualities inside a concave function. See Pakes and McGuire (1994) for details.

Since consumer surplus over time may be driven more by an industry’s innovation rate than by its pricing behavior, providing an endogenous long-run innovation rate that varies across market structures and regulatory controls seems particularly important for policy work. In our model the long-run rate of innovation is an equilibrium outcome that depends on consumer preferences, firms’ costs, and any regulatory stipulations in effect.

Another odd implication is that consumers’ utility rankings of the inside goods depends on the level of the outside good’s quality. For example, a consumer indifferent between two products that differ by one quality level will strictly prefer the higher of the two products if the outside alternative improves. As such, the relative market shares depend on the outside good’s absolute quality: as the outside good increases, holding inside goods fixed, the relative market share of the higher quality product increases. Since our model uses a linear specification over the quality index, utility rankings and relative shares are independent of the outside good’s quality.¹⁷

Finally, we note that our relative state-space is truly a normalization: the game using the relative state space is an exact transformation of a dynamic game that is initially expressed in absolute terms. Our focus on durable goods forces us to develop a model that is, from the consumer’s perspective, consistent with respect to shifts in relative qualities when the baseline good’s quality changes. A consumer purchases a product with quality q_{jt} knowing that this product has expected utility γq_{jt} next period as well (since depreciation is zero). The quality index must enter utility linearly to correctly sum these flows across time in a relative state space reformulation. Such consistency issues are best addressed by initially writing the model in absolute levels.¹⁸ Conceptually, developing the model in levels is advantageous since agents’ primitives are, in most cases, more naturally defined over absolute levels.

In essence, the EP bounds approach defines quality relative to the outside good and generate an upper bound by manipulating the behavior of lead firms, whereas we define

¹⁷The quality index may itself be a concave transformation of measured quality, to reflect diminishing marginal utility for measured quality.

¹⁸Goettler, Parlour, and Rajan (2005, 2006) also take the approach of initially specifying their dynamic asset trading game in absolute levels and then exactly transforming the game to a relative state.

quality relative to the frontier and generate a lower bound by truncating the degree to which firms and outside options can be inferior. Since the industry leaders generate most of the sales, profits, and surplus, assumptions regarding severe laggards are more innocuous than assumptions restricting the benefits to innovation by frontier firms.¹⁹

3.2 Approximation of Δ_t

A challenge in solving our model is that Δ_t is a high-dimensional simplex. We approximate this continuous state variable with a discretization that restricts $\Delta_t \in \{\Delta^d\}_{d=1}^D$. To be precise, let Δ'_{t+1} denote the (unapproximated) transition implied by (9) and (10), and let $\rho_d(\Delta'_{t+1})$ denote the distance between Δ'_{t+1} and the d^{th} distribution of our discretization. Several candidate distance metrics are available: the Kullback-Leibler divergence measure, sum of squared errors of PDFs or CDFs, and the mean, among others. Since we are using the approximation to obtain firms' and consumers' continuation values the distance metric should be based on moments of the distribution that are most relevant to the future profitability, pricing, and investment behavior. For logit based demand systems the mean is the most relevant moment.²⁰ We therefore define

$$\rho_d(\Delta'_{t+1}) = \left| \sum_k k \Delta'_{k,t+1} - \sum_k k \Delta_k^d \right|, \text{ for all } d \in (1, \dots, D), \quad (15)$$

where the summation is over the discrete qualities from \underline{q} to \bar{q} tracked by Δ .

Now let d_1 and d_2 denote the superscripts of the two distributions closest to Δ'_{t+1} ,

¹⁹While we developed this alternative bounds approach in the context of durable goods, its merits apply equally to the non-durable case. In future research, we will assess the effect of using our specification for the non-durable case explored in Pakes and McGuire (1994). We also note that some features of the Ericson-Pakes specification that we have excluded may be adopted. For example, the outside good's quality (for non-owners) may be allowed to stochastically improve according to an exogenous process exactly as in Ericson and Pakes. A bound on inside goods' qualities relative to the outside good can also be imposed. We would argue, however, that this bound should be generated by modifying the innovation technology $f(\cdot|\cdot)$ to reflect the difficulty of advancing too quickly, rather than by modifying the demand system to eliminate the benefit of higher quality.

²⁰Fixing consumers' conditional choice probabilities and firms' relative qualities, we generate random ownership distributions and regress the resulting profits on moments of the random Δ s. The mean is easily the best predictor of a Δ 's profitability, with an R^2 of .995.

and define the stochastic transition of the discretized Δ as

$$\Delta_{t+1} = \begin{cases} \Delta^{d_1} & \text{with probability } \frac{\rho_{d_1}}{\rho_{d_1} + \rho_{d_2}} \\ \Delta^{d_2} & \text{with probability } \frac{\rho_{d_2}}{\rho_{d_1} + \rho_{d_2}} \end{cases}. \quad (16)$$

Our discretization is a multidimensional version of the stochastic transition used by Benkard (2004) to approximate production experience. With multiple dimensions, one could consider transitioning to more than the two closest points, whereas with one dimension using only the two closest is obvious. To retain the suitability of using only the closest two distributions, we generate $\{\Delta^d\}_{d=1}^D$ from a single family of distributions parameterized by a scalar.²¹ The discrete grid of this scalar is chosen such that the mean qualities implied by the $\{\Delta^d\}_{d=1}^D$ are .1 apart and range from $\bar{q} - 11$ to $\bar{q} - 1$. The fact that profitability and prices are driven primarily by the mean quality of the ownership distribution suggests that this set of $\{\Delta^d\}_{d=1}^D$ is sufficiently rich to capture the tradeoffs associated with dynamic demand.²²

This approximation retains the key feature of dynamic demand—lowering price today reduces expected future demand. Since Δ_{t+1} is used only in computing continuation values, the effect of the approximation on the equilibrium will depend on the degree of curvature in value functions with respect to Δ and on the coarseness of the discretization.

3.3 Solving and Simulating Industry Equilibrium

We compute the equilibrium using a Gauss-Jacobi scheme to update the value and policy functions.²³ Starting at iteration $k = 0$, we initialize the consumer value function \hat{V}^0 and the firms' value functions W_j^0 each to zero, and policy functions (x_j^{*0}, p_j^{*0}) to yield innovation rates of 50 percent and prices 20 percent higher than marginal costs.²⁴

²¹We use the logit, which has CDF for the k^{th} quality level of $z \exp(q_k)/(1 + z \exp(q_k))\kappa$ where z is the scalar parameter and $\kappa = z \exp(\bar{q})/(1 + z \exp(\bar{q}))$ is a normalization constant.

²²This choice for $\{\Delta^d\}_{d=1}^D$ is also computationally efficient since the two closest Δ^d s may be obtained using an indexing formula instead of a search.

²³For robustness we also compute the equilibrium using Gauss-Seidel with random orderings of the states, as discussed in Section 4.2.1.

²⁴The consumers smoothed value function (which integrates over the idiosyncratic ε) is $\hat{V}(\omega, \Delta, \tilde{\omega}) = \log \sum_{j=0, \dots, J} \hat{V}_j(\omega, \Delta, \tilde{\omega})$, where the product-specific \hat{V}_j functions are defined in equation (13).

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

1. For each $\tilde{\omega} \in \Delta$, evaluate the consumer's value function \hat{V}^k given the firms' policy functions $\{x_j^{*k-1}, p_j^{*k-1}\}_{j=1}^J$ and the next period approximation parameters ρ'^{k} from the previous iteration.
2. For each $j \in J$, evaluate firm j 's value function W_j^k given the other firms' policy functions from the previous iteration $\{x_{j'}^{*k-1}, p_{j'}^{*k-1}\}_{j' \neq j}^J$.
3. Update the consumer value functions $\hat{V}^{k+1} \leftarrow \hat{V}^k$, the firms value functions $W_j^{k+1} \leftarrow W_j^k, \forall j$, and their policy functions $\{x_j^{*k+1}, p_j^{*k+1}\} \leftarrow \{x_j^{*k}, p_j^{*k}\}$.
4. Check for convergence in the sup norm of all agents' value functions. If convergence is not achieved, return to step (1).

To simulate the converged model we first specify an initial state for the industry (ω_0, Δ_0) . Then for each simulated period $t = 0, \dots, T$, we implement each firm's optimal price and investment according to the equilibrium policy functions, process the stochastic evolution of ownership as described in equation (16), and process the stochastic innovation outcomes according to $f(\cdot|\cdot)$.

4 Empirical Application

This paper has two components—a theory component that develops a general model of durable goods competition and an empirical component that applies the model to the CPU industry.

In the empirical application, we account for important asymmetries between Intel and AMD by allowing them to differ in their costs of production and innovation and brand equity. For the purposes of illustrating some theoretical properties of the model, we consider the symmetric case in which firms have identical brand intercepts, innovation efficiencies, and marginal costs. In the symmetric setting differences in firm behavior are entirely due to the stochastic nature of innovation and evolution of the industry, as opposed to being driven by exogenous asymmetries.

A natural question to ask is why have these asymmetries developed? To answer such a question, one would need a model that endogenizes the evolution of brand equity, marginal costs, and innovation costs, which is well beyond our current objectives. We take these asymmetries to be exogenous and investigate their implications for pricing and investment strategies, and their impact on profits and surplus.

4.1 Data

Using data from 1993 to 2004, we construct empirical moments for the PC processor industry. Over this period the industry is essentially a duopoly, with Intel and AMD controlling about 95 percent of the market.

We obtained information on PC processor unit shipments, manufacturer prices, and product quality measures, by processor, from a variety of sources. We obtained quarterly global unit shipment data from In-Stat/MDR, an industry research firm that specialized in microprocessors. Lastly, we used a PC processor speed benchmark from the CPU Scorecard (www.cpuscorecard.com) to obtain a single index of quality that is comparable across firms and product generations.²⁵

Our model restricts each firm to offer only one product. Since AMD and Intel each offer several products in each period, we must average the prices and qualities. Figure 1 presents the log qualities and prices of each firm’s offerings from 1993 Q1 to 2004 Q4.²⁶ Intel clearly dominates in the early period with much higher quality. AMD’s introduction of the K6 processor in early 1997 narrows the gap, but parity is not achieved until the introduction of the AMD Athlon in mid-1999. The change in Intel’s log-quality from over the 48 quarters corresponds to approximately an average of 11 percent improvement per quarter. This rate of improvement translates into a doubling in CPU speed every 6.7 quarters, which is consistent with the 18 to 24 month elapsed-time for doubling transistor counts on integrated circuits that has become known as “Moore’s Law”.

²⁵The list of processor speed ratings from the CPU Scorecard does not contain all the processors in the data set: 74 of 217 processors did not have benchmarks (38 from AMD and 36 from Intel). To fill in the missing values, we impute the missing benchmark based on the available ratings.

²⁶All prices are converted to base year 2000 dollars.

The correlation between Intel's price and its quality advantage is evident by comparing the last plot in Figure 1 with Intel's price in each of the first three plots. Average prices were generally above \$300 before the K6 introduction and around \$200 after the Athlon was introduced. The correlation between average prices and the difference between Intel and AMD's quality is .63 for Intel and $-.57$ for AMD. Such correlations are consistent with the theoretical model presented in Section 2.

The first two price graphs show (local) peaks at the end of 1999, despite the two firm's being roughly tied on the quality dimension. These high prices reflect the high demand due to the tech bubble and to the fact that each firm recently increased the speed of its chips. Our model accounts for variations in prices due to the recent upgrades, but does not contain aggregate demand shocks which would be necessary to match precisely the variance in prices, for example.

Figure 2 illustrates the dominance of Intel, as evidenced by a market share that barely dips below 80 percent. The steady expansion of total CPU shipments from 1993 through 1999, followed by the flattening of shipments until growth resumes in 2004 is also evident. The slow period from 2000 through 2003 primarily reflects the shock of the bursting of the late nineties tech bubble. Since our model does not implement aggregate shocks, we will not concern ourselves with this aspect of the data.

Figure 3 reveals another striking asymmetry in this industry: AMD invests in R&D less than one-fourth the amount invested by Intel, yet is able to offer similar, sometimes even higher, quality products beginning in 1999. This asymmetry warrants estimating separate investment efficiencies for each firm.

4.2 Estimation

To estimate the model's parameters we use a method of simulated moments (MSM) estimator which minimizes the distance between a set of unconditional moments of our data and their simulated counterparts implied by our model. Hall and Rust (2003) refer to this type of estimator as a simulated minimum distance (SMD) estimator because it minimizes a weighted distance between actual and simulated moments. The esti-

mator may also be viewed as taking the indirect inference approach of Smith (1993), Gouriéroux, Monfort, and Renault (1993), and Gallant and Tauchen (1996) in which the moments to match are derived from an auxiliary model that is easier to evaluate than the structural model of interest. Regardless of the label used, the estimator is in the class of generalized method of moments (GMM) estimators introduced by Hansen (1982) and augmented with simulation by Pakes and Pollard (1989).

The general idea behind the SMD estimator is to assume that the data are observations from a stationary distribution and to search over the space of structural parameters to find the model which has a stationary distribution that yields moments matching the actual moments. For example, the average prices in the actual and simulated data ought to be similar. The estimator is obtained as the minimum of a quadratic form of the difference between the actual and simulated moments. We use enough simulations that the variance in the estimator is due entirely to the finite sample size. As such, the efficient weight matrix is the inverse of the covariance matrix of the actual data's moments.²⁷

4.2.1 A Simulated Minimum Distance Estimator

In this subsection we present the assumptions and details of our estimator, such as which moments to match. Our presentation follows Hall and Rust (2003).

The model presented in Section 2 generates a stochastic process for $\mu_t = \{\omega_t, \Delta_t, p_t, x_t, s_t\}$, where ω denotes qualities relative to the frontier, Δ is the ownership distribution, p denotes prices, x denotes investments, and s denotes market shares.

²⁷Since we obtain the efficient weight matrix directly from the data, we do not need a two-step GMM estimator to obtain efficiency.

The transition density, f_μ , for this Markov process is given by

$$\begin{aligned}
f_\mu(\omega_{t+1}, \Delta_{t+1}, p_{t+1}, x_{t+1}, s_{t+1} | \omega_t, \Delta_t, p_t, x_t, s_t, \theta) &= \prod_{j=1}^J f(\omega_{j,t+1} - \omega_{j,t} | \omega_t, x_t) \\
&\times g(\Delta_{t+1} | \Delta_t, s_t) \\
&\times I\{p_{t+1} = p(\omega_{t+1}, \Delta_{t+1})\} \\
&\times I\{x_{t+1} = x(\omega_{t+1}, \Delta_{t+1})\} \\
&\times I\{s_{t+1} = s(\omega_{t+1}, \Delta_{t+1})\}
\end{aligned} \tag{17}$$

where θ denotes the vector of K parameters to be estimated. Note that f_μ is degenerate since prices, investments, and market shares are deterministic functions of the state variables ω_{t+1} and Δ_{t+1} . The model would need to be modified, perhaps by adding aggregate shocks, if we were to use maximum likelihood since the data would almost surely contain observations having zero likelihood. This degeneracy, however, is not a problem for the SMD estimator we define below, because it is based on predicting moments of the distribution μ_t , not particular realizations of μ_t given μ_{t-1} .²⁸

For each candidate value of θ encountered, we solve for equilibrium and simulate the model S times for T periods each, starting at the initial state (ω_0, Δ_0) which we observe in the data. These $S \times T$ simulated periods each have three stochastic outcomes—each firms' investment outcome and the random transition of Δ_t . The set of i.i.d. $U(0,1)$ draws for these outcomes, denoted $\{\{U_t^n\}_{t=1}^T\}_{n=1}^S$, is held fixed throughout the estimation procedure to preserve continuity of the estimator's objective function. The set of simulated industry outcomes is denoted $\{\{\mu_t(\theta, U_{<t}^n, \omega_0, \Delta_0)\}_{t=1}^T\}_{n=1}^S$, where the subscript in $U_{<t}^n$ indicates that μ_t^n depends on only the first $t - 1$ realizations of U^n .

The SMD estimator is the parameter value that yields the best fit of simulated moments to moments in the data. But which moments should we match? Gallant and Tauchen (1996) suggest using the score of an auxiliary model which closely approximates the distribution of the data. If the auxiliary model nests the structural model then the estimator is as efficient as maximum likelihood. Hall and Rust (2003) use simple

²⁸One reason we do not attempt to predict each period's realization of μ_t given μ_{t-1} is that we do not observe Δ_t in each period.

statistics such as means and covariances. We use a mixture of simple moments and estimates of approximations to policy functions, as embodied by a function m which maps either real or simulated data into the L moments we match. The specification of m is provided in the next subsection. Identification requires $L \geq K$.

The vector of moments using actual data is denoted $m_T \equiv m(\{\mu_t^{actual}\}_{t=1}^T)$ and the simulated moment vector is the average over the S simulations:

$$m_{S,T}(\theta) = \frac{1}{S} \sum_{n=1}^S m(\{\mu_t(\theta, U_{<t}^n, \omega_0, \Delta_0)\}_{t=1}^T). \quad (18)$$

The simulated minimum distance estimator $\hat{\theta}_T$ is then defined as

$$\hat{\theta}_T = \underset{\theta \in \Theta}{\operatorname{argmin}} (m_{S,T}(\theta) - m_T)' A_T (m_{S,T}(\theta) - m_T) \quad (19)$$

where A_T is an $L \times L$ positive definite weight matrix.

Assumption 1. *For any $\theta \in \Theta$ the process $\{\mu_t(\theta, U_{<t}^n, \omega_0, \Delta_0)\}$ is ergodic with unique invariant density $\Psi(\mu|\theta)$ given by:*

$$\Psi(\mu'|\theta) = \int f_\mu(\mu'|\mu, \theta) d\Psi(\mu|\theta). \quad (20)$$

A valid concern with using moments based on simulated equilibrium outcomes is that the equilibrium may not be unique. Two-stage approaches in which policy functions are first estimated nonparametrically, as in Bajari, Benkard, and Levin (2007), permit the model to have multiple equilibria. The assumption they invoke is that the data are all from the same equilibrium, which is clearly a weaker assumption than Assumption 1. Unfortunately, we do not have sufficient data to use a two-stage approach. We note that other researchers have proceeded with SMD estimation of equilibrium models without establishing uniqueness. For examples, see Gowrisankaran and Town (1997), Epple and Seig (1999), and Xu (2007).

As in Pakes and McGuire (1994), we use value function iteration to numerically solve

for equilibrium. Whereas general equation solvers may be computationally superior, value function iteration has the conceptual advantage of restricting the set of equilibria to those that are limits of finite horizon games.²⁹ Although this equilibrium refinement does not guarantee uniqueness, it likely reduces the potential for multiplicity.

We have also established that the same equilibrium is obtained when we tweak the value function algorithm in various ways. For example, we consider Gauss-Jacobi updating and Gauss-Seidel updating with randomly determined orderings of the states.³⁰ Finally, we note that the objective function defining the SMD estimator appears to be smooth, which would not be the case if the algorithm were jumping to different equilibria with small changes in the model.

Assumption 2. *The structural model presented in Section 2 is correctly specified. As such, there exists a $\theta^* \in \Theta$ for which each simulated sequence $\{\mu_t^n\}, n = 1, \dots, S$ from the initial state (ω_0, Δ_0) has the same probability distribution as the observed sequence $\{\mu_t\}$.*

This assumption enables us to use the standard GMM formula for the asymptotic covariance matrix of $\hat{\theta}_T$. We could alternatively relax this assumption, and bootstrap the covariance matrix.

Define the functions $E[m|\theta]$, $\nabla E[m|\theta]$, and $\nabla m_{S,T}$ as:

$$\begin{aligned} E[m|\theta] &= \int m(\mu) d\Psi(\mu|\theta) \\ \nabla E[m|\theta] &= \frac{\partial}{\partial \theta} E[m|\theta] \\ \nabla m_{S,T} &= \frac{\partial}{\partial \theta} m_{S,T}(\theta). \end{aligned} \tag{21}$$

²⁹Initializing firms' and consumers' value functions to zero implies a terminal payoff of zero in the last period. The first iteration then solves for optimal strategies of a one period game. The second iteration solves for optimal strategies of a two-period game, and so on. When determining a current iteration's optimal actions, the algorithm actually holds other players' policy functions fixed at the previous iteration's values. The resulting update is therefore not really an equilibrium, unless the policy functions have indeed converged. One could modify the algorithm to literally compute, say, the two-period equilibrium in its second iteration, but at great computational cost. As the algorithm progresses, the implied policy functions converge to an equilibrium of the infinite horizon game.

³⁰The ordering has an effect because Gauss-Seidel uses updated values as soon as they become available within an iteration.

Assumption 3. θ^* is identified; that is, if $\theta \neq \theta^*$ then $E[m|\theta] \neq E[m|\theta^*] = E[m(\{\mu_t^{actual}\}_{t=1}^T)]$. In addition, $\text{rank}(\nabla E[m|\theta]) = K$ and $\lim_{T \rightarrow \infty} A_T = A$ with probability 1 where A is an $L \times L$ positive definite matrix.

The optimal weight matrix is $\Omega(m, \theta^*)^{-1} \equiv E[(m(\mu) - E[m(\mu)])(m(\mu) - E[m(\mu)])']^{-1}$, the inverse of the covariance matrix of the moment vector, where the expectation is taken with respect to the ergodic distribution of μ given $\theta = \theta^*$. Using $A_T = [\text{cov}(\{\mu_t^{actual}\}_{t=1}^T)]^{-1}$ as a consistent estimate of the optimal weight matrix, the estimator $\hat{\theta}_T$ has the property

$$\sqrt{T}(\hat{\theta}_T - \theta^*) \implies N\left(0, (1 + 1/S) (\nabla E[m|\theta^*]' \Omega(m, \theta^*)^{-1} \nabla E[m|\theta^*])^{-1}\right). \quad (22)$$

We choose S to be sufficiently high (10,000) that simulation error has a negligible effect.

4.2.2 Auxiliary Models

The structural model in Section 2 implies that prices, investments, and market shares are functions of the state variables ω_t and Δ_t . A natural choice of auxiliary model would therefore be semi-parametric approximations to these policy functions. Data restrictions and known differences between our model and the actual CPU industry lead us to depart somewhat from this choice.

First, we do not have a good measure of Δ so we are unable to condition on this state variable. Second, we have only 48 quarterly observations, so we restrict ourselves to parsimonious (i.e., linear) approximations to the policy functions. Finally, the model assumes the market size M is fixed, whereas the data exhibit an upward trend in sales, revenues, and R&D expenditures, as illustrated in Figure 2 and Figure 3.³¹ Hence, instead of estimating an equation for investment levels, our auxiliary model matches investment per unit revenue, which is stationary.

Our moment vector, m_T , consists of the following twelve moments:

³¹Allowing market size to grow over time would render the model computationally intractable. An implication of this misspecification is that our model underestimates the future benefits to innovation. Unless growth rates are very high, this effect should be minor.

- coefficients from regressing each firm’s price on a constant and $\omega_{Intel} - \omega_{AMD}$,
- coefficients from regressing Intel’s share of sales on a constant and $\omega_{Intel} - \omega_{AMD}$,
- mean innovation rates for each firm, defined as $\frac{1}{T}(q_T - q_0)/\delta$,
- mean $\omega_{Intel} - \omega_{AMD}$ and Mean $|\omega_{Intel} - \omega_{AMD}|$, and
- mean investment per unit revenue for each firm.

An obvious difference between our model and the CPU industry is that both Intel and AMD offer multiple versions of their products, whereas our model permits only one product per firm.³² We therefore average prices and qualities across the products offered by each firm. The resulting average price series are plotted for both Intel and AMD in Figure 1, along with the price of each firm’s frontier product. The figure also presents each firm’s frontier quality and the difference in their average qualities.

In the estimation of the model, we are able to match the correlation of prices with $\omega_{Intel} - \omega_{AMD}$, but we underestimate this covariance (primarily) because the variance in average prices exceeds that of our simulated prices. This outcome is not surprising, since a multiproduct firm will tend to vary its highest quality products’ prices more than a single product firm would vary its one price. Comparing the frontier and average prices in Figure 1 reveals that the average price’s variation is determined largely by the variance in each firm’s frontier price. For comparison we also estimate our model using a modification of the price equations in which the mean prices and correlations are matched instead of the regression estimates.

We denote this latter specification the “Alt Aux.” model, and the former the “Baseline Aux.” model. As reported in Section 4.2.4, the structural estimates are quite similar across the two auxiliary models.

We estimate the covariance matrix of the moment vector m_T using a bootstrap procedure with 1000 replications. Its inverse is an estimate of the optimal weighing matrix, which we use for A_T .

³²Extending the model to allow each firm to offer multiple products would be a considerable undertaking.

Experimentation with the structural model indicates that the moments provided by the auxiliary models are quite sensitive to the structural parameters, which is encouraging from an identification standpoint. Being a highly nonlinear model, all the structural parameters influence all the auxiliary moments. However, the connection between some parameters and moments are particularly tight. For example, the auxiliary pricing equations respond sharply to changes in the price and quality coefficients and AMD’s fixed effect. The market share equation also plays an important role in identifying these three preference parameters. The intercepts in the auxiliary pricing equations are sensitive to the marginal costs. The spillover parameter is identified primarily by the mean absolute difference in quality across firms—a high spillover keeps the qualities similar. Finally, the investment efficiencies are chosen to match the firms’ observed R&D expenditures per unit revenue.

A typical concern in empirical industry studies is the endogeneity of prices. We assume that all unobserved demand shocks affect the overall demand for microprocessors, but not consumers’ preferences for one brand over another. This assumption is based on the direct observability of product quality through the performance benchmarks. As such we base the market share moments on “inside” shares only (i.e., each firm’s share of total cpu sales).

4.2.3 Parameterizations

We estimate $K = 8$ parameters in $\theta = (\gamma, \alpha, \xi_{AMD}, a_{0,Intel}, a_{0,AMD}, a_1, mc_{Intel}, mc_{AMD})$, where the a parameters relate to the innovation process detailed below. Table 1 lists the parameter values of the remaining model parameters, which we now discuss.

We choose the log-quality grid size δ to be .1823, which corresponds to 20 percent quality improvements from one grid point to the next. We use fifteen grid points for Δ to track the ownership distribution, which implies $\bar{\delta}_c = 14 * .1823 = 2.552$. Our choice of δ and $\bar{\delta}_c$ reflects the following considerations: i) a fine enough grid (i.e., low enough δ) that in equilibrium firms will have varying degrees of differentiation, ii) a $\bar{\delta}_c$ that is sufficiently high that consumers rarely reach the lowest grid point before upgrading, and

iii) a number of grid points that is computationally manageable. We choose $\bar{\delta}_f$ to be five δ steps so that the leader may be up to 149 percent higher quality than the laggard, which exceeds the observed maximum quality difference.

Following Pakes and McGuire (1994) we specify the probabilities of successful and failed investments, respectively, as $f(1|x) = ax/(1 + ax)$ and $f(0|x) = 1/(1 + ax)$, with a “cost” of investment of $c = \$1$ for each investment dollar.³³ We modify this innovation setup by allowing a to vary for the laggard according to the degree to which the firm is behind. In particular, we specify $a_j(q_t) = a_{0,j}(1 + a_1(\frac{\bar{q}_t - q_{jt}}{\delta})^2)$, which reflects the increased difficulty of advancing the frontier, relative to catching up to the frontier.

We set β to .975, which corresponds to an annual discount factor around .90 since we view a period as being three months. We choose the period length to be three months since our data is quarterly and firms planning horizons are often on a quarterly basis.

We set the market size M to be 400 million consumers. Determining the appropriate market size is difficult because the CPU market is global and a significant share of demand comes from corporations. The effect of varying M is easy to predict: investment is increasing in M because the benefit of innovation scales linearly with M , but the cost of innovation is fixed. The pricing policy function, however, should be unaffected by M .

We fix $\xi_{Intel} = 0$. With three choices in each period, two mean utilities can be estimated, which suggests we should estimate ξ_{Intel} as well as ξ_{AMD} . However, the share of the “no purchase” option is declining over time due to the trend in unit sales and assumption of fixed M . Hence, we prefer to fix $\xi_{Intel} = 0$ and estimate ξ_{AMD} off the relative sales of Intel and AMD.³⁴

4.2.4 Estimates and Model Fit

We report the model’s fit in Table 2 and the parameter estimates in Table 3 for each of the two auxiliary models we consider. Comparing the “Observed” column with the

³³The functional form $ax/(1 + ax)$ for the probability of successful innovation provides closed form solutions for optimal investment, as discussed in Appendix A.

³⁴We note, however, that the percent improvement for upgrade purchases implied by the estimated model is 280 percent, which seems quite plausible. Since this moment helps identify the fixed effects but is already well matched, adding it to estimate ξ_{Intel} would yield a value close to zero.

“Simulated” column in Table 2 reveals that the model fits the twelve moments reasonably well. In the third column for each auxiliary model, we report a pseudo-t for each moment’s fit by dividing the difference between the actual and simulated values by the standard error of the actual moment.

Moments that have t-values below three are generally considered to be well-matched. The moments we have difficulty fitting in the baseline auxiliary model relate to the pricing equations. In particular, we under-predict the sensitivity of firm’s prices to the difference in qualities. This sensitivity could be increased by increasing the quality coefficient γ . However, increasing γ would worsen the fit of the coefficient on $\omega_{Intel} - \omega_{AMD}$ in the share equation, which is already too high.

In the alternative auxiliary model, which substitutes correlations between each firm’s price and $\omega_{Intel} - \omega_{AMD}$ for the regression coefficient, all the moments have t-values below 2.5. Using the correlation instead of the regression coefficient improves the fit because the variance of prices in the data exceed that implied by our model. As discussed in the previous section, the relatively low variation in prices is due to firm’s selling just one product. We include the alternative auxiliary model to illustrate that the model fits the data extremely well, once this discrepancy is taken into account.³⁵

In both specifications the model slightly overpredicts the innovation rates, but both firms’ rates are within two standard errors of their observed values. The model accurately predicts the average quality difference of 1.2 δ steps between the leader and laggard, and the average absolute difference of 1.4 δ steps. R&D as a share of revenue is also predicted with a high degree of accuracy. The standard errors of these two moments are small, so the t-values are above 1.7 in the baseline model despite being close.

Table 3 provides the structural estimates and their standard errors for both of the auxiliary models we consider. All the parameters are statistically significant given the relatively small standard errors. The estimated marginal costs are similar to the average

³⁵Since we have more moments than estimated parameters, we can formally test the model’s specification as being that which generated the data. The objective function’s value of 55.7 in the base model and 26.2 using the alternative model, with only 4 over-identifying restrictions, leads to rejecting the model. Rarely do structural models of this sort pass such specification tests, since the real-world is typically too complicated for a tractable model to mimic perfectly.

manufacturing costs reported by In-Stat of \$39.29 for AMD and \$43.19 for Intel. Intel’s estimated cost may be higher because of the significant advertising of their processors.

4.3 Results

We now use these parameter values as the basis for six different industry scenarios, which correspond to the column headers in Table 4: 1) Intel-AMD duopoly, 2) symmetric duopoly, 3) monopoly, 4) myopic pricing duopoly, 5) myopic pricing monopoly, and 6) social planner. Scenario 1 is the baseline model using the estimates reported in column 1 of Table 3. Scenario 2 modifies the model by using Intel’s firm-specific values for both firms, so that we can get a sense of what the duopoly outcome would be like if both firms were on equal footing. Scenario 3 uses Intel’s parameters for the monopolist. Scenarios 4 and 5 highlight the importance of accounting for the dynamic nature of demand by computing equilibrium when firms are myopic with respect to the pricing decision. Under myopic pricing we solve for the equilibrium when firms and consumers know that firms are behaving myopically with respect to the effect of current prices on future demand.³⁶ Finally, scenario 6 considers the social planner with one product who maximizes the sum of discounted profits and discounted consumer surplus.³⁷

For each scenario we solve for optimal policies and simulate 10,000 industries each for 300 periods, starting from a state in which demand is moderate (average $\tilde{\omega}$ is 10δ relative to frontier at 14δ). If the scenario involves two firms then AMD’s quality (or the laggard, if symmetric) is one δ -step behind Intel’s quality (or the initial leader if the industry is symmetric). We then analyze the simulated data to characterize the equilibrium behavior of firms and consumers and to identify observations of particular

³⁶After price is myopically chosen to maximize current profits, investment is optimally chosen taking into account the dynamic tradeoffs.

³⁷We have recently computed the social planner’s optimal behavior when he controls two products, but have yet to update the full set of results. The planner benefits only slightly from having a second product available. Currently, we weight firm profits equally with consumer surplus. Being a global market, however, a domestic social planner would potentially include only domestic consumers and firms in its objective function. For example, if all firms were domestic but only half the consumers were domestic, then the weight on firms in the objective function should be twice the weight on consumer surplus.

interest. Finally, we consider various counterfactual experiments to further illustrate the properties of the model and its implications for policy analysis.

Much of the information we provide in this results section is merely designed to instill confidence that the model yields equilibrium outcomes that make sense. Particular findings that we wish to emphasize are set apart as “observations”.

Firm Behavior in Equilibrium: We first characterize the firms’ optimal price and investment policy functions in the scenarios 1 and 3 (the baseline duopoly and monopoly settings). The complexity of the state space, due to the distribution of ownership, prohibits a state-by-state inspection of the policy functions. For illustrative purposes, however, we pick two values the ownership distribution—one with low demand and one with high demand—and plot each firm’s value function, pricing function, investment function, innovation rate, and resulting market shares (given prices and consumers’ purchase behavior). The resulting ten plots appear in Figure 4. The first row depicts the ownership distribution of each respective column. In each of the lower ten plots the horizontal axis is the difference in quality (measured in δ steps) between AMD and Intel. Negative values indicate that AMD is the laggard and positive values indicate AMD is the quality leader.

The first column of plots corresponds to an ownership distribution that has a high mean quality of currently owned products which implies few consumers are ready to upgrade. The second column of plots corresponds to an ownership distribution with a low mean quality of currently owned products implying many consumers are ready to upgrade. The title of each plot reports the monopolist’s corresponding value for comparison.³⁸

As expected, the market shares of both firms are substantially higher in column two than column one, as are the value functions (particularly Intel’s), since the ownership distribution corresponds to higher demand. Intel’s prices are higher in column 2 since the leader charges high prices when demand is high. AMD’s prices, however, are only

³⁸The monopolist only has one number since quality difference (i.e., the x-axis) is undefined with only one firm.

marginally higher when demand is high. The monopolist's value and price are also both higher in the high demand state.

Interestingly, the difference in investment behavior across these two states depends on industry structure. The monopolist invests more in the low demand state, since investment is more crucially needed to induce future upgrades. In the duopoly, however, both firms invest more (at each level of quality difference) in the high demand state, which reflects the desire to have a quality advantage (or less of a disadvantage) when consumers are primed to upgrade.

For both Intel and AMD, investment is generally increasing in their own quality relative to their competitor. This relationship between investment and quality advantage primarily reflects the effect of spillovers. If a leader's advantage increases, the competitor's investment efficiency increases substantially which increases its probability of innovating. The leader must therefore invest more to try to preserve his advantage. Similarly, when a laggard becomes more competitive its efficiency declines due to the weaker spillover benefit. The laggard invests more to at least partially make up for this reduced efficiency.

To further illustrate the nature of equilibrium, we provide time series plots of the simulated industry. Two hundred simulated periods for the estimated model appear in Figure 5. The variation in prices over time is evident in both the price plot of Figure 5 and the full distribution of realized margins reported in Table 5. Note from the table that the variation is substantially greater in the duopoly than in the monopoly, and greater when firms optimally account for the dynamic nature of demand.

Finally, Table 6 reports the simulated outcomes separately for Intel and AMD. The total profits are discounted lifetime profits, which therefore correspond to market capitalization. The values of \$433 billion for Intel and \$38 billion for AMD are broadly consistent with market valuations for Intel and AMD over the period of our data, though perhaps a little high.

Figure 6 depicts the difference between optimal pricing by a durable good monopolist and pricing by a myopic monopolist that ignores the dynamic nature of demand. The

top plot corresponds to the payoff (per consumer) facing the myopic monopolist as a function of price. The price that maximizes the current period’s gross profits (which excludes investments) is around \$300. The lower plot corresponds to the true payoff facing the durable good monopolist as a function of current price. The payoff is expressed in a per-period equivalent (by multiplying by $(1 - \beta)$). This per-period payoff is lower for two reasons: investment costs are included and the dynamic monopolist accepts lower current period profits in order to preserve future profits. This preservation is achieved by charging a much higher price, around \$450, in this case.

Consumer Behavior in Equilibrium: We now characterize consumers’ policy functions. The consumer’s only decision is when and what to buy. Since the “when” part corresponds to the current ownership \tilde{q} , we plot in Figure 7 the choice probabilities for each ownership vintage, averaged across states encountered in the duopoly simulations of scenario 1 (the baseline). As expected, the consumer is more likely to upgrade her product, to either the frontier or non-frontier offering, the lower is her current vintage relative to the frontier offering. Consumers with vintages within two δ -steps (i.e., 44%) of the frontier upgrade around 12 percent of the time, compared to more than 50 percent of the time for owners of the lowest quality vintage.

As consumers implement their policies function, they generate a sequence of ownership distributions across time. Figure 8 depicts the average ownership distribution for the baseline duopoly and monopoly cases. Because monopolists charge higher prices, consumers are less likely to upgrade from a given vintage to the frontier in the monopoly case. In the duopoly consumers also have the option to upgrade to the non-frontier product. Both these forces cause the monopolist’s ownership distribution to have more mass on the older vintages, compared to the duopolists’ ownership distribution.

Importantly, Figure 8 reveals that consumers almost never reach the lower bound on vintage, which implies our bounding approach (as discussed in Section 3.1), indeed has little effect on equilibrium behavior. If consumers reached this bound often, then we would simply lower this bound by increasing $\bar{\delta}_c$.

Multiplying the choice probabilities in Figure 7 by the ownership distribution in Figure 8 yields the portion of purchasers from each ownership vintage, as presented in Figure 9. In the duopoly most purchasers upgrade from products that are two to five δ steps below the frontier, which implies the frontier is 44 percent to 149 percent faster than their current product. Of course, some of the consumers upgrade to the non-frontier firm’s offering. As reported in Table 4, the average upgrade percent improvement is 100 percent. This percent improvement for upgrades is perhaps somewhat low, but given that the estimation did not attempt to fit this moment, we find it acceptable. In the monopoly, the higher prices induce consumers to upgrade only when the percent improvement over their current product is reasonably high. Most upgrades in the monopoly are to products that are three to seven δ steps below the frontier, which corresponds to percent improvements of roughly 73 and 258 percent, respectively. The mean upgrade percent in the monopoly 150 percent.

Observations of Interest: Having established that consumers’ and firms’ policy functions match our intuition, we now present and discuss features of equilibrium behavior and outcomes of particular interest.

Observation 1. *Margins (defined as $(p - mc)/mc$) and profits are significantly higher when firms correctly account for the dynamic nature of demand. The differences are larger for monopoly than duopoly.*

From Table 4 we see that monopoly profits are 34 percent higher and margins are 102 percent higher when the monopolist accounts for the dynamic nature of demand, compared to “myopic pricing” which ignores the decline in future demand due to current sales. Industry profits for the duopoly are 14 percent higher and margins are 34 percent higher when the firms account for the dynamic nature of demand, compared to myopic pricing.

This result highlights the importance of accounting for the dynamic nature of demand when analyzing pricing behavior in durable goods markets. Standard practice in

the empirical industrial organization and marketing literatures is to observe prices and use first-order conditions from a static profit maximization to infer marginal costs. Observation 1 suggests that marginal cost estimates computed in this manner for durable goods will be too high. That is, prices are high because the firm does not want to reduce future demand, not because its marginal costs are high.

Table 6 reports the separate outcomes for Intel and AMD under both optimal and myopic pricing. The stark asymmetry of this industry is matched by our model.

Observation 2. *(i) Equilibrium investment is too low, relative to the socially optimal level, for both the monopoly and duopoly market structures.*

(ii) Innovation is slightly higher with a monopoly than with a duopoly.

The per-firm investment levels (measured in millions of dollars) reported in Table 4 for the duopoly, monopoly, and social planner (with one firm to control) are, respectively, 1206, 2584, and 9787. The resulting innovation rates for the industry’s frontier product in each market structure are 0.680 for the duopoly, 0.694 for the monopolist, and 0.896 for the planner.

The finding that innovation by a monopolist exceeds that of a duopoly may come as a surprise to many readers. This finding is driven by two features of the model—the monopolist must innovate to induce consumers to upgrade, and the monopolist is able to extract much of the potential surplus from these upgrades because of its substantial pricing power. If the good’s durability were reduced, by introducing physical depreciation, the monopolist’s innovation would fall since the “competition-with-itself” would decline.

Observation 3. *(i) The AMD-Intel duopoly attains 84.8 percent of the planner’s social surplus, whereas the monopoly attains 83.4 percent—a difference of $(84.8 - 83.4)/83.4 = 1.7$ percent.*

(ii) Consumers’ share of social surplus is 90.2 percent in the AMD-Intel duopoly, compared to 88.1 percent in the monopoly.

(iii) *Consumer surplus is \$432 billion per year in the AMD-Intel duopoly, compared to \$415 billion per year in the monopoly—a difference of 4 percent.*

(iv) *The benefits of competition are substantially higher if Intel were facing a symmetric competitor: consumer surplus and social surplus are 8.2 percent and 6.7 percent higher in a symmetric duopoly than in a monopoly, and duopoly margins are less than half the monopolist's.*

(v) *The social inefficiencies of duopoly and monopoly are due to both higher prices and lower investment.*

Firms' profits are calculated in the usual manner as the discounted sum of per period profits. However, the appropriate measure of consumer surplus is less obvious. We compute consumer surplus assuming mean utility from the outside alternative (for non-owners) would remain at zero forever in the absence of innovation by these two firms. If this were not the case, then an exogenous improvement in the outside alternative (for nonowners) would need to be included in the model. While such a modification could easily be accommodated, our stance is that the role of an outside good in models of industry evolution ought to be minimized.³⁹

Consumer surplus for a consumer in period t who currently owns a product with quality \tilde{q}_t is the expected utility flow in that period, divided by the price coefficient to convert utils to dollars. In the static logit model the expected utility is computed using the “inclusive value” formula $\log \sum_j \exp(u_{jt} - \varepsilon_{jt})$. In the dynamic setting, however, we cannot use this formula since the choice probabilities are functions of the discounted continuation values and our measure of surplus uses only the utility flows, not utility flows plus discounted continuation values. Hence, we define

³⁹If this alternative is truly a substitute for the inside goods, then its price and quality transitions ought to react strategically to the inside goods, in which case it becomes an inside good itself.

$$\begin{aligned}
CS(\tilde{q}_t, q_t) &= \frac{1}{\alpha} \mathbf{E} [u_{jt}] \\
CS(\tilde{q}_t, q_t) &= \frac{1}{\alpha} \sum_{j \in \{0, \dots, J\}} s_{jt|\tilde{q}} (\gamma q_{jt} - \alpha p_{jt} + \xi_j - \mathbf{E} [\varepsilon_{jt} | \text{choose } j]) \\
&= \frac{1}{\alpha} \sum_{j \in \{0, \dots, J\}} s_{jt|\tilde{q}} (\gamma q_{jt} - \alpha p_{jt} + \xi_j - \log s_{jt|\tilde{q}})
\end{aligned} \tag{23}$$

where the quality levels are in absolute terms (i.e., not relative to the frontier), $s_{jt|\tilde{q}}$ is the conditional choice probability and $\log s_{jt|\tilde{q}}$ is the expected value of ε_{jt} given that j is chosen.⁴⁰

Aggregate discounted consumer surplus over a simulation run is the discounted sum of the per period surplus, integrated over the distribution of consumer product ownership:⁴¹

$$CS = M \sum_{t=0}^T \beta^t \sum_{\tilde{q}=q_t}^{\tilde{q}_t} CS(\tilde{q}, q_t) \cdot \Delta_{\tilde{q},t}. \tag{24}$$

Table 4 reports the aggregate discounted CS and industry profits for each of the scenarios we consider. The duopoly value for CS is \$4.32 trillion, which corresponds to \$432 billion per year (using an annual discount factor of .9). This value is 4.1 percent higher than the CS generated by the monopoly. In terms of social surplus, the monopolist is able to make-up for part of this reduced CS through its significantly higher industry profits. Nonetheless, the monopoly's social surplus is still about 1.7 percent below the duopoly social surplus.

Observation 4. *Unexpected 10 percent increases in price for one period yield:*

(i) *Intel's price increase causes its share to fall by 18.1 percent, AMD's share to increase by 4.1 percent, and the no-purchase share to increase by 3.7 percent.*

(ii) *AMD's price increase causes its share to fall by 12.7 percent, Intel's share to increase by 0.4 percent, and the no-purchase share to increase by 0.4 percent.*

(iii) *Intel's price increase as a monopolist causes its share to fall by 23.7 percent, and the no-purchase share to increase by 4.2 percent.*

⁴⁰We use a distribution for ε that is shifted by Euler's constant so that its mean is zero.

⁴¹Consumer surplus may be computed directly from the value functions as $\frac{M}{\alpha} \sum_{\tilde{q}=q_t}^{\tilde{q}_t} \hat{V}(q_0, \Delta_0, \tilde{q}) \cdot \Delta_{\tilde{q},0}$.

(iv) A firm's markup (in both duopoly and monopoly) is positively related to the elasticity of demand, in direct contrast with the "inverse elasticity rule" of static pricing.

These elasticities are averages from 10,000 simulations in which the unexpected, single-period price shock occurred in period 100 of each simulation. The stochastic nature of the model implies that the state at which this price shock is implemented varies across simulations. In Figure 10 we plot the elasticity and optimal markup at each of the states encountered in these simulations. For both duopoly firms and for Intel as a monopolist, the markup is increasing in elasticity. This contrasts with the inverse relationship of static pricing, as depicted by performing the price shock exercise for the monopolist that prices myopically.

Hence, the static relationship between markups and (short-run) elasticities is reversed due to the dynamic nature of demand for durable goods. Intel (and AMD) could indeed increase current sales substantially by lowering price today, but this would lower its future demand too much to warrant the change.

Observation 5. *As depicted in Figure 11, consumer surplus, margins, innovation, and profits all increase monotonically in the rate of time preference upon which consumers base decisions.*⁴²

If consumers do not value the future benefit of a durable good, their willingness to pay for the good decreases, which causes firms to reduce prices and earn lower profits. Innovation is also reduced, since consumers are not willing to pay as much for their upgrade purchases.

This result contrasts with the finding in Nair (2007) that optimal prices for a durable goods monopolist are decreasing in the consumers' rate of time preference. The difference in our findings is due to the fact that we increase the consumer's discounted utility flow

⁴²Consumer surplus is constructed assuming that consumers truly value the future using the discount rate .975, but are myopic, to varying degrees, in their decision making. Using the varying rate of time preference for decision making as the true discount factor merely steepens the slope of the consumer surplus plot.

from the durable good when β increases, whereas Nair (2007) fixes the discounted utility flow at the estimated values. His result applies when consumers use different discount factors for the valuation of a product's discounted utility flows and the determination of when to purchase a product.

4.4 Counterfactual Analysis

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

In Figure 12 we plot the margins, innovation rates, consumer surplus, and social surplus when the "access denied" portion of the market varies from zero to one (in .1 increments). We find that barring AMD from 30 percent of the market decreases consumer surplus, social surplus, and innovation all by about 2 percent. The innovation rate begins to increase when the share of the market from which AMD is barred surpasses .8, and eventually passes the duopoly level of innovation when Intel is a monopolist. This pattern indicates that as Intel gains exclusive access to a large share of the market, its investment increases to reap the greater benefit from higher quality that results from its pricing power.

5 Conclusions

This paper presents a dynamic model of durable goods oligopoly with endogenous innovation. The model entails two methodological contributions. First we propose an alternative approach to bounding the state space in dynamic oligopoly models of the Ericson and Pakes (1995) type. Under this alternative specification, the innovation rate

for the industry is endogenous instead of being determined by an outside good's exogenous rate of innovation. This feature makes our model particularly suitable for studying industries in which innovation is a significant source of consumer surplus. Second, we develop a simple approximation for the ownership distribution, which is an endogenous state variable that summarizes the state of demand when products are durable goods.

We estimate the model using data from the computer processor industry and show that accounting for the durable nature of products in an equilibrium setting has important implications for firm and consumer behavior and market outcomes. Because increased sales today lowers future demand, firms set prices higher when correctly accounting for this dynamic aspect of demand.

We compute profits and consumer surplus under alternative market structures and find that consumer surplus is 4.1 percent higher with AMD duopoly than without AMD competing against Intel. Social surplus is 1.7 percent higher with competition from AMD. Interestingly, the monopolist innovates more than the duopoly, as its market power enables it to better extract the potential gains to trade resulting from innovations.

Several possibilities remain for future research. For example, natural questions arise regarding the firm's choice of leasing versus selling its durable product. Under what conditions should a firm sell versus lease? Is it more advantageous for the leader or laggard to lease? Does leasing reduce innovation and social welfare?

Another extension of interest is to consider switching costs for consumers when upgrading from one firm's product to a different firm's product. Conceptually, such an extension is trivial, though the state space would grow substantially since the ownership distribution would need to be firm-specific.

Finally, allowing for multiproduct firms would enable researchers to study whether competition is enhanced or diminished when leading firms offer inferior products to compete more directly with the top offerings of lagging firms.

Appendix A: Optimal Prices and Investments

As discussed in Section 2.3, we use a sequential approach to solve for the simultaneously chosen prices and investments. Following Pakes and McGuire (1994), we specify the probability of successful innovation to be $f(1|x) = a_j x / (1 + a_j x)$ where a_j denotes the firm's investment efficiency. This specification yields a closed form solution to the first order condition in (11). Let

$$EW^+(p_{jt}) = \sum_{\substack{q_{-j,t+1} \\ \Delta_{t+1}}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}) h_f(q_{-j,t+1}|q_t, \Delta_t) g_f(\Delta_{t+1}|\Delta_t, q_t, p_t, q_{t+1}) f(\tau_{jt} = \delta|x_{jt})$$

$$EW^-(p_{jt}) = \sum_{\substack{q_{-j,t+1} \\ \Delta_{t+1}}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}) h_f(q_{-j,t+1}|q_t, \Delta_t) g_f(\Delta_{t+1}|\Delta_t, q_t, p_t, q_{t+1}) f(\tau_{jt} = 0|x_{jt})$$

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

$$x_{jt}^*(p_{jt}) = \frac{1}{a_j} \left(\frac{c}{\beta a_j (EW^+(p_{jt}) - EW^-(p_{jt}))} \right)^{-1/2} - 1. \quad (25)$$

To determine the optimal price, consider the derivative of the firm's value function with respect to price, $\frac{\partial W}{\partial p_{jt}} = 0$, which implies

$$\frac{\partial \pi_j(p_t, q_t, \Delta_t)}{\partial p_{jt}} + \beta \sum_{\tau_{jt}, q_{-j,t+1}, \Delta_{t+1}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}) h_f(q_{-j,t+1}|q_t, \Delta_t) \frac{\partial g_f(\Delta_{t+1}|\Delta_t, q_t, p_t, q_{t+1})}{\partial p_{jt}} f(\tau_{jt}|x_{jt}^*(p_{jt})) = 0 \quad (26)$$

where the partial derivative $\frac{\partial x^*(p)}{\partial p}$ may be ignored due to the Envelope theorem. Recall the important dynamic trade-off—a higher price today implies that more people will be available in the next period to purchase the product. The second term of this first-order condition captures this benefit to raising price, and leads to forward-looking firms pricing higher than myopic firms who ignore this dynamic aspect of demand.

We use Brent's method to solve for the optimal price. For each candidate price we use $x^*(p_{jt})$, the optimal investment level given this price, to evaluate the probability of a successful innovation. While we have yet to prove that the optimal price is uniquely

determined, inspection of the first-order-condition as a function of p_{jt} at many states indicates that this appears to be the case. The pair $(p_{jt}^*, x_{jt}^*(p_{jt}^*))$ is the optimal set of controls at this state.

Appendix B

Proof of Proposition 1: We prove the proposition for the case of a finite horizon, using backwards induction, since this approach enables us to impose rational expectations regarding future outcomes.

Consider the finite game with T periods in which a consumer starting at state $(q_1, \Delta_1, \tilde{q}_1, \varepsilon_1)$ maximizes expected discounted utility

$$V^T(q_1, \Delta_1, \tilde{q}_1, \varepsilon_1) = \max_{\{y_t(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) \in (0, 1, \dots, J)\}_{t=1}^T} \mathbb{E} \left[\sum_{t=1}^T \beta^t (\gamma q_{y_t, t} - \alpha p_{y_t, t} + \xi_{y_t} + \varepsilon_{y_t, t}) \right] \quad (27)$$

where $q_{0,t} = \tilde{q}_t$ and $p_{0,t} = 0$ in each period, $y_t(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t)$ is the consumer's policy function, and the expectation is taken with respect to information available at time t . In this game each firm j maximizes expected discounted net profits

$$W_j^T(q_{j1}, q_{-j,1}, \Delta_1) = \max_{\{p_{jt}(q_t, \Delta_t), x_{jt}(q_t, \Delta_t)\}_{t=1}^T} \mathbb{E} \left[\sum_{t=1}^T \beta^t (M s_{jt}(p_t, q_t, \Delta_t)(p_{jt} - mc_j) - cx_{jt}) \right] \quad (28)$$

where M is the market size, $s_{jt}(\cdot)$ is the market share for firm j as defined in equation (8), p_t is the vector of J prices, x_{jt} is investment by firm j , and mc_j is firm j 's constant marginal cost of production.

In period T firms and consumers play a standard static differentiated products game given the state of the industry, as described by (q_T, Δ_T) . Since consumers' utility functions are linear in the quality index, consumers' choices are insensitive to shifts in all qualities (q_t and \tilde{q}) by some constant \hat{q} . The market share function therefore satisfies $s_{jt}(p_t, q_t, \Delta_t) = s_{jt}(p_t, q_t - \hat{q}, \Delta_t)$, which implies firms' prices are insensitive to shifts in all qualities. The period T value functions V^T and W^T , for consumers and firms therefore satisfy

$$\begin{aligned} \text{Firms: } W_j^T(q_{jt}, q_{-j,t}, \Delta_t) &= W_j^T(q_{jt} - \hat{q}, q_{-j,t} - \hat{q}, \Delta_t) \\ \text{Consumers: } V^T(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) &= \gamma \hat{q} + V^T(q_t - \hat{q}, \Delta_t, \tilde{q}_t - \hat{q}, \varepsilon_t). \end{aligned} \quad (29)$$

Note that each consumer's utility shifts by $\gamma\hat{q}$ when all qualities shift by \hat{q} .

Now consider equilibrium outcomes in period $T - 1$ taking as given the period T equilibrium payoffs. Each consumer solves

$$\begin{aligned}
V^{T-1}(q_{T-1}, \Delta_{T-1}, \tilde{q}_{T-1}, \varepsilon_{T-1}) &= \max_{y \in (0, 1, \dots, J)} \gamma q_{y, T-1} - \alpha p_{y, T-1} + \xi_y + \varepsilon_{y, T-1} \\
&+ \beta \sum_{q_T} \int_{\varepsilon_T} V^T(q_T, \Delta_T, \tilde{q}_T, \varepsilon_T) dF_\varepsilon(\varepsilon_T) \\
&\prod_{j=1}^J f(q_{jT} - q_{j, T-1} | x_{j, T-1}, q_{T-1})
\end{aligned} \tag{30}$$

where $\tilde{q}_T = \max(q_{y, T-1}, \bar{q}_T - \delta_c)$ is the transition of \tilde{q} accounting for the maximum allowed difference between the frontier product's quality \bar{q}_T and each consumers \tilde{q} , and the deterministic transition to Δ_T is based on consumers' choices, as detailed in equation (9). Since each consumer is small relative to M , her actions do not affect the transition of Δ .

Each firm j solves

$$\begin{aligned}
W_j^{T-1}(q_{j, T-1}, q_{-j, T-1}, \Delta_{T-1}) &= \max_{\substack{p_{j, T-1} \\ x_{j, T-1}}} M s_{j, T-1}(p_{T-1}, q_{T-1}, \Delta_{T-1})(p_{j, T-1} - mc_j) - cx_{j, T-1} \\
&+ \beta \sum_{q_T} W_j^T(q_{j, T}, q_{-j, T}, \Delta_T) \prod_{j=1}^J f(q_{jT} - q_{j, T-1} | x_{j, T-1}, q_{T-1})
\end{aligned} \tag{31}$$

In these equations defining V^{T-1} and W^{T-1} , the products' future qualities are uncertain. Rational expectations regarding this uncertainty is achieved by using the firm's investments in period $T - 1$ to determine the distribution of q_T .⁴³

Now consider these same maximizations at a state with all qualities shifted by \hat{q} :

$$\begin{aligned}
V^{T-1}(q_{T-1} - \hat{q}, \Delta_{T-1}, \tilde{q}_{T-1} - \hat{q}, \varepsilon_{T-1}) &= \max_{y \in (0, 1, \dots, J)} \gamma(q_{y, T-1} - \hat{q}) - \alpha p_{y, T-1} + \xi_y + \varepsilon_{y, T-1} \\
&+ \beta \sum_{q_T} \int_{\varepsilon_T} V^T(q_T - \hat{q}, \Delta_T, \tilde{q}_T - \hat{q}, \varepsilon_T) dF_\varepsilon(\varepsilon_T) \\
&\prod_{j=1}^J f(q_{jT} - \hat{q} - (q_{j, T-1} - \hat{q}) | x_{j, T-1}, q_{T-1} - \hat{q})
\end{aligned} \tag{32}$$

⁴³Recall that $f(\cdot | x_{jt}, q_t)$ is the probability distribution of j 's investment outcome, which is restricted to be either no improvement in quality, or improvement by one δ -step.

and

$$\begin{aligned}
W_j^{T-1}(q_{j,T-1} - \hat{q}, q_{-j,T-1} - \hat{q}, \Delta_{T-1}) &= \max_{\substack{p_{j,T-1} \\ x_{j,T-1}}} Ms_{j,T-1}(p_{T-1}, q_{T-1} - \hat{q}, \Delta_{T-1})(p_{j,T-1} - mc_j) \\
&\quad - cx_{j,T-1} + \beta \sum_{q_T} W_j^T(q_{j,T} - \hat{q}, q_{-j,T} - \hat{q}, \Delta_T) \\
&\quad \prod_{j=1}^J f(q_{jT} - \hat{q} - (q_{j,T-1} - \hat{q}) | x_{j,T-1}, q_{T-1} - \hat{q}).
\end{aligned} \tag{33}$$

Substitute the right-hand sides of (29) into (33) and (32). Then note that $f(q_{jT} - \hat{q} - (q_{j,T-1} - \hat{q}) | x_{j,T-1}, q_{T-1} - \hat{q}) = f(q_{jT} - q_{j,T-1} | x_{j,T-1}, q_{T-1})$ by algebra and the assumption that the “spillover” aspect of investment outcomes depends on quality differences between the investing firm and the frontier product. As such, firms’ investment choices are unaffected by the \hat{q} shift. Consumers’ and firms’ discounted continuation values are therefore insensitive to the \hat{q} shift. Since current flow utility is insensitive to the quality shift (by linearity), consumers’ period $T - 1$ choices (i.e., $s_{j,T-1}$) must be insensitive to the shift, which further implies firms’ $T - 1$ prices are insensitive to the shift. Implementing these equivalences converts (33) into (31), exactly, and converts (32) into (30), except for a $-(\gamma\hat{q} + \beta\gamma\hat{q})$ term which does not affect the consumer’s choice. The modified (32) is

$$\begin{aligned}
V^{T-1}(q_{T-1} - \hat{q}, \Delta_{T-1}, \tilde{q}_{T-1-\hat{q}}, \varepsilon_{T-1}) &= \max_{y \in (0,1,\dots,J)} \gamma(q_{y,T-1} - \hat{q}) - \alpha p_{y,T-1} + \xi_y + \varepsilon_{y,T-1} \\
&\quad + \beta \sum_{q_T} \int_{\varepsilon_T} (-\gamma\hat{q} + V^T(q_T, \Delta_T, \tilde{q}_T, \varepsilon_T)) dF_\varepsilon(\varepsilon_T) \\
&\quad \prod_{j=1}^J f(q_{jT} - q_{j,T-1} | x_{j,T-1}, q_{T-1})
\end{aligned} \tag{34}$$

By induction, the optimal consumer policies $y_t(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t)$ and firm policies $p_t(q_{jt}, q_{-jt}, \Delta_t)$ and $x_t(q_{jt}, q_{-jt}, \Delta_t)$ are insensitive to shifts in all qualities, for all t . The firm’s value functions W^t are also insensitive to \hat{q} shifts and the consumers’ value function V^t is shifted by $\gamma\hat{q} \sum_{t'=0}^{T-t} \beta^{t'}$.

To complete the proof, choose $\hat{q} = \bar{q}_t$, the quality of the frontier product in period t . □

References

- [1] Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005), “Competition and Innovation: an Inverted-U Relationship,” *The Quarterly Journal of Economics*, 120(2), 701–728..
- [2] Bajari, P., L. Benkard, and J. Levin (2007), “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, forthcoming.
- [3] Benkard, L. (2004), “A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft,” *Review of Economic Studies*, 71(3), 581–611.
- [4] Bond, E. and L. Samuelson (1984), “Durable Good Monopolies with Rational Expectations and Replacement Sales,” *RAND Journal of Economics*, 15(3), 336–345.
- [5] Bulow, J. (1982), “Durable Goods Monopolists,” *Journal of Political Economy*, 90, 314–332.
- [6] Carranza, J. (2004), “Demand for Durable Goods and the Dynamics of Quality,” mimeo, University of Wisconsin-Madison.
- [7] Chen, J., S. Esteban, and M. Shum (2008), “How Much Competition is a Secondary Market?” mimeo, UC Irvine.
- [8] Coase, R. (1972), “Durability and Monopoly,” *Journal of Law and Economics*, 15, 143–149.
- [9] Desai, P.S. and D. Purohit (1999), “Competition in Durable Goods Markets: The Strategic Consequences of Leasing and Selling,” *Marketing Science*, 18(1), 42–58.
- [10] Doraszelski, U. and M. Satterthwaite (2007), “Computable Markov-Perfect Industry Dynamics: Existence, Purification, and Multiplicity,” Harvard University, mimeo.
- [11] Dubé, J. P. , G. Hitsch, and P. Chintagunta (2007) “Dynamic Standards Competition and Tipping: The Case of 32/64 Bit Video Game Consoles,” manuscript, Chicago GSB.

- [12] Ericson, S. and A. Pakes (1995), “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62(1), 53–82.
- [13] Esteban, S. and M. Shum (2007), “Durable-Goods Oligopoly with Secondary Markets: the Case of Automobiles,” *RAND Journal of Economics*.
- [14] Fudenberg, D. and J. Tirole (1998), “Upgrades, Trade-ins, and Buybacks,” *RAND Journal of Economics*, 29(2), 237–252.
- [15] Gallant, R. and G. Tauchen (1996), “Which Moments to Match?” *Econometric Theory*, 12, 657–681.
- [16] Gilbert, R. (2005), “Looking for Mr. Schumpeter: Where are We in the Competition-Innovation Debate?” forthcoming, National Bureau of Economic Research.
- [17] Goettler, R., C. Parlour, and U. Rajan (2005), “Equilibrium in a Dynamic Limit Order Market,” *Journal of Finance*, vol. 60, no. 5, 2149–2192.
- [18] Goettler, R., C. Parlour, and U. Rajan (2006), “Microstructure Effects and Asset Pricing,” Carnegie Mellon University, mimeo #2006-E97.
- [19] Gordon, B. R. (2008), “A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry,” working paper, Columbia University.
- [20] Gouriéroux, C., A. Monfort, and E. Renault (1993), “Indirect Inference,” *Journal of Applied Econometrics*, 8, S85–S118.
- [21] Gowrisankaran, G. and M. Rysman (2007), “Dynamics of Consumer Demand for New Durable Consumer Goods,” working paper, Boston University.
- [22] Gowrisankaran, G. and R.J. Town (1997), “Dynamic Equilibrium in the Hospital Industry,” *Journal of Economics and Management Strategy*, 6(1), 45–74.
- [23] Gul, F. (1987), “Noncooperative Collusion in Durable Goods Oligopoly,” *RAND Journal of Economics*, 18(2), 248–254.

- [24] Hall, G. and J. Rust (1993), “Simulated Minimum Distance Estimation of a Model of Optimal Commodity Price Speculation with Endogenously Sampled Prices,” mimeo, Yale University.
- [25] Hansen, L. (1982), “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, 50, 1929–1954.
- [26] Hendel, I. and A. Lizzeri (1999), “Interfering with Secondary Markets,” *RAND Journal of Economics*, 30(1), 1–21.
- [27] Kleiman, E. and T. Ophir (1966), “The Durability of Durable Goods,” *Review of Economic Studies*, 33(2), 165–178.
- [28] Lee, I. H. and J. Lee (1998), “A Theory of Economic Obsolescence,” *Journal of Industrial Economics*, 46(3), 383–401.
- [29] Levinthal, D. and D. Purohit (1989), “Durable Goods and Product Obsolescence,” *Marketing Science*, 8(1), 35–56.
- [30] Macieira, J. (2007) , “Extending the Frontier: A Structural Model of Investment and Technological Competition in the Supercomputer Industry,” Virginia Tech.
- [31] Markovich, S. (forthcoming), “Snowball: A Dynamic Oligopoly Model with Network Externalities,” *Journal of Economic Dynamics and Control*.
- [32] Markovich, S. and Moenius, J. (2005), “Winning While Losing: Competition Dynamics in the Presence of Indirect Network Effects,” Working paper, Northwestern University.
- [33] Melnikov, O. (2001), “Demand for Differentiated Durable Products: The Case of the US Computer Printer Industry,” mimeo, Yale University.
- [34] Nair, H. (2007), “Intertemporal Price Discrimination with Forward-Looking Consumers: Application to the US Market for Console Video-Games,” *Quantitative Marketing and Economics*, 5(3), 239–292.

- [35] Pakes, A. and D. Pollard (1989), “Simulation and Asymptotics of Optimization Estimators,” *Econometrica*, 57, 1027-1057.
- [36] Pakes, A. and P. McGuire (1994), “Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model,” *RAND Journal of Economics*, 25(4), 555–589.
- [37] Rust, J. (1986), “When is it Optimal to Kill off the Market for Used Durable Goods?” *Econometrica*, 54(1), 65–86.
- [38] Rust, J. (1987), “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55(5), 999–1033.
- [39] Smith, A. (1993) “Estimating Nonlinear Time Series Models using Vector Autoregressions: Two Approaches,” *Journal of Applied Econometrics*, 8, 63–84.
- [40] Song, M. (2003), “A Dynamic Analysis of Cooperative Research in the Semiconductor Industry,” mimeo University of Rochester.
- [41] Song, I. and Chintagunta, P. (2003), “A Micromodel of New Production Adoption with Heterogeneous and Forward-Looking Consumers: Application to the Digital Camera Category,” *Quantitative Marketing and Economics*, 1, 371–407.
- [42] Stokey, N. (1981), “Rational Expectations and Durable Goods Pricing,” *Bell Journal of Economics*, 12, 112–128.
- [43] Swan, P. (1970), “Durability of Consumption Goods,” *American Economic Review*, 60(5), 884–894.
- [44] Sieper, E. and P. Swan (1973), “Monopoly and Competition in the Market for Durable Goods,” *Review of Economic Studies*, 40(3), 333–351.
- [45] Waldman, M. (1996), “Durable Goods Pricing When Quality Matters,” *Journal of Business*, 69(4), 489–510.

- [46] Waldman, M. (2003), “Durable Goods Theory for Real World Markets,” *Journal of Economics Perspectives*, 17(1), 131–154.
- [47] Xu, Yi (2007), “A Structural Empirical Model of R&D Investment, Firm Heterogeneity, and Industry Dynamics,” Working paper, New York University.

Table 1: Base Case Parameters

Parameter	Value
Intel fixed effect, ξ_{Intel}	0
Discount factor, β	0.975
Log-quality stepsize, δ	0.1823 (20% quality increments)
$\bar{q} - \tilde{q}$ bound, $\bar{\delta}_c$	2.5522 (14 δ -steps)
$\bar{q} - q_j$ bound, $\bar{\delta}_f$	0.9115 (5 δ -steps)
Innovation pdf, $f(\tau = \delta x, q)$	$a(\cdot)x/(1 + a(\cdot)x)$, x in \$millions
Innovation spillover, $a(\cdot)$	$a_{0,j}(1 + a_1(\frac{\bar{q}_t - q_{jt}}{\delta})^2)$
Investment cost, c	1
Market size, M	400 million

Table 2: Empirical vs. Simulated Moments

Auxiliary Moment	Baseline Aux. Model			Alt. Aux. model		
	Observed	Simulated	pseudo-t	Observed	Simulated	pseudo-t
Intel price equation:						
constant	228.102	228.840	0.079	270.844	274.073	0.282
$\omega_{Intel} - \omega_{AMD}$	197.335	61.239	3.141	0.633	0.637	0.056
AMD price equation:						
constant	180.554	139.712	3.826	149.185	137.738	1.277
$\omega_{Intel} - \omega_{AMD}$	-139.870	-15.300	5.038	-0.572	-0.748	2.083
Intel share equation:						
constant	0.830	0.831	0.230	0.830	0.821	1.368
$\omega_{Intel} - \omega_{AMD}$	0.074	0.114	2.311	0.074	0.116	2.447
Mean Innovation Rates:						
Intel	0.571	0.676	1.587	0.571	0.731	2.416
AMD	0.602	0.672	0.881	0.602	0.714	1.419
Mean Quality Differences:						
$\omega_{Intel} - \omega_{AMD}$	1.217	1.112	0.506	1.217	1.500	1.353
$ \omega_{Intel} - \omega_{AMD} $	1.412	1.587	0.978	1.412	1.716	1.700
Mean R&D / Revenue:						
Intel	0.114	0.120	1.717	0.114	0.111	0.760
AMD	0.203	0.226	2.524	0.203	0.197	0.758
Objective Function :		55.8			26.2	

The “Alt. Aux. Model” differs from the baseline by using the correlations between each firm’s price and $\omega_{Intel} - \omega_{AMD}$ instead of the regression coefficient.

The pseudo-t is $|Observed - Simulated|/Observed_Std_Error$.

Table 3: Parameter Estimates

Parameter	Baseline Aux.		Alt. Aux.	
	Estimate	Std. Err.	Estimate	Std. Err.
Price, α	0.01017	0.001184	0.00938	0.000198
Quality, γ	0.38560	0.037586	0.40607	0.020751
AMD Fixed Effect, ξ_{AMD}	-2.49436	0.111348	-2.62150	0.046973
Intel Innovation, $a_{0,Intel}$	0.00088	0.000044	0.00100	0.000112
AMD Innovation, $a_{0,AMD}$	0.00197	0.000432	0.00185	0.000055
Spillover, a_1	1.30389	0.025923	1.36845	0.033894
Intel Marginal Cost, mc_{Intel}	57.76092	1.807299	71.12646	0.843711
AMD Marginal Cost, mc_{AMD}	30.34652	1.411450	20.73857	0.937190

Table 4: Industry Measures under Various Scenarios

	Intel-AMD	Symmetric	Monopoly	Myopic Pricing		Social Planner
	Duopoly	Duopoly		Duopoly	Monopoly	
Industry Profits (\$millions)	472069	535865	559302	412513	416265	-391252
Consumer Surplus	4321806	4490887	4152413	4126303	3895274	6044158
Social Surplus	4793875	5026752	4711715	4538816	4311540	5652905
Period Profits per Consumer	28.017	32.660	32.309	24.728	24.330	-24.468
Margins	3.146	2.197	4.461	2.312	2.212	0.000
Price	221.738	184.675	315.451	177.961	185.516	57.761
Frontier Innovation Rate	0.680	0.635	0.694	0.591	0.523	0.896
Industry Investment	1205.781	1267.145	2584.089	806.987	1270.218	9787.185
Average $ q_{1t} - q_{2t} /\delta$	1.702	1.259	0.000	1.429	0.000	0.000
Frontier Quality $t = 1$	2.735	2.735	2.735	2.735	2.735	2.735
Frontier Quality $t = 100$	15.018	14.191	15.272	13.404	12.207	18.907
Mean Quality Upgrade	1.006	0.587	1.523	0.709	0.691	0.553
Firm 1 market share	0.172	0.162	0.150	0.206	0.214	0.426
Firm 2 market share	0.029	0.143	0.000	0.026	0.000	0.000

Reported values are based on 10,000 simulations of 300 periods each.

Symmetric duopoly uses Intel’s firm specific parameters for both firms.

Under “myopic pricing” firms choose price ignoring its effect on future demand.

Profits, investment, and surplus numbers are reported in millions of dollars.

Profits and surplus are discounted back to period 0, except “Period Profits per Consumer”.

Social surplus is the sum of consumer surplus and industry profits.

The social planner and monopolist are restricted to offering one product.

Margins are computed as $(p - mc)/mc$. Price and margins are share-weighted averages.

Table 5: Margin Distribution

	Duopoly	Monopoly	Myopic Pricing	
			Duopoly	Monopoly
Min	2.261	2.480	2.154	2.117
5%	2.743	3.565	2.237	2.117
10%	2.824	3.750	2.247	2.117
25%	2.965	4.073	2.277	2.163
50%	3.138	4.602	2.313	2.209
75%	3.355	4.906	2.350	2.264
90%	3.549	4.924	2.378	2.302
95%	3.656	4.963	2.405	2.321
Max	4.046	6.231	2.614	2.571
Mean	3.146	4.461	2.312	2.212
Stdev	0.481	0.658	0.041	0.120

Distribution of sales-weighted average price margins $((p_{jt} - mc_j)/mc_j)$ under the different industry scenarios.

Table 6: Intel and AMD outcomes with optimal pricing and myopic pricing

	Optimal Pricing		Myopic Pricing	
	Intel	AMD	Intel	AMD
Total Profits (\$millions)	433596	38473	382466	30047
Period Profits	10345	862	9181	710
Period Profits per Consumer	25.863	2.154	22.952	1.776
Market Share	0.172	0.029	0.206	0.026
Margin	3.092	3.492	2.184	3.332
Investment	2016.584	395.866	1264.725	349.613
Innovation Rate	0.679	0.677	0.590	0.592

Figure 1: CPU Prices and Qualities: 1993 Q1 to 1994 Q4

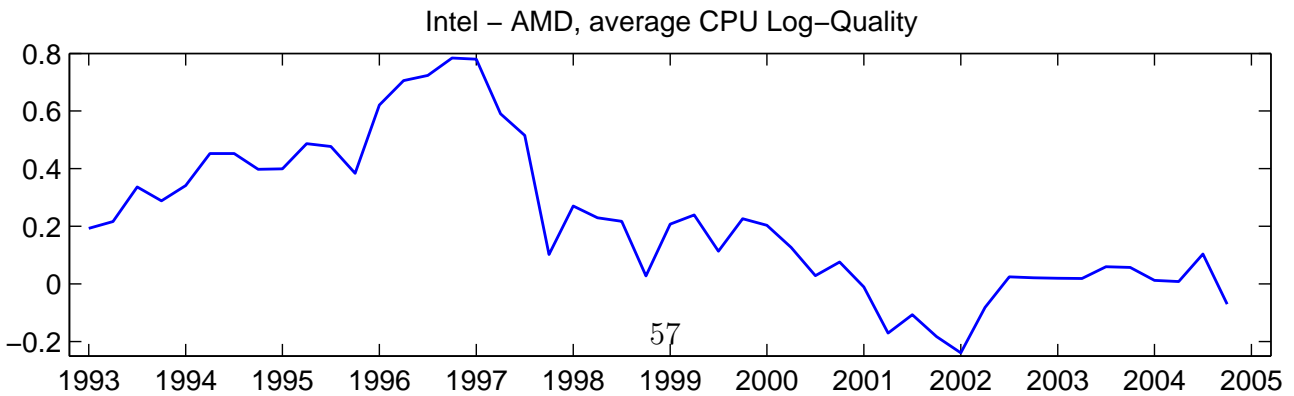
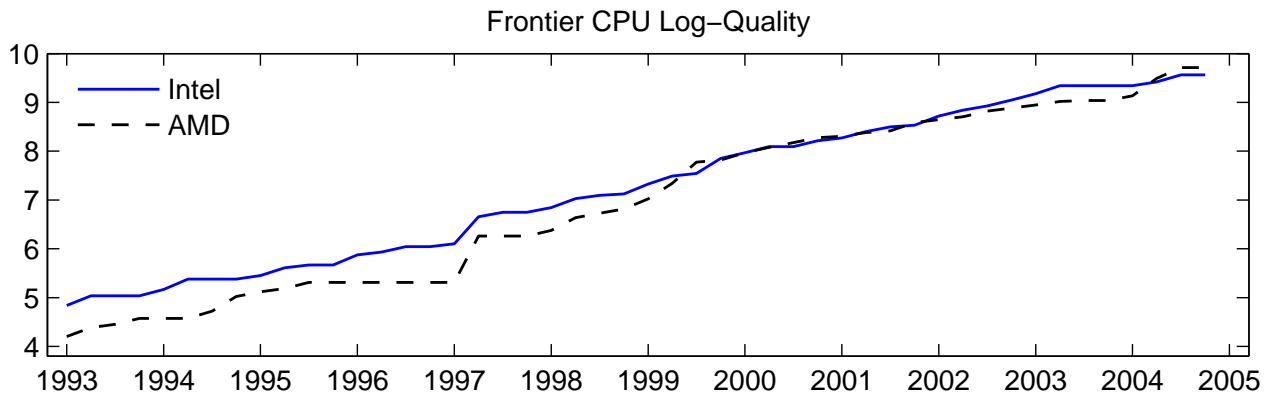
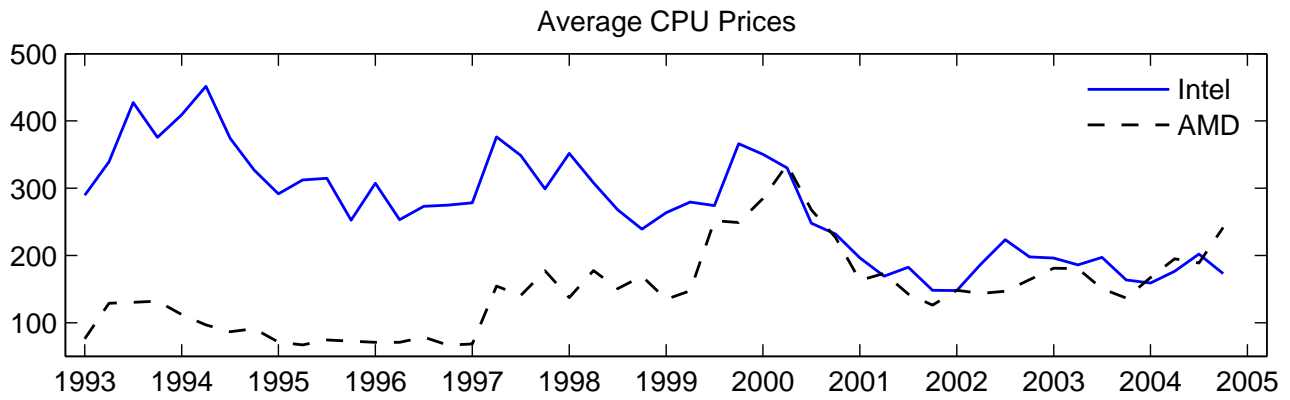
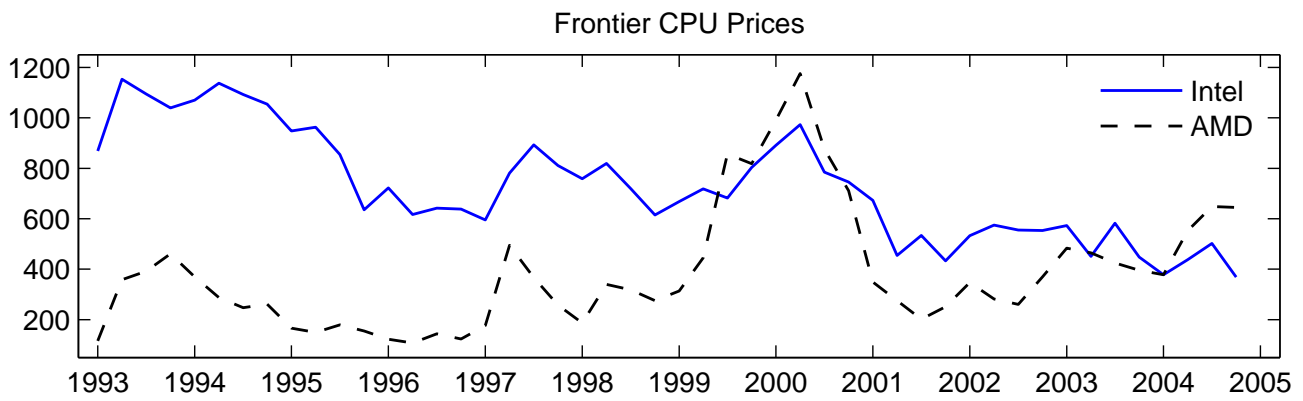


Figure 2: CPU Shipments and Market Shares

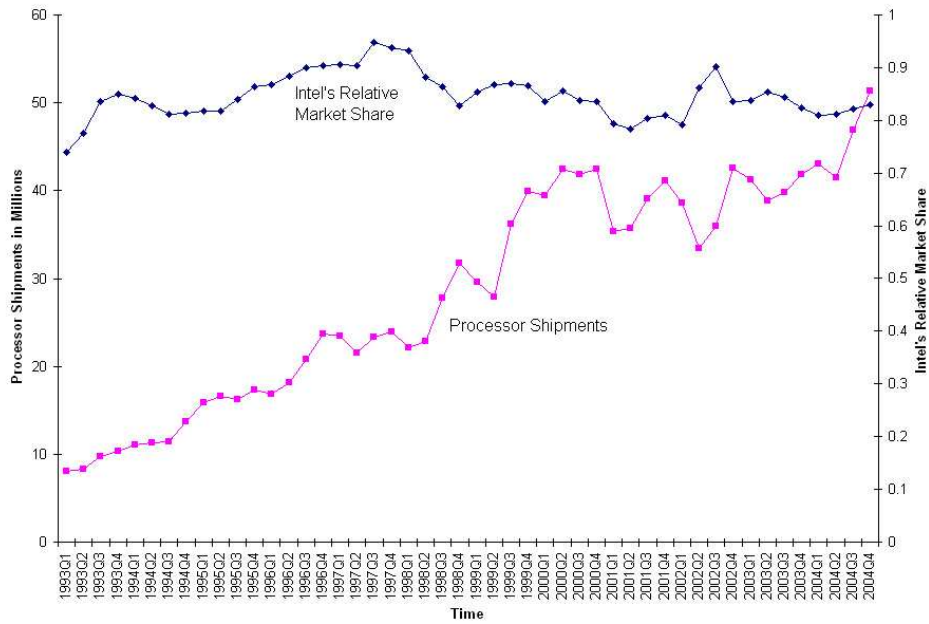


Figure 3: R&D Spending

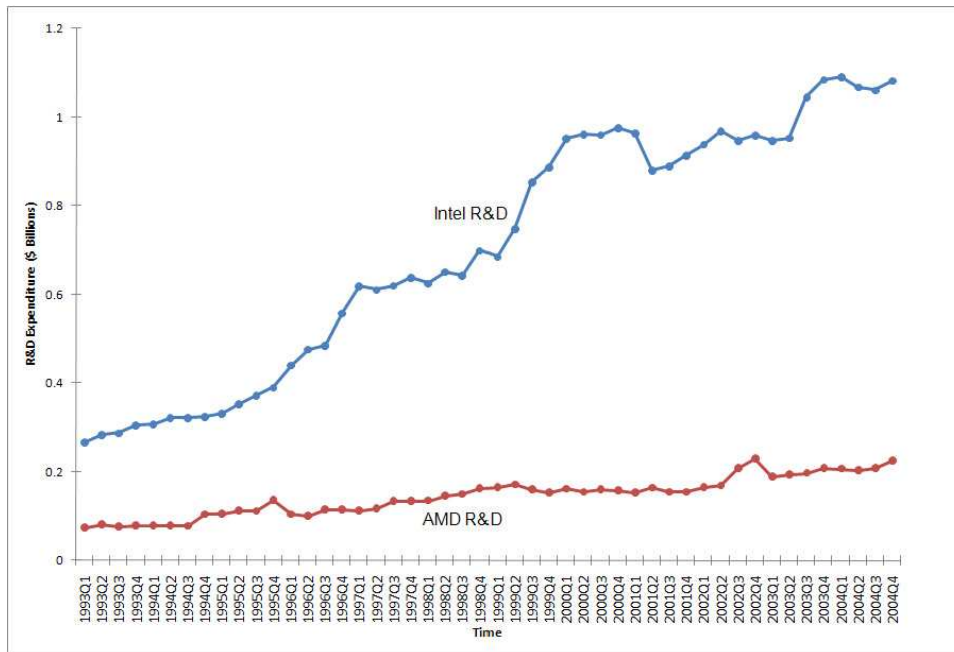


Figure 4: Value and Policy Functions: Column 1 has ownership with many recent purchases; Column 2 has ownership with mostly old vintages. Value functions are reported in \$ billions and investments are in \$ millions.

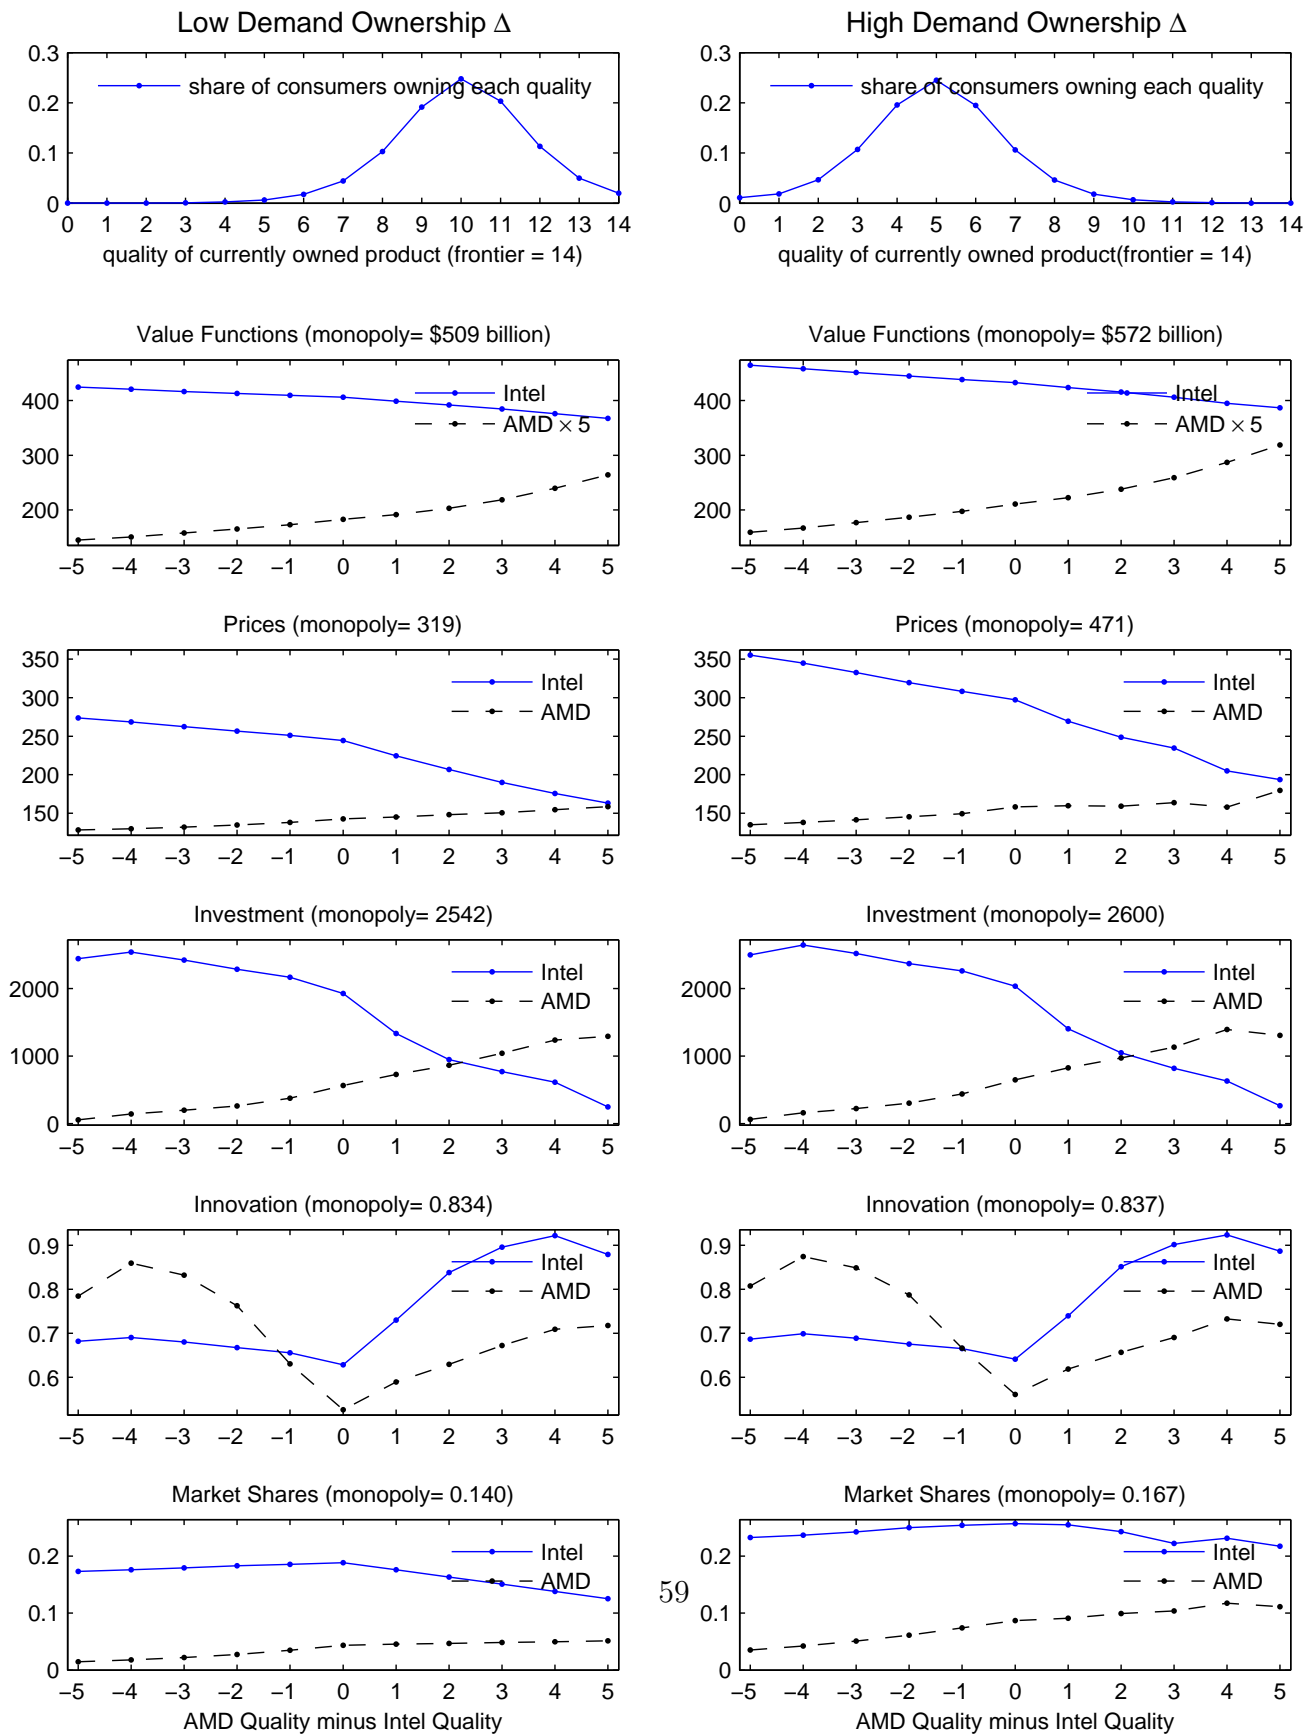


Figure 5: Simulated Time Series

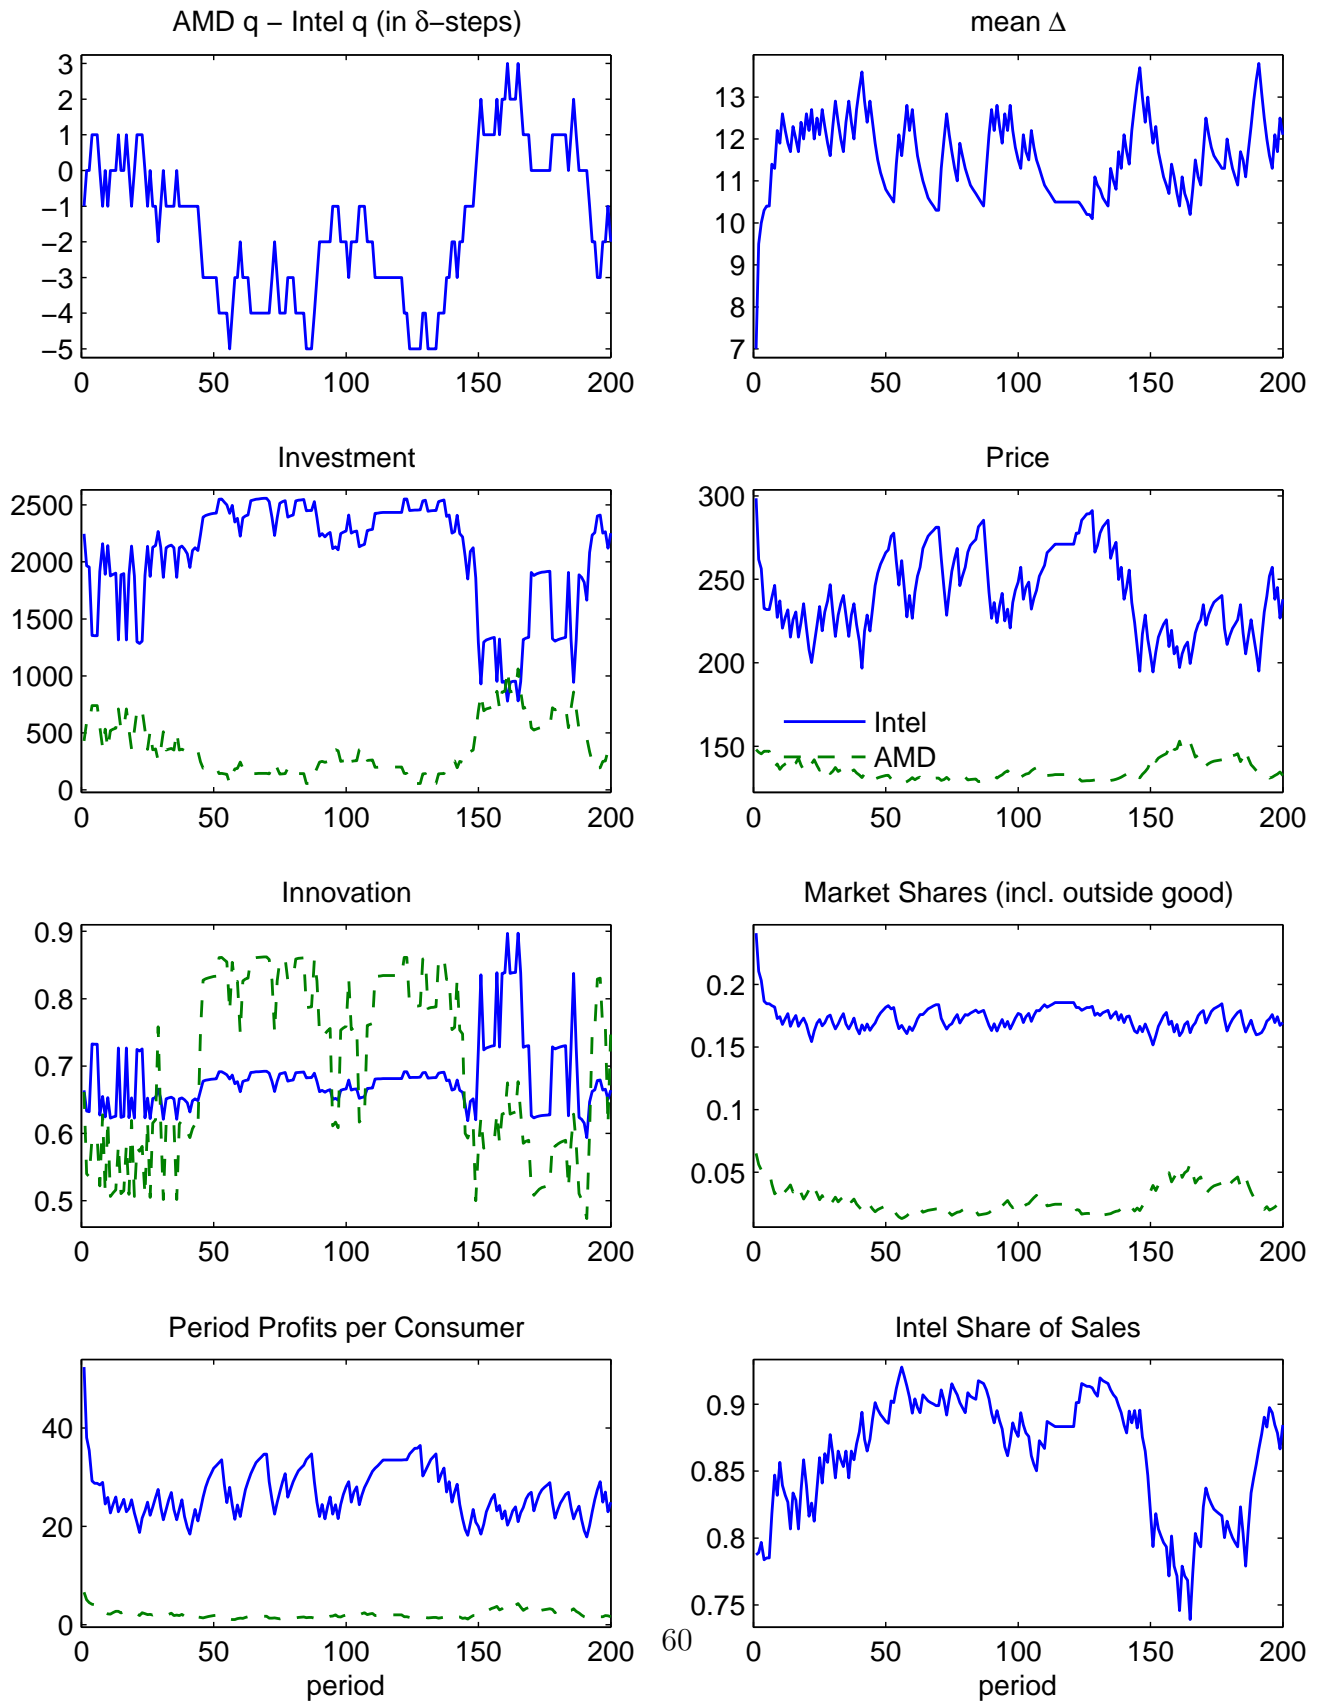


Figure 6: Myopic vs. Optimal pricing for Monopoly

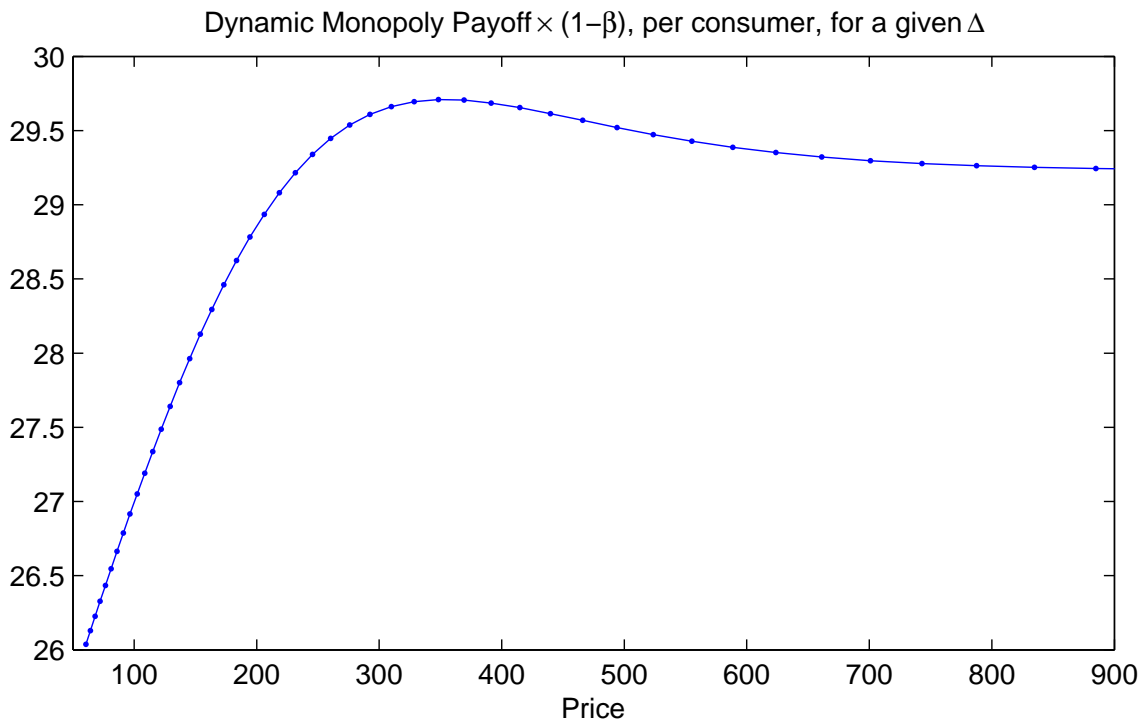
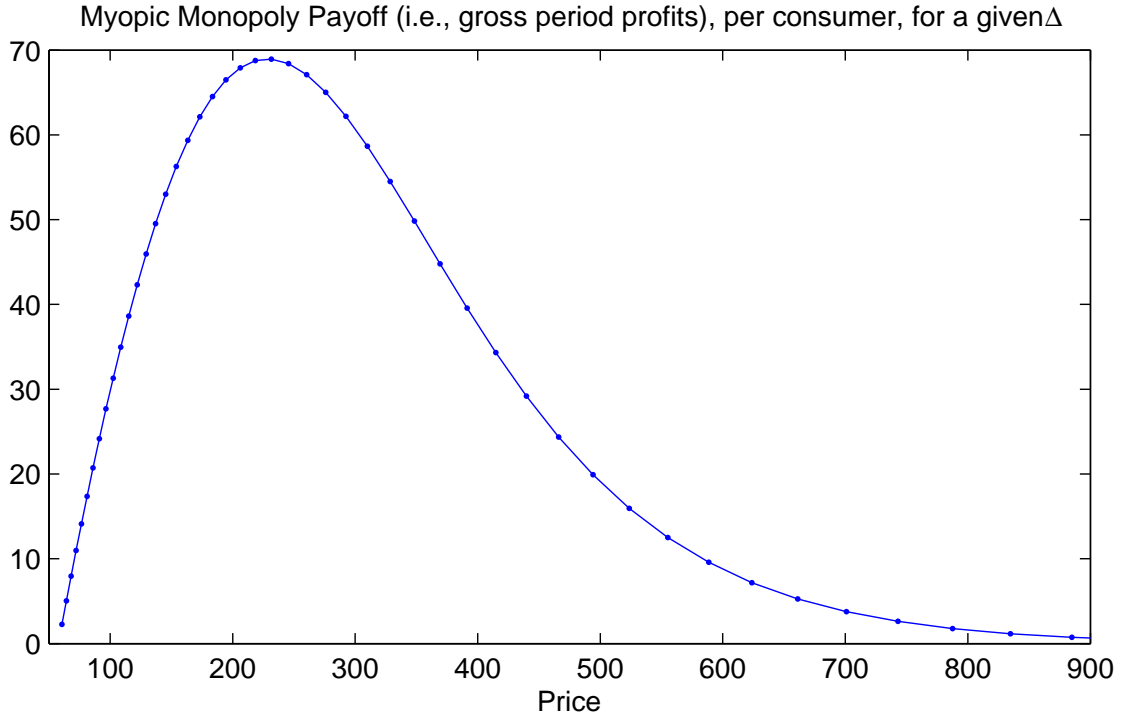


Figure 7: Purchase Probabilities by Vintage, Duopoly Case

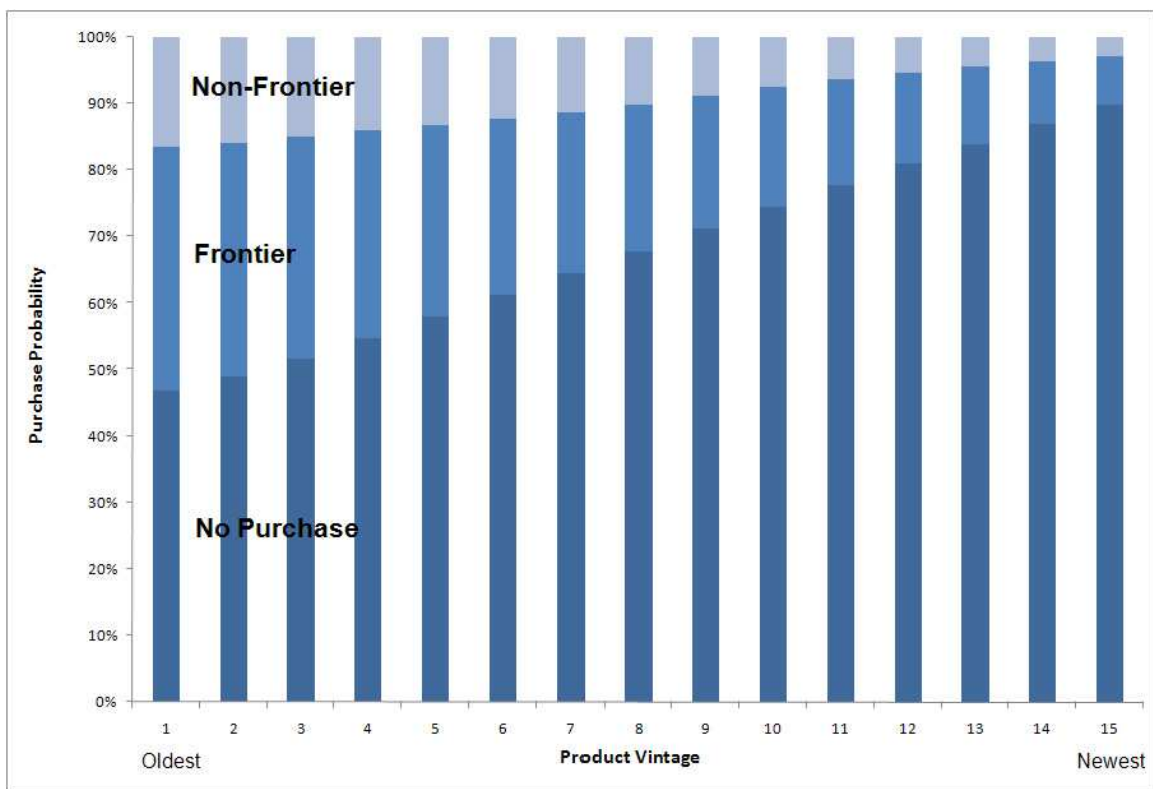


Figure 8: Ownership Distribution, Duopoly and Monopoly Cases

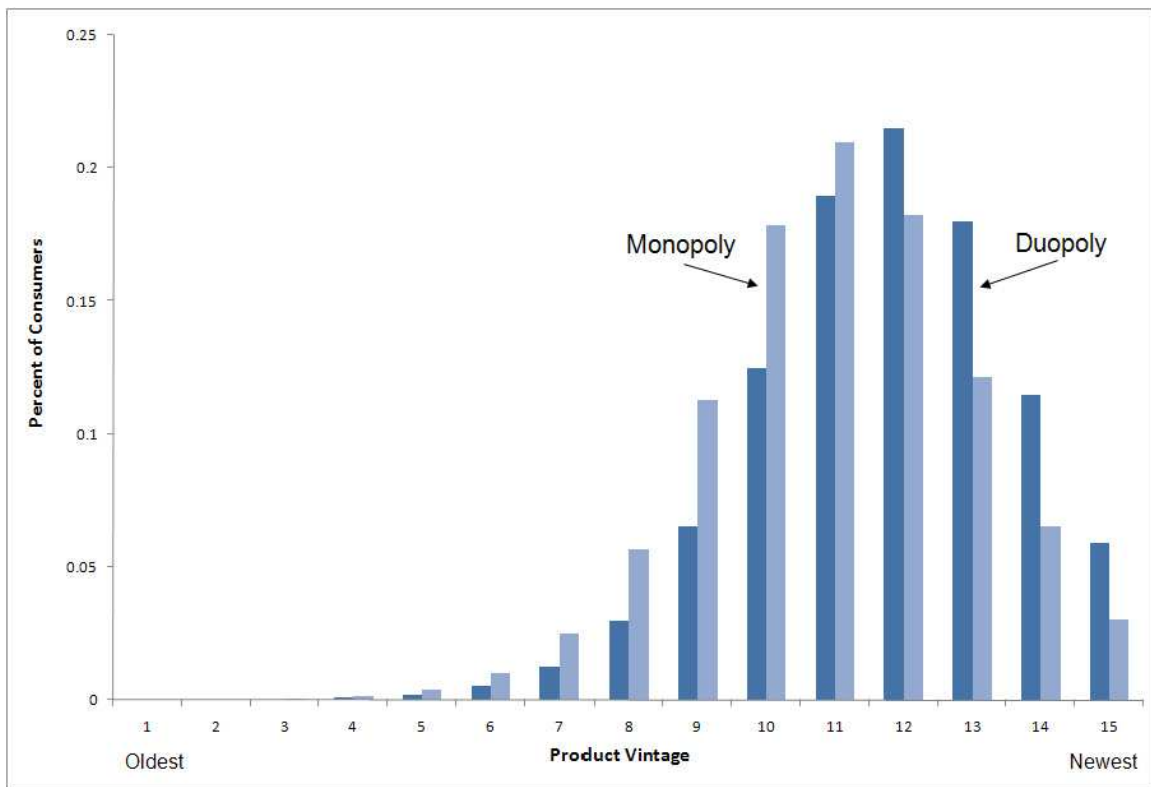


Figure 9: Each Vintage's Share of Purchases

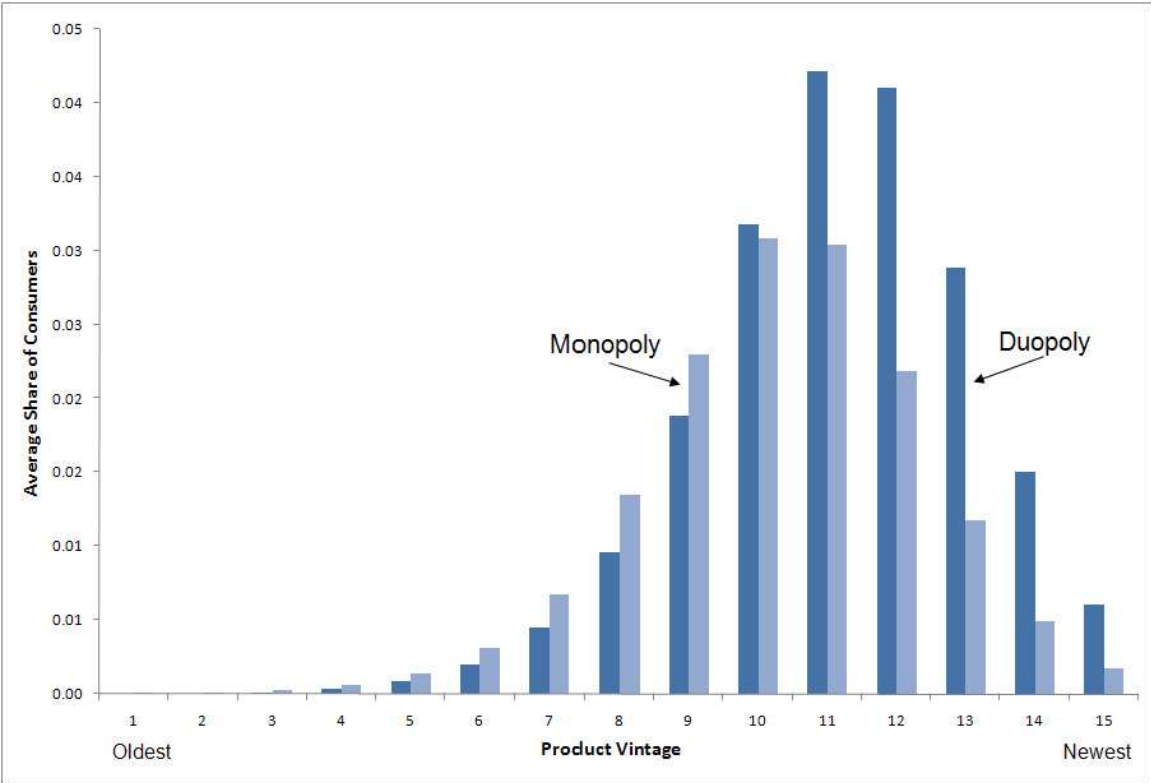


Figure 10: Elasticities and Markups for duopoly, monopoly, and myopic pricing monopoly

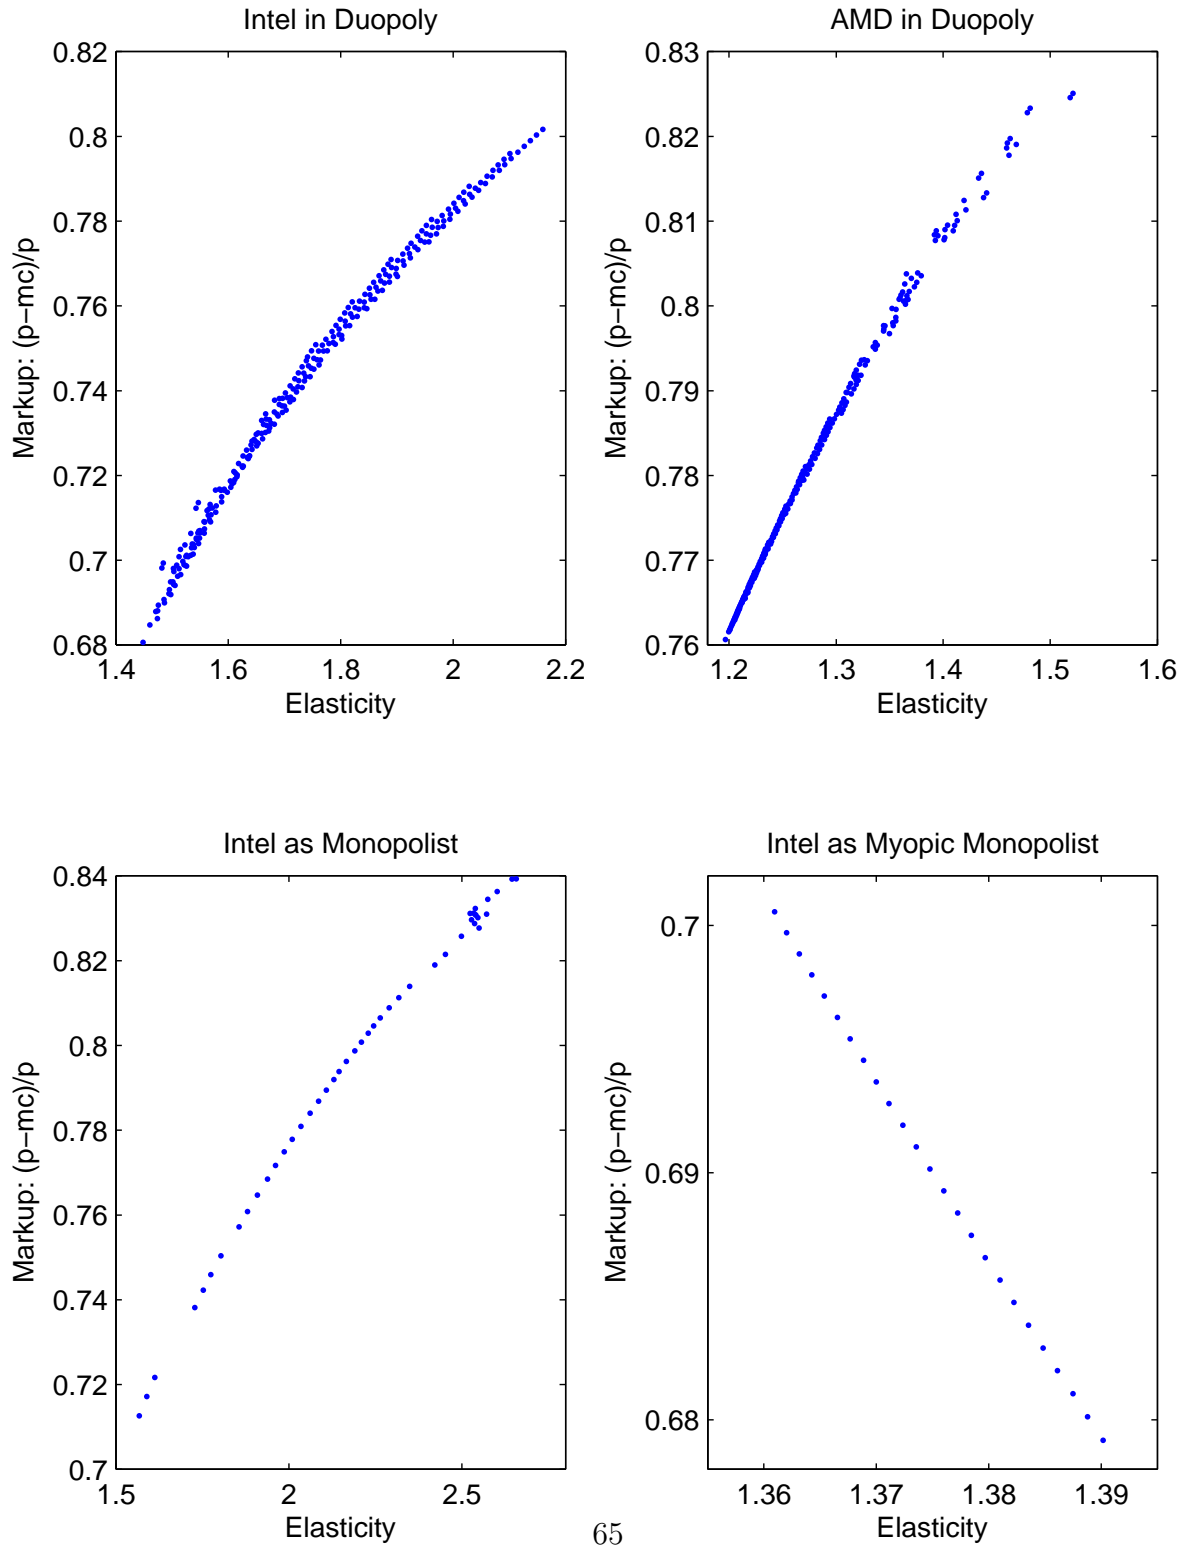


Figure 11: Comparison of Outcomes Varying Consumer Discount Factor

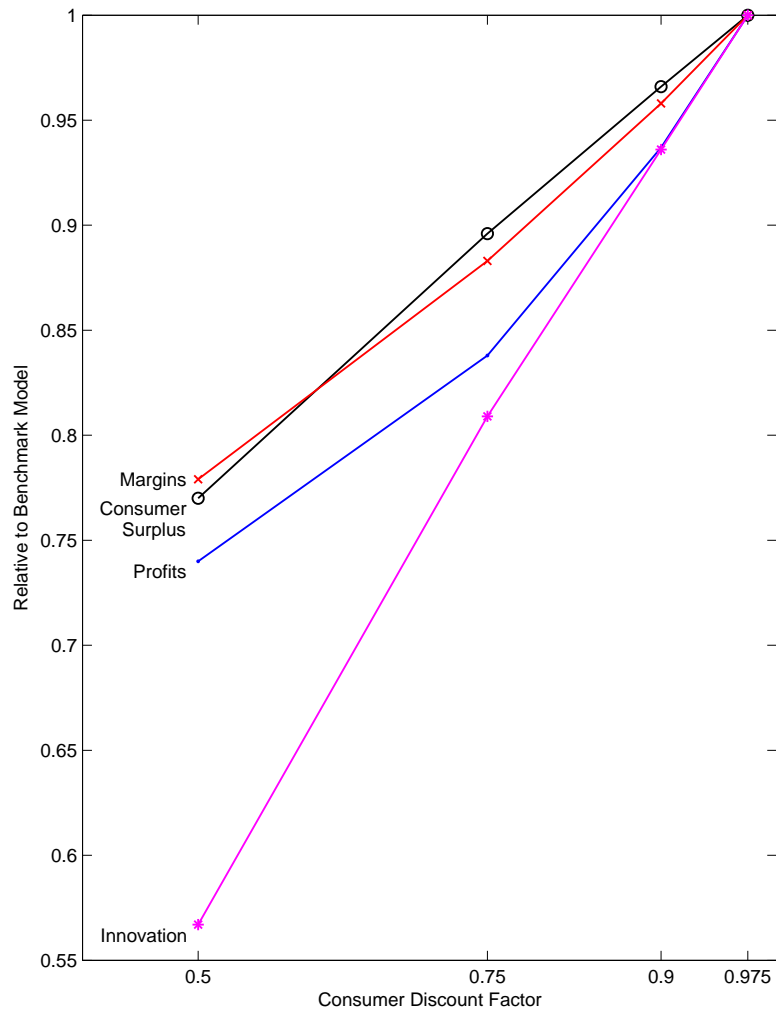


Figure 12: Market Restriction

