Competition and Innovation in the Microprocessor Industry: Does AMD spur Intel to innovate more? *

Ronald Goettler†
University of Chicago

Brett Gordon‡
Columbia University

June 23, 2009

Abstract

We propose and estimate a model of dynamic oligopoly with durable goods and endogenous innovation to examine the relationship between market structure and the evolution of quality. 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 long-run rate of innovation. We estimate the model for the PC microprocessor industry and perform counterfactual simulations to measure the benefits of competition. Consumer surplus is 2.5 percent higher ($5 billion per year) with AMD than if Intel were a monopolist. Innovation, however, would be higher without AMD present. Counterfactuals reveal that consumer surplus can actually increase as the market moves toward monopoly, which suggests policymakers ought to consider the dynamic trade-off of lower current consumer surplus from higher prices for higher future surplus from more innovation. Comparative statics reveal that competition does induce higher innovation if consumers have sufficiently high preferences for quality and low price sensitivity.

Keywords: dynamic oligopoly, competition and innovation, 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. A previous version of this paper circulated as "Durable Goods Oligopoly with Innovation: Theory and Empirics."

†E-mail: goettler@uchicago.edu

‡E-mail: brg2114@columbia.edu

1 Introduction

Economists have long sought to understand the relationship between market structure and innovation to inform policy governing antitrust, patent regulation, and economic growth. However, after more than 60 years of research, the conclusions of both the theoretical and empirical literatures are decidedly mixed. The original theoretical hypothesis, proposed by Schumpeter (1942), posits a positive relationship between market concentration and innovation. Fellner (1951) and Arrow (1962) argue for a linear negative relationship, and Scherer (1967) proposes a model yielding an inverted-U relationship. The empirical literature has found varying degrees of support for each of these hypothesis, leading Cohen and Levin (1989), in their review of the literature, to state "the empirical results bearing on the Schumpeterian hypotheses are inconclusive." They argue the lack of an empirical consensus is primarily due to data measurements errors and the difficulty of controlling for various industry-specific factors.

More recently, Aghion, et al. (2005) estimate an inverted-U relationship using a panel of UK firms, controlling for industry effects and year effects. Below we replicate their Figure 1 to illustrate that although a statistically significant relationship exists between the Lerner index and citation-weighted patent counts, the substantial variation in the data around the curve suggests one might hesitate to rely on it when evaluating a particular merger.[1]

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

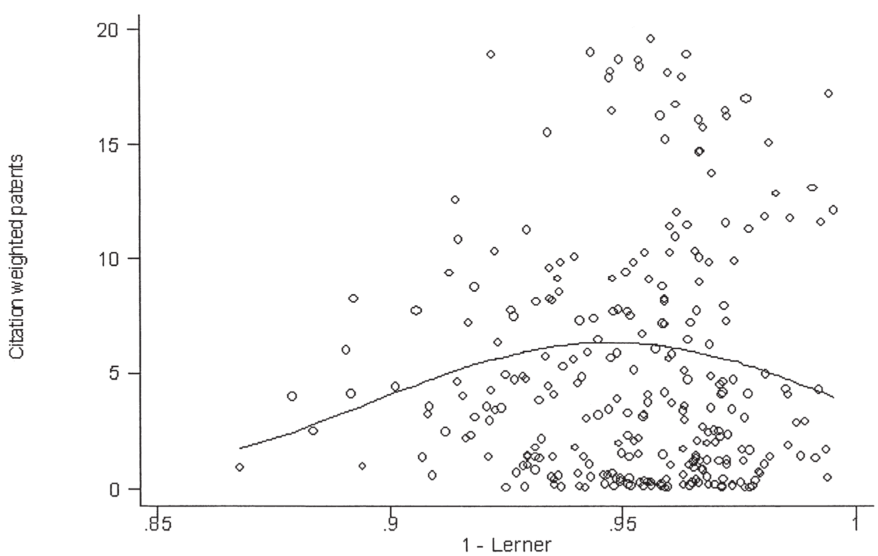

FIGURE I

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.

[1] A potential explanation for this variation lies in the theoretical result of Vives (2008) who finds an ambiguous relationship between the Lerner index and expenditures on cost-reducing innovations.

In this paper we pursue a complementary approach to the reduced-form empirical studies. Rather than attempting to characterize the relationship between market structure and innovation across multiple industries, we focus on understanding this relationship in a particular industry.[2] We construct a structural model of dynamic oligopoly with endogenous innovation and use it to assess the effect of competition on industry outcomes. Our work is a natural extension to the early industry simulation models of Nelson and Winter (1982) and Grabowski and Vernon (1987) and to the dynamic oligopoly model of Ericson and Pakes (1995). To our knowledge we are the first to investigate the effect of competition on innovation using structural empirical methods.

We estimate the model for the PC microprocessor industry and perform counterfactuals to measure the effect of market structure on innovation, profits, and consumer surplus. Lacking variation in the number of firms, our identification of the effect of competition on innovation relies on estimating consumer preferences and firms' innovation efficiencies, from which we can derive the benefits and costs of innovation for any exogenously specified market structure.[3] This approach accords with Dorfman and Steiner (1954), Needham (1975), and Lee (2005), who find that consumer preferences and firm competencies are key determinants of R&D.

The microprocessor industry is well-suited for our study for at least three reasons. First, processor innovations are easily measured via changes on performance benchmarks. Second, the industry is close to being a pure duopoly with Intel and Advanced Micro Devices (AMD) comprising 95 percent of sales. Finally, the industry is of extreme importance to the economy and is in the midst of an antitrust lawsuit in which AMD alleges Intel's anti-competitive practices have restricted its access to consumers.[4]

However, the analysis is complicated by the fact that microprocessors, like most high-tech goods, are durable. Sellers therefore face a dynamic trade-off: selling more today reduces future demand. One strategy firms use to mitigate this consequence of durability is to continually improve product quality to induce upgrade purchases 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. Our model of dynamic oligopoly with durable goods and endogenous innovation sheds light on these issues.

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 later. Since consumers' choices depend on the products they currently own (if any), the distribution of currently owned products affects aggregate demand. We explicitly model the endogenous evolution of this distribution and its effect on equilibrium behavior. In particular, we show that accounting for

[2]This path of inquiry is potentially more relevant given federal antitrust investigations, such as Intel vs. AMD, and Microsoft vs. Netscape, in which policymakers must concentrate on the unique aspects of competition and innovation in specific industries.

[3]Our model can be extended to allow for endogenous market structures. Counterfactuals would then correspond to exogenous shifts in industry or policy features, such as entry costs or restrictions on firm size.

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

product durability has significant implications for prices, innovation, profits, and consumer surplus.

Importantly, our model can generate either a positive or a negative relationship between competition and innovation. Thus, we allow the data to guide our conclusions. We find innovation would be 6.4 percent higher if Intel were a monopoly. This result 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 invests more is that its pricing power enables it to extract more of the surplus generated by innovations.[5] Consumer surplus, however, would be 2.5 percent ($5 billion per year) lower without AMD, since prices would be substantially higher. In a counterfactual in which we vary the share of the market from which AMD is excluded, we find that consumer surplus is actually higher when AMD is partially excluded. This finding supports the FTC's recent consideration of the dynamic trade-off of lower current consumer surplus from higher prices for higher future surplus from more innovation.[6]

To gain further insight, we perform comparative statics by varying consumer preferences for quality and price (vertical competition) and the scale of consumers' idiosyncratic preferences (horizontal differentiation). We find that equilibrium innovation rates increase monotonically as preferences for quality increase and as distaste for price declines. In contrast, equilibrium innovation is U-shaped in horizontal differentiation. When products are near perfect substitutes, firms innovate rapidly to ensure survival. As substitutability declines, innovation to ensure survival drops, but eventually picks up again in response to firms' increased local market power. Our main empirical result that innovation is lower with competition holds for all levels of horizontal differentiation, and for all configurations of vertical competition except when consumers highly value quality and are relatively insensitive to price.

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 account 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.

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

[5] In a nondurable goods setting, Macieira (2007) finds competition increases innovation in the supercomputer industry.

[6] Gilbert (2005) reports that the Federal Trade Commission increasingly cites the potential effect of competition on innovation as a concern, despite the absence of conclusive theoretical or empirical evidence.

it also implies the industry's long-run 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 long-run innovation rate is endogenous.

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 (2008) and Markovich and Moenius (2008) 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 prices in the automobile industry, taking durability as fixed. Chen, Esteban, and Shum (2008) assess the competitive effect of secondary markets for automobiles. Lee (2008) estimates supply and demand in the two-sided video game market. 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.[7]

Our work also connects to the large theoretical literature on durable goods, which is reviewed by Waldman (2003). Early work focuses 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 investigates the robustness of 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 microprocessors). Firms do, however, choose innovation rates, which determine the rate at which goods become obsolete. Though similar, durability and obsolescence have a poten-

[7] Our 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 consider only 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 (forthcoming) and Gowrisankaran and Rysman (2007) estimate models that allow for replacement purchases. By estimating supply as well as demand, we are able to compute equilibrium outcomes under counterfactual scenarios of interest.

tially important difference: durability entails commitment since the good is produced with a given durability, whereas obsolescence depends on future innovations.

The second area, beginning with Coase (1972) and followed by Stokey (1981) and Bulow (1982), among others, considers the monopolist's time inconsistency problem: the firm would like to commit to a fixed price over time but after selling to today's buyers, will then want to lower the price to sell to those consumers who were unwilling to buy at the supposedly fixed price.[8] 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, we assume firms commit to fixed prices within periods, but not across periods.[9]

Thus far, this literature has focused 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 $j \in \{1, \ldots, J\}$ sells a single product with time-varying log-quality denoted $q_{jt} \in \{0, \delta, 2\delta, \ldots\}$.[10] In each period firms simultaneously choose prices p_{jt} and investment x_{jt}.[11] 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). Hence, the distribution of currently owned products affects current demand. We denote this endogenous distribution Δ_t.

We do not consider secondary markets since computers and microprocessors are rarely resold.[12]

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

[9] An interesting comparative static may be to see how industry outcomes vary with period length. Such a comparative static, however, is not trivial to construct since it involves changing the scale of several parameters simultaneously and tweaking the innovation process to maintain *ceteris paribus*.

[10] We could normalize q_{jt} to be positive integers, but we can interpret the model more easily 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. In the conclusion we conjecture on the robustness of our results to the multiproduct setting.

[11] The 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.

[12] One could add a secondary market in which adverse selection results in an (exogenous) transaction cost. The ownership distribution Δ_t would then convey the set of used goods available for trade.

We also model consumers as owning one microprocessor at a time.[13]

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. All agents observe the vector of firms' qualities $q_t = (q_{1t}, \ldots, q_{Jt})$ and the ownership distribution Δ_t. These two state variables comprise the state space of payoff relevant variables for firms. The consumer's state space consists of the quality of her currently owned product $\tilde{q}_t$, the firms' current offerings q_t, and the ownership distribution Δ_t. This latter state variable is relevant to the consumer since it affects firms' current and future prices and investment levels.[14]

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 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.[15] For conciseness we omit the oft used i subscript for consumers.

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 available to the consumer 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, a variety of factors could determine the utility from the no-purchase option. In the context of CPUs, the outside good for "non-owners" may consist of using computers at schools and libraries or using old computers received from 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 $\max(q_t)$.

[13] Hendel (1999) estimates a static model of multi-unit purchases of personal computers and Cho (2008) estimates a dynamic model of computer replacement by a telecommunications firm using many computers.

[14] We assume consumers observe the ownership distribution merely as a convenient way to impose rational expectations of future prices and qualities. Rationality requires consumers act as if they condition on the ownership distribution since it influences innovation and future prices through firms' policy functions.

[15] As explained by Rust (1996) the *independence from irrelevant alternatives* (IIA) property generally fails to hold in dynamic contexts since the attributes of all the products typically enter the continuation values even when the utility flows depend only on the characteristics of the product chosen.

To ensure that our choice of $\bar{\delta}_c$ does not drive the model's behavior, we check that consumers upgrade frequently enough that the quality of their most recent purchase is rarely below $\underline{q}_t$.[16]

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 consumers with $\tilde{q} \leq \underline{q}_t$ as owning the lower bound itself. Hence, $\Delta_t = (\Delta_{\underline{q}_t,t}, \ldots, \Delta_{k,t}, \ldots, \Delta_{\bar{q}_t,t})$, where $\Delta_{k,t}$ is the fraction of consumers whose outside option (i.e., current product) has quality q_{kt}.

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

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

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 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 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 $\tilde{q}_t$. With $J+1$ choices in each period, we normalize ξ_0 to be zero and simply exclude it from u_{0t}. We expect $\xi_j < 0$ for $j > 0$ since in a given period most consumers do not upgrade. We interpret the highest of these ξ_j as an upgrade switching cost and the differences between ξ_j and this maximum to be brand preferences. One could alternatively specify the ξ_j to enter the utility flow for the no-purchase option, but in our empirical application having future utility depend solely on the product's

[16] The assumption that non-owners' utility increases in the frontier quality is unnecessary in markets with sufficiently few non-owners (e.g., automobiles and appliances). However, if many sales are to new adopters and the outside good does not improve along with the frontier, then $\bar{\delta}_c$ will need to be chosen high enough that most consumers will have purchased the product before the frontier achieves a quality advantage of $\bar{\delta}_c$.

quality seems more natural.

Each consumer is small relative to the market so that her actions do not affect the evolution of Δ_t. We also assume consumers have identical preferences. 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. Although 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 its ability to introduce higher quality products into the market. Firms choose a level $x_j \in \mathbb{R}_+$ to invest in R&D . The R&D outcome, denoted $\tau_{jt} = q_{j,t+1} - q_{jt}$, is probabilistic and stochastically increasing in the level of current investment. We restrict $\tau_{jt} \in \{0, \delta\}$ and denote its probability distribution $f(\cdot|x, q_t)$.[17] 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 to the frontier than when advancing it.

As with most industries, innovation in the microprocessor industry is influenced by factors other than investment by its firms. For example, scientific discovery from public and quasi-public research institutions clearly contribute to innovation. Such external forces are implicitly captured in our model through the innovation efficiency of investments by "inside" firms.

The innovation literature distinguishes between process and product innovations, as does Ericson and Pakes (1995). We focus on product innovations for a few reasons. First, quality improvements are frequent, easily observed, and directly affect consumers' choices. By comparison, innovations that lower costs, such as increasing the size of silicon wafers (from which individual processors are cut) are infrequent and difficult to observe. Microprocessor firms also invest in capacity, but third-party microprocessor foundries are able to provide additional capacity as needed.[18] As such, improving product quality appears to be the most important investment for firms in this industry.

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.[19] 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 $a_j(q_t) = a_{0,j} \max(1, a_1(\frac{\bar{q}_t - q_{jt}}{\delta})^{1/2})$, which reflects the increased difficulty of advancing the frontier relative to catching up to it.[20]

[17] Proportional improvements in quality, as implied by a log-quality grid, suit the microprocessor industry.

[18] Interestingly, in 2008 AMD announced plans to go "fabless" with all production occurring at fabrication plants on a contractual basis.

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

[20] We experimented with linear and convex spillovers as well, but this concave function fit the best. The parameters

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

$$\pi_j(p_t, q_t, \Delta_t) = M s_{jt}(p_t, q_t, \Delta_t)(p_{jt} - mc_j(q_t)) \ , \tag{4}$$

where M is the (fixed) market size, $s_{jt}(\cdot)$ is the market share for firm j, and p_t is the vector of J prices. Firm j's constant marginal costs are given by

$$mc_j(q_t) = mc_0 + mc_{slope}\,(\bar{q}_t - q_{j,t}) \ , \tag{5}$$

where $mc_{slope} < 0$ implies production costs are lower for non-frontier firms.[21]

Each firm maximizes its expected discounted profits, yielding the Bellman equation

$$\begin{aligned} W_j(q_{jt}, q_{-j,t}, \Delta_t) \quad = \quad & \max_{p_{jt}, x_{jt}} \ \pi_j(p_t, q_t, \Delta_t) - c x_{jt} + \beta \sum_{\tau_{jt},\, q_{-j,t+1},\, \Delta_{t+1}} W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}) \\ & f(\tau_{jt}|x_{jt})\ h_f(q_{-j,t+1}|q_t, \Delta_t)\ g_f(\Delta_{t+1}|\Delta_t, q_t, p_t, q_{t+1}) \ , \end{aligned} \tag{6}$$

where c is the unit cost of investment, $h_f(\cdot|\cdot)$ is the firm's beliefs about its competitors' future 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.[22]

Following Rust (1987) we assume consumers' $\{\varepsilon_{jt}\}$ are multivariate extreme-value and integrate over future ε_{jt} to obtain the product-specific value function

$$\begin{aligned} \hat{V}_j(q_t, \Delta_t, \tilde{q}_t) \quad = \quad & u_{jt} - \varepsilon_{jt} + \beta \sum_{q_{t+1},\, \Delta_{t+1}} \log\left(\sum_{j' \in \{0,\ldots,J\}} \exp\left\{ \hat{V}_{j'}(q_{t+1}, \Delta_{t+1}, \tilde{q}_{t+1}) \right\} \right) \\ & h_c(q_{t+1}|q_t, \Delta_t)\ g_c(\Delta_{t+1}|\Delta_t, q_t, p_t, q_{t+1}) \ . \end{aligned} \tag{7}$$

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

$$s_{jt|\tilde{q}} = \frac{\exp\{\hat{V}_j(q_t, \Delta_t, \tilde{q}_t\}}{\sum_{k \in \{0,\ldots,J\}} \exp\{\hat{V}_k(q_t, \Delta_t, \tilde{q}_t)\}} \ . \tag{8}$$

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

$$s_{jt} = \sum_{\tilde{q} \in \{\underline{q}_t, \ldots, \bar{q}_t\}} s_{jt|\tilde{q}}\ \Delta_{\tilde{q},t} \ . \tag{9}$$

These market shares translate directly into the law of motion for the distribution of ownership.[23]

other than a_1 were insensitive to the form of the spillover.

[21] In our application $|mc_{slope}|$ is small enough that marginal costs are always positive.

[22] Expressing 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.

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

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

$$\Delta_{k,t+1}(\cdot) = s_{0t|k}\Delta_{kt} + \sum_{j=1,\dots,J} s_{jt} I(q_{jt} = k) \ , \tag{10}$$

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 $t+1$, we shift the interim Δ_{t+1} that results from equation (10) via

$$\Delta_{t+1} = \Gamma\left(\Delta_{t+1}\right). \tag{11}$$

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

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}:

$$\begin{aligned} -c + \beta \sum_{\tau_{jt},\, q_{-j,t+1},\, \Delta_{t+1}} & W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}) \\ h_f(q_{-j,t+1}|q_t, \Delta_t)\; g_f(\Delta_{t+1}|\Delta_t, q_t, p_t, q_{t+1})\; f(\tau_{jt}|x_{jt}) \frac{\partial f(\tau_{jt}|x_{jt})}{\partial x_{jt}} & \quad = \quad 0 \ . \end{aligned} \tag{12}$$

Given q_t and outcomes for $(\tau_{jt},\ q_{-j,t+1})$, the transition for Δ_{t+1} depends only on prices. Given our choice for f we analytically compute optimal investment as a function of price, $x^*_{jt}(p_{jt})$, as detailed in Appendix A.

Finally, we note that physical depreciation of goods could be added to the model by supposing that currently owned products decline each period by one (or more) quality grid-steps with some fixed probability. In our application of the model to the CPU industry, however, physical depreciation is zero.

2.3 Equilibrium

We consider pure-strategy Markov-Perfect Nash Equilibrium (MPNE) of this dynamic oligopoly game in which all state variables are indeed Markov. Our MPNE extends that of 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 determine 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, an 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 (10) and (11).[24]

The functional form of the investment transition function satisfies the unique-investment-choice (UIC) admissibility criterion in Doraszelski and Satterthwaite (2007). To guarantee existence of equilibrium one must show that a pure-strategy equilibrium exists in both investment choices and prices. Ericson and Pakes (1995) and the 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). We have verified numerically that the firm's objective function is concave in its own price and that the best-response functions yield unique equilibrium prices at each state holding behavior at other states fixed.

3 Computation

This section discusses the details behind the computation of the MPNE we 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. We relegate the algorithmic details of computing and simulating the equilibrium to Appendix B.

[24]Symmetry 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 of 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:*

$$\begin{array}{rrcl} \textit{Firms:} & W_j(q_{jt}, q_{-j,t}, \Delta_t) & = & W_j(q_{jt} - \bar{q}_t, q_{-j,t} - \bar{q}_t, \Delta_t) \\ \textit{Consumers:} & V(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) & = & \frac{\gamma \bar{q}_t}{1-\beta} + V(q_t - \bar{q}_t, \Delta_t, \tilde{q}_t - \bar{q}_t, \varepsilon_t) \ . \end{array} \tag{13}$$

The proof, which appears in Appendix C, 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 tracks 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 the Bellman equations. When integrating over future states, some of the possible states involve frontier quality increasing (always by δ units). The consumer's continuation value for such an 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 $t + 1$. We can then express the consumer's product-specific value function in equation (7) using the relative state space as

$$\hat{V}_j(\omega_t, \Delta_t, \tilde{\omega}_t) = \gamma\omega_{jt} - \alpha p_{jt} + \beta \sum_{I_{\bar{q}_t},\, \omega_{t+1},\, \Delta_{t+1}} \log\left(\sum_{j' \in \{0,\ldots,J\}} \exp\left\{ \frac{\gamma\delta I_{\bar{q}_t}}{1-\beta} + \hat{V}_{j'}(\omega_{t+1}, \Delta_{t+1}, \tilde{\omega}_{t+1}) \right\} \right)$$
$$h_c(I_{\bar{q}_t}, \omega_{t+1}|\omega_t, \Delta_t)\; g_c(\Delta_{t+1}|\Delta_t, \omega_t, p_t, I_{\bar{q}_t}) \;, \tag{14}$$

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 (6) using the relative state space becomes

$$\begin{aligned} W_j(\omega_{jt}, \omega_{-j,t}, \Delta_t) = \max_{p_{jt}, x_{jt}} \; & \pi_j(p_t, \omega_t, \Delta_t) - cx_{jt} + \\ & \beta \sum_{\tau_{jt},\, \omega_{-j,t+1},\, I_{\bar{q}_t},\, \Delta_{t+1}} W_j(\omega_{jt} + \tau_{jt} - I_{\bar{q}_t}, \omega_{-j,t+1} - I_{\bar{q}_t}, \Delta_{t+1}) \\ & h_f(I_{\bar{q}_t}, \omega_{-j,t+1}|\omega_t, \Delta_t)\; g_f(\Delta_{t+1}|\Delta_t, \omega_t, p_t, I_{\bar{q}_t})\; f(\tau_{jt}|x_{jt}) \;, \end{aligned} \tag{15}$$

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.[25] In particular, we verify that investment behavior when a laggard is maximally inferior does not suggest 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 (2007). 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.[26] This concavity implies the derivative of

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

[26]The 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.

market share with respect to a firm's own quality goes to zero regardless of its competitors' qualities. Since investment is costly, a relative quality above which investment is zero will exist, thereby establishing an upper bound. Firms exiting when relative quality gets sufficiently low establishes the lower bound.

The EP approach to bounding the state space has a few potential drawbacks. Perhaps the most significant is that the exogenous innovation rate of the outside good solely determines the industry's long-run innovation rate. Improvements in the outside good provide a continual need for inside firms to invest to remain competitive. If the outside good never improves, the equilibrium has no investment and no innovation in the long run. Since an industry's innovation rate, rather than its pricing behavior, may be the primary determinant of discounted consumer surplus, providing an endogenous long-run innovation rate that varies across market structures and regulatory controls seems 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 implication of EP is that consumers' utility rankings of inside goods depends on the outside good's quality. For example, a consumer indifferent between two products that differ in quality will strictly prefer the higher quality product 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.[27]

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 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. Initially writing the model in absolute levels best addresses such consistency issues.[28] 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 generates an upper bound by manipulating the behavior of lead firms, whereas we define quality relative to the frontier and generate a lower bound by truncating the degree to which firms and outside options can be inferior. Since industry leaders generate most of the sales, profits, and surplus, assumptions regarding severe laggards are more innocuous than assumptions restricting the benefits to innovation

[27]The quality index may itself be a concave transformation of measured quality, such as the logarithm, to reflect diminishing marginal utility for measured quality.

[28]Goettler, Parlour, and Rajan (2005, 2008) also take the approach of initially specifying their dynamic asset trading game in absolute levels and then transforming the game to use relative states.

by frontier firms.[29]

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 (10) and (11) 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 most relevant to future profitability, pricing, and investment. For logit demand systems the mean is the most relevant moment.[30] We therefore define

$$\rho_d(\Delta'_{t+1}) = \left| \sum_k k\Delta'_{k,t+1} - \sum_k k\Delta^d_k \right|, \quad \text{for all } d \in (1, \ldots, D), \tag{16}$$

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}, and define the stochastic transition of the discretized Δ as

$$\Delta_{t+1} = \begin{cases} \Delta^{d_1} \text{ with probability } \frac{\rho_{d_1}}{\rho_{d_1}+\rho_{d_2}} \\ \Delta^{d_2} \text{ with probability } \frac{\rho_{d_2}}{\rho_{d_1}+\rho_{d_2}} \end{cases}. \tag{17}$$

Our discretization is a multidimensional version of the stochastic transition Benkard (2004) uses 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.[31] We choose the discrete grid of this scalar such that the mean qualities 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 trade-offs associated with dynamic demand.[32]

[29] Although 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 of Pakes and McGuire (1994). We also note that our approach could allow for features of the EP specification that we currently exclude. For example, the outside good's quality (for non-owners) may be allowed to stochastically improve according to an exogenous process exactly as in EP.

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

[31] We 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.

[32] This choice for $\{\Delta^d\}_{d=1}^D$ is also computationally efficient since one may obtain the two closest Δ^ds using an indexing formula instead of a search.

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.

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.[33] To illustrate theoretical properties of the model, in Section 5 we present comparative statics for the symmetric case in which firms have identical brand intercepts and innovation efficiencies.

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 combine 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 and sales-weighted blended unit production costs from In-Stat/MDR, an industry research firm that specializes in microprocessors. We use a PC processor speed benchmark from the CPU Scorecard (www.cpuscorecard.com) to construct a single index of quality comparable across firms and product generations.[34] All prices and costs are converted to base year 2000 dollars.

Aggregate computer ownership and penetration rates come from the Homefront study created by Odyssey, a consumer research firm that specializes in technology products.[35] They survey respondents about the characteristics (e.g., processor speed and manufacturer) of their primary or most recently purchased PC. These data allow us to construct semi-annual observations of the ownership distribution for those consumers who have purchased at least one computer. Figure 1 illustrates this distribution, and its evolution over time, for Intel microprocessors. AMD's distribution looks similar. We interpolate these semi-annual distributions to obtain quarterly data that we then combine with the quarterly penetration rate of PCs in U.S. households to obtain a measure of the state variable Δ. We assume consumers who have yet to purchase a PC by period t derive utility from the outside

[33] To investigate why these asymmetries developed, one would need a model that endogenizes the evolution of brand equity and innovation efficiencies, which is well beyond our current objectives.

[34] We impute missing benchmark scores for 77 processors based on the 177 scores that are reported by cpuscorecard.com.

[35] The firm conducts semi-annual telephone surveys using a nationally representative sample of 1500 to 2500 households. The households do not belong to a pre-chosen panel, and the firm draws a new sample for each wave of the study. The survey data is available neither at the household level nor in panel form. To increase the accuracy of the relevant sample, approximately 500 additional PC-owning households are oversampled.

good equivalent to a PC that is 7.8 percent the speed of the frontier ($\bar{q}_t$), which translates to fourteen δ-steps.[36] We define $\bar{\Delta}_t$ to be the mean quality represented by this measure of Δ_t.

Our model restricts each firm to offer only one product. Since AMD and Intel offer multiple products each period, we must average prices and qualities.[37] Since we do not observe quantities of each processor, we use simple averages of the speeds and prices of processors offered by each firm in each period. The resulting average price series for both Intel and AMD and their respective frontier product prices comprise the first row of plots in Figure 2.[38] The next row of plots presents each firm's frontier quality and the difference in their average qualities from 1993 Q1 to 2004 Q4. 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 over the 48 quarters corresponds to an approximate average of 11 percent improvement per quarter. This rate of improvement implies a doubling in CPU speed every 6.7 quarters, which is consistent with the 18 to 24 months per doubling of transistor counts on integrated circuits that has become known as "Moore's Law."

The correlation between Intel's average price and its quality advantage is evident in the plots. Average prices were generally above \$300 before the K6 was introduced and around \$200 after the Athlon was introduced. The correlation between average prices and the difference between Intel and AMD's average quality is .67 for Intel and $-.56$ for AMD.[39] The significant correlation (.44) between Intel's market share, appearing in the last plot of Figure 2, and its quality advantage is also evident. Such correlations are consistent with our model.

The plot of price divided by log-quality in the bottom of Figure 2 illustrates the rapid innovation in this industry. Since this ratio captures changes in both quality and pricing, we plot our average prediction of its evolution as a visual summary of the model's fit.

Innovation rates are difficult to infer visually from the quality levels, so we report histograms of quarterly changes in each firm's frontier and average log-quality measures in Figure 3. The mean quarterly innovation rates are between 10.4 and 11.3 percent for Intel and AMD's average and frontier qualities.[40] Fewer than one-fifth of the quarters have innovation rates (average or frontier) exceeding the 20 percent steps we impose in our specification, and gains much higher than 20 percent are rare. Moreover, only 3 of the 48 quarters have changes in the log-quality difference between Intel and

[36] For comparison, the 80286 processor (three generations before the Pentium) is 8.6 percent the speed of the Pentium. In essence, we assume non-adopters have access to PCs that are roughly three generations old.

[37] Extending the model to allow each firm to offer multiple products would be a considerable undertaking. If each firm offers two products instead of one, the product-market choice jumps from one price to two prices and one quality (that of the low offering). Since this optimization is solved at each state for each iteration of the value function and may have many local minima, the increase in computational demands is substantial.

[38] The first two price graphs show (local) peaks at the end of 1999 despite the two firms offering similar qualities. These high prices reflect the high demand from the tech bubble and the fact that each firm had recently increased quality. Our model accounts for variations in prices due to recent quality improvements but does not contain the aggregate demand shocks needed to match precisely the variance in prices. We leave the addition of serially correlated aggregate demand shocks for future research since their inclusion substantially complicates the computation of equilibrium.

[39] These correlations are .59 and -.55 respectively, using each firm's frontier product instead of averaging over the product lines.

[40] Note that the mix of products offered can change in such a manner that average quality declines. Indeed, AMD has two quarters in which average quality declines, albeit by less than 1 percent.

AMD that exceed our log-quality step-size of .1823. Hence, the model's inability to deliver large innovations in a single quarter does not conflict with observed innovation in this industry.

Finally, quarterly R&D investment levels, obtained from firms' annual reports, steadily rise over time. Interestingly, AMD invests less than one-fourth the amount Intel does, yet is able to offer similar, sometimes even higher, quality products beginning in 1999. Spillovers offer a potential explanation for this asymmetry, since AMD is usually in the position of playing catch-up. However, we find the model fits the data significantly better when we allow AMD's investment efficiency $a_{0,AMD}$ to be higher than Intel's $a_{0,Intel}$.

4.2 Estimation

To estimate the model's parameters, we use a method of simulated moments (MSM) estimator that minimizes the distance between a set of unconditional moments of our data and their simulated counterparts from 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. One may also view the estimator 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).

For each candidate value of the K-vector θ, 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.[41] The simulated minimum distance estimator $\hat{\theta}_T$, which we detail in Appendix D, is

$$\hat{\theta}_T = \underset{\theta \in \Theta}{\operatorname{argmin}}\ (m_{S,T}(\theta) - m_T)' A_T (m_{S,T}(\theta) - m_T)\ , \tag{18}$$

where m_T is the L-vector of observed moments, $m_{S,T}(\theta)$ is the vector of simulated moments, and A_T is an $L \times L$ positive definite weight matrix. 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.[42]

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. Assuming the data come from the same equilibrium is weaker than our assumption that the model has a unique equilibrium. Unfortunately, we do not have sufficient data to use a two-stage

[41] The only parameter values for which we have been unable to compute an equilibrium are those that are so extreme that the computer generates floating point underflows or overflows. Such parameter values are well beyond the relevant range for estimation.

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

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).

We use value function iteration to numerically solve for equilibrium. Although 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.[43] Though this equilibrium refinement does not guarantee uniqueness, it likely reduces the potential for multiplicity.[44]

4.2.1 Auxiliary Model

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 (i.e., moments to match) would therefore be semi-parametric approximations to these policy functions. We have only 48 quarterly observations, so we restrict ourselves to linear approximations.[45]

One difference between our model and the real world requires care when choosing moments to match. For stationarity, we assume market size M is fixed, whereas the data exhibit an upward trend in sales, revenues, and R&D expenditures. As such, we choose moments that are stationary in both the data and the model. For example, instead of matching the share of the outside good (which declines over time), we match the difference between the frontier quality and the average quality of consumers' current vintages. This moment relates to the outside good's share since a higher share implies less frequent upgrading, which leads to a wider gap between frontier quality and consumers' current vintages. We also match investment per unit revenue, which is stationary in the data, instead of the upward trending R&D.

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

- coefficients from regressing each firm's price on a constant, $\omega_{Intel,t} - \omega_{AMD,t}$, and $\omega_{own,t} - \bar{\Delta}_t$, where $\bar{\Delta}_t$ is the mean quality currently owned (see Section 4.1),
- coefficients from regressing Intel's share of sales on a constant and $\omega_{Intel,t} - \omega_{AMD,t}$,
- mean $(\bar{q}_t - \bar{\Delta}_t)$ (which captures the rate at which consumers upgrade),

[43] Initializing 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.

[44] We have established that we obtain the same equilibrium when we tweak the value function algorithm: we consider Gauss-Jacobi updating and Gauss-Seidel updating with randomly determined orderings of the states. 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.

[45] Gallant and Tauchen (1996) suggest using the score of an auxiliary model that closely approximates the distribution of the data. If the auxiliary model nests the structural model, the estimator is as efficient as maximum likelihood. Hall and Rust (2003) use simple statistics, such as means and covariances. We use a mixture of simple moments and estimates of approximations to policy functions.

- mean innovation rates for each firm, defined as $\frac{1}{T}(q_T - q_0)/\delta$,
- mean $(\omega_{Intel,t} - \omega_{AMD,t})$ and mean $|\omega_{Intel,t} - \omega_{AMD,t}|$, and
- mean investment per unit revenue for each firm.[46]

These moments and their fitted values appear in Table 2 in the next section.

In the estimation we are able to match the correlation of prices with $\omega_{Intel,t} - \omega_{AMD,t}$, 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 2 reveals that variation in average price is determined largely by variation in the frontier price.[47]

We use 10,000 bootstrap replications to estimate the covariance matrix of the moment vector m_T. Its inverse is an estimate of the optimal weighing matrix, which we use for A_T.

4.2.2 Identification

Experimentation with the structural model indicates the auxiliary model's moments are quite sensitive to the structural parameters, which is encouraging from an identification standpoint. Being a nonlinear model, all the structural parameters influence all the auxiliary moments. However, the connection between some parameters and moments are particularly tight.

The demand-side parameters $(\alpha, \gamma, \xi_{Intel}, \xi_{AMD})$ are primarily identified by the six moments in the price equations, the two moments in the Intel share equation, and the mean ownership quality relative to the frontier quality. The auxiliary pricing equations respond sharply to changes in any of these four parameters.[48] The market share equation is primarily sensitive to γ and $\xi_{Intel} - \xi_{AMD}$. The mean $(\bar{q}_t - \hat{\Delta}_t)$ decreases if consumers upgrade more quickly, and is therefore like an outside share equation that identifies the levels of ξ. As such, we interpret ξ_{Intel} as a switching cost and $\xi_{Intel} - \xi_{AMD}$ as a brand effect.

The supply-side parameters $(a_{0,Intel}, a_{0,AMD}, a_1)$, which govern the investment process, are primarily identified by observed innovation rates, quality differences, and investment levels. The investment efficiencies are chosen such that the observed investment levels (per unit revenue) yield innovation at the observed rates. The spillover parameter a_1 is chosen to match the mean absolute difference and the mean difference in quality across firms—a high spillover keeps the qualities

[46] R&D and revenue data correspond to firm-wide activity. In the absence of R&D expenditures for different aspects of their businesses, we assume Intel and AMD invest in their business units proportional to the revenue generated by each unit. For both firms microprocessors comprise the bulk of revenues. According to Intel's 2003 annual report, the "Intel Architecture Business" (i.e., its microprocessor unit) delivered 87 percent of its consolidated net revenue.

[47] We considered using frontier quality and price instead of averages. Consumers' choices (i.e., market shares) then appear to be insensitive to quality and price since the frontier values fluctuate more than the qualities and prices of the mid-level products most consumers buy.

[48] The intercepts in the auxiliary pricing equations are sensitive to the marginal costs, so we could, in fact, estimate marginal costs instead of using the observed costs.

similar.[49]

To check our intuition for identification, we simulate 50 data sets using our model and parameter estimates as the true data generating process and then seek to recover the parameter values. Our estimator converged in each of the runs, using starting values randomly chosen within 30 percent of the true values. Table 1 reports the results of this exercise. The mean of the 50 estimates is close to the true parameter vector: all the means are within 2.66 standard errors of the truth, with five of the seven parameters being within 1.25 standard errors.

The ability of our estimator to recover consumer preferences and firms' innovation parameters is important for our empirical strategy of identifying the effect of competition on innovation. We do not observe variation in the number of firms. Hence, our conclusions regarding the effect of competition on innovation rely on estimating the costs and benefits of innovation, as determined by the structural supply and demand side parameters.[50]

4.2.3 Estimates and Model Fit

We estimate $\theta = (\gamma, \alpha, \xi_{Intel}, \xi_{AMD}, a_{0,Intel}, a_{0,AMD}, a_1, mc_0, mc_{slope})$. In a first-stage we use average unit cost data to estimate marginal costs specified by equation 5. We then use the simulated minimum distance estimator in Section 4.2 to estimate the remaining parameters.

We first fix a few parameters. We set the log-quality grid size δ to .1823, which corresponds to 20 percent improvements between grid points. We set $\bar{\delta}_c$ to 2.552, which corresponds to a maximum of fifteen δ-steps between consumers' $\tilde{q}$ and the frontier.[51] We choose $\bar{\delta}_f$ to be six δ steps so that the leader may be nearly 200 percent higher quality than the laggard, which far exceeds the observed maximum quality difference. We assume each period is three months and set β to .975.[52] We set the market size M to 400 million consumers.[53]

We report the model's fit in Table 2 and the parameter estimates in Table 3. Comparing the "Observed" column with the "Simulated" column in Table 2 reveals that the model fits the fifteen moments reasonably well. In the third column 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 four moments we have difficulty fitting relate to the pricing equations. In particular, we under-predict

[49] In principle, one could estimate these parameters in a first stage. Our attempts to do so indicated that more data are needed for this approach due to the highly stochastic nature of innovation.

[50] One could consider variation in firms' relative qualities as a form of market structure variation and investigate its relationship with innovation. In our data such relationships are insignificant. More data could potentially reveal a significant relationship, but our sense is that such variation in competition is "too local" to induce significant responses in investment behavior.

[51] Our choice of δ and $\bar{\delta}_c$ reflects the following considerations: i) the ability to replicate "Moore's Law" when firms innovate in 40 to 60 percent of the periods (so that $q_{Intel} - q_{AMD}$ varies), ii) a sufficiently high $\bar{\delta}_c$ that consumers rarely reach the lowest grid point before upgrading, and iii) a computationally manageable number of grid points.

[52] Our quantity data are quarterly, and firms' pricing and product releases are roughly quarterly.

[53] Determining the appropriate market size is difficult because the CPU market is global and a significant share of demand comes from corporations. Our choice for M appears sensible in that the model's implied market capitalizations for Intel and AMD are similar to the observed values.

the sensitivity of prices to the difference in qualities and to the difference between each firm's quality and the average vintage. Increasing the quality coefficient γ increases this sensitivity, but worsens the fit of the coefficient on $\omega_{Intel} - \omega_{AMD}$ in the share equation.[54] Although the simulated pricing constant for Intel is significantly above the actual constant, the mean predicted price for Intel is roughly the mean observed price.[55]

The model slightly overpredicts the innovation rates, but both firms' rates are within one standard error of their observed values. The model slightly underpredicts the average quality difference of 1.05 δ steps between the leader and laggard and slightly overpredicts the average absolute difference of 1.27 δ steps, though both predictions are within 1.5 standard errors of the observed values. R&D per unit revenue is also predicted accurately for each firm.

Table 3 provides the structural estimates and their standard errors. All the parameters are statistically significant given the relatively small asymptotic standard errors.

4.3 Empirical Results

We now use these parameter values as the basis for the seven industry scenarios in Table 4: 1) Intel-AMD duopoly, 2) symmetric duopoly, 3) monopoly, 4) symmetric duopoly with no spillovers, 5) myopic pricing duopoly, 6) myopic pricing monopoly, and 7) social planner. Scenario 1 is the baseline model using the estimates in column 1 of Table 3. Scenario 2 modifies the model by using Intel's firm-specific values for both firms since AMD's low ξ hampers its ability to compete. Scenario 3 uses Intel's parameters for the monopolist. Scenario 4 illustrates the effect of innovation spillovers. Scenarios 5 and 6 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 firms are behaving myopically with respect to the effect of current prices on future demand.[56] Finally, scenario 7 considers the social planner who maximizes the sum of discounted profits and discounted consumer surplus. The planner sets prices and investment for two products, but the outcome is nearly identical to the case of one product since the planner essentially invests in only the lead product.[57]

For each scenario we solve for optimal policies and simulate 10,000 industries each for 300 periods,

[54] In Section 4.2.1, we explain that average prices for a multiproduct firm fluctuate more than the price of a single-product firm. When Intel releases a new CPU, its price is very high, which moves the (unweighted) average price by more than a single-product firm would change its price after an improvement of the same magnitude. The multiproduct firm initially charges a high price for the new product as it first sells to the few consumers willing to pay top dollar for the new frontier. Adding consumer heterogeneity and multiple products to capture such price discrimination would be a natural, albeit difficult, extension to our model.

[55] Since 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 69.85 with eight 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.

[56] After price is myopically chosen to maximize current profits, investment is optimally chosen taking into account the dynamic trade-offs.

[57] 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.

starting from the state of the first quarter of our data. 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 interest. Finally, we consider various counterfactual experiments to further illustrate the properties of this industry and its implications for policy analysis.

Much of the information we provide in this section is merely designed to instill confidence that the model yields sensible outcomes. Particular findings 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 scenarios 1 and 3 (the baseline duopoly and monopoly settings). The size 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 for the ownership distribution—representing low and high demand states—and plot each firm's value function, pricing function, investment function, innovation rate, and equilibrium market shares. These plots appear in Figure 4 below their respective plot of the low-demand and high-demand ownership distribution. 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.[58]

As expected, the market shares and prices of both firms are substantially higher in column two than column one, as are the value functions (particularly Intel's), since the ownership distribution in column two corresponds to higher demand. The monopolist's value, price, and market share are also higher in the high demand state.[59]

The difference in investment behavior across these two demand 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 slightly 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.

[58] The monopolist only has one number for each plot since quality difference (i.e., the x-axis) is undefined with only one firm.

[59] The kink where qualities are tied in the price and market share curves, which are more pronounced in the low-demand state, reflect the fact that the ownership Δ is defined relative to the frontier. As such, when Intel's advantage over AMD increases from an x-axis value of -1 to -2, Intel's quality relative to the ownership Δ is unchanged, so the increase in price is small. However, when Intel's disadvantage relative to AMD increases from an x-axis value of 1 to 2, Intel also experiences a decline in its quality relative to Δ, so its price decline is relatively large.

As quality leaders both Intel and AMD invest (and hence innovate) more as their quality advantage increases, except for a slight dip when they become maximally ahead.[60] The non-monotonic innovation outcomes across the full range of relative qualities are due to the spillover aspect of innovation. Recall that a firm's innovation efficiency increases when it becomes a laggard and continues to increase as its relative quality falls. Finally, we note that although innovation is lowest for each firm when qualities are the same, the probability of the frontier increasing peaks at this point since the frontier advances when *either* firm is successful.

We acknowledge that the degree of variation in investment and innovation across demand states and across quality configurations in the policy function plots is unrealistic. We face a significant tradeoff when choosing the period length of our Markov model: pricing decisions are best modeled using a short time horizon whereas investment decisions typically reflect long planning horizons. For this study, accurately measuring consumer preferences is more important than having the correct model of innovation since the benefits to investment (as measured by market power) vary significantly across market structures, whereas the innovation process is arguably the same. Since consumer preferences are more accurately estimated using a relatively short period length, we maintain that firms act on a quarterly basis.[61]

Finally, Table 5 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 $223 billion for Intel and $20 billion for AMD are broadly consistent with market valuations for Intel and AMD over the period of our data.

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 5(a) the choice probabilities for each ownership vintage, averaged across states encountered in the duopoly simulations of scenario 1 (the baseline). As expected, the lower a consumer's current vintage relative to the frontier, the more likely she is to upgrade. Consumers with vintages within two δ-steps (i.e., 44%) of the frontier upgrade less than 10 percent of the time, compared to one-third of the time for owners of the lowest quality vintage.[62]

As consumers implement their policy functions, they generate a sequence of ownership distributions across time. Figure 5(b) depicts the average ownership distribution for the baseline duopoly and

[60] When a firm is maximally ahead, the incentive to innovate is reduced. These states are rarely encountered in the simulations, and all our results are robust to further relaxing the bound on quality differences.

[61] Two modifications to our model are possibly worth considering to smooth out fluctuations: adding convex adjustment costs to investment, and having firms choose a fixed investment level prior to playing a dynamic pricing game. Unfortunately, both approaches are computationally difficult to implement. The former approach requires adding previous investment as a state variable for each firm, and the latter approach nests the dynamic pricing game into the payoff of an investment level game.

[62] Restricting consumers to only upgrade to products with strictly higher quality is trivial to implement. The firms' policy functions are barely affected by such a restriction, except when the quality coefficient is extremely low (which we encounter in comparative statics reported below). When the quality coefficient is low, firms reduce investment since generating upgrade sales beyond those generated by the logit error becomes too costly. If strict upgrades are enforced, then firms innovate regardless of the quality coefficient since innovation is the only way to generate sales.

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 forces cause the monopolist's ownership distribution to have more mass on the older vintages compared to the duopolists' ownership distribution.

Importantly, Figure 5(b) 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 5(a) by the ownership distribution in Figure 5(b) yields the portion of purchasers from each ownership vintage.[63] In the duopoly most purchasers upgrade from products three to seven δ steps below the frontier, which implies the frontier is 73 percent to 258 percent faster than their current product. Of course, some of the consumers upgrade to the non-frontier firm's offering. As we report in Table 4, the average upgrade improvement is 152 percent. In the monopoly the higher prices induce consumers to upgrade only when the percent improvement over their current product is quite high. Most upgrades in the monopoly are to products that are five to eight δ steps below the frontier, which corresponds to improvements of 149 and 330 percent, respectively. The mean upgrade in the monopoly is 244 percent.

Observations of Interest: Having established the sensibility of consumers' and firms' policy functions, we now discuss features of equilibrium behavior and outcomes of particular interest for the microprocessor industry. Here we evaluate the model using the parameter estimates, whereas in the next section we explore the model's properties at different parameter values.

Observation 1. *Regarding the effect of competition on innovation, we find*

i. Innovation is higher with a monopoly than with the Intel-AMD duopoly. The difference is more pronounced when comparing monopoly to a symmetric duopoly pitting Intel against another Intel.

ii. Equilibrium investments for monopoly and duopoly market structures are below the socially optimal levels chosen by the planner.

The industry 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, 1234, 1571, and 5644. The resulting innovation rates for the industry's frontier product in each market structure are 0.643 for the duopoly, 0.684 for the monopolist, and 0.885 for the planner. The symmetric duopoly's innovation rate is only 0.558.

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

[63]The resulting graph is similar to a scaled down version of Figure 5(b).

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.

The absence of technology spillovers in the monopoly is a potential factor in the monopolist's higher innovation compared to a duopoly in which firms mimic, to some degree, each others innovations. However, as we report in Table 4, the innovation rate in the symmetric duopoly with no spillovers (scenario 5) is actually lower than the innovation rate in the symmetric duopoly with spillovers (scenario 3). The direct effect of removing the spillover is to increase the incentive to innovate since it cannot be copied. The direct effect, however, is dominated by the equilibrium effect of one firm dominating the industry, as evidenced by the high average quality difference. The absence of a threat from the weak laggard induces the leader to reduce investment. Note, however, that the presence of the laggard nonetheless keeps margins much lower than in the monopoly case. Hence, high margins, rather than an absence of spillovers, provide the incentive for rapid innovation by the monopolist.

Importantly, our model yields higher innovation with competition when evaluated using different parameter values. As such, our model indeed lets the data speak on this fundamental question. In the next section we present comparative statics that reveal when competition fosters higher innovation.

Of course, policymakers are more concerned with surplus and profits than with innovation alone. Firms' profits are calculated in the usual manner as the discounted sum of per-period profits. See Appendix E for details on computing consumer surplus.[64]

Observation 2. *Regarding the effect of competition on surplus, we find*

i. The AMD-Intel duopoly attains 84.09 percent of the planner's social surplus, whereas the monopoly attains 83.98 percent—a negligible difference.

ii. Consumers' share of social surplus is 89.2 percent in the AMD-Intel duopoly, compared to 87.1 percent in the monopoly.

iii. Consumer surplus is $201 billion per year in the AMD-Intel duopoly, compared to $196 billion per year in the monopoly—a difference of 2.5 percent.

iv. The benefits of competition are higher if Intel were facing a symmetric competitor: consumer surplus is 4.0 percent higher and social surplus is 2.5 percent higher in a symmetric duopoly than in a monopoly.

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

[64]We compute consumer surplus (CS) 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. Although such a modification could easily be accommodated, its specification would be ad hoc. We also acknowledge that measuring consumer surplus for a product that has transformed our world on so many levels is an almost futile effort. As such, we focus on differences in surplus across scenarios rather than levels.

Table 4 reports the aggregate discounted CS and industry profits for each of the scenarios we consider. The duopoly value for CS is \$2.01 trillion, which corresponds to \$201 billion per year (using an annual discount factor of .9). This value is 2.5 percent higher than the CS generated by the monopoly. In terms of social surplus, the monopolist is able to make up for nearly all of this reduced CS with higher industry profits.

Given that prices are 44 percent higher in the monopoly, one might be surprised that consumer surplus is only 2.5 percent lower and social surplus is unchanged. This finding reflects the compounding benefits of the monopolist's faster innovation. A quality improvement increases consumer surplus in all future periods in addition to the current period, whereas a higher price only decreases current surplus. Hence, a small increase in the innovation rate can offset a large increase in prices.

Next we report demand elasticities.

Observation 3. *Using unexpected 10 percent increases in prices for one period, we find:*

i. Own-price elasticities are 1.94 for Intel and 1.47 for AMD.[65]

ii. Cross-price elasticities are .04 for Intel (when p_{AMD} increases) and .29 for AMD.

iii. The own-price elasticity is 2.58 for Intel when AMD is absent.

iv. Some of the lost sales due to the price hike are recouped in subsequent periods: the average percent of lost sales recouped in the four subsequent periods are 13.80 for Intel, .01 for AMD, and 12.2 percent for Intel as a monopolist. Over ten periods the firms recoup, respectively, 21.0, .01, and 20.8 percent.

v. Markups (in both duopoly and monopoly) are positively related to elasticity, in direct contrast with the "inverse elasticity rule" of static pricing, except for firms with quality vastly inferior to the frontier.

Only the last of these observations requires explanation. The elasticities are averages from 10,000 simulations in which the unexpected, single-period price shock occurred in period 101 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 6 we plot the elasticity and optimal markup at each of the states encountered in these simulations. For both duopoly firms when they are leaders, and for Intel as a monopolist, the markup is increasing in elasticity. By contrast, the price shock exercise for the myopically-pricing monopolist yields the familiar inverse relationship of static pricing.

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

The series of lines in each of the duopoly plots correspond to the lowering of marginal costs as the firm falls further behind the leader. As the laggard falls further behind, the effect of its current

[65] These elasticities are somewhat lower than the 2.32 to 2.74 range reported in Prince (2008) for personal computer purchases.

sales on future profitability diminishes and the familiar inverse relationship between elasticity and markup is restored.

This breakdown of the inverse-elasticity rule, however, does not imply that as consumers become more price sensitive firms raise prices. Indeed, in the comparative statics of the next section, we see that as price sensitivity increases, due to either a higher α or lower variance of ϵ, prices rise. The breakdown reflects the fact that dynamic considerations lead firms to price higher and logit demand systems have elasticities that increase in price.

Recently AMD 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 that has exclusive access to a portion of the market is restricted to offer the same price in both sub-markets.

Observation 4. *Excluding AMD from an increasing portion of the market yields:*

i. prices and innovation steadily rising, and

ii. consumer and social surplus initially higher, but eventually lower than when AMD competes for all consumers.

In Figure 7 we plot the price, 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 (share-weighted) price and innovation both rise steadily as AMD is increasingly barred from the market. Consumer surplus is actually higher when AMD is barred from a portion of the market (peaking when the restriction is 50 percent). Although the surplus gains are small, this finding highlights the importance of accounting for innovation when considering antitrust policy. The decrease in current-period consumer surplus from higher prices can be more than offset by higher innovation leading to higher future surplus. Unfortunately, identifying when this dynamic trade-off calls for leniency is not a simple task. We hope the model and methods we present here provide stepping stones for such considerations.

5 Comparative Statics and Theoretical Findings

We now turn to results based on our model that go beyond our particular empirical application. In addition to being of interest themselves, these observations also serve to assess the robustness of our empirical findings.

Observation 5. *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 54 percent higher and margins are 128 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 18 percent higher and margins are 22 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 5 suggests that marginal cost estimates computed in this manner for durable goods will be too high. That is, prices are high because the firm does not want to reduce future demand, not because its marginal costs are high.

Observation 6. *As Figure 8 depicts, consumer surplus, margins, innovation, and profits all increase monotonically in the rate of time preference upon which consumers base decisions.*[66]

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 Nair (2007), who finds optimal prices for a durable goods monopolist decreasing in consumers' rate of time preference. The difference in our findings is due to the fact that we increase the consumer's discounted utility flow from the durable good when β increases, whereas Nair (2007) fixes the discounted utility flow at the estimated values. In essence, his counterfactual corresponds to shrinking the length of the time period whereas our counterfactual corresponds to increasing the rate of time preference.[67]

The primary empirical finding of this paper is that Intel would innovate more if it were not competing against AMD. We now illustrate that the relationship between competition and innovation hinges on consumer preferences, which is consistent with Dorfman and Steiner (1954) and Lee (2005) who find that price and quality preferences primarily determine R&D intensity.

Observation 7. *Comparative statics in γ and α (the quality and price coefficients) over the ranges of .2 to .6 for γ and .004 to .01575 for α reveal*

i. Innovation is increasing in γ and decreasing in α for both monopoly and duopoly (symmetric and asymmetric).

[66] Consumer 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.

[67] The discount factor equals $e^{-TimePreference \times PeriodLength}$, which approaches 1 as either the rate of time preference approaches infinity or the period length shrinks to 0. Coase (1972) conjectures that a durable-good monopolist's profit goes to zero as period length shrinks to 0. The conjecture is often misleadingly stated as pertaining to consumers' discount factors going to 1, which implies profits go to zero as the "time rate of preference" goes to infinity. To change the period length in our model, one must change the scale of all per-period utility flows (i.e., the quality coefficient and the logit error variance). The innovation process, however, cannot be modified in a simple manner to preserve *ceteris paribus* as needed for a comparative static.

ii. The effect of competition on innovation (measured by the symmetric duopoly's innovation minus the monopoly's innovation) is increasing in γ and decreasing in α, except where both γ is low and α is high, as Figure 10 depicts.

iii. Innovation is higher for the duopoly than the monopoly when γ is high and α is low.

The finding that innovation is increasing as consumers care more about quality (γ higher) is not surprising. The part of (i) regarding price sensitivity is also intuitive: innovation declines as α increases since extracting surplus from innovations becomes more difficult. Figure 9 depicts the result in (i) for the case of symmetric duopoly.

Regarding the effect of competition on innovation, we find it intuitive that the slope of the innovation surface over the (γ, α) plane is steeper for duopoly than monopoly. This higher slope leads to the duopoly eventually becoming more innovative than the monopoly as γ increases and α falls.

The degree of horizontal differentiation (i.e., local market power or substitutability) varies across industries. To assess the generality of our result that durable goods monopolists innovate more than duopolists, we therefore consider the comparative static in which the standard deviation of consumers' idiosyncratic utility (ε) varies over a wide range (from .048 to 10). The pure vertical model corresponds to a standard deviation of zero and the standard logit model has a standard deviation of $\pi/\sqrt{6}$, roughly 1.28.[68]

Observation 8. *Varying the degree of local market power by changing var(ε), we find*

i. As the variance of ε increases, starting from the standard logit value of 1.28, innovation in the monopoly decreases, whereas innovation in the duopoly increases.

ii. As the variance of ε decreases, starting from the standard logit value of 1.28, innovation in the monopoly decreases and innovation in the duopoly first increases until the laggard gives up, at which point innovation declines.

iii. Monopoly has more innovation than duopoly, except for very low values of var(ε).

iv. Prices and consumer surplus decline in both the monopoly and duopoly as the variance of ε declines, except when the duopolists rapidly increase innovation as the industry nears the case in which only one firm effectively survives.

Figure 6 presents four graphs depicting the findings of this comparative static. The upper-right figure illustrates that as the variance of ε approaches zero, the quality difference between the firms shrinks until this variance crosses a threshold around .48, at which point the equilibrium is characterized by one of the firms essentially giving up, yielding quality differences near the maximum allowed. Relaxing the maximal difference between qualities would lead the duopoly to offer the same outcomes as the monopoly since a duopoly with only one effective firm is essentially a monopoly.

[68] Decreasing the error's standard deviation by a given percent is identical to simultaneously increasing the demand parameters (γ, α, ξ) by the same percent.

Hence, we conclude that for all levels of local market power, as captured by var(ε), the monopoly has weakly more innovation than the duopoly. The case of high preferences for quality and low price sensitivity, from the previous observation, remains the only scenario in which we find competition leads to more innovation.

The intuition behind finding (i) is that more variance in ε reduces the monopolist's competition with its future self, and so innovation is less important as a means of generating upgrade purchases. For duopolists the additional variance provides an increase in local market power, which enables them to extract more of the gains from offering higher qualities.

Theoretical papers often consider both firm count and substitutability as candidate measures of competitiveness. We focus on firm count since this measure may be influenced by regulatory policy. Nonetheless, we can use this comparative static to relate to others' findings. The first plot of Figure 6 reveals a U-shaped relation between substitutability and innovation by duopolists.[69] We find the U-shape to be intuitive: when substitutability is very high firms innovate to ensure their survival, and when substitutability is sufficiently low firms increase innovation in response to the substantial gains in local market power. In a nondurable setting in which firms invest to lower costs, Vives (2008) finds increasing substitutability leads to more R&D for linear and constant elasticity demand systems but no change for logit demand.[70] The differences in our findings from Vives (2008) is consistent with the overall sensitivity of results to model assumptions throughout the literature on competition and innovation. This sensitivity of theoretical findings suggests a need for empirical case-studies to identify the relationship between competition and innovation in particular industries of interest.

6 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 long-run innovation rate for the industry is endogenous instead of 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 extend the EP framework to accommodate durable goods by developing a simple approximation for the ownership distribution, which is an endogenous state variable that summarizes the state of demand.

We estimate the model using data from the microprocessor 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. We compute profits and consumer surplus under alternative market structures and find that consumer surplus is 2.5 percent higher with AMD duopoly

[69] We ignore the left-most portion of the curve over which only one firm is essentially active.

[70] We report Vives' results for the case of price competition and restricted entry since this case is closest to ours. He also considers quantity competition and entry.

than without AMD competing against Intel, though social surplus is unchanged. In support of Schumpeter's hypothesis, the monopolist innovates more than the duopoly, as its market power enables it to better extract the potential gains to trade resulting from innovations.

Extending our model to allow for multiproduct firms would enable researchers to study whether competition and consumer surplus are enhanced or diminished when leading firms offer inferior products to compete more directly with the top offerings of lagging firms.[71]

Another extension could compare our results for durable goods with those for nondurable goods to shed light on the particular role of durability. Durable-goods monopolists do face competition, albeit from themselves. As such, our finding that innovation declines when moving from one to two firms does not rule-out an inverted U-shaped relationship between innovation and the number of firms selling nondurables.

Appendix A: Optimal Prices and Investments

As discussed at the end of Section 2.2, 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 (12). 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

$$x^*_{jt}(p_{jt}) = \frac{1}{a_j}\left(\frac{c}{\beta a_j(EW^+(p_{jt}) - EW^-(p_{jt}))}\right)^{-1/2} - 1\ . \tag{19}$$

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

$$\begin{aligned} \frac{\partial \pi_j(p_t,q_t,\Delta_t)}{\partial p_{jt}} + \beta \sum_{\tau_{jt},\ q_{-j,t+1},\ \Delta_{t+1}} & W_j(q_{jt} + \tau_{jt}, q_{-j,t+1}, \Delta_{t+1}) \\ h_f(q_{-j,t+1}|q_t, \Delta_t)\ \frac{\partial g_f(\Delta_{t+1}|\Delta_t,q_t,p_t,q_{t+1})}{\partial p_{jt}}\ & f(\tau_{jt}|x^*_{jt}(p_{jt})) \quad = \quad 0\ , \end{aligned} \tag{20}$$

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

[71] As a partial investigation of the effect of moving to a multiproduct setting, we compare profit gains from quality improvements in the multiproduct and single-product static logit model. The model has two consumer types (preferences for quality) and is calibrated to match shares and prices of Intel and AMD. Moving from single to multiple products per firm increases the incentive to innovate (i.e., profit gain) for both monopolists and duopolists. However, the effect is larger for the monopolists, which suggests our finding that competition reduces innovation in the microprocessor industry will be robust to this extension.

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. Although we not proven that the optimal price is uniquely determined, inspection of the firms' best-response functions at many states suggests prices are uniquely determined. The pair $(p^*_{jt}, x^*_{jt}(p^*_{jt}))$ is the optimal set of controls at this state.

Appendix B: Solving and Simulating Industry Equilibrium

We compute the equilibrium using a Gauss-Jacobi scheme to update the value and policy functions.[72] Starting at iteration $k = 0$, we initialize the consumer value function $\hat{V}^0$ and the firms' value functions W^0_j each to zero, and policy functions (x^{*0}_j, p^{*0}_j) to yield innovation rates of 50 percent and prices 20 percent higher than marginal costs.[73]

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

1. For each $\tilde{\omega} \in \Delta$, evaluate the consumer's value function $\hat{V}^k$ given the firms' policy functions $\left\{x^{*k-1}_j, p^{*k-1}_j\right\}_{j=1}^{J}$ and the next period approximation parameters $\rho'^{,k}$ from the previous iteration.
2. For each $j \in J$, evaluate firm j's value function W^k_j given the other firms' policy functions from the previous iteration $\left\{x^{*k-1}_{j'}, p^{*k-1}_{j'}\right\}_{j' \neq j}^{J}$.
3. Update the consumer value functions $\hat{V}^{k+1} \leftarrow \hat{V}^k$, the firms' value functions $W^{k+1}_j \leftarrow W^k_j, \forall j$, and their policy functions $\left\{x^{*k+1}_j, p^{*k+1}_j\right\} \leftarrow \left\{x^{*k}_j, p^{*k}_j\right\}$.
4. Check for convergence in the sup norm of all agents' value functions with a tolerance of 1e-10. 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 $t = 0, \ldots, T$, 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 (17), and process the stochastic innovation outcomes according to $f(\cdot|\cdot)$.

Appendix C: Transforming to a Relative State-Space

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

$$V^T(q_1, \Delta_1, \tilde{q}_1, \varepsilon_1) = \max_{\{y_t(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) \in (0,1,\ldots,J)\}_{t=1}^{T}} \mathrm{E}\left[\sum_{t=1}^{T} \beta^t \left(\gamma q_{y_t,t} - \alpha p_{y_t,t} + \xi_{y_t} + \varepsilon_{y_t,t}\right)\right], \tag{21}$$

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

$$W^T_j(q_{j1}, q_{-j,1}, \Delta_1) = \max_{\{p_{jt}(q_t, \Delta_t),\ x_t(q_t, \Delta_t)\}_{t=1}^{T}} \mathrm{E}\left[\sum_{t=1}^{T} \beta^t \left(M s_{jt}(p_t, q_t, \Delta_t)(p_{jt} - mc_j) - c x_{jt}\right)\right], \tag{22}$$

[72] For robustness we also compute the equilibrium using Gauss-Seidel with random orderings of the states.

[73] The consumer's smoothed value function (which integrates over the idiosyncratic ε) is $\hat{V}(\omega, \Delta, \tilde{\omega}) = \log \sum_{j=0,\ldots,J} \hat{V}_j(\omega, \Delta, \tilde{\omega})$, using the product-specific $\hat{V}_j$ functions of equation (14).

where M is market size, $s_{jt}(\cdot)$ is the market share for firm j as defined in equation (9), 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 therefore satisfy

$$\begin{array}{rrcl} \text{Firms:} & W_j^T(q_{jt}, q_{-j,t}, \Delta_t) & = & W_j^T(q_{jt} - \hat{q}, q_{-j,t} - \hat{q}, \Delta_t) \\ \text{Consumers:} & V^T(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) & = & \gamma\hat{q} + V^T(q_t - \hat{q}, \Delta_t, \tilde{q}_t - \hat{q}, \varepsilon_t) \ . \end{array} \tag{23}$$

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

$$\begin{aligned} V^{T-1}(q_{T-1}, \Delta_{T-1}, \tilde{q}_{T-1}, \varepsilon_{T-1}) \quad = \quad & \max_{y \in (0,1,\ldots,J)} \gamma q_{y,T-1} - \alpha p_{y,T-1} + \xi_y + \varepsilon_{y,T-1} \\ & + \beta \sum_{q_T} \int_{\epsilon_T} V^T(q_T, \Delta_T, \tilde{q}_T, \varepsilon_T) \, dF_\varepsilon(\varepsilon_T) \\ & \prod_{j=1}^{J} f(q_{jT} - q_{j,T-1} | x_{j,T-1}, q_{T-1}) \ , \end{aligned} \tag{24}$$

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 consumer's $\tilde{q}$, and the deterministic transition to Δ_T is based on consumers' choices, as detailed in equation (10). Since each consumer is small relative to M, her actions do not affect the transition of Δ.

Each firm j solves

$$\begin{aligned} W_j^{T-1}(q_{j,T-1}, q_{-j,T-1}, \Delta_{T-1}) \quad = \quad & \max_{\substack{p_{j,T-1} \\ x_{j,T-1}}} M s_{j,T-1}(p_{T-1}, q_{T-1}, \Delta_{T-1})(p_{j,T-1} - mc_j) - c x_{j,T-1} \\ & + \beta \sum_{q_T} W_j^T(q_{j,T}, q_{-j,T}, \Delta_T) \prod_{j=1}^{J} f(q_{jT} - q_{j,T-1} | x_{j,T-1}, q_{T-1}) \ . \end{aligned} \tag{25}$$

In these equations defining V^{T-1} and W^{T-1}, the products' future qualities are uncertain. Rational expectations regarding this uncertainty are achieved by using the firm's investments in period $T-1$ to determine the distribution of q_T.[74]

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

$$\begin{aligned} V^{T-1}(q_{T-1} - \hat{q}, \Delta_{T-1}, \tilde{q}_{T-1} - \hat{q}, \varepsilon_{T-1}) \quad = \quad & \max_{y \in (0,1,\ldots,J)} \gamma(q_{y,T-1} - \hat{q}) - \alpha p_{y,T-1} + \xi_y + \varepsilon_{y,T-1} \\ & + \beta \sum_{q_T} \int_{\epsilon_T} V^T(q_T - \hat{q}, \Delta_T, \tilde{q}_T - \hat{q}, \varepsilon_T) \, dF_\varepsilon(\varepsilon_T) \\ & \prod_{j=1}^{J} f(q_{jT} - \hat{q} - (q_{j,T-1} - \hat{q}) | x_{j,T-1}, q_{T-1} - \hat{q}) \end{aligned} \tag{26}$$

and

$$\begin{aligned} W_j^{T-1}(q_{j,T-1} - \hat{q}, q_{-j,T-1} - \hat{q}, \Delta_{T-1}) \quad = \quad & \max_{\substack{p_{j,T-1} \\ x_{j,T-1}}} M s_{j,T-1}(p_{T-1}, q_{T-1} - \hat{q}, \Delta_{T-1})(p_{j,T-1} - mc_j) \\ & - c x_{j,T-1} + \beta \sum_{q_T} W_j^T(q_{j,T} - \hat{q}, q_{-j,T} - \hat{q}, \Delta_T) \\ & \prod_{j=1}^{J} f(q_{jT} - \hat{q} - (q_{j,T-1} - \hat{q}) | x_{j,T-1}, q_{T-1} - \hat{q}) \ . \end{aligned} \tag{27}$$

[74] Recall 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.

Substitute the right-hand sides of (23) into (27) and (26). 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 (27) into (25), exactly, and converts (26) into (24), except for a $-(\gamma\hat{q}+\beta\gamma\hat{q})$ term that does not affect the consumer's choice. The modified (26) is

$$\begin{aligned} V^{T-1}(q_{T-1}-\hat{q},\Delta_{T-1},\tilde{q}_{T-1-\hat{q}},\varepsilon_{T-1}) \quad = \quad & \max_{y\in(0,1,\ldots,J)} \gamma(q_{y,T-1}-\hat{q}) - \alpha p_{y,T-1} + \xi_y + \varepsilon_{y,T-1} \\ & +\beta \sum_{q_T} \int_{\epsilon_T} \left(-\gamma\hat{q} + V^T(q_T,\Delta_T,\tilde{q}_T,\varepsilon_T)\right) \, dF_\varepsilon(\varepsilon_T) \\ & \prod_{j=1}^{J} f(q_{jT}-q_{j,T-1}|x_{j,T-1},q_{T-1}) \; . \end{aligned} \tag{28}$$

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.

□

Appendix D: A Simulated Minimum Distance Estimator

Our presentation of the assumptions and details of our estimator 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. The transition density, f_μ, for this Markov process is given by

$$\begin{aligned} f_\mu(\omega_{t+1},\Delta_{t+1},p_{t+1},x_{t+1},s_{t+1}|\omega_t,\Delta_t,p_t,x_t,s_t,\theta) \quad = \quad & \prod_{j=1}^{J} f(\omega_{j,t+1}-\omega_{j,t}|\omega_t,x_t) \\ \times \quad & g(\Delta_{t+1}|\Delta_t,s_t) \\ \times \quad & I\{p_{t+1}=p(\omega_{t+1},\Delta_{t+1})\} \\ \times \quad & I\{x_{t+1}=x(\omega_{t+1},\Delta_{t+1})\} \\ \times \quad & I\{s_{t+1}=s(\omega_{t+1},\Delta_{t+1})\} \; , \end{aligned} \tag{29}$$

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}.

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 $\left\{\{U_t^n\}_{t=1}^T\right\}_{n=1}^S$, is held fixed throughout the estimation procedure to preserve continuity of the estimator's objective function. The set of simulated industry outcomes is denoted $\left\{\{\mu_t(\theta,U_{<t}^n,\omega_0,\Delta_0)\}_{t=1}^T\right\}_{n=1}^S$, where the subscript in $U_{<t}^n$ indicates that μ_t^n depends on only the first $t-1$ realizations of U^n.

The 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:

$$m_{S,T}(\theta) = \frac{1}{S}\sum_{n=1}^{S} m\left(\{\mu_t(\theta, U^n_{<t}, \omega_0, \Delta_0)\}_{t=1}^{T}\right), \tag{30}$$

where the initial state (ω_0, Δ_0) corresponds to the first quarter of our data.

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

$$\hat{\theta}_T = \underset{\theta\in\Theta}{\operatorname{argmin}}\ (m_{S,T}(\theta) - m_T)' A_T (m_{S,T}(\theta) - m_T)\ , \tag{31}$$

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

We make the following assumptions.

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

$$\Psi(\mu'|\theta) = \int f_\mu(\mu'|\mu, \theta) d\Psi(\mu|\theta). \tag{32}$$

Assumption 2. *The structural model presented in Section 2 is correctly specified. As such, a $\theta^* \in \Theta$ exists for which each simulated sequence $\{\mu^n_t\}, n = 1, \ldots, S$ 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 $\mathrm{E}[m|\theta]$, $\nabla\mathrm{E}[m|\theta]$, and $\nabla m_{S,T}$ as

$$\begin{array}{lcl} \mathrm{E}[m|\theta] & = & \int m(\mu) d\Psi(\mu|\theta) \\ \nabla\mathrm{E}[m|\theta] & = & \frac{\partial}{\partial\theta}\mathrm{E}[m|\theta] \\ \nabla m_{S,T} & = & \frac{\partial}{\partial\theta} m_{S,T}(\theta). \end{array} \tag{33}$$

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

The optimal weight matrix is $\Omega(m, \theta^*)^{-1} \equiv \mathrm{E}[(m(\mu) - \mathrm{E}[m(\mu)])(m(\mu) - \mathrm{E}[m(\mu)])']^{-1}$, 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 $A_T = [cov(\{\mu^{actual}_t\}_{t=1}^T)]^{-1}$ as a consistent estimate of the optimal weight matrix, the estimator $\hat{\theta}_T$ has the property

$$\sqrt{T}(\hat{\theta}_T - \theta^*) \Longrightarrow N\left(0, (1 + 1/S)\left(\nabla\mathrm{E}[m|\theta^*]'\Omega(m, \theta^*)^{-1}\nabla\mathrm{E}[m|\theta^*]\right)^{-1}\right). \tag{34}$$

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

Appendix E: Deriving Consumer Surplus

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

$$
\begin{aligned}
CS(\tilde{q}_t, q_t) &= \tfrac{1}{\alpha}\mathrm{E}\left[u_{jt}\right] \\
&= \tfrac{1}{\alpha}\sum_{j\in\{0,\ldots,J\}} s_{jt|\tilde{q}}\left(\gamma q_{jt} - \alpha p_{jt} + \xi_j - \mathrm{E}\left[\varepsilon_{jt}|\text{choose } j\right]\right) \\
&= \tfrac{1}{\alpha}\sum_{j\in\{0,\ldots,J\}} s_{jt|\tilde{q}}\left(\gamma q_{jt} - \alpha p_{jt} + \xi_j - \log s_{jt|\tilde{q}}\right),
\end{aligned}
\tag{35}
$$

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.[75]

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:[76]

$$
CS = M\sum_{t=0}^{T}\beta^t\sum_{\tilde{q}=\underline{q}_t}^{\bar{q}_t} CS(\tilde{q}, q_t)\cdot\Delta_{\tilde{q},t}. \tag{36}
$$

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] Arrow, K. (1962), "Economic Welfare and the Allocation of Resources for Invention," in R. Nelson, ed., *The Rate and Direction of Inventive Activity* Princeton, NJ, Princeton University Press, 609-625.

[3] Bajari, P., L. Benkard, and J. Levin (2007), "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, forthcoming.

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

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

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

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

[8] Chen, J., S. Esteban, and M. Shum (2008), "How Much Competition is a Secondary Market?" mimeo, UC Irvine.

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

[10] Cho, S. (2008) "The Effectiveness of Parametric Approximation: A Case of Mainframe Computer Investment," *Journal of Economic Research*, 13, 125–148.

[11] Cohen, W. and R. Levin (1989), "Empirical Studies of Innovation and Market Structure," in R. Schmalensee and R. Willig, eds., *Handbook of Industrial Organization, Volume 2*, North-Holland, Amsterdam, 1059–1107.

[12] Dasgupta, P and J. Stiglitz (1980), "Industrial Structure and the Nature of Innovative Activity," *Economic Journal*, 90, 266-293.

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

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

[13] 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.

[14] Dixit, A. and J. Stiglitz (1977), "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67, 297-308.

[15] Doraszelski, U. and A. Pakes (2007) "A Framework for Applied Dynamic Analysis in IO," in M. Armstrong and R. Porter, eds., *Handbook of Industrial Organization, Volume 3*, North-Holland, Amsterdam, 1887–1966.

[16] Doraszelski, U. and M. Satterthwaite (2007), "Computable Markov-Perfect Industry Dynamics: Existence, Purification, and Multiplicity," Harvard University, mimeo.

[17] Dorfman, R. and P. O. Steiner (1954), "Optimal Advertising and Optimal Quality," *American Economic Review*, 44, 826–836.

[18] 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.

[19] Epple, D. and H. Sieg (1999), "Estimating Equilibrium Models of Local Jurisdictions," *Journal of Political Economy*, 107(4), 645–681.

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

[21] Esteban, S. and M. Shum (2007), "Durable-Goods Oligopoly with Secondary Markets: the Case of Automobiles," *RAND Journal of Economics*, 38(2), 332–354.

[22] Fellner, W. (1951), "The Influence of Market Structure on Technological Progress," *Quarterly Journal of Economics*, 65, 556–577.

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

[24] Gallant, R. and G. Tauchen (1996), "Which Moments to Match?" *Econometric Theory*, 12, 657–681.

[25] Gilbert, R. (2005), "Looking for Mr. Schumpeter: Where are We in the Competition-Innovation Debate?" forthcoming, National Bureau of Economic Research.

[26] Goettler, R., C. Parlour, and U. Rajan (2005), "Equilibrium in a Dynamic Limit Order Market," *Journal of Finance*, 60(5), 2149–2192.

[27] Goettler, R., C. Parlour, and U. Rajan (forthcoming), "Informed Traders and Limit Order Markets," *Journal of Financial Economics*.

[28] Gordon, B. R. (forthcoming), "A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry," *Marketing Science*.

[29] Gouriéroux, C., A. Monfort, and E. Renault (1993), "Indirect Inference," *Journal of Applied Econometrics*, 8, S85–S118.

[30] Gowrisankaran, G. and M. Rysman (2007), "Dynamics of Consumer Demand for New Durable Consumer Goods," working paper, Boston University.

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

[32] Grabowski, H. G. and J. M. Vernon (1987), "Pioneers, Imitators, and Generics – A Simulation Model of Schumpeterian Competition," *Quarterly Journal of Economics*, 102, 491–525.

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

[34] Hall, G. and J. Rust (2003), "Simulated Minimum Distance Estimation of a Model of Optimal Commodity Price Speculation with Endogenously Sampled Prices," mimeo, Yale University.

[35] Hansen, L. (1982), "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1929–1954.

[36] Hendel, I. (1999), "Estimating Multiple-Discrete Choice Models," *Review of Economic Studies*, 66(2), 423–446.

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

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

[39] Lee, C. Y. (2005), "A New Perspective on Indusry R&D and Market Structure," *Journal of Industrial Economics*, 53(1), 101–122.

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

[41] Lee, R. S. (2008), "Vertical Integration and Exclusivity in Platform and Two-Sided Markets," working paper, NYU.

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

[43] Macieira, J. (2007) , "Extending the Frontier: A Structural Model of Investment and Technological Competition in the Supercomputer Industry," Virginia Tech.

[44] Markovich, S. (2008), "Snowball: A Dynamic Oligopoly Model with Network Externalities," *Journal of Economic Dynamics and Control*, 32, 909–938.

[45] Markovich, S. and Moenius, J. (2008), "Winning While Losing: Competition Dynamics in the Presence of Indirect Network Effects," Working paper, Northwestern University.

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

[47] Nair, H. (2007), "Intertemporal Price Discrimination with Forward-Looking Consumers: Application to the US Market for Console Video-Games," *Quantitative Marketing and Economics*, 5(3), 239–292.

[48] Needham, D. (1975), "Market Structure and Firms' R&D Behavior," *Journal of Industrial Economics*, 23(4), 241–255.

[49] Nelson, R.R. and S.G. Winter (1982), "The Schumpeterian tradeoff revisited," *American Economic Review*, 72, 114–132.

[50] Pakes, A. and D. Pollard (1989), "Simulation and Asymptotics of Optimization Estimators," *Econometrica*, 57, 1027–1057.

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

[52] Prince, J. (2008), "Repeat Purchase and Rapid Quality Improvement: Structural Estimation of the Demand for Personal Computers," *Journal of Economics and Management Strategy*, 17(1), 1–33.

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

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

[55] Rust, J. (1996), "Estimation of Dynamic Structural Models, Problems and Prospects: Discrete Decision Processes," in C. Sims, ed., *Advances in Econometrics: Sixth World Congress, Vol. II*, Cambridge University Press, 119–170.

[56] Scherer, F. (1967), "Market Structure and the Employment of Scientists and Engineers," *American Economic Review*, 57, 524-531.

[57] Schumpeter, J (1942), *Capitalism, Socialism, and Democracy*, New York, Harper.

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

[59] Smith, A. (1993) "Estimating Nonlinear Time Series Models using Vector Autoregressions: Two Approaches," *Journal of Applied Econometrics*, 8, 63–84.

[60] Song, M. (2006), "A Dynamic Analysis of Cooperative Research in the Semiconductor Industry," mimeo University of Rochester.

[61] 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.

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

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

[64] Vives, X. (2008), "Innovation and Competitive Pressue," *The Journal of Industrial Economics*, 56(3), 419–469.

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

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

[67] Xu, Yi (2007), "A Structural Empirical Model of R&D Investment, Firm Heterogeneity, and Industry Dynamics," Working paper, New York University.

Table 1: Monte Carlo Results

Parameter	"truth"	mean $\hat{\theta}$	std.dev. $\hat{\theta}$	std.err. $\hat{\theta}$	t-stat
Price, α	0.0139	0.014099	0.001127	1.59E-04	1.248821
Quality, γ	0.266193	0.266086	0.019364	2.74E-03	0.039028
Intel Fixed Effect, ξ_{Intel}	-0.70526	-0.69612	0.091838	1.30E-02	0.703466
AMD Fixed Effect, ξ_{AMD}	-3.00061	-3.0331	0.086395	1.22E-02	2.659806
Intel Innovation, $a_{0,Intel}$	0.001376	0.001356	0.000141	1.99E-05	0.993388
AMD Innovation, $a_{0,AMD}$	0.003538	0.003594	0.000335	4.74E-05	1.181202
Spillover, a_1	3.949222	3.74234	0.577757	8.17E-02	2.531993

The t-stat column reports ("truth" $-$ mean $\hat{\theta}$) / (std.err. $\hat{\theta}$). The "std.dev. $\hat{\theta}$" column is a parametric bootstrap for standard errors of the estimates reported in Table 3.

Table 2: Empirical vs. Simulated Moments

Auxiliary Moment	Observed	Simulated	pseudo-t
Intel price equation:			
constant	44.131	137.516	-3.294
$\omega_{Intel,t} - \omega_{AMD,t}$	141.645	26.256	2.562
$\omega_{Intel,t} - \bar{\Delta}_t$	136.843	47.019	3.850
AMD price equation:			
constant	134.973	117.376	0.374
$\omega_{Intel,t} - \omega_{AMD,t}$	-135.905	-23.183	-4.009
$\omega_{AMD,t} - \bar{\Delta}_t$	31.429	8.461	0.890
Intel share equation:			
constant	0.831	0.833	-0.306
$\omega_{Intel,t} - \omega_{AMD,t}$	0.079	0.086	-0.371
Potential Upgrade Gains:			
Mean $(\bar{q}_t - \bar{\Delta}_t)$	0.895	0.868	0.651
Mean Innovation Rates:			
Intel	0.590	0.633	-0.839
AMD	0.610	0.635	-0.558
Mean Quality Differences:			
$\omega_{Intel,t} - \omega_{AMD,t}$	1.046	0.786	1.414
$\lvert\omega_{Intel,t} - \omega_{AMD,t}\rvert$	1.274	1.442	-1.091
Mean R&D / Revenue:			
Intel	0.114	0.116	-0.579
AMD	0.203	0.200	0.366
Objective Function :		69.85	

The pseudo-t is $|Observed - Simulated|/Observed\_Std\_Error$.

Table 3: Parameter Estimates

Parameter	Estimate	Asymptotic Std. Errors	Par. Bootstrap Std. Errors
Price, α	0.01390	0.00142	0.001127
Quality, γ	0.26615	0.03170	0.019364
Intel Fixed Effect, ξ_{intel}	-0.70589	0.01781	0.091838
AMD Fixed Effect, ξ_{AMD}	-3.00158	0.14494	0.086395
Intel Innovation, $a_{0,Intel}$	0.00138	0.00025	0.000141
AMD Innovation, $a_{0,AMD}$	0.00354	0.00073	0.000335
Spillover, a_1	3.95201	0.04922	0.577757
Stage-1 Marginal Cost Equation			
Constant, mc_0	44.7636	1.0853	
slope, mc_{slope}	-26.8630	5.0932	

The stage-1 marginal cost regression uses $\max(0, \omega_{competitor,t} - \omega_{own,t})$ as the regressor. The parametric bootstrap standard errors are from column 3 of the monte carlo results in Table 1.

Table 4: Industry Measures under Various Scenarios

	Intel-AMD Duopoly	Symmetric Duopoly	Monopoly	No Spillover Duopoly	Myopic Pricing Duopoly	Myopic Pricing Monopoly	Social Planner
Industry Profits ($billions)	243.6	268.2	290.2	267.4	206.3	188.7	-225.7
Consumer Surplus	2005.9	2035.4	1956.4	2038.0	1903.4	1919.7	2900.8
Social Surplus	2249.5	2303.7	2246.6	2305.4	2109.7	2108.4	2675.1
CS as share of monopoly CS	1.0253	1.0404	1	1.0417	0.9729	0.9813	1.4827
SS as share of monopoly SS	1.0013	1.0254	1	1.0262	0.9391	0.9385	1.1907
SS as share of planner SS	0.8409	0.8612	0.8398	0.8618	0.7886	0.7882	1
Period Profits per consumer	14.34	15.88	17.22	15.77	12.40	10.93	-14.11
Price	167.63	133.36	241.25	136.39	120.25	130.98	38.40
Margins	2.964	2.264	4.390	2.830	2.425	1.926	0
Frontier Innovation Rate	0.643	0.558	0.684	0.521	0.570	0.551	0.885
Industry Investment ($millions)	1234	1126	1571	903	618	894	5644
Quality Increase by $t = 100$	11.594	10.068	12.340	9.922	10.139	9.934	15.985
Avg. $\|q_{1t} - q_{2t}\|/\delta$	1.657	1.468	n.a.	5.799	2.149	n.a.	5.663
Avg. Percent Quality Upgrade	152.4	80.7	244.0	102.1	102.7	106.6	87.1
Firm 1 market share	0.118	0.106	0.107	0.117	0.141	0.152	0.288
Firm 2 market share	0.021	0.097	n.a.	0.082	0.026	n.a.	0.024

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 and surplus are reported in billions of dollars; investments are 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 "no spillover" duopoly uses symmetric firms, both with Intel's parameters. The social planner offers two products. The monopolist offers one product. Margins are computed as $(p - mc)/mc$. Price and margins are share-weighted averages. Firm 1 is Intel; firm 2 is AMD, except the symmetric duopoly and planner use two Intels. The social planner reported above sells two products, but the results are nearly identical for a single-product planner.

Table 5: Intel and AMD Outcomes with Optimal and Myopic Pricing

	Optimal Pricing		Myopic Pricing	
	Intel	AMD	Intel	AMD
Discounted Profits ($billions)	223.2	20.4	190.3	16.1
Period Profits	5.27	.46	4.52	.44
Period Profits per Consumer	13.18	1.16	11.30	1.10
Market Share	0.118	0.021	0.141	0.026
Margin	3.128	2.138	2.612	1.697
Investment ($billions)	1.016	0.218	0.273	0.346
Innovation Rate	0.641	0.637	0.553	0.573

Figure 1: Ownership Distribution of Intel CPUs

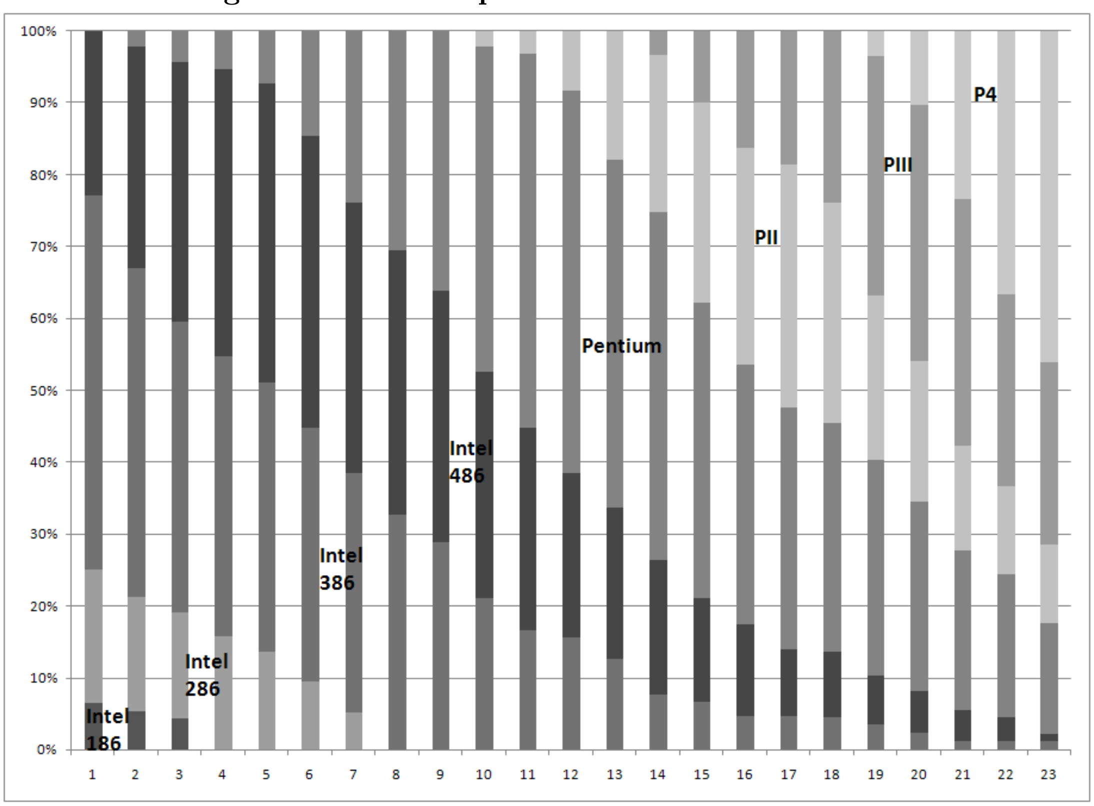

Figure 2: CPU Prices, Qualities, and Shares: 1993 Q1 to 2004 Q4

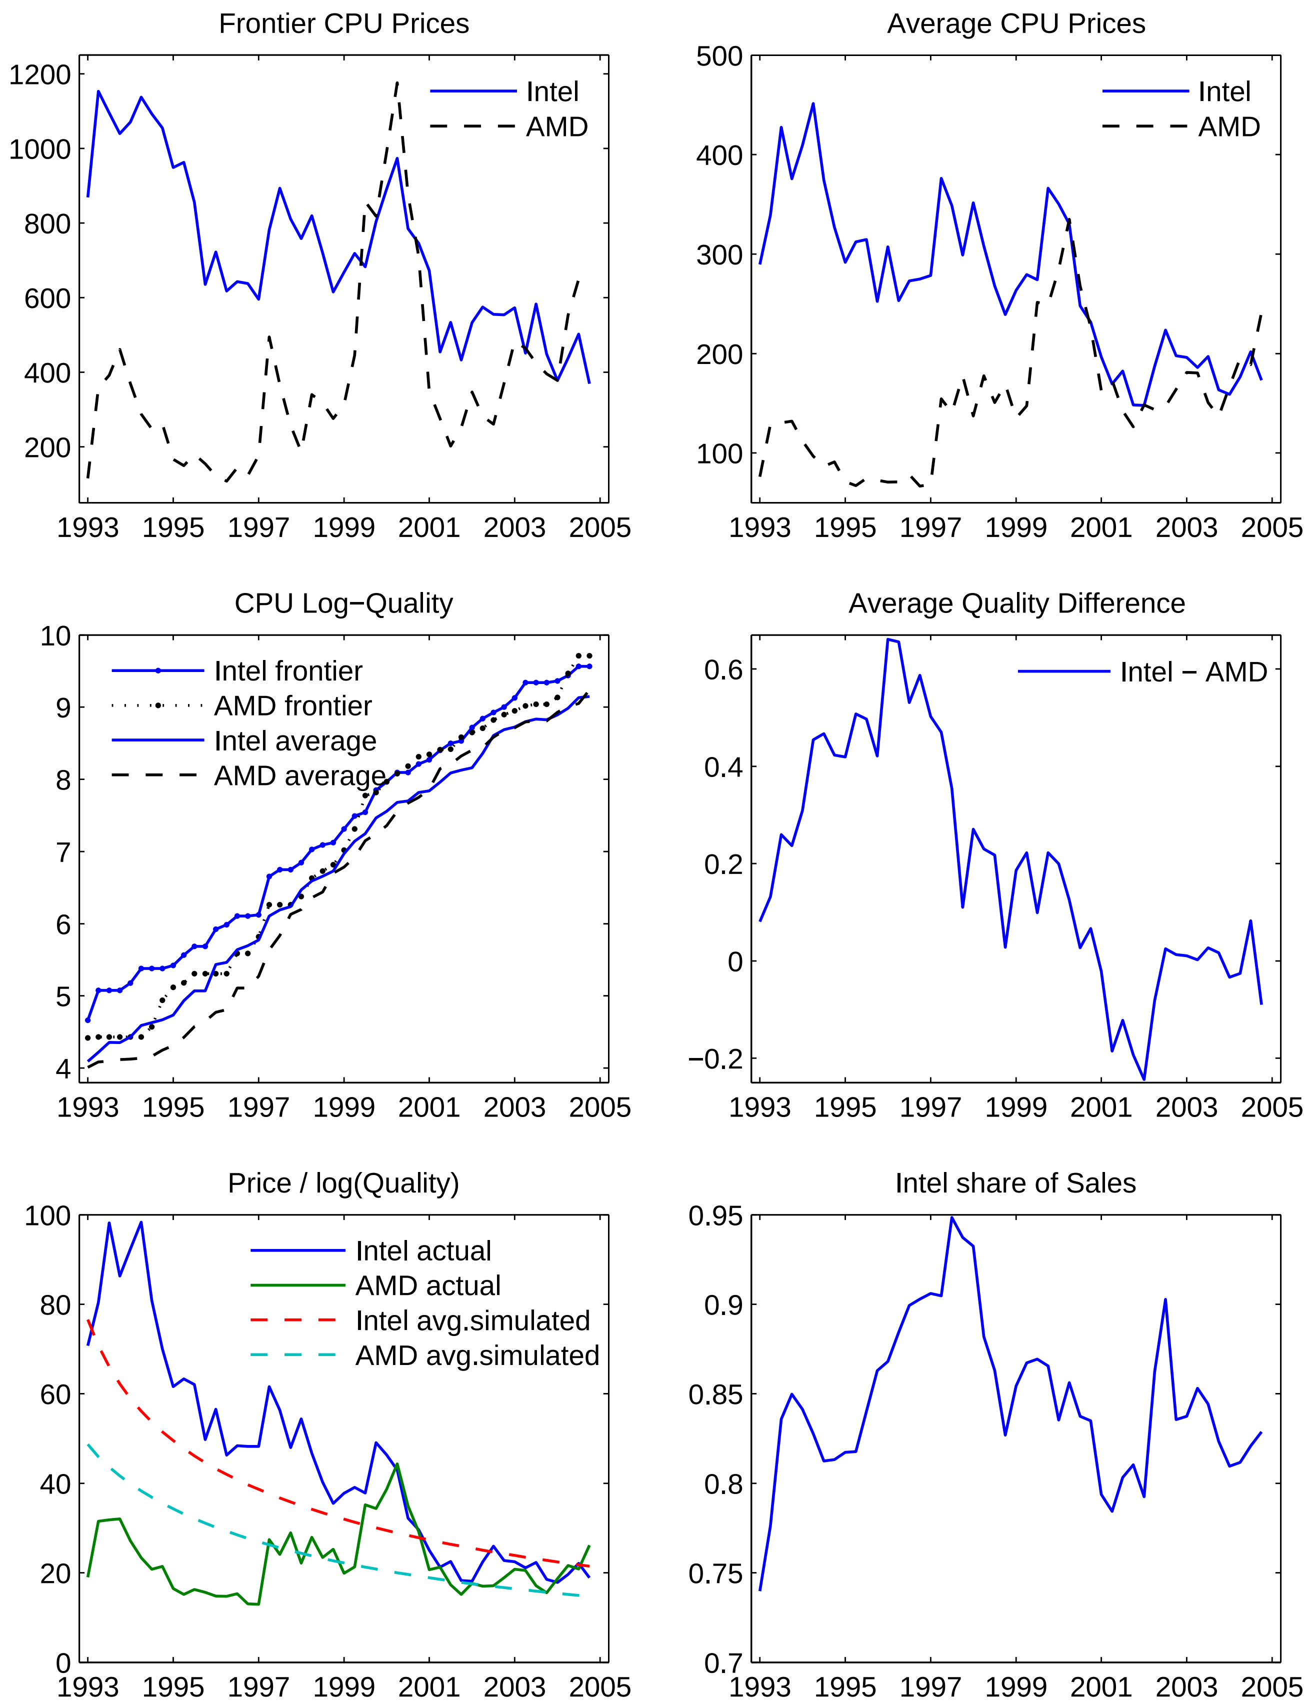

Figure 3:

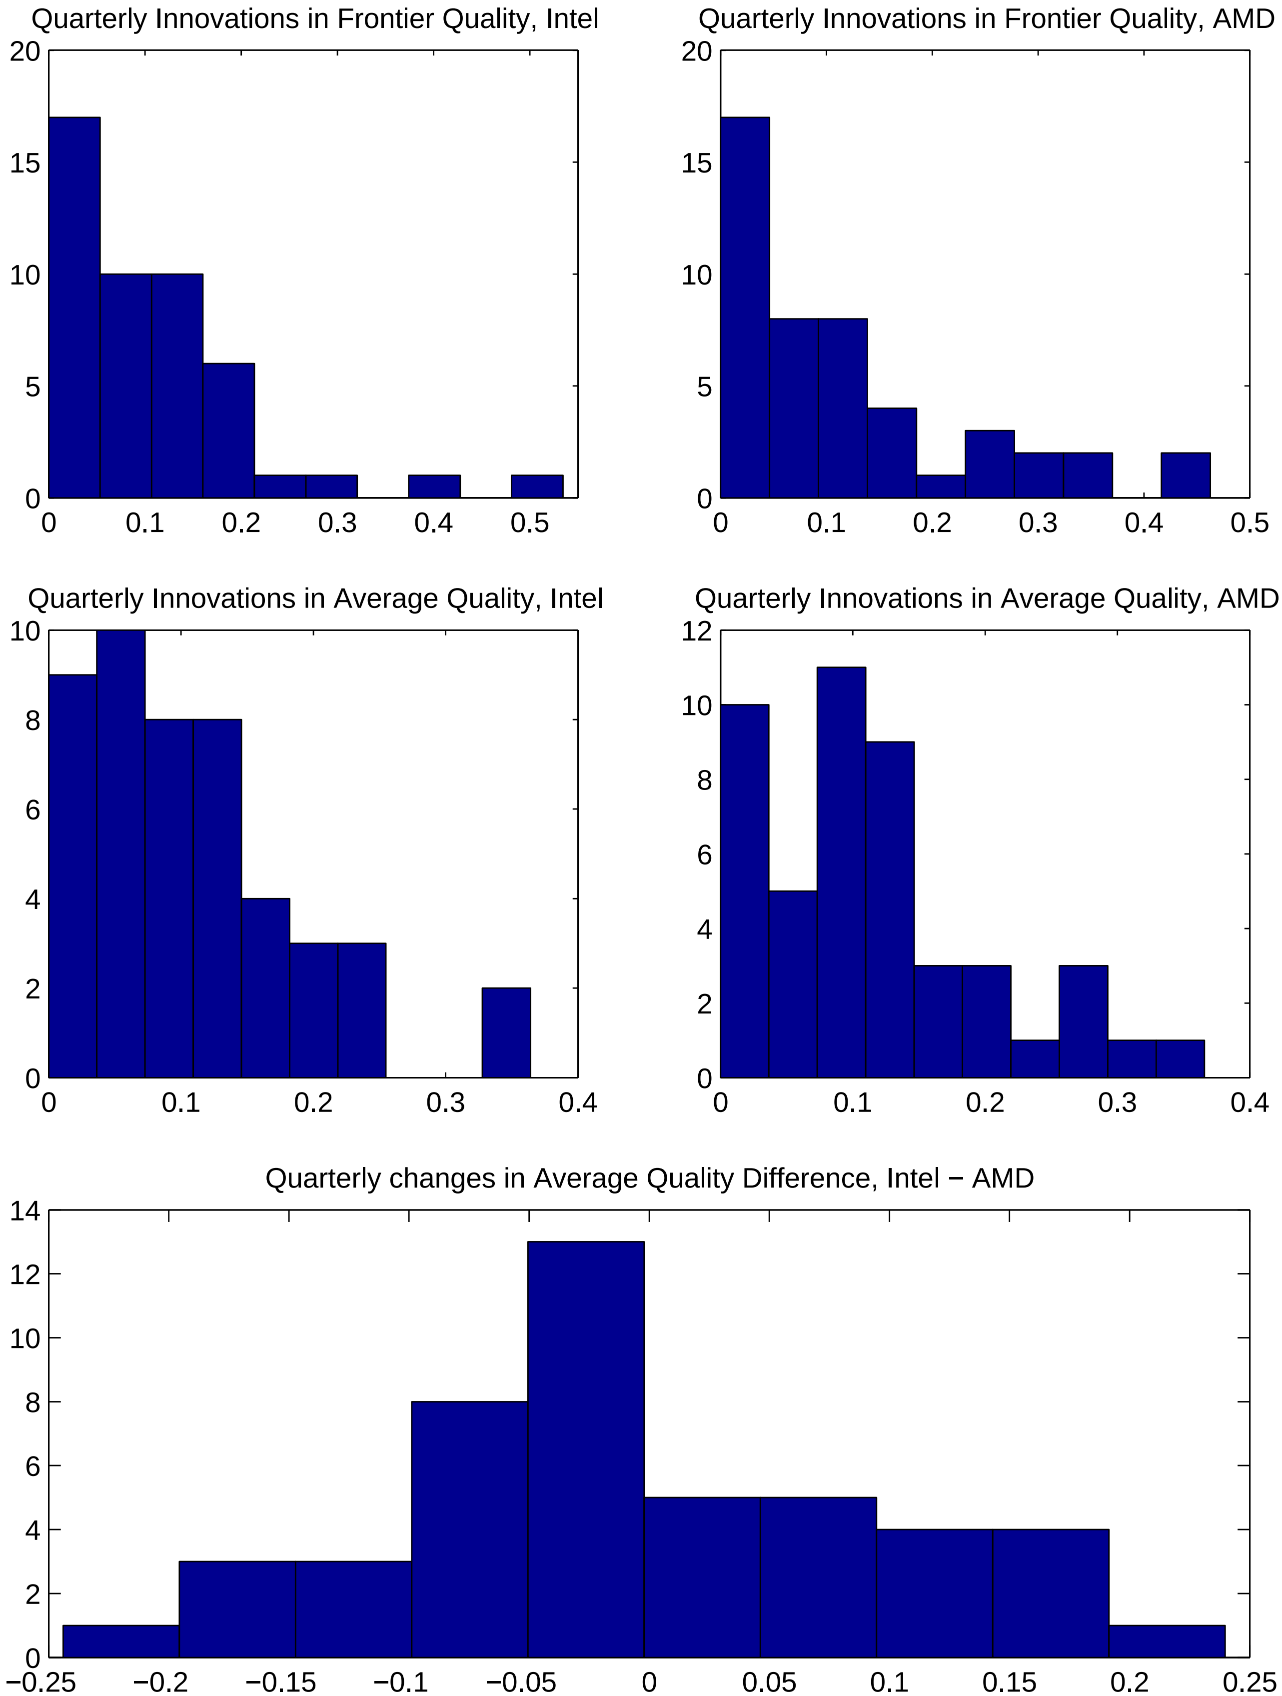
Quarterly Innovations in Frontier Quality, Intel
20
15
10
5
0
0
0.1
0.2
0.3
0.4
0.5
Quarterly Innovations in Frontier Quality, AMD
20
15
10
5
0
0
0.1
0.2
0.3
0.4
0.5
Quarterly Innovations in Average Quality, Intel
10
8
6
4
2
0
0
0.1
0.2
0.3
0.4
Quarterly Innovations in Average Quality, AMD
12
10
8
6
4
2
0
0
0.1
0.2
0.3
0.4
Quarterly changes in Average Quality Difference, Intel − AMD
14
12
10
8
6
4
2
0
−0.25
−0.2
−0.15
−0.1
−0.05
0
0.05
0.1
0.15
0.2
0.25

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.

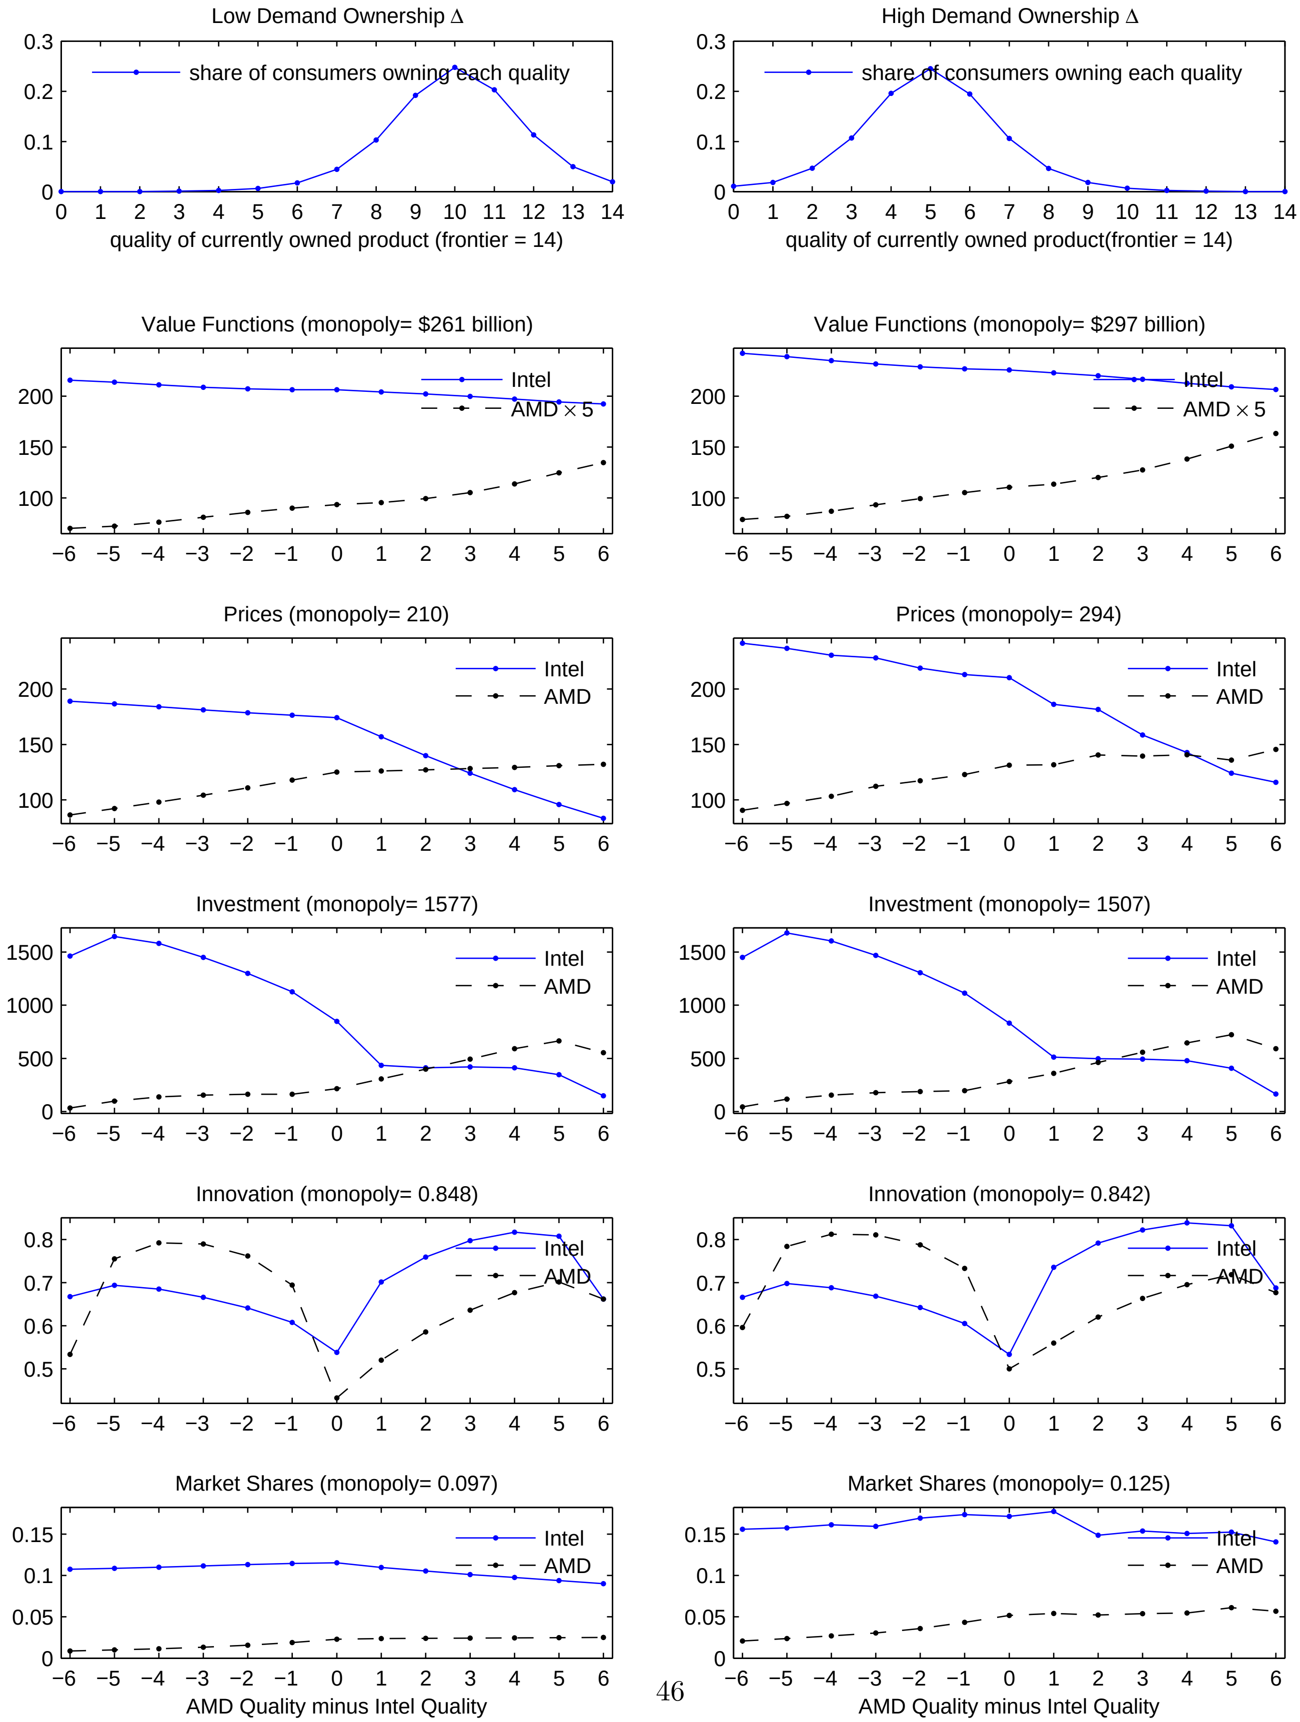

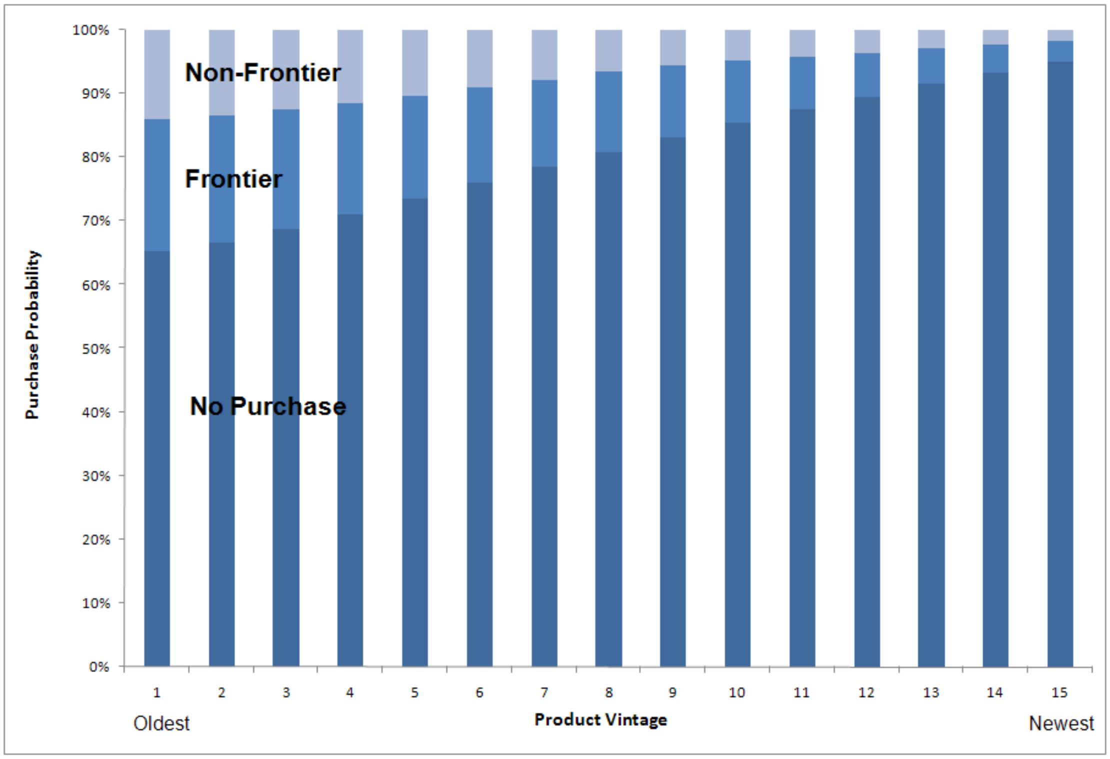

(a) **Purchase Probabilities by Vintage, Duopoly Case**

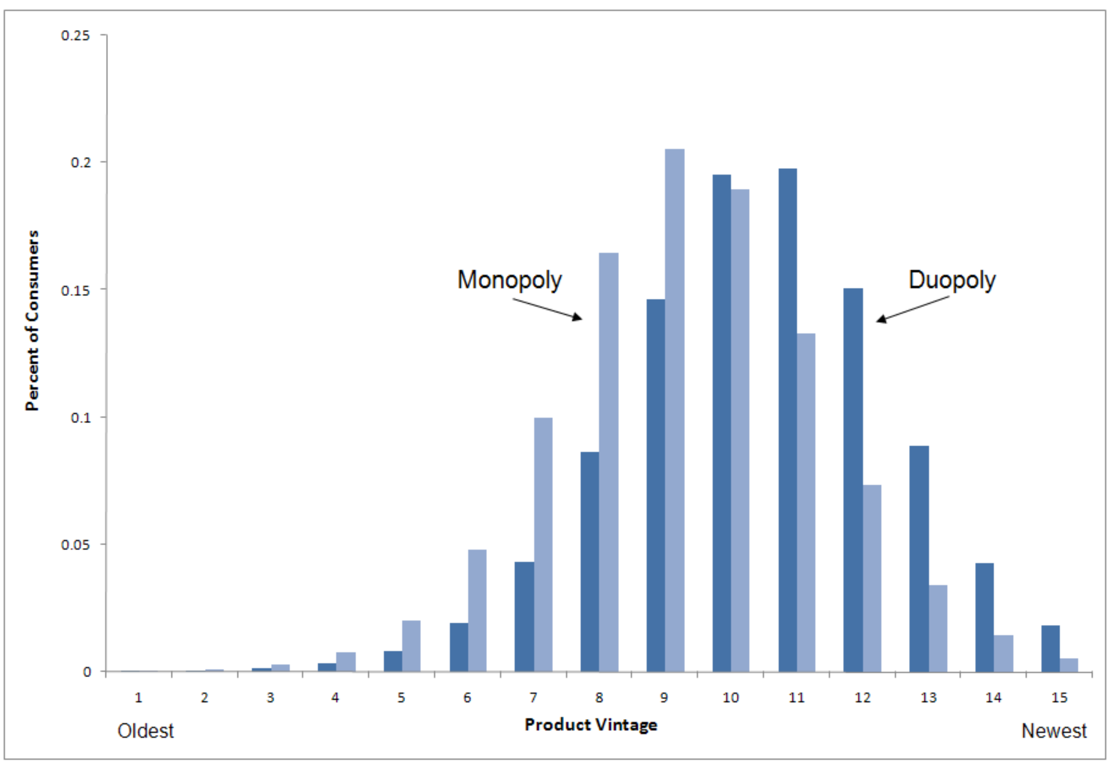

(b) **Ownership Distribution, Duopoly and Monopoly Cases**

Figure 5:

Figure 6: Elasticities and Markups for duopoly, monopoly, and myopic pricing monopoly

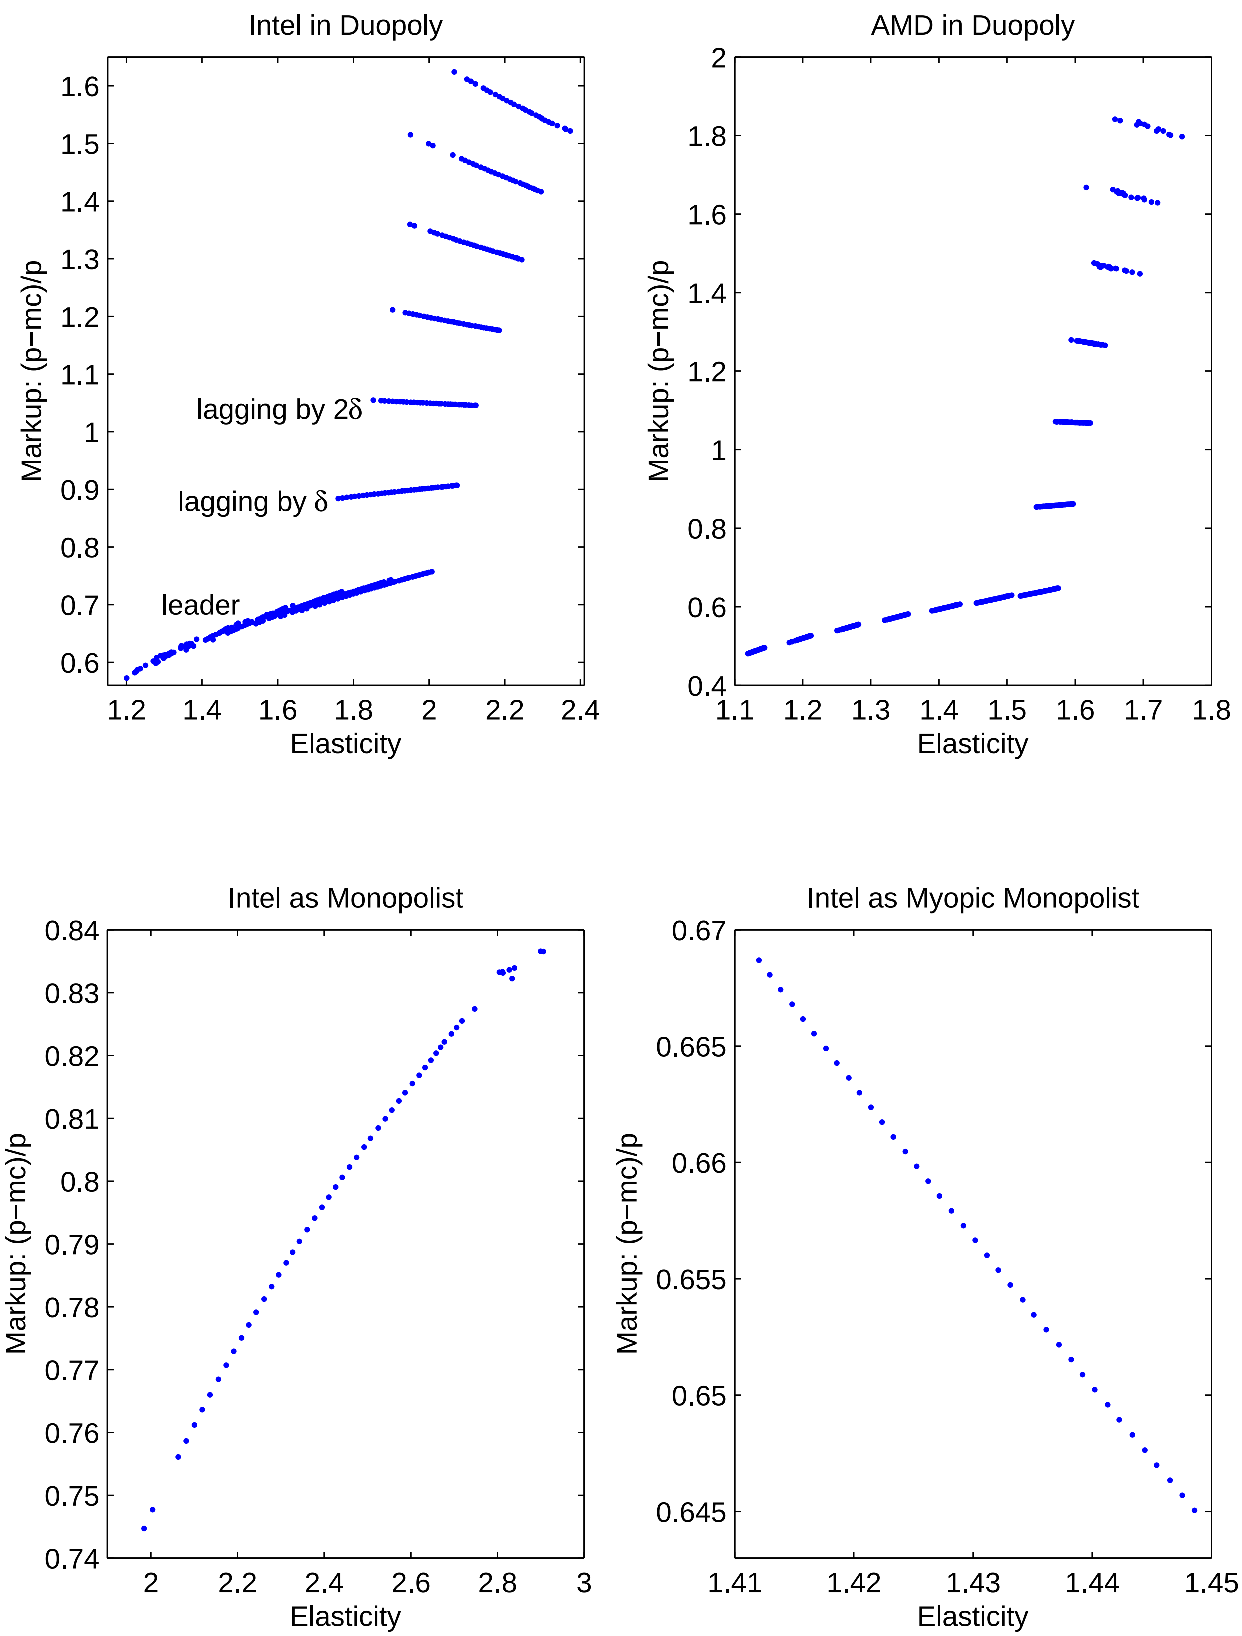

Figure 7: Market Restriction

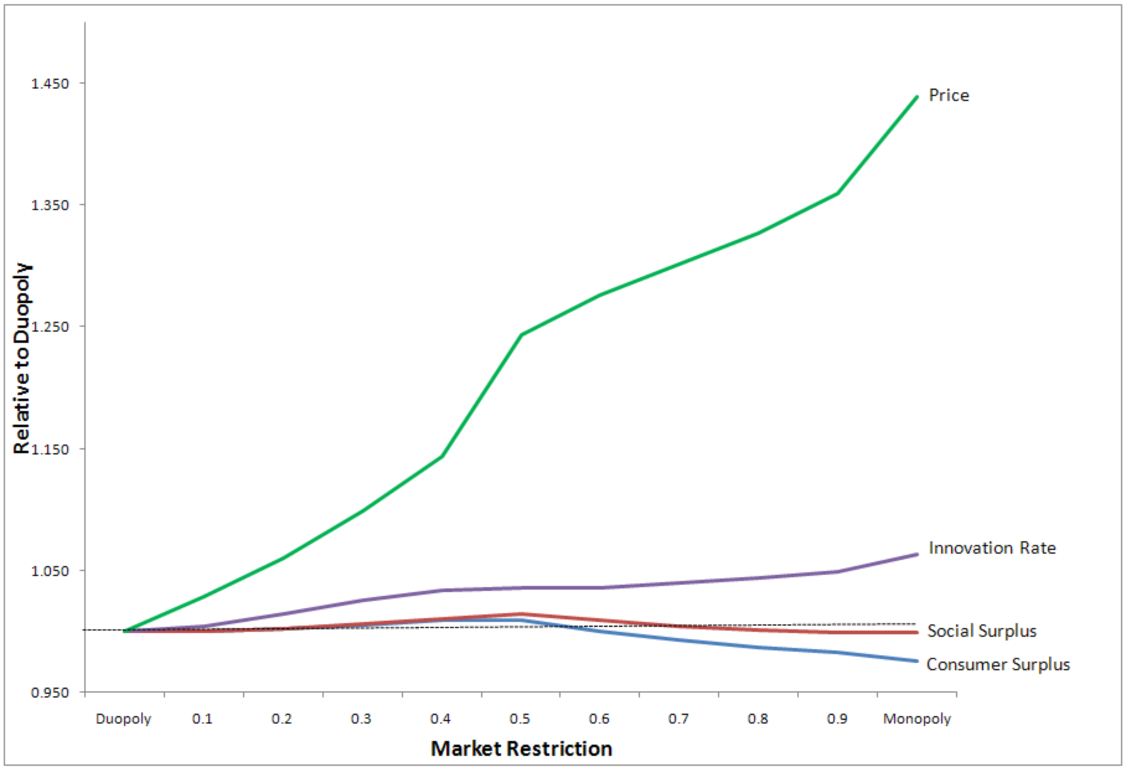

Figure 8: Comparison of Outcomes Varying Consumer Discount Factor

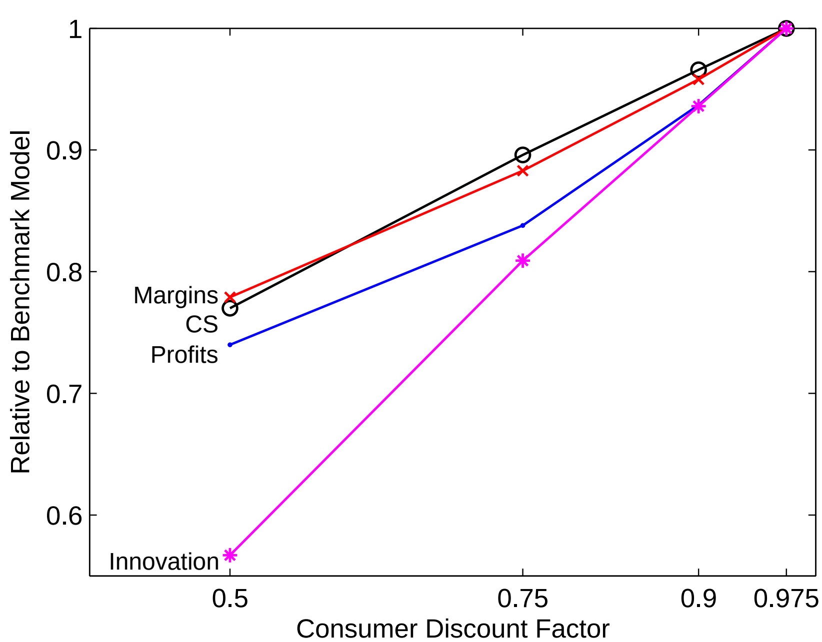

Figure 9: Symmetric Duopoly Innovation as (α, γ) Vary

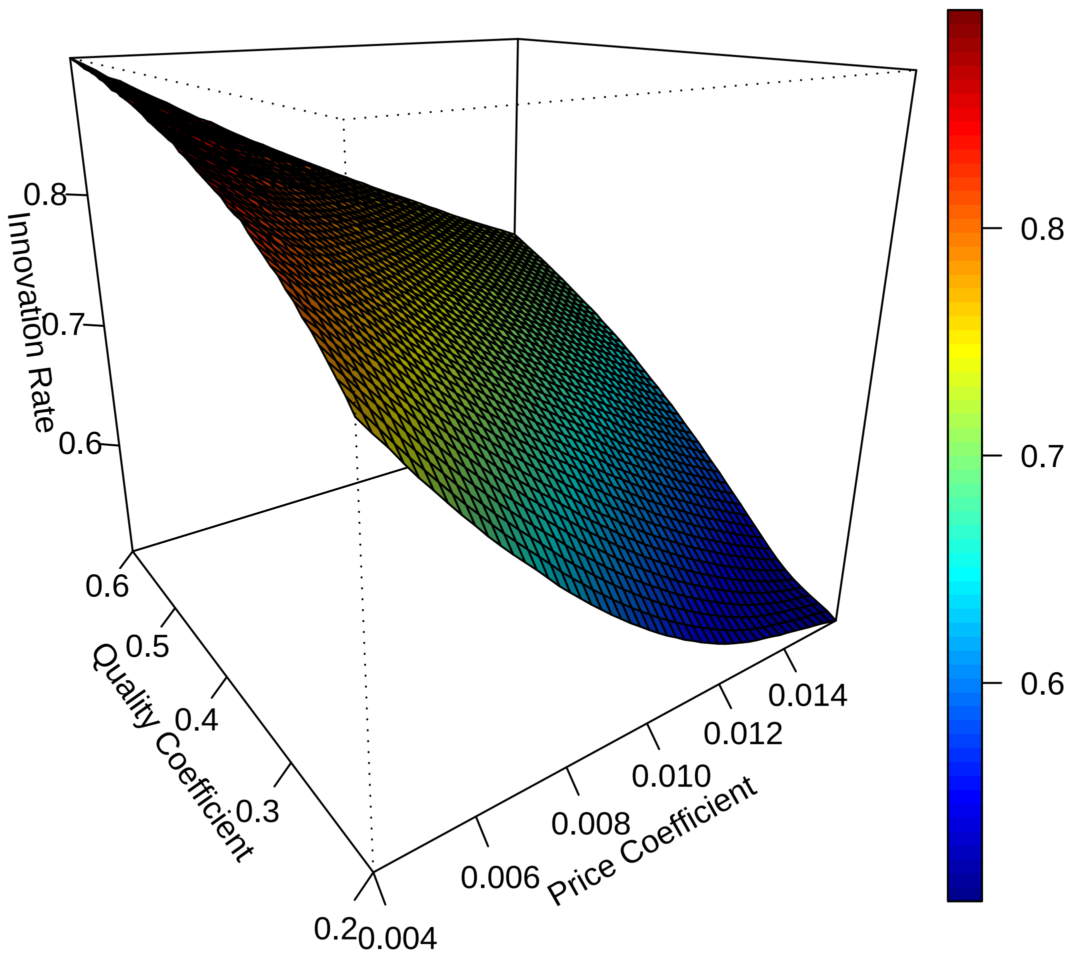

Figure 10: Symmetric Duopoly Minus Monopoly Innovation as (α, γ) Vary

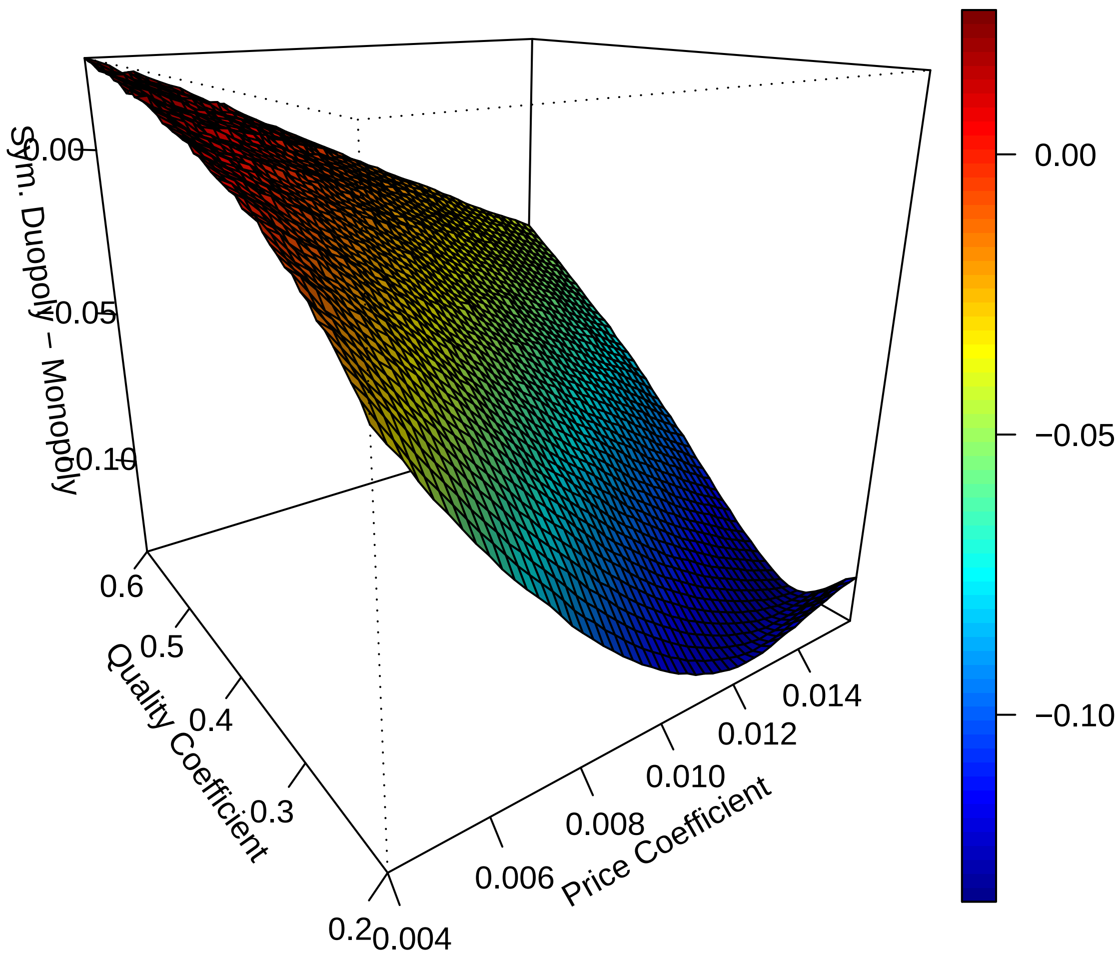

Figure 11: Industry Outcomes as var(ε) Varies