Does AMD spur Intel to innovate more?*

Ronald Goettler[†]
University of Chicago

Brett R. Gordon[‡]
Columbia University

December 16, 2011

Abstract

We estimate an equilibrium model of dynamic oligopoly with durable goods and endogenous innovation to examine the effect of competition on innovation in the PC microprocessor industry. Firms make dynamic pricing and investment decisions while consumers make dynamic upgrade decisions, anticipating product improvements and price declines. Consistent with Schumpeter, we find the rate of innovation in product quality would be 4.2 percent higher without AMD present, though higher prices would reduce consumer surplus by $12 billion per year. Comparative statics illustrate the role of product durability and provide implications of the model for other industries.

Keywords: competition and innovation, dynamic oligopoly, durable goods, estimation of dynamic games, microprocessors, CPUs

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

*Both authors contributed equally to this research and are listed alphabetically. We would like to thank Ana Aizcorbe, Ron Borkovsky, Günter Hitsch, Ali Hortaçsu, John Rust, Stephen Ryan, Holger Sieg, Gabriel Weintraub, three anonymous referees, and the editor for helpful comments. All remaining errors are our own. Goettler acknowledges financial support from the True North Communications, Inc. Faculty Research Fund at the University of Chicago Booth School of Business. 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. The original theoretical hypothesis, proposed by Schumpeter (1942), posits a positive relationship between market concentration and innovation. Arrow (1962) argues for a negative relationship, and Scherer (1967) proposes a model yielding an inverted-U relationship. The empirical literature has found mixed support for each of these hypotheses, partly due to the difficulty of controlling for industry-specific factors, leading Cohen and Levin (1989) to state, "The empirical results bearing on the Schumpeterian hypotheses are inconclusive." Despite the absence of conclusive theoretical or empirical evidence, the Federal Trade Commission increasingly cites the potential negative effect of competition on innovation as a concern (Gilbert 2006).

In this paper, we pursue a complementary approach to the reduced-form empirical studies in Cohen and Levin's review and continued by others, such as Blundell, Griffith, and Van Reenen (1999) and Aghion et al. (2005). Rather than attempt to characterize the relationship between market structure and innovation across industries, we focus on understanding this relationship in a particular industry. We construct and estimate a structural model of dynamic oligopoly with endogenous innovation to assess the effect of competition on innovation, profits, and consumer surplus in the personal computer (PC) microprocessor industry. Because microprocessors are durable, firms must compete with the stock of used goods and consumers must account for the evolution of prices and qualities when timing their purchases. We model product durability and show that its effect on equilibrium innovation can limit welfare losses due to market power. Understanding the effect of product durability on firm behavior is important since durable goods comprise 55 percent of all manufactured goods (Economic Report of the President, 2011).

We study the microprocessor industry for three primary reasons. First, the industry is important to the economy: Jorgenson, Ho, and Samuels (2010) report that the computer-equipment manufacturing industry generated 25 percent of U.S. productivity growth from 1960 to 2007. Second, recent antitrust lawsuits claim Intel's anti-competitive practices, such as rewarding PC manufacturers who exclusively use Intel microprocessors, have restricted AMD's access to consumers. Intel settled these claims in 2009 with a \$1.25 billion payment to AMD but is still under investigation by government authorities in the United States, Europe, and Asia.[1] Finally, most studies rely on indirect measures of innovation, such as patents, whereas innovations in microprocessors are directly measured via improved performance on benchmark tasks.

Several industry features and stylized facts motivate our model. First and foremost, the market is essentially a duopoly, with AMD and Intel selling 95 percent of the PC central-processing units

[1] S. Lohr and J. Kanter, "AMD-Intel Settlement Won't End Their Woes," New York Times, November 13, 2009.

(CPUs).[2] Accordingly, we cannot treat firms as being small relative to the industry, as in Hopenhayn (1992) and Klette and Kortum (2004), and instead model their strategic interaction through Markov-perfect Nash equilibrium. Second, AMD and Intel invest substantially in R&D—respectively, 20 and 11 percent of revenues on average over the 1993 to 2004 span of our data. Innovation is rapid, with new products being released nearly every quarter and CPU performance doubling roughly every seven quarters. Quarterly innovations, however, vary: the standard deviation in quarterly performance gains is slightly higher than the average gain. Finally, AMD and Intel extensively cross-license each other's technologies, which leads to an industry structure in which neither firm gets too far ahead and technological leadership changes hands. To capture these supply-side features, we model innovation in an AMD-Intel duopoly as stochastic gains on a quality ladder in which success is more likely with higher investments and for laggards who benefit from innovation spillovers.

Consumer behavior also guides our model. As microprocessors are durable, replacement drives demand: 82 percent of PC purchases in 2004 were replacements (Computer Industry Almanac, 2005). A short-term increase in innovation widens the quality gap between currently-owned products and new offerings, boosting demand and raising prices and sales. After the upgrade boom, prices and sales fall as replacement demand drops. Firms must continue to innovate to rebuild replacement demand, because microprocessors do not physically depreciate. We model this upgrade cycle and the timing of consumers' purchases given beliefs about future prices and innovation. Because Intel and AMD tend to revise prices and product offerings quarterly, our infinite-horizon, discrete-time model has 3-month periods.[3]

To identify the effect of competition on innovation, we estimate consumer preferences and firms' innovation efficiencies, which determine benefits and costs of innovation, and solve for equilibrium under various competitive scenarios. This approach accords with Dorfman and Steiner (1954), Needham (1975), and Lee (2005), who find consumer preferences and firm competencies are key determinants of R&D. We estimate preferences and innovation efficiencies using a minimum distance estimator to match simulated moments from our model's equilibrium to observed aggregate moments, such as average prices and innovation rates, constructed from quarterly CPU prices, qualities, market shares, and innovation.[4] We then compare outcomes across counterfactual simulations with AMD

[2]Cyrix Corporation (acquired in 1997 by National Semiconductor), Transmeta Corporation, and VIA Technologies were fringe players trying to break into the market during the 1990s and early 2000s, but none succeeded. The AIM Alliance of Apple Computer, IBM, and Motorola supplied the PowerPC microprocessor for Apple, which garnered a 2 percent share of sales in 2003.

[3]We assume a firm's ability to commit to prices is exogenously specified by the period length: firms commit to fixed prices within, but not across periods. Thus, we do not address the time-inconsistency problem of wanting to commit today to a high price, but then wanting to lower it later after some consumers buy at the high price (Coase 1972, Stokey 1981, Bulow 1982, and Bond and Samuelson 1984). Assessing the effect of period length on industry outcomes would be interesting, though difficult to implement since it involves changing the scale of several parameters simultaneously and tweaking the innovation process to maintain *ceteris paribus*.

[4]Several studies estimate demand for durable goods, taking product quality as exogenous (Melnikov 2001, Song and Chintagunta 2003, Nair 2007, Gordon 2009, Gowrisankaran and Rysman 2009, and Carranza 2010). Our econometric model differs from these in its use of supply-side equilibrium restrictions to help identify the structural parameters.

either removed or strengthened to be an equal competitor to Intel.[5] Importantly, our model can generate either a positive or a negative relationship between competition and innovation, depending on parameter values. The data therefore guide our conclusions.

We find the rate of innovation in product quality would be 4.2 percent higher if Intel were a monopolist, consistent with Schumpeter. Without AMD, higher margins spur Intel to innovate faster to generate upgrade sales. This result, however, depends on the degree of competition from past sales. If first-time purchasers were to arrive sufficiently faster than we observe, innovation in an Intel monopoly would be lower, not higher, since upgrade sales would be less important.

Consumer surplus would be 4.2 percent lower ($12 billion per year) in an Intel monopoly since the surplus gains from higher innovation are smaller than the losses from the 50 percent increase in prices. Like Coase's (1972) conjecture and the ensuing literature, we show that product durability can limit welfare losses from market power.[6] We hypothetically vary depreciation and market growth to show, respectively, that lowering durability or its importance increases the surplus loss from removing AMD. Unlike Coase, though, the mechanism in our model involves innovation as well as pricing.

We also evaluate the effect of Intel's alleged anti-competitive practices, by performing counterfactual simulations in which we vary the share of the market from which AMD is foreclosed. The industry innovation rate peaks when AMD is foreclosed from half the market and consumer surplus peaks with 40 percent foreclosure. This latter result reveals that the surplus gains from faster innovation can exceed losses due to higher prices. We therefore find support for the FTC's recent emphasis on the dynamic trade-off between lower current consumer surplus from higher prices and higher future surplus from more innovation.

To further understand the relationship between competition and innovation, we perform additional comparative statics by varying (i) consumer preferences for quality and price, (ii) product substitutability, and (iii) the degree of innovation spillovers that enable firms to innovate more efficiently when catching up to the frontier.

We find that equilibrium innovation rates increase monotonically as preferences for quality increase and as price sensitivity declines, for both duopoly and monopoly. As explained in section 5, duopoly innovation is more sensitive to preferences. Consequently, industry innovation is higher in the duopoly than in the monopoly when quality preferences are high and price sensitivity is low.

Innovation spillovers reduce incentives for leaders to innovate but also ensure laggards do not fall so far behind that they give up trying to remain competitive, as they do in our model without spillovers. We show duopoly innovation increases as spillovers decrease, as long as the laggard never concedes leadership. With no spillovers or large spillovers, monopoly innovation is higher than duopoly innovation, but with moderate spillovers duopoly innovation is higher.

[5] Our model can be extended to yield an endogenous number of firms. Counterfactuals in the number of firms would then correspond to exogenous shifts in entry and exit costs.

[6] Carlton and Gertner (1989) show that competition from past sales of durable goods limits the welfare loss of mergers that increase market power. See the review article by Waldman (2007).

As a whole, our comparative statics demonstrate that competition's effect on innovation depends on industry characteristics that likely vary across industries and perhaps across time within an industry. Such variation might be one reason cross-industry studies have difficulty identifying robust relationships.

Our work relates to the literatures on endogenous growth theory and dynamic oligopoly. A series of papers in endogenous growth theory (Aghion and Howitt 1992, Aghion et al. 2001, and Aghion et al. 2005) examines the relationship between competition and innovation. In addition to providing suggestive evidence of an inverted-U relationship between the Lerner index and patent-production in UK industries, Aghion et al. (2005) develop a model of technological innovation that generates this relationship. We demonstrate that the durability of goods and nonzero investment by frontier firms in our model generate implications that differ from those in Aghion et al. (2005).

Vives (2008) also investigates the effect of competition on innovation by firms selling nondurable goods. He finds firms innovate less when facing more competitors and innovate more when competition increases via greater product substitutability. We find, with durable goods, that the effect of more competitors on innovation depends on consumer preferences and the strength of innovation spillovers.

Our work is a natural extension of 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, hereafter EP). The EP framework has been applied to a variety of industries, as summarized by Doraszelski and Pakes (2007), but none of the studies considers durable goods with forward-looking consumers and endogenous innovation. Given the prominence of durable goods in our economy (e.g., airplanes, automobiles, consumer electronics, etc.) and the importance of innovation for economic growth, filling this gap is a major contribution of our paper.

We incorporate durable goods into the EP framework as applied to differentiated products by Pakes and McGuire (1994).[7] In our model, firms make dynamic pricing and investment decisions while taking into account the dynamic behavior of consumers. In turn, when considering to buy now or later, consumers account for the fact that firms' strategies lead to higher-quality products and lower prices. Since consumers' choices depend on the products they currently own, the distribution of currently-owned products affects aggregate demand. We model the endogenous evolution of this distribution and its effect on equilibrium behavior. Prices, innovation, profits, and consumer surplus are all substantially higher when firms correctly account for the dynamic nature of demand arising from durability. We find that ignoring the dynamic nature of demand for durable goods leads to a reversal of the effect of competition on innovation.

[7]The theoretical literature on durable goods, reviewed by Waldman (2003), focuses on monopoly and perfect competition, whereas we consider the more empirically relevant market structure of oligopoly. This literature also focuses on endogenous product durability (i.e., the rate of depreciation), whereas we study endogenous obsolescence due to innovation. Though similar, durability and obsolescence have an important difference: durability entails commitment since the good is produced with a given durability, whereas obsolescence depends on future innovations.

In the EP framework, the industry's long-run innovation rate equals the exogenous rate at which the outside good's quality improves because returns to innovation are assumed to go to zero when a firm's quality is sufficiently higher than the outside good, regardless of competitors' qualities. We relax this assumption to obtain an endogenous long-run innovation rate that depends on consumer preferences and firms' technologies. Endogenous innovation is important for policy work because the compounding effects of innovation on consumer surplus can dominate pricing effects.[8]

In the following section, we describe aspects of the microprocessor industry that motivate our model and empirical strategy. In section 3, we present our model of firm and consumer behavior. In section 4, we estimate the model using the microprocessor data and discuss implications specific to that industry. In section 5, we perform a series of comparative statics to further illustrate the model's properties and its implications for other industries. Section 6 concludes.

2 Data and Industry Background

Intel co-founder Gordon Moore predicted in 1965 that the number of transistors per integrated circuit would double every two years, thereby doubling performance. Panel (a) of Figure 1 depicts "Moore's law" over the 48 quarters in our data from 1993 through 2004 by plotting the log-quality of the frontier CPU for Intel and AMD, where quality is measured using processor speed benchmarks from www.cpuscorecard.com and www.cpubenchmark.net.[9] The mean quarterly percentage change in CPU performance from 1993 to 2004 is 10.2 percent for Intel and 11 percent for AMD. Nearly one-fifth of the quarters have gains exceeding 20 percent and more than one-fifth of the quarters have no improvements in frontier quality. Accordingly, We model firms as innovating with uncertain outcomes to climb a quality ladder.

The largest performance gains result from major redesigns of the microprocessor die, such as Intel's progression from the 386 to the 486 to the Pentium and AMD's progression from the K5 to the K6 to the Athlon. Smaller gains arise from other design changes, such as adding a math coprocessor to the 486SX to create the 486DX. From 1993 to 2004, AMD and Intel sold processors from 10 and 20 different die designs, respectively. As a firm gains experience manufacturing a given design, the yield of usable dies from each silicon wafer increases, which lowers unit costs. With experience, the firm also increases processor speed. An average of 8.2 processor speeds were offered for each die-design.

Since few consumers purchase frontier CPUs, we average the log-qualities of each firm's CPU

[8] Goettler and Gordon (2011) use a dynamic oligopoly model similar to the one in this paper to investigate the relationship between various measures of competition and innovation when goods are nondurable.

[9] We splice two benchmarks to construct a single index of quality comparable across product generations since no single benchmark spans our dataset. The growth of mobile computing and server farms in recent years has led consumers and firms to focus on power consumption as well as execution speed. Over our period, however, desktops comprised over 80 percent of CPU sales and performance per unit time, not per watt, was the focus.

offerings in each quarter and plot the difference in average log-qualities in panel (b) of Figure 1.[10] Intel's initial quality advantage is moderate in 1993 to 1994, then becomes large when it releases the Pentium. AMD's introduction of the K6 processor in 1997 narrows the gap, but parity is not achieved until sales of the AMD Athlon gained traction in mid-2000.

Unit shipments, manufacturers' average selling prices (ASP), and production costs are provided by In-Stat/MDR, a market research firm specializing in the microprocessor industry. ASPs in panel (d) are lower and less variable than frontier-product prices in panel (c). We assume retail CPU prices are the same as manufacturer prices since consumers tend to buy CPUs as part of a PC and the PC manufacturing sector is competitive, with margins below 5 percent.[11] All prices and costs are converted to base year 2000 dollars.

The covariation in Intel's share of sales, its quality advantage, and its ASP is evident by comparing their plots, vertically arranged on the right-hand side of Figure 1. Over our sample, the correlation between Intel's ASP and its quality advantage is .66, and the correlation between AMD's ASP and Intel's advantage is $-.34$. The correlation between Intel's share and its quality advantage is .39. These correlations are consistent with the model we present in section 3 and help identify its parameters, as discussed in section 4.1.2.

CPU prices also depend on competition from CPUs bought in the past. To measure such competition, we average the log-quality of currently-owned CPUs, as reported in consumer surveys conducted by Odyssey, a consumer research firm specializing in technology products.[12] This average quality trails the quality of frontier CPUs in panel (c) for two reasons: consumers rarely purchase the frontier product and only upgrade every 3.3 years (Gordon, 2009). The correlation of each firm's price with its quality relative to the average quality currently owned is .69 for Intel and .37 for AMD.

Although prices and production costs of a given processor fall over time, more complicated chip designs lead to stationary prices and unit costs, as depicted in panels (d) and (e) in Figure 1. The significant correlation of .48 between each firm's unit costs (sales-weighted blended unit production costs) and its quality relative to its competitor motivates our model for costs in the next section.

Finally, quarterly R&D investment levels, obtained from firms' annual reports, are a relatively constant share of revenue. Although AMD's investment share of revenue is nearly double Intel's share, AMD's investment level is about one-fourth the level of Intel. Nonetheless, AMD is able to

[10]Ideally we would use sales of each CPU to construct average log-quality, but we only observe quantities at the die-design level. In each quarter, we equally allocate a die's sales across the CPUs with its design.

[11]In 2002, 30% of PCs were sold by unbranded "white box" manufacturers (J. Spooner, "Dell Eyes 'White-Box' Market," CNET News, August 20, 2002).

[12]The semi-annual Homefront surveys by Odyssey provide a national sample of 1500 to 2500 households reporting the processor speed and manufacturer of their primary or most recently purchased PC. We interpolate these semi-annual ownership distributions yielding quarterly data that we combine with the quarterly penetration rate of PCs in U.S. households to obtain the ownership distribution across all consumers, including those who have yet to adopt. We assume consumers who have yet to purchase a PC have public access to a PC with a processor 7.8 percent the speed of the frontier. For comparison, the 80286 processor (three generations before the Pentium) is 8.6 percent the speed of the Pentium.

offer similar, sometimes even higher, quality products beginning in 1999. To explain this asymmetry, our model in the next section allows for innovation spillovers since AMD is usually in the position of playing catch-up.

3 Model

We present a dynamic model of differentiated-products oligopoly for a durable good. Although we interpret some model details in the context of microprocessors, the model applies to any durable good. We abstract away from the role of computer manufacturers because consumers can choose either firm's microprocessors regardless of their choice of other computer components (e.g., disk drive, memory, video card, monitor, etc.).

Time, indexed by t, is discrete with an infinite horizon. Each firm $j \in \{1, \ldots, J\}$ sells a single product and invests to improve its quality. If successful, quality improves next period by a fixed proportion; otherwise it is unchanged.[13] Consequently, we denote log-quality $q_{jt} \in \{\ldots, -2\delta, -\delta, 0, \delta, 2\delta, \ldots\}$.

A key feature of demand for durable goods is that the value of the no-purchase option is endogenous because it depends on past choices. Consumers decide each period whether to buy a new product or to continue using the one they already own. This feature generates a dynamic trade-off for pricing: selling more in the current period reduces demand in future periods because recent buyers are unlikely to buy again in the near future. The distribution of currently-owned products, denoted Δ_t, therefore affects current demand.

Firms and consumers are forward-looking and take into account the optimal dynamic behavior of other agents 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 simultaneously choosing prices p_{jt} and investment x_{jt}. 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. We assume consumers observe Δ_t merely as a convenient way to impose rational expectations of future prices and qualities. Rationality requires consumers to act as if they condition on the ownership distribution since it influences innovation and future prices through firms' policy functions.

We restrict firms to sell only one product because the computational burden of allowing multi-product firms is prohibitive—the state space grows significantly and the optimization within each state becomes substantially more complex. Accounting for multiple products would be important if our focus were on price discrimination or product-line pricing and quality choices (Aizcorbe and

[13] Borkovsky (2008) studies the timing of new releases, and Holmes, Levine, and Schmitz (2011) explore the effect of switchover disruptions on the incentives to innovate.

Kortum 2005, Gordon 2009, and Nosko 2010). Our demand model captures the market features that are most relevant for our focus on endogenous innovation: consumers upgrade when the offered qualities are sufficiently higher than their currently-owned quality, and consumers expect innovations to raise future quality and lower future prices per unit quality.

We do not consider entry and exit since they rarely occur in the CPU industry. We also do not consider secondary markets since computers and microprocessors are rarely resold. With resale, the ownership distribution would convey the set of used goods available for trade, as in the model with car resale in Chen, Esteban, and Shum (2011).

3.1 Consumers

We model consumers as owning one microprocessor at a time.[14] Utility for a consumer i from firm j's new product with quality q_{jt} is given by

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

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_{ijt} captures idiosyncratic variation, which is i.i.d. across consumers, products, and periods.[15]

We assume brand preference only affects utility at the time of purchase and normalize the brand preference for the no-purchase option to be zero. Utility from the no-purchase option is then

$$u_{i0t} = \gamma \tilde{q}_{it} + \varepsilon_{i0t} \ . \tag{2}$$

In principle, the model has two outside alternatives: for consumers with previous purchases, $\tilde{q}_{it}$ is the quality of their most recent purchase, and for non-owners $\tilde{q}_{it}$ is the quality available through other means, such as public access.[16]

To facilitate bounding the state space, we assume $\tilde{q}_{it}$ is within $\bar{\delta}_c$ of the industry's frontier product. That is, $\tilde{q}_{it} \geq \underline{q}_t \equiv \bar{q}_t - \bar{\delta}_c$, where $\bar{q}_t \equiv \max(q_t)$. To ensure our choice of $\bar{\delta}_c$ does not affect equilibrium behavior, we check that consumers upgrade frequently enough that the quality of their most recent purchase rarely matches $\underline{q}_t$.

Since the ownership distribution only has mass at vintages weakly above $\underline{q}_t$, we define the ownership state variable $\Delta_t = (\Delta_{\underline{q}_t,t}, \ldots, \Delta_{k,t}, \ldots, \Delta_{\bar{q}_t,t})$, where $\Delta_{k,t}$ is the fraction of consumers whose outside option has quality $\tilde{q}_{it} = q_{kt}$.

Each consumer maximizes her expected discounted utility, yielding a value function V that satis-

[14] Cho (2008) estimates a dynamic model of computer replacement by a telecommunications firm using many computers.

[15] As explained in Rust (1996), the independence from irrelevant alternatives (IIA) property of logit demand fails to hold in dynamic contexts since the attributes of all the products enter the continuation values.

[16] For CPUs, the outside good for non-owners might consist of using computers at schools and libraries or using old computers received from family or friends who have upgraded.

fies Bellman's equation. Omitting i and t subscripts for conciseneness and using the prime superscript to denote next period values,

$$V(q,\Delta,\tilde{q},\varepsilon) = \max_{y\in(0,\dots,J)} \; u_y + \beta \sum_{q',\Delta'} \int V(q',\Delta',\tilde{q}'_y,\varepsilon')\; f_\varepsilon(\varepsilon')d\varepsilon'\; h_c(q'|q,\Delta,\varepsilon)\; g_c(\Delta'|\Delta,q,q',\varepsilon), \tag{3}$$

where y denotes the optimal choice in the current period, $h_c(\cdot|\cdot)$ is the consumer's beliefs about future product qualities, $g_c(\cdot|\cdot)$ is the consumer's beliefs about the transition kernel for Δ', and f_ε is the density of ε. The evolution of $\tilde{q}$ is trivial: if $y=0$ then $\tilde{q}' = \max(\tilde{q},\underline{q}')$; otherwise $\tilde{q}' = q_y$. Each consumer is small relative to the market so that her actions do not affect the evolution of Δ to Δ'.

Following Rust (1987), we assume ε are multivariate extreme-value and integrate over ε to obtain the smoothed Bellman equation

$$\bar{V}(q,\Delta,\tilde{q}) = \log\left(\sum_{j\in\{0,\dots,J\}} \exp\left\{u_j - \varepsilon_j + \beta\sum_{q',\Delta'}\bar{V}(q',\Delta',\tilde{q}'_j)\; h_c(q'|q,\Delta)\; g_c(\Delta'|\Delta,q,p,q')\right\}\right), \tag{4}$$

from which we construct product-specific value functions:

$$v_j(q,\Delta,\tilde{q}) = u_j - \varepsilon_j + \beta\sum_{q',\Delta'}\bar{V}(q',\Delta',\tilde{q}'_j)\; h_c(q'|q,\Delta)\; g_c(\Delta'|\Delta,q,p,q')\;. \tag{5}$$

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

$$s_{j|\tilde{q}} = \frac{\exp\{v_j(q,\Delta,\tilde{q}\})}{\sum_{k\in\{0,\dots,J\}} \exp\{v_k(q,\Delta,\tilde{q})\}}\;. \tag{6}$$

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

$$s_j = \sum_{\tilde{q}\in\{\underline{q},\dots,\bar{q}\}} s_{j|\tilde{q}}\;\Delta_{\tilde{q}}\;. \tag{7}$$

These market shares translate directly into the law of motion for Δ, which tracks the ownership of products between $\underline{q}_t$ and $\bar{q}$. Assuming $\bar{q}$ is unchanged between the current and next periods, the share of consumers owning a product of quality k at the start of the next period is the share who retain product k plus the share of consumers who bought a new product from any firm offering quality k. If a firm advances the quality frontier with a successful R&D outcome then Δ' shifts because the ownership distribution is defined relative to the frontier quality. We relegate the notational details of the evolution of Δ to a footnote.[17]

[17] Assuming $\bar{q}$ is unchanged between the current and next periods and letting $\mathcal{I}(\cdot)$ denote an indicator function

$$\Delta'_k(\Delta,q,p|\bar{q}'=\bar{q}) = s_{0|k}\Delta_k + \sum_{j=1,\dots,J} s_j\mathcal{I}(q_j=k)\;.$$

3.2 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 and chooses an investment level $x_j \in \mathbb{R}_+$.[18] We restrict the innovation outcome $\tau_j = q'_j - q_j$ to be either 0 or δ, with the probability of success given by

$$\chi_j(\tau = \delta | x, q) = \frac{a^j(q)x}{1 + a^j(q)x}, \tag{9}$$

which yields a closed form for optimal investment (Pakes and McGuire, 1994). The investment efficiency $a^j(q) = a_{0,j} \max(1, a_1(\frac{\bar{q} - q_j}{\delta})^{1/2})$ is higher for firms below the frontier ($\bar{q} - q_j > 0$), assuming a positive innovation spillover a_1. This spillover implies an increased difficulty of advancing the frontier relative to catching up to it.[19] Linear and convex spillovers yield similar results to the concave $a^j(q)$ we use. The probability of failure is $\chi_j(\tau = 0|x, q) = 1 - \chi_j(\tau = \delta | x, q)$.

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

$$\pi_j(p, q, \Delta) = M s_j(p, q, \Delta)(p_j - mc_j(q)) \ , \tag{10}$$

where M is the fixed market size, $s_j(\cdot)$ is the market share for firm j from equation (7), and p is the vector of J prices. In section 4.1.3 we discuss the possibility of reallocating mass to the lowest vintage in Δ to capture one effect of consumers entering the market, while retaining fixed M to ensure stationarity. Firm j's constant marginal costs are given by

$$mc_j(q) = \lambda_0 + \lambda_1 \left(\bar{q} - q_j\right) \ , \tag{11}$$

where $\lambda_1 < 0$ implies production costs are lower for non-frontier firms. In our application, $|\lambda_1|$ is small enough that marginal costs are always positive.

Each firm maximizes its expected discounted profits, for which the Bellman equation is

$$W_j(q_j, q_{-j}, \Delta) = \max_{p_j, x_j} \pi_j(p, q, \Delta) - x_j + \beta \sum_{\tau_j, q'_{-j}, \Delta'} W_j(q_j + \tau_j, q'_{-j}, \Delta') \chi_j(\tau_j | x_j, q) h_{f_j}(q'_{-j} | q, \Delta) g_{f_j}(\Delta' | \Delta, q, p), \tag{12}$$

If a firm advances the quality frontier then Δ' shifts: the second element of Δ' 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 $y = (y_1, y_2, \ldots, y_L)$ as $\Gamma(y) = (y_1 + y_2, y_3, \ldots, y_L, 0)$. If the quality frontier advances at the end of the current period, we shift the interim Δ' in the above equation via $\Gamma(\cdot)$. Hence

$$\Delta'(\Delta, q, p) = \mathcal{I}(\bar{q}' = \bar{q}) \Delta'(\Delta, q, p | \bar{q}' = \bar{q}) + \mathcal{I}(\bar{q}' > \bar{q}) \Gamma\left(\Delta'(\Delta, q, p | \bar{q}' = \bar{q})\right). \tag{8}$$

[18] Pillai (2009) finds innovation in the microprocessor industry depends in part on innovations by upstream manufacturers of semiconductor equipment. We implicitly assume these external forces do not vary over time.

[19] If recent investments have failed to increase quality, the firm is more likely to be a laggard. The spillover therefore mimics, to a degree, the effect of including a state variable for cumulative R&D investments since the previous innovation. Actually including such a state variable significantly raises the computational burden.

where $h_{f_j}(\cdot|\cdot)$ is firm j's beliefs about competitors' future quality levels, and $g_{f_j}(\cdot|\cdot)$ is its beliefs about the transition kernel for Δ, which is based on beliefs about consumers' choices given prices and qualities.

Firms simultaneously choose prices and investments to satisfy the first-order conditions

$$\frac{\partial W}{\partial p_j} = \frac{\partial \pi_j(p,q,\Delta)}{\partial p_j} + \beta \sum_{\tau_j, q'_{-j}, \Delta'} W_j(q_j+\tau_j, q'_{-j}, \Delta') h_{f_j}(q'_{-j}|q,\Delta) \frac{\partial g_{f_j}(\Delta'|\Delta,q,p)}{\partial p_j} \chi_j(\tau_j|x_j,q) = 0 \quad (13)$$

and

$$\frac{\partial W_j}{\partial x_j} = -1 + \beta \sum_{\tau_j, q'_{-j}, \Delta'} W_j(q_j+\tau_j, q'_{-j}, \Delta') h_{f_j}(q'_{-j}|q,\Delta) g_{f_j}(\Delta'|\Delta,q,p) \frac{\partial \chi_j(\tau_j|x_j,q)}{\partial x_j} = 0. \quad (14)$$

Recall the important dynamic trade-off—a higher price today implies more people will be available in the next period to purchase the product. The presence of $\frac{\partial g_{f_j}(\Delta'|\Delta,q,p)}{\partial p_j}$ in $\frac{\partial W}{\partial p_j}$ captures this benefit of raising price and leads to forward-looking firms pricing higher than myopic firms that ignore this dynamic aspect of demand.

3.3 Equilibrium

We consider pure-strategy Markov-Perfect Nash Equilibrium (MPNE) of this dynamic oligopoly game. 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 about the probability of future states.

The equilibrium specifies that (1) firms' and consumers' equilibrium strategies depend only on current state variables, which comprise all payoff-relevant variables, (2) consumers have rational expectations about firms' policy functions, which determine future qualities and prices, and the evolution of the ownership distribution, and (3) each firm has rational expectations about 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, $h_c^*(q'|q,\Delta,\tilde{q}) = \prod_{j=1}^{J} \chi_j(\tau = q'_j - q_j \, |x_j^*, q)$, $h_{f_j}^*(q'_{-j}|q,\Delta) = \prod_{k\neq j}^{J} \chi_j(\tau = q'_k - q_k \, |x_k^*, q)$, and g_c^* and $g_{f_j}^*$ are derived from the law of motion for Δ as described in footnote 17.

In some of the counterfactuals and comparative statics, we impose symmetry, which implies $W_j^* = W^*$, $x_j^* = x^*$, $p_j^* = p^*$, $h_{f_j}^* = h_f^*$, and $g_{f_j}^* = g_f^*$ for all j. Symmetry also requires firm-specific parameters—brand intercepts ξ_j and investment efficiencies $a_{0,j}$—to be the same across firms.

Besanko et al. (2010) and Borkovsky, Doraszelski, and Kryukov (2010) document the existence of multiple equilibria in dynamic oligopoly models based on EP. To reduce multiplicity, we focus on equilibria that are limits to finitely repeated games: we use backwards induction to solve for an equilibrium of the T-period game and then let $T \to \infty$. For each T and for each state, we solve the system of first-order conditions in equations (13) and (14). Our numerical algorithm for computing equilibrium to the infinite horizon game corresponds to value function iteration with (a) initial values of zero for $\bar{V}$ and W and (b) equilibrium strategies being played within each state for each iteration, as opposed to merely playing best responses to strategies from the previous iteration. This refinement yields a unique equilibrium if the subgame within each state at each iteration has a unique equilibrium. Inspections of best-response functions at various states during convergence suggests this refinement indeed yields a unique equilibrium.

We relegate the algorithmic details of computing and simulating the MPNE to Appendix A. One issue worth highlighting is that to evaluate firms' first-order conditions, we must solve a fixed-point in Δ' such that consumers' current beliefs about Δ' match the Δ' in equation (8) that results from the choice probabilities in equation (6).

3.4 Bounding the State Space

Product qualities q_t increase without bound. To numerically solve for equilibrium, we transform the state space to one that is finite by measuring all qualities relative to the current period's maximum quality $\bar{q} = \max(q)$. Our ability to implement this transformation without altering the dynamic game itself hinges on the following proposition.

Proposition 1. *Shifting q and $\tilde{q}$ by $\bar{q}$ has no affect on firms' payoffs and shifts consumers' payoffs in each state by $\frac{\gamma\bar{q}}{1-\beta}$, the discounted value of the reduced utility in each period. More formally,*

$$\begin{array}{lrcl} \textit{Firms:} & W_j(q_j - \bar{q}, q_{-j} - \bar{q}, \Delta) & = & W_j(q_j, q_{-j}, \Delta) \\ \textit{Consumers:} & V(q - \bar{q}, \Delta, \tilde{q} - \bar{q}, \varepsilon) + \frac{\gamma\bar{q}}{1-\beta} & = & V(q, \Delta, \tilde{q}, \varepsilon)\ . \end{array} \tag{15}$$

The proof, which appears in Appendix B, rests on the following properties of the model: (1) log-quality q 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 $\chi_j(\cdot)$, which is independent of quality levels; and (3) Δ 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.

To facilitate writing the value functions in terms of a relative state space, we define $\omega = q - \bar{q}$ and $\tilde{\omega} = \tilde{q} - \bar{q}$ as analogs to the original state variables. We also define the indicator variable $I_{\bar{q}} = 1$ if $\bar{q}' > \bar{q}$ to indicate an improvement in the frontier product. We can then express the consumer's product-specific value function in equation (5) using the relative state space as

$$v_j(\omega, \Delta, \tilde{\omega}) = \gamma\omega_j - \alpha p_j + \xi_j + \beta \sum_{q', \Delta'} \left(\frac{\gamma\delta I_{\bar{q}}}{1-\beta} + \bar{V}(\omega', \Delta', \tilde{\omega}') \right) \, h_c(I_{\bar{q}}, \omega' | \omega, \Delta) \, g_c(\Delta' | \Delta, \omega, p, I_{\bar{q}}), \quad (16)$$

where the outside alternative's p_0 and ξ_0 are zero and, in a slight abuse of notation, $h_c(I_{\bar{q}}, \omega'|\omega, \Delta)$ and $g_c(\Delta'|\Delta, \omega, p, I_{\bar{q}})$ are the analogs of the consumer's transition kernels for q' and Δ' in the original state space. $\frac{\gamma\delta I_{\bar{q}}}{1-\beta}$ is the discounted value of one δ-step of quality each period, which must be explicitly added when an improvement in frontier quality causes $\tilde{\omega}'$ to drop by δ even though $\tilde{q}'$ is unchanged. Since v_j is the product-specific value function, $\tilde{\omega}' = \omega'_j$.

Firm j's value function in equation (12) using the relative state space becomes

$$\begin{aligned} W_j(\omega_j, \omega_{-j}, \Delta) \quad = \quad & \max_{p_j, x_j} \; \pi_j(p, \omega, \Delta) - x_j + \beta \sum_{\tau_j, \, \omega'_{-j}, \, I_{\bar{q}}, \, \Delta'} W_j(\omega_j + \tau_j - I_{\bar{q}}, \omega'_{-j} - I_{\bar{q}}, \Delta') \\ & h_{f_j}(I_{\bar{q}}, \omega'_{-j}|\omega, \Delta) \; g_{f_j}(\Delta'|\Delta, \omega, p, I_{\bar{q}}) \; \chi_j(\tau_j|x_j, \omega) \; , \end{aligned} \quad (17)$$

where ω'_{-j} 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 $h_{f_j}(I_{\bar{q}}, \omega'_{-j}|\omega, \Delta)$, $\chi_j(\tau_j|x_j, \omega)$, and $g_{f_j}(\Delta'|\Delta, \omega, p, I_{\bar{q}})$ as the analogs of the firm's transition kernels for competitors' qualities and Δ'.

Finally, we invoke a knowledge-spillover argument to bound the difference between each firm's quality and the frontier quality. We denote the maximal difference in firms' qualities $\bar{\delta}_f$ and modify the transition kernels $\chi_j(\cdot)$ and $h_{f_j}(\cdot)$ accordingly. We choose $\bar{\delta}_f < \bar{\delta}_c$ since, in most markets, quality differences among new products are less than the quality gap between the frontier and products from which consumers have yet to upgrade. We also choose $\bar{\delta}_f$ to be sufficiently large that firms never reach the bound in equilibria computed during estimation. Note that if firms were permitted to exit, quality differences would be bounded automatically by the exiting of firms with sufficiently low relative quality.

Our bounding approach differs from the EP framework for differentiated products, as detailed in Pakes and McGuire (1994) and Doraszelski and Pakes (2007).[20] In EP, the industry's long-run innovation rate is solely determined by the exogenous innovation rate of the outside good. Improvements in the outside good provide a continual need for inside firms to invest to remain competitive. If the outside good never improves, the equilibrium has no investment and no innovation in the long run. In our model, the long-run rate of innovation is an equilibrium outcome that depends on consumer preferences, firms' costs, and the regulatory environment.

[20] The standard normalization in discrete choice models subtracts the mean utility of the outside good from all options. EP, however, subtracts the outside good's quality from firms' qualities inside a concave function. Concavity implies the derivative of market share with respect to a firm's own quality goes to zero regardless of competitors' qualities. Since investment is costly, a relative quality above which investment is zero will exist, thereby establishing an upper bound. Firms exit when relative quality gets sufficiently low, which establishes the lower bound.

In essence, EP 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 by frontier firms.

4 Empirical Application

This paper has two components: a theory component that develops a dynamic oligopoly model with durable goods 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 brand fixed-effects and costs of production and innovation. In section 5, we present comparative statics for the symmetric case in which firms have identical brand intercepts and innovation efficiencies, to illustrate broader implications of the model.

4.1 Estimation

We estimate the cost parameters $\lambda = (\lambda_0, \lambda_1)$ in equation (11) in a first stage using linear regression, yielding $\hat{\lambda}$. To estimate the dynamic parameters $\theta = (\gamma, \alpha, \xi_{Intel}, \xi_{AMD}, a_{0,Intel}, a_{0,AMD}, a_1)$, 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) in the data. The simulated minimum distance estimator $\hat{\theta}_T$, which we detail in Appendix C, is

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

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. Hence, the efficient weight matrix is the inverse of the covariance matrix of the actual data's moments. We use 10,000 bootstrap replications to estimate this covariance matrix. Since we obtain the efficient weight matrix directly from the data, we do not need a two-step GMM estimator to obtain efficiency.

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. Their assumption that the data arise 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 approach. As discussed in section 3.3, we only consider equilibria that are limits of finite horizon games to reduce the concern of multiple equilibria.

4.1.1 Moments to Match

We match a combination of simple moments and coefficients from linear approximations to firms' policy functions. 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. We therefore choose moments that are stationary in both the data and the model. For example, we match investment per unit revenue, which is stationary in the data, instead of the trending investment levels.

Our moment vector, m_T, consists of the following 15 moments:

- average prices and the coefficients (other than the constant) from regressing each firm's price on a constant, $q_{Intel,t} - q_{AMD,t}$, and $q_{own,t} - \bar{\Delta}_t$, where $\bar{\Delta}_t = \sum_{k=\underline{q}_t}^{\bar{q}_t} k\Delta_{kt}$ is the mean log-quality currently owned in period t,

- coefficients from regressing Intel's share of sales on a constant and $q_{Intel,t} - q_{AMD,t}$,

- mean $(\bar{q}_t - \hat{\Delta}_t)$, where $\hat{\Delta}_t$ is the same as $\bar{\Delta}_t$ except non-owners are excluded. This moment captures the rate at which consumers upgrade: if consumers upgrade quickly, all else equal, the average difference between $\bar{q}_t$ and $\hat{\Delta}_t$ will be low.

- mean innovation rates for each firm, defined as $\frac{1}{T}(q_T - q_0)/\delta$,

- mean $(q_{Intel,t} - q_{AMD,t})$ and share of quarters with $q_{Intel,t} \geq q_{AMD,t}$, and

- mean investment per unit revenue for each firm.[21]

Recall that q and Δ measure log-quality which implies quality differences are proportional. These moments and their fitted values appear in Table 1.

[21] 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, its microprocessor unit delivered 87 percent of its consolidated net revenue.

4.1.2 Identification

Experimentation with the structural model reveals the moments we seek to match are sensitive to the structural parameters. Being a nonlinear model, all the structural parameters influence all the moments, though the connections between some parameters and moments are more direct.

The demand-side parameters $(\alpha, \gamma, \xi_{Intel}, \xi_{AMD})$ are primarily identified by the pricing moments, the Intel share equation moments, and the mean ownership quality relative to the frontier quality. The pricing moments respond sharply to changes in any of these four parameters. 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 akin to an outside share equation that identifies the levels of ξ. We interpret ξ_{Intel} as a hassle cost of upgrading one's computer 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 difference in quality across firms—a high spillover keeps the qualities similar.

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. Consequently, our conclusions regarding the effect of competition on innovation rely on estimating the costs and benefits of innovation, as determined by the structural parameters governing supply and demand.

One could consider variation in firms' relative qualities as a form of market structure variation and investigate its relationship with innovation. In our data, innovation since the previous quarter is positively related to that quarter's difference in firms' qualities. We do not use these moments, however, since the p-values are .12 and .18 for AMD and Intel, respectively. We note in our discussion of firms' policy functions in section 4.2.1, however, that the innovation policies exhibit this same positive correlation.

4.1.3 Estimates

We use the simulated minimum distance estimator in equation 18 to estimate the dynamic parameters θ given the first-stage marginal cost estimates $\hat{\lambda}$. We first fix a few model-setup parameters. We set δ to .1823, which yields quality gains of 20 percent between rungs on the quality ladder. We set $\bar{\bar{\delta}}_c$ to 5.287, which corresponds to a maximum of 29 δ-steps between consumers' $\tilde{q}$ and the frontier. Our choice of δ and $\bar{\bar{\delta}}_c$ reflects the following considerations: (i) the ability to replicate "Moore's Law" when firms innovate in 40 to 60 percent of the periods, (ii) a sufficiently high $\bar{\bar{\delta}}_c$ that consumers rarely reach the lowest grid point before upgrading, and (iii) computation time. We choose $\bar{\bar{\delta}}_f$ to be eight δ steps, which exceeds the observed maximum quality difference of 5.2 δ-steps. Since our quantity data are quarterly and firms' pricing and product releases are roughly quarterly, we assume each period is three months and set β to .975. We set the market size M to 400 million consumers, such that the model's implied market capitalizations for Intel and AMD are similar to their observed values.

The market size for microprocessors is arguably growing over time as new computer applications are developed and as complimentary components (e.g., memory, disk drives, and monitors) become better and cheaper. Market expansion corresponds to adding new consumers with vintage $\underline{q}$ and increasing M accordingly. Unfortunately, increasing M results in a non-stationarity that is computationally burdensome. Instead, we adjust Δ' to reflect the composition effect of market expansion by adding a mass of consumers, equal to 2.6 percent of M, to the lowest vintage in Δ' and re-normalize to maintain a fixed M. This arrival rate matches the average quarterly growth from 1993 to 2003 in computer ownership by U.S. households according to the U.S. Census Current Population Survey (CPS) Computer Ownership Supplement. The high demand from a mass of consumers with $\tilde{q} = \underline{q}$ in each period raises equilibrium prices and, since inducing upgrades becomes less critical for sustained demand, lowers innovation rates.

We report the model's fit in Table 1 and the parameter estimates in Table 2. The model fits the 15 moments reasonably well, despite having only seven parameters. As is typical with structural econometric models, the data formally reject our model using a J-stat test since the real world is too complicated for a tractable model to mimic perfectly.

Table 2 provides the structural estimates and their standard errors. All the parameters are statistically significant given the relatively small asymptotic standard errors. Dividing the estimated quality coefficient by the price coefficient implies consumers are willing to pay \$21 for a δ increase in log-quality per period, which translates, for example, to \$51 for a 20 percent faster CPU to be used for 16 quarters ($\delta\gamma(1-\beta^{16})/[(1-\beta)\alpha]$). Dividing Intel's fixed effect by the price coefficient implies upgrading to a new computer is associated with a hassle cost of \$48. Dividing $\xi_{Intel} - \xi_{AMD}$ by the price coefficient implies consumers are willing to pay \$194 for the Intel brand over the AMD brand. The model needs this strong brand effect to explain the fact that AMD's share never rises above 22 percent in the period during which AMD had a faster product. Intel and AMD's innovation efficiencies are estimated to be .001 and .0019, respectively, as needed for AMD to occasionally be the technology leader while investing much less. Intel's price elasticity for current sales with respect to an unexpected one-period price change is 2.16, compared to 1.77 for AMD. These elasticities are lower than the range reported in Prince (2008) for PC purchases, perhaps reflecting the importance of the CPU to the PC's performance.

4.2 Empirical Results

We use the baseline parameter estimates to compare seven industry scenarios in Table 3: (1) AMD-Intel 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 all the estimates in column 1 of Table 2. 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 price myopically by solving $\frac{\partial \pi_j(p,q,\Delta)}{\partial p_j} = 0$ instead of the dynamic first-order condition in equation (13). 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 quickly transitions to states with investment in only the frontier.

For each scenario, we solve for optimal policies and simulate 10,000 industries each for 300 periods, starting from the initial state in our data. We then analyze the simulated data to characterize the equilibrium behavior of firms and consumers and to identify observations of interest. Finally, counterfactual experiments illustrate implications of the model for policy analysis.

Our characterizations of equilibrium behavior in sections 4.2.1 and 4.2.2 instill confidence that the model yields sensible outcomes. We set apart findings of interest as "observations" in section 4.2.3.

4.2.1 Firm Behavior in Equilibrium

Figure 2 presents value functions, pricing, innovation, market shares, and period profits for the monopoly and the symmetric duopoly at select states. Figure 3 presents these equilibrium outcomes for the same monopoly and for the symmetric duopoly without spillovers. We evaluate the duopoly without spillovers as a theoretical exercise to illustrate the properties of the model. We suspect most industries exhibit some degree of innovation spillovers and explore further the effect of spillovers in section 5.

We present the symmetric duopoly case, for which leader-laggard policy differences reflect only quality differences (not different firm-specific parameters). In both figures, the x-axis in the first two columns of plots is the ownership distribution state variable $\bar{\Delta}$. Demand is high when consumers' average quality $\bar{\Delta}$ is low. Accordingly, value functions and prices both decline as $\bar{\Delta}$ increases for the monopolist (in column 1) and the duopolists (in column 2). In the second column, outcomes are separately presented for the leader and laggard when their qualities differ by 4δ and for the firms when they are tied. As expected, values, prices, and market shares are highest for the leader and lowest for the laggard, with the tied firms in between.

In the third column of both figures, we fix Δ at its most frequent value in simulations of the symmetric duopoly with spillovers and vary the leader's quality advantage on the x-axis. The laggard is 8δ behind at the leftmost value and tied for the lead at the rightmost value. Accordingly, the leader's value function, prices, and shares decline as its advantage shrinks, whereas the laggard's value, prices, and shares increase as it catches up to the leader.

The value functions, prices, market shares, and period profits match our intuition. The outcome of greatest interest is the innovation rate. As the ownership distribution becomes newer, the monopolist slightly increases innovation, whereas the duopolists slightly decrease innovation, both with and without the spillover. In the duopoly, returns to investment are driven more by business stealing than by the building of future demand. The business-stealing motive is greater for duopolists when consumers are primed to upgrade, as indicated by a low Δ.

Innovation by the 4δ leader and tied firms are much higher in Figure 3 without spillovers than in Figure 2 with spillovers. The reason is that without spillovers, the laggard struggles to catch up and indeed gives up completely once he falls behind by 7δ. Equilibrium is thus characterized by high innovation initially as firms battle to be the reigning leader, after which innovation drops to zero for the laggard and below .5 for the leader. The innovation plot in the last column of Figure 3 depicts this storyline. Without spillovers, the leader increases innovation as the laggard gains, peaking when the laggard is δ behind. With spillovers, the leader decreases innovation as the laggard catches up and firms take turns being leaders. Losing the current battle does not permanently lower profits when spillovers enable a return to leadership, which reduces the incentive to fight. With spillovers, the difference in value functions between the leader and laggard is $75 billion, compared to over $300 billion without spillovers, despite the similar differences in period profits reported in the bottom row of each figure.

The differences in innovation policies with and without spillovers yield dramatically different distributions of states visited, as depicted by the histograms in the top rows of figures 2 and 3. Without spillovers, the firms tend to be at their maximal degree of differentiation; with spillovers, they tend to be tied or off by one step. The ownership distributions encountered also differ: consumers tend to own older vintages without spillovers since they upgrade less often in response to the leader's higher prices (given its large quality advantage).

The change in period profits when a firm's relative quality changes by one step—the x-axis in the bottom-right panel—represents the immediate impact of innovation on a firm's net cash flow. The substantial difference in innovation with and without spillovers, despite the similar immediate effect on profits, suggests innovation is driven primarily by long-run considerations.

4.2.2 Consumer Behavior in Equilibrium

In Figure 4 we plot the choice probabilities at each ownership vintage, averaged across states encountered in the AMD-Intel duopoly. The lower a consumer's vintage relative to the frontier, the more likely she is to upgrade. As reported in Table 3, when consumers upgrade in the AMD-Intel duopoly, the average improvement in quality is 261 percent, compared to 410 percent in the monopoly.

As consumers implement their policy functions, they generate a sequence of ownership distributions across time. Figure 5 depicts the average ownership distributions for the AMD-Intel duopoly and the monopoly. Because monopolists charge higher prices, consumers are less likely to upgrade from a given vintage to the frontier in the monopoly case. In addition, consumers in the duopoly usually have an option to upgrade to a non-frontier product. Both of these forces cause the ownership distribution to be older in the monopoly. Figure 5 also suggests that consumers rarely reach the lowest vintage. Indeed, for the oldest distribution encountered in the simulations, only .00001 percent of consumers are at the lowest vintage, which ensures the lower bound has no effect on equilibrium behavior.

4.2.3 Observations Specific to the Microprocessor Industry

Having established the sensibility of consumers' and firms' policy functions, we now compare the estimated model with counterfactual models of the microprocessor industry. Here we evaluate the model and counterfactuals at the parameter estimates, whereas in section 5 we present comparative statics to more broadly assess the model's implications for the effect of competition on outcomes across industries characterized by different consumer preferences, depreciation rates, and innovation spillovers.

We first assess the importance of accounting for the dynamic nature of demand by comparing outcomes when Intel and AMD price myopically by solving $\frac{\partial \pi_j(p,q,\Delta)}{\partial p_j} = 0$ instead of the dynamic first-order condition in equation (13).

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

In Table 3 monopoly profits are 76 percent higher and margins are 156 percent higher when the monopolist accounts for the dynamic nature of demand (scenario 3), compared to myopic pricing (scenario 6). Industry profits for the AMD-Intel duopoly (scenario 1) are 28 percent higher and margins are 58 percent higher when firms account for the dynamic nature of demand, compared to myopic pricing (scenario 5). These higher margins induce firms to innovate more rapidly: the duopoly innovation rate is 34 percent higher with optimal (dynamic) pricing and the monopoly innovation rate is 42 percent higher.

Accounting for dynamic demand is more important for the monopolist because competition is solely with itself, whereas the duopolists are primarily concerned with each other. Moreover, duopolists are less concerned about the effect of current pricing on future demand since future demand is a shared resource.

This result highlights the importance of accounting for dynamic demand when analyzing the pricing of durable goods. Standard practice in the empirical industrial organization and marketing literatures is to observe prices and use first-order conditions from a static profit maximization to infer marginal costs. Observation 1 suggests marginal cost estimates computed in this manner for durable goods will be too high. Prices are high, in part, because firms want to preserve future demand, not only because marginal costs are high. Since the incentive to preserve future demand is increasing in market concentration, this over-estimation of costs will be greatest for concentrated markets.

In the next three observations, we compare market outcomes under alternative market structures for the microprocessor industry. The monopoly counterfactual corresponds to a world in which AMD never existed, not a world in which Intel merges with AMD, since no such merger would ever be pursued. As such, the monopolist sells and invests in one product, not two.

Observation 2. *Regarding the effect of competition on innovation in the CPU industry, we find*

i. The rate of innovation in product quality is 4.2 percent higher with a monopoly than with the AMD-Intel duopoly. The difference is more pronounced when comparing the monopoly to a symmetric duopoly pitting Intel against another Intel, with or without spillovers.

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

iii. In the counterfactuals with firms pricing myopically, the AMD-Intel duopoly innovates more than monopoly.

The average industry investment levels, reported in millions of dollars in Table 3, for the duopoly, monopoly, and social planner are, respectively, \$830 million, \$1,672 million, and \$6,672 million per period. The resulting innovation rate for the industry's frontier product is 0.599 for the duopoly, 0.624 for the monopoly, and 0.869 for the planner. The symmetric duopoly's innovation rate is only 0.501 because the intense price competition when Intel faces a copy of itself reduces investment returns.

The finding that innovation by a monopoly exceeds that of a duopoly reflects two features of the model: the monopoly must innovate to induce consumers to upgrade, and the monopoly is able to extract much of the potential surplus from these upgrades because of its substantial pricing power. If competition with itself were reduced by a steady flow of new consumers into the market, the monopoly would reduce innovation below that of the duopoly. We illustrate this result with a comparative static in the next section.

Observation 2.iii illustrates the importance of correctly accounting for durability when evaluating incentives to innovate, since the effect of competition on innovation is reversed when firms (or researchers) do not account for the dynamic nature of demand for durable goods.

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 other's innovations. As we report in Table 3, 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 innovations cannot be copied. The direct effect, however, is dominated by the equilibrium effect of one firm eventually dominating the industry, as evidenced by the innovation policies and histogram of q difference in Figure 3. The absence of a threat from the weak laggard, who eventually gives up and stops innovating, induces the leader to reduce investment. The presence of the laggard nonetheless keeps margins lower than in the monopoly. Thus, market power, not an absence of spillovers, provides the incentive for rapid innovation by the monopolist, compared to the duopolist. In the next section, we consider spillovers of varying degrees between the estimated level and no spillover.

Importantly, our model yields higher innovation with competition when evaluated using different values for price and quality preferences. As such, our model indeed lets the data speak on this fundamental question. In the next section, we present comparative statics to show that competition fosters higher innovation when consumers highly value quality and are relatively insensitive to price.

Of course, policymakers are more concerned with surplus and profits than with innovation *per se*. We compute firms' profits as the discounted sum of per-period profits and consumer surplus directly from the value functions: $CS = \frac{M}{\alpha} \sum_{\tilde{q}=\underline{q}_t}^{\bar{q}_t} \bar{V}(q_0, \Delta_0, \tilde{q}) \cdot \Delta_{\tilde{q},0}$. We 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.

Observation 3. *Regarding the effect of competition on surplus, we find*

i. The AMD-Intel duopoly generates 4.2 percent more consumer surplus than the monopoly.

ii. The AMD-Intel duopoly generates 1 percent less social surplus than the monopoly. The duopoly and monopoly generate 92.9 and 94 percent, respectively, of the planner's social surplus.

iii. Consumers' share of social surplus is 88 percent in the AMD-Intel duopoly, compared to 83.4 percent in the monopoly.

Table 3 reports the aggregate discounted CS and industry profits for each of the scenarios we consider.[22] The AMD-Intel duopoly CS of \$2.98 trillion corresponds to \$298 billion per year, using an annual discount factor of .9. Although both the AMD-Intel and symmetric duopolies generate more CS than the monopoly, higher industry profits enable the monopolist to generate more social surplus than the duopolies.

[22]The compounding effect of the monopoly's higher innovation rates implies the consumer surplus gain in duopoly relative to monopoly is larger the shorter the time horizon. Using the 48-quarter horizon of our data, the gain in CS when moving from the monopoly to the AMD-Intel duopoly is 7.1%, instead of 4.2% using 300 quarters.

As noted in observation 3.iii, CS comprises more than 83 percent of social surplus whether the industry is a monopoly or duopoly. Moreover, consumers are the primary benefactors of innovation opportunities, regardless of market structure, as evidenced by comparisons with (unreported) counterfactuals in which firms are barred from innovation. Monopoly profits are 2.6 percent higher with innovation than without innovation, whereas CS in the monopoly is 64.2 percent higher with innovation. Duopoly profits are actually 13 percent lower with innovation, whereas CS in the duopoly is 65.7 percent higher with innovation.

To put the 4.2 percent CS gain due to competition from AMD in perspective, the CS gain from an increase in frontier quality by one δ-step is \$55.4 billion.[23] The \$121 billion higher CS under duopoly, compared to monopoly, therefore equals the CS gain from 2.2 innovations, which is roughly one year's worth of innovations (under either monopoly or duopoly). Again, we see that the difference in CS between duopoly and monopoly are small relative to the overall gains from innovation.

Recently Intel paid AMD \$1.25 billion to settle claims that Intel's anti-competitive practices foreclosed AMD from many consumers. 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 Intel has exclusive access. Let ζ denote this portion. Period profits for j = Intel are then

$$\hat{\pi}_j(p,q,\Delta) = M\left[\zeta \hat{s}_j(p_j,q_j,\Delta) + (1-\zeta)s_j(p,q,\Delta)\right](p_j - mc_j(q)), \tag{19}$$

where $\hat{s}_j(p_j,q_j,\Delta)$ is Intel's market share in the sub-market in which it competes only with the outside good (i.e., Δ). We assume Intel sets the same price in each sub-market and consumers are randomly assigned to each market each period.

Observation 4. *As AMD is excluded from an increasing portion of the market,*

i. margins monotonically rise and innovation exhibits an inverted-U with a peak at $\zeta = .5$,

ii. consumer surplus rises initially, peaking at $\zeta = .4$, then declines, eventually falling below the consumer surplus with no foreclosure.

In Figure 6 we plot margins, innovation rates, consumer surplus, and social surplus when the foreclosed portion of the market varies from zero to one. Not surprisingly, share-weighted margins rise monotonically as AMD is increasingly barred from the market. Industry innovation peaks at 4.8 percent higher than the estimated AMD-Intel duopoly innovation rate when AMD is barred from half the market, but then drifts down to the 4.2 percent higher innovation of the monopoly. Consumer surplus is actually higher when AMD is barred from a portion of the market, peaking at 40 percent

[23] We evaluate the gain in CS from Intel advancing when AMD is two steps behind and Δ is at its most common value. The gains are less than the upper-bound $M\delta\gamma/(\alpha(1-\beta))$ =\$61.5 billion since not all consumers upgrade immediately to the improved product. The gain converges to \$61.5 billion as the Δ distribution shifts to older vintages.

foreclosure. Although the CS gains are small, this finding highlights the importance of accounting for innovation in antitrust policy: the decrease in consumer surplus from higher prices can be more than offset by the compounding effects of higher innovation rates.

5 Comparative Statics

We present comparative statics in preferences for price and quality, depreciation of the good's quality, the magnitude of innovation spillovers, and product substitutability. In addition to being of interest themselves as characterizations of outcomes for a wide array of durable-goods markets, these results assess the robustness of our earlier findings specific to the microprocessor industry. We also relate our findings to Aghion et al. (2005).

5.1 Comparative Static in Consumer Preferences

Our primary empirical result is that Intel would innovate more if it were not competing against AMD. We now illustrate in Figure 7 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 5. *Comparative statics in the quality and price coefficients, γ and α, reveal*

i. *Innovation is increasing in γ and decreasing in α for both monopoly and duopoly.*
ii. *The effect of competition on innovation is increasing in γ and decreasing in α, except where both γ is low and α is high, as Figure 7 depicts.*
iii. *Innovation is higher for the duopoly than the monopoly when γ is high and α is low.*
iv. *Consumer surplus is higher for the monopoly than the duopoly when γ is low and α is high. This result illustrates again that higher innovation in the monopoly can more than offset higher prices, yielding higher consumer surplus than obtained in the duopoly.*

We first note that our estimates of .2764 and .0131 for γ and α in the microprocessor industry are far from the region of preferences for which competition increases innovation.

Part (i) is intuitive: firms innovate faster when consumers are willing to pay more for quality, either due to a higher coefficient on quality or a lower coefficient on price. We plot these monotonic relationships in the top two panels of Figure 7 for the duopoly and monopoly, respectively.

Part (ii) is less obvious. The higher slope of the duopoly in the top panel relative to the monopoly in the second panel implies the duopoly eventually innovates more than the monopoly as γ increases and α falls, both of which increase consumers' willingness to pay for quality. But why is duopoly innovation more sensitive to preferences than is monopoly innovation? To gain insight, first express average industry innovation in the duopoly as a weighted average:

$$Pr(q_{1t} = q_{2t})\, \mathrm{E}\left[1 - (1 - \chi_j(\tau_1 = 1|x^*_{1t}(q_t, \Delta_t)))^2\right] + \sum_{k=1}^{\delta_f} Pr(q_{2t} = q_{1t} - k)\, \mathrm{E}\left[\chi_j(\tau_1 = 1|x^*_{1t}(q_t, \Delta_t)\right], \tag{20}$$

where E integrates over Δ, $Pr(\cdot)$ are probability weights, and firms are ordered within each period such that q_{1t} is the frontier quality. The first term reflects the mechanical benefit of having two firms: when firms are tied, innovation by either firm advances the frontier. By comparison, the average monopoly innovation rate is $\mathrm{E}\left[\chi_j(\tau_1 = 1|x^*_{1t}(\Delta_t))\right]$. Differentiating these innovation rates with respect to γ (or similarly α) reveals that two factors contribute to innovation being more sensitive to preferences in the duopoly than in the monopoly: the effect of γ on firms' investment policies $x^*(\cdot)$ and on which states are encountered (i.e., the probability weights). As detailed in Appendix D, both channels lead to the duopoly increasing innovation faster than the monopoly as γ increases (or, similarly, as α decreases). Hence, innovation in the duopoly eventually exceeds innovation in the monopoly as consumers are willing to pay more for quality.

5.2 Comparative Static in Depreciation

Although quality does not depreciate in our empirical application to microprocessors, augmenting the model to accommodate depreciation is easy, as detailed in Appendix E. Figure 8 presents a comparative static relating depreciation and innovation.[24]

Observation 6. *As depreciation increases:*

i. *Innovation declines faster in the duopoly than in the monopoly.*
ii. *Margins increase faster in the monopoly than in the duopoly.*
iii. *Consumer surplus declines faster in the monopoly than in the duopoly.*
iv. *Discounted profits increase faster in the monopoly than in the duopoly.*

Two forces affect equilibrium behavior when depreciation increases. First, the ownership distribution ages more quickly, which reduces the need for firms to innovate to induce upgrade purchases. Second, consumers expect to use each purchase over fewer periods (since they upgrade more quickly), which reduces the discounted utility derived from each purchase. These forces have opposing effects on prices, with the latter effect shifting the pricing policy function lower and the former effect lowering the Δ values at which we evaluate the policy function. The higher prices from lowering Δ dominates

[24] The highest depreciation rate we consider in Figure 8 is a 20 percent reduction (one δ-step) per quarter. Since we assume the outside good's quality is within $\bar{\delta}_c$ of the frontier, a monopolist selling a good with 100 percent depreciation each period (i.e., a nondurable) will stop innovating upon reaching this bound. Clearly, with nondurables, the bound can affect equilibrium strategies, as discussed in Goettler and Gordon (2011).

the shift in the pricing policy function, causing a moderate net increase in prices as depreciation increases. The margins plotted in Figure 8 reflect this rise.

In both the monopoly and duopoly, discounted lifetime profits increase by greater proportions than the margins increase since firms sell more units and investment costs decline. Consumer surplus falls in both the monopoly and duopoly, since consumers pay higher prices for goods that are less durable and of lower quality (since innovation declines). The faster decline of consumer surplus in the monopoly implies the surplus gain from competition is increasing in the depreciation rate. Since depreciation reduces durability, this result demonstrates the role of durability in limiting welfare losses from market power.

One might expect duopoly innovation to decline by less than monopoly innovation as depreciation increases, since the monopolist faces competition only from the durability of its own products, whereas duopolists face competition from past units sold as well as each others' current offerings. The difference in discounted utility derived from competing offerings that differ in quality by one step, however, shrinks as the unit's expected time in use declines due to depreciation. This lower difference implies a reduced competitive gain from innovation, which reduces the business-stealing incentive to innovate in the duopoly.

5.3 Comparative Static in the Market Growth Rate

Observation 7. *As M grows due to entry by new consumers, innovation increases in the duopoly and does so at a faster rate than in the monopoly.*

We present the market growth comparative static in two steps since computing the equilibrium when the market grows each period is computationally impractical due to the resulting nonstationarity (see section 4.1.3). In the top panel of Figure 9, we increase the proportion of consumers who enter the market each period with $\tilde{q} = \underline{q}$, while re-normalizing market size to keep M fixed. This reallocation of consumers to $\underline{q}$ reduces competition from past sales (i.e., Δ) in both monopoly and duopoly. Since competition from past sales is the monopolist's *only* competition, innovation in the monopoly decreases. Duopolists, on the other hand, are primarily concerned with competition from each other since their qualities tend to be closer to each other than to consumers' vintages. Accordingly, duopolists increase innovation in response to the increased demand from the reallocation.

In the lower panel, we present the comparative static for equilibrium innovation as M increases. Innovation increases faster in the duopoly than in the monopoly. Combining the two effects, innovation increases in the duopoly as M grows due to entry by new consumers, and does so at a faster rate than in the monopoly. Since the two components of market growth by entry of new consumers have opposing effects on monopoly innovation, the sign of the net effect in the monopoly is not obvious. Regardless, for sufficiently high market growth due to entry by new consumers, innovation is higher in the duopoly than in the monopoly.

Increasing the market growth rate is similar to increasing depreciation in that both cause the ownership distribution to age faster. Increasing the market growth rate, however, does not lower the expected utility derived over each purchase's useful life, as occurs with increased depreciation. The competitive gain from innovation, due to consumers deriving higher utility over many periods when purchasing the better of two products, is therefore not lowered by market growth.

5.4 Comparative Static in the Innovation Spillover

To investigate further the effect of spillovers, initially noted in our discussion of the policy plots in figures 2 and 3, we present, in Figure 10, symmetric-duopoly outcomes when the spillover effect varies from its estimated value (in the AMD-Intel duopoly) to no spillover.

Observation 8. *As the spillover declines from its estimated value to zero,*

i. Innovation, consumer surplus, and social surplus steadily rise until the spillover is sufficiently small that a severely lagging firm concedes the market by ceasing to innovate.[25] *At this point, innovation plunges and both surplus measures decline.*

ii. Margins and the difference in firms' qualities increase gradually at first and then sharply when the spillover is sufficiently small that the laggard concedes the market by ceasing to innovate.

iii. For moderate spillovers (30 to 40 percent of the estimated spillover), the duopoly has innovation, consumer surplus, and social surplus exceeding those in the monopoly.

One might expect the surplus measures with no spillover to be lower than with the estimated spillover (of $a_1 = 3.94$) since margins are 43 percent higher and innovation is 13 percent lower with no spillover. However, these measures are averages over the 300 periods simulated for each of the 10,000 simulated industries, and the no-spillover duopoly is initially a fierce battle for supremacy. During this initial period, innovation is extremely high and margins are relatively low, since the firms have similar qualities. Hence the early periods, which receive greater weight in the discounted sum of utility flows, deliver substantial discounted surplus. That is, the surplus is heavily front-loaded in the no-spillover industry.

For comparison, the monopolist's values relative to the estimated-spillover symmetric duopoly are 2.34 for margins, .949 for consumer surplus, 1.004 for social surplus, and 1.246 for innovation. The peak duopoly innovation rate, when a_1 is 40 percent of the estimated spillover, exceeds innovation in the monopoly by 11.2 percent. Consumer surplus and social surplus are maximized by the duopoly with 35 percent of the estimated spillover, yielding surplus gains relative to the monopolist of 20 percent and 10 percent, respectively.

[25] Although the laggard ceases to innovate, it remains in the industry at the maximum quality disadvantage $\bar{\delta}_f$. The laggard's presence restricts the leader's market power, thereby reducing its innovation below that of a monopolist.

Scott (1984) and Levin et al. (1985) find that differences in technological opportunity and appropriability conditions explain much of the variation in R&D intensity across industries. Since innovation spillovers reduce a firm's ability to appropriate profits from its innovations, Observation 8 supports their finding.

The inverted-U relationship between spillovers and innovation, combined with the monotonic relationship between spillovers and margins, implies an inverted-U relationship between margins and innovation. Thus innovation spillovers may be added to the list of possible structural causes for the inverted-U relationship between margins and innovation identified by Aghion et al. (2005).

5.5 Comparative Static in Product Substitutability

Product substitutability is governed by the variance of the idiosyncratic utility shock ε. As $1/var(\varepsilon)$ approaches zero, firms enjoy local monopolies with no product substitutability. As $1/var(\varepsilon)$ approaches infinity, products become perfect substitutes and the market eventually yields a "winner-take-all" outcome. Figure 11 depicts market outcomes as we vary product substitutability.

Observation 9. *Regarding the effect of product substitutability on innovation, we find*

ii. *Innovation in the monopoly exhibits an inverted-U as substitutability increases.*

i. *Innovation in the duopoly increases as substitutability increases, until $var(\varepsilon)$ becomes too small for firms with similar qualities to coexist. Beyond this "shakeout" threshold, the laggard eventually concedes the market as evidenced by the sharp increase in the quality difference.*

iii. *Duopoly innovation is higher than monopoly innovation when substitutability is near the shakeout threshold.*

Vives (2008) concludes that duopolists facing logit demand increase product innovation when substitutability increases, which matches Observation 9.i until the shakeout threshold is reached. Vives (2008) also shows R&D on cost-reducing technologies is unaffected by product substitutability in the logit model for nondurable goods.

5.6 Relating to Aghion et al. (2005)

In the model of Aghion et al. (2005), a monopolist would never innovate. As such, Aghion et al. (2005) vary competition not by the number of firms, but by the degree to which duopolists collude when tied. As competition increases, due to less collusion, a firm with inferior technology decreases investment and firms at the same technology level increase investment. The former is the Schumpeterian effect since the lower profits in the tied state reduce the laggard's incentive to innovate, whereas the latter effect is the "escape-the-competition" effect since each tied firm increases investment as the profit gap widens between a tied firm and a leader firm. When competition is low, the industry spends more time in the tied state since the laggard innovates rapidly, leading to a tied state, and tied firms innovate slowly. Consequently, the escape-the-competition effect dominates when competition is low, which causes an increase in competition to increase average industry innovation over time. When competition is high, the industry tends to be in the unlevel state, leading the Schumpeterian effect on laggard's innovation to reduce average industry innovation when competition increases further.

Our model with durable goods differs from the nondurable-goods model of Aghion et al. (2005) in several ways. Their firms' technology levels differ by at most one innovation step, whereas our firms can differ by multiple steps. Industry profits in their model depend only on technology differences, which implies a firm with a quality advantage will never innovate. A leader in our model, however, innovates to increase its quality advantage, relative to both its competitor and to the stock of used durables.

To better relate to Aghion et al. (2005), we transform our model to use nondurable goods and consider a measure of competition more similar to their ability-to-collude measure: product substitutability as measured by the variance of the idiosyncratic utility shock ε.[26] More intense competition via greater product substitutability (i.e., lower variance of ε) raises the profit gain when a tied firm becomes a leader, and reduces the profit gain when a laggard becomes tied for the lead. This pattern matches the effect of increased competition via less collusion. As illustrated in Figure 12, this nondurable version of our model yields Proposition 1 of Aghion et al. (2005), which states that innovation by the laggard decreases with greater competition, and innovation by the tied (neck-in-neck) firms increases. The declining share of periods in which firms are tied also matches their Proposition 4, which states that the expected technological gap between firms increases with competition. Our nondurable model, however, does not generate the inverted-U of their Proposition 2, since the leader in our model innovates faster as competition increases, whereas the leader in their model never innovates.

We apply this same measure of competition to our durable goods model and find that their Proposition 1 no longer holds: both the laggard and tied firms increase innovation when product market competition increases. In short, we show that their Proposition 1 depends on whether the good is durable and their Proposition 4 depends on whether the leader invests.

6 Conclusion

In this paper, we estimate a dynamic model of durable goods oligopoly with endogenous innovation and use it to assess the effect of competition on innovation in the PC microprocessor industry. Consumers are better off under a duopoly due to lower margins: consumer surplus is higher with

[26] In the nondurable model, consumers are myopic and there is no ownership distribution. We also assume no outside good is available since constant industry revenue in Aghion et al. (2005) implies firms essentially compete only within an industry. For Figure 12, firms differ by no more than one quality step, though the results are qualitatively the same when greater differences are allowed.

AMD competing against Intel than without AMD. However, in support of Schumpeter's hypothesis, industry innovation is higher with Intel as a monopolist.

Two forces drive innovation: competition between firms for the technological frontier and competition with the installed base to induce consumers to upgrade. Duopolists face both forces, whereas a monopolist only faces the latter. This latter effect highlights the importance of product durability, as the absence of depreciation necessitates innovation to induce upgrades. Another important distinction in durable goods markets is between replacement and first-time purchases. Market growth—the addition of potential first-time purchasers—can mitigate competition with the installed base. Rapid market growth reduces the innovation incentives for a monopolist who can exploit demand from first-time buyers, leading the duopoly to innovate faster than the monopolist. Our finding that Intel would innovate more rapidly as a monopolist could therefore be reversed if markets for microprocessors grew more rapidly.

Extending our model to allow for endogenous product durability, as in Rust (1986) for a monopolist, would be interesting. The monotonic increase in duopoly profits as depreciation exogenously increases in Observation 6, combined with the observed near-perfect durability of microprocessors, suggests a prisoner's dilemma: firms would like the industry to sell less durable goods, but they cannot commit to doing so.

One of our goals is to demonstrate the value of structural empirical methods to investigate the effect of competition on innovation. We hope future work will adopt this approach to examine other industries in depth to complement insights from cross-sectional studies of this important issue.

Appendix A: Solving and Simulating Industry Equilibrium

Besanko et al. (2010) and Borkovsky, Doraszelski, and Kryukov (2010) document the existence of multiple equilibria in dynamic oligopoly models based on Ericson and Pakes (1995). To reduce multiplicity, we focus on equilibria that are limits to finitely repeated games: we use backwards induction to solve for an equilibrium of the T-period game and then let $T \to \infty$. For each T and for each state, we solve the system of first-order conditions in equations (13) and (14). Our numerical algorithm for computing equilibrium to the infinite horizon game therefore corresponds to value function iteration with (a) initial values of $\bar{V}^0 = 0$ and $W^0 = 0$ and (b) equilibrium strategies being played within each state for each iteration, as opposed to merely playing best responses to strategies from the previous iteration.

Given $\chi_j(\tau_j = 1|x, q) = a^j(q)x/(1 + a^j(q)x)$, we can simplify the investment first-order condition in equation (14) to

$$x_j - \left(\frac{a^j(q)}{1 - (\beta a^j(q)(EW^+(p_j) - EW^-(p_j)))^{-1/2}} - a^j(q) \right)^{-1} = 0, \tag{21}$$

where

$$EW^+(p_j) = \sum_{q'_{-j}, \Delta'} W_j(q_j + \delta, q'_{-j}, \Delta')\; h_{f_j}(q'_{-j}|q, \Delta)\; g_{f_j}(\Delta'|\Delta, q, p)$$

and

$$EW^{-}(p_j) = \sum_{q'_{-j},\Delta'} W_j(q_j + 0, q'_{-j}, \Delta')\ h_{f_j}(q'_{-j}|q, \Delta)\ g_{f_j}(\Delta'|\Delta, q, p)$$

are the expected continuation values conditional on positive and negative innovation outcomes, respectively. The dependence of these expectations on p_j is through the effect of price on the ownership transition to Δ'.

For each iteration $k = 1, 2, \ldots$, we follow these steps:

1. Simultaneously solve firms' first-order conditions in equations (13) and (21) for $\{p_j^*, x_j^*\}_{j=1}^J$ at each state given continuation values determined by W^{k-1} and $\bar{V}^{k-1}$ for firms and consumers, respectively. Since the FOC depend on consumers' current choices which in turn depend on their rational expectations of Δ', for each conjectured $\{p_j, x_j\}_{j=1}^J$ we solve for the fixed point in Δ' such that consumers' expectations for Δ' are realized (i.e., equation (8) in footnote 17 is satisfied). Denote Δ'^* this fixed point when firms play $\{p_j^*, x_j^*\}_{j=1}^J$.
2. Let W^k equal the discounted payoffs given firms' current policies $\{p_j^*, x_j^*\}_{j=1}^J$, continuation ownership distribution Δ'^*, and continuation values based on W^{k-1}.
3. For each consumer state, which adds $\tilde{\omega}$ to the industry state vector, evaluate the consumer's smoothed value function $\bar{V}^k$ given firms' current policies $\{p_j^*, x_j^*\}_{j=1}^J$, continuation ownership of Δ'^*, and continuation values based on $\bar{V}^{k-1}$.
4. Check for convergence in the sup norms of $|\bar{V}^k - \bar{V}^{k-1}|$ and $|W^k - W^{k-1}|$ 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). In our first quarter, AMD is 1.7 δ-steps behind Intel. Hence we start 70 percent of the simulations with AMD two steps behind Intel and 30 percent with AMD one step behind. 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 evolution of ownership given consumers' equilibrium choice probabilities, and process the stochastic innovation outcomes according to $\chi_j(\cdot|\cdot)$.

A challenge in solving our model is that Δ is a high-dimensional simplex. We approximate this continuous state variable with a discretization that restricts $\Delta \in \{\Delta^d\}_{d=1}^D$. Let $\rho_d(\Delta')$ denote the distance between the exact continuation Δ', as implied by equations in Footnote 17, 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. 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. We therefore define

$$\rho_d(\Delta'_{t+1}) = \left| \sum_k k\Delta'_{k,t+1} - \sum_k k\Delta_k^d \right|, \quad \text{for all } d \in (1, \ldots, D), \tag{22}$$

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

We generate $\{\Delta^d\}_{d=1}^D$ using a distribution parameterized by a scalar and choose the scalar's discrete grid such that mean qualities are .25 apart and range from 9 to 29, relative to $\bar{q}$ fixed at 30. 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.

In computing equilibrium, we use a cubic spline to interpolate between Δ grid points since solving the firms' first-order conditions requires differentiable continuation values. When simulating the model, rather than interpolating policy functions, we force Δ to remain on the grid by randomizing Δ_{t+1} to be one of the two closest Δ_d, with probabilities proportional to the distances between these closest Δ_d and the exact continuation Δ.

Appendix B: 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

$$V^T(q_1, \Delta_1, \tilde{q}_1, \varepsilon_1) = \max_{\{y_t(q_t,\Delta_t,\tilde{q}_t,\varepsilon_t)\in(0,1,\ldots,J)\}_{t=1}^T} \mathrm{E}\left[\sum_{t=1}^{T} \beta^t \left(\gamma q_{y_t,t} - \alpha p_{y_t,t} + \xi_{y_t} + \varepsilon_{y_t,t}\right)\right] , \tag{23}$$

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

$$W_j^T(q_{j1}, q_{-j,1}, \Delta_1) = \max_{\{p_{jt}(q_t,\Delta_t),\ x_t(q_t,\Delta_t)\}_{t=1}^T} \mathrm{E}\left[\sum_{t=1}^{T} \beta^t \left(M s_{jt}(p_t, q_t, \Delta_t)(p_{jt} - mc_j) - x_{jt}\right)\right] , \tag{24}$$

where M is market size, $s_{jt}(\cdot)$ is the market share for firm j as defined in equation (7), 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

$$\begin{array}{rrcl} \text{Firms:} & W_j^T(q_{jt}, q_{-j,t}, \Delta_t) & = & W_j^T(q_{jt} - \hat{q}, q_{-j,t} - \hat{q}, \Delta_t) \\ \text{Consumers:} & V^T(q_t, \Delta_t, \tilde{q}_t, \varepsilon_t) & = & \gamma\hat{q} + V^T(q_t - \hat{q}, \Delta_t, \tilde{q}_t - \hat{q}, \varepsilon_t) . \end{array} \tag{25}$$

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

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

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 (17). Since each consumer is small relative to M, her actions do not affect the transition of Δ.

Each firm j solves

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

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. Recall that $\chi_j(\cdot|x_{jt}, q_t)$ 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.

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

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

and

$$\begin{aligned} W_j^{T-1}(q_{j,T-1}-\hat{q}, q_{-j,T-1}-\hat{q}, \Delta_{T-1}) \quad = \quad & \max_{\substack{p_{j,T-1}\\ x_{j,T-1}}} M s_{j,T-1}(p_{T-1}, q_{T-1}-\hat{q}, \Delta_{T-1})(p_{j,T-1}-mc_j) \\ & -x_{j,T-1} + \beta \sum_{q_T} W_j^T(q_{j,T}-\hat{q}, q_{-j,T}-\hat{q}, \Delta_T) \\ & \prod_{j=1}^{J} \chi_j(q_{jT}-\hat{q}-(q_{j,T-1}-\hat{q})|x_{j,T-1}, q_{T-1}-\hat{q}) \ . \end{aligned} \tag{29}$$

Substitute the right-hand sides of (25) into (29) and (28). Then note that $\chi_j(q_{jT}-\hat{q}-(q_{j,T-1}-\hat{q})|x_{j,T-1}, q_{T-1}-\hat{q}) = \chi_j(q_{jT}-q_{j,T-1}|x_{j,T-1}, q_{T-1})$ 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 (29) into (27), exactly, and converts (28) into (26), except for a $-(\gamma\hat{q}+\beta\gamma\hat{q})$ term that does not affect the consumer's choice. The modified (28) is

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

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.

$\square$

Appendix C: 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 3 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

$$\begin{aligned} f_\mu(\omega_{t+1}, \Delta_{t+1}, p_{t+1}, x_{t+1}, s_{t+1} | \omega_t, \Delta_t, p_t, x_t, s_t, \theta) &= \prod_{j=1}^{J} \chi_j(\omega_{j,t+1} - \omega_{j,t} | \omega_t, x_t) \\ &\times \; g(\Delta_{t+1} | \Delta_t, s_t) \\ &\times \; I\{p_{t+1} = p(\omega_{t+1}, \Delta_{t+1})\} \\ &\times \; I\{x_{t+1} = x(\omega_{t+1}, \Delta_{t+1})\} \\ &\times \; I\{s_{t+1} = s(\omega_{t+1}, \Delta_{t+1})\} \,, \end{aligned} \tag{31}$$

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 firm's 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 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:

$$m_{S,T}(\theta) = \frac{1}{S} \sum_{n=1}^{S} m\left(\{\mu_t(\theta, U_{<t}^n, \omega_0, \Delta_0)\}_{t=1}^T\right), \tag{32}$$

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

$$\hat{\theta}_T = \underset{\theta \in \Theta}{\operatorname{argmin}} \, (m_{S,T}(\theta) - m_T)' A_T (m_{S,T}(\theta) - m_T) \,, \tag{33}$$

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 $\{\mu_t(\theta, U_{<t}^n, \omega_0, \Delta_0)\}$ is ergodic with unique invariant density $\Psi(\mu|\theta)$ given by*

$$\Psi(\mu'|\theta) = \int f_\mu(\mu'|\mu, \theta) d\Psi(\mu|\theta). \tag{34}$$

Assumption 2. *The structural model presented in section 3 is correctly specified. As such, a $\theta^* \in \Theta$ exists for which each simulated sequence $\{\mu_t^n\}, n = 1, \ldots, S$ 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

$$\begin{array}{lcl} \mathrm{E}[m|\theta] & = & \int m(\mu)d\Psi(\mu|\theta) \\ \nabla\mathrm{E}[m|\theta] & = & \frac{\partial}{\partial\theta}\mathrm{E}[m|\theta] \\ \nabla m_{S,T} & = & \frac{\partial}{\partial\theta}m_{S,T}(\theta). \end{array} \tag{35}$$

Assumption 3. *θ^* is identified; that is, if $\theta \neq \theta^*$ then $E[m|\theta] \neq E[m|\theta^*] = E[m(\{\mu_t^{actual}\}_{t=1}^T)]$. In addition, $rank(\nabla E[m|\theta]) = K$ 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 $A_T = [cov(\{\mu_t^{actual}\}_{t=1}^T)]^{-1}$ as a consistent estimate of the optimal weight matrix, the estimator $\hat{\theta}_T$ has the property

$$\sqrt{T}(\hat{\theta}_T - \theta^*) \Longrightarrow N\left(0, (1+1/S)\left(\nabla\mathrm{E}[m|\theta^*]'\Omega(m,\theta^*)^{-1}\nabla\mathrm{E}[m|\theta^*]\right)^{-1}\right). \tag{36}$$

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

Appendix D: Decomposing the Effect of the Quality Coefficient on Industry Innovation

As discussed in Observation 5, two factors contribute to innovation being more sensitive to preferences (γ) in the duopoly than in the monopoly: the effect of γ on firms' investment policies $x^*(\cdot)$ and on which states are encountered.

To evaluate the effect of γ on firms' investment policies, we condition on typical Δ values for the monopoly and duopoly and, in Figure 13, plot the effect of γ on firms' investment levels and innovation rates, the marginal effect of investment on the probability of an innovation ($\frac{\partial\chi_j(x^*)}{\partial x}$), the marginal effect of γ on investment ($\frac{\partial x^*}{\partial\gamma}$), and the marginal effect of γ on innovation ($\frac{\partial\chi_j(x^*)}{\partial\gamma} = \frac{\partial\chi_j(x^*)}{\partial x}\frac{\partial x^*}{\partial\gamma}$). This last marginal effect, plotted in panel 6, is higher for the leader and for each tied firm than for the monopolist, and is much higher for *either* tied firm innovating. Hence changes in innovation policy functions contribute to the industry innovation increasing more in the duopoly than in the monopoly as γ increases. Interestingly, for low γ, the monopolist's investment is more sensitive to γ than duopolists' investments, but the marginal effect of investment on innovation is sufficiently higher for the duopolists that the ultimate effect on innovation is greater in the duopoly.

The effect of γ on industry innovation through investment policies $x^*(\cdot)$, however, can potentially be mitigated by the equilibrium transiting to states with lower duopoly innovation or higher monopoly innovation. Quality preferences can effect the evolution of two state variables: the ownership distribution Δ and, in the duopoly, the quality difference $q_{1t} - q_{2t}$. The effect of γ on realized ownership distributions is small and has little impact on innovation in both the duopoly and monopoly since firms' innovation policies are relatively insensitive to Δ, as illustrated by the innovation policies in panels 7 and 8 of Figure 2. The quality difference in panel 5 of Figure 13, however, widens as γ increases, which has two distinct effects on innovation. First, the leader innovates more with a wider lead, as illustrated by the innovation policy in panel 9 of Figure 2. Second, the share of periods with tied firms declines. Since the frontier advances when *either* tied firm innovates, one might expect this factor to put downward pressure on innovation as γ increases. However, as seen in panel 1 of Figure 13, the gap between the probability of either tied firm innovating and an individual tied firm innovating increases in γ. This larger gap offsets the lower probability of being tied, as revealed by the nearly flat product of the gap and the probability in panel 5. The net effect of γ's influence on states encountered is therefore the positive effect of higher innovation by the duopoly leader as the quality difference widens.

Appendix E: Adding Depreciation to the Model

To allow for physical depreciation of the durable good, as needed for Observation 6, we modify the transition kernels for Δ and $\tilde{q}$. Let ϕ denote the probability that each consumer's $\tilde{q}_t$ declines by one δ-step. The post-depreciation ownership share for each vintage k in Δ' is then

$$\Delta'_k = \begin{cases} \Delta'_k(\Delta, q, p) + \phi\Delta'_{k+\delta}(\Delta, q, p) & \text{if } q_k = \underline{q}' \\ (1-\phi)\Delta'_k(\Delta, q, p) + \phi\Delta'_{k+\delta}(\Delta, q, p) & \text{if } \underline{q}' < q_k < \bar{q}' \\ (1-\phi)\Delta'_k(\Delta, q, p) & \text{if } q_k = \bar{q}' \end{cases}, \tag{37}$$

where $\Delta'(\cdot)$ is given in Footnote 17, the first line enforces the lower bound for $\tilde{q}'$, and the third line acknowledges that the frontier does not gain mass from a higher vintage. The consumer's continuation values in equations (3), (4), and (5) must also integrate over the realization of this stochastic depreciation for the consumer's $\tilde{q}'$.

References

[1] Aghion, P. and P. Howitt (1992), "A Model of Growth through Creative Destruction," *Econometrica*, 60(2), 323–351.

[2] Aghion, P., C. Harris, P. Howitt, and J. Vickers (2001), "Competition, Imitation, and Growth with Step-by-Step Innovation," *Review of Economic Studies*, 68(3), 467–492.

[3] 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.

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

[5] 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.

[6] Bajari, P., L. Benkard, and J. Levin (2007), "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75, 1331–1370.

[7] Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010), "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics," *Econometrica*, 78, 453–508.

[8] Blundell, R., R. Griffith, and J. Van Reenen (1999) "Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms," *Review of Economic Studies*, 66, 529–554.

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

[10] Borkovsky, R. (2008), "The Timing of Version Releases in R&D-intensive Industries: A Dynamic Duopoly Model," mimeo University of Toronto.

[11] Borkovsky, R., Doraszelski, U., and Y. Kryukov (2010), "A Dynamic Quality Ladder Model with Entry and Exit: Exploring the Equilibrium Correspondence Using the Homotopy Method," mimeo University of Toronto.

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

[13] Carlton, D. and R. Gertner (1989), "Market Power and Mergers in Durable-Good Industries," *Journal of Law and Economics*, 32, 203–226.

[14] Carranza, J. (2010), "Product innovation and adoption in market equilibrium: The case of digital cameras," *International Journal of Industrial Organization*, 28(6), 604–618.

[15] Chen, J., S. Esteban, and M. Shum (2011), "How Much Competition is a Secondary Market?" mimeo, UC Irvine.

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

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

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

[19] *Computer Industry Almanac* (2005), "Worldwide PC market" report, Computer Industry Almanac, Arlington Heights, IL, http://www.c-i-a.com/wwpc.htm.

[20] 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, Vol. 3*, North-Holland, Amsterdam, 1887–1966.

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

[22] Economic Report of the President (2011), "Table B-12, Gross domestic product (GDP) by industry, value added, in current dollars and as a percentage of GDP, 1979–2009," United States Government Printing Office, Washington.

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

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

[25] Gilbert, R. (2006), "Looking for Mr. Schumpeter: Where are We in the Competition-Innovation Debate?" in A. Jaffe, J. Lerner, and S. Stern, eds., *Innovation Policy and the Economy, Vol. 6*, NBER Books.

[26] Goettler, R. and B. R. Gordon (2011), "Competition and Product Innovation in Dynamic Oligopoly," working paper, Columbia Business School.

[27] Gordon, B. R. (2009), "A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry," *Marketing Science*, 28(5), 846–867.

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

[29] Gowrisankaran, G. and M. Rysman (2009), "Dynamics of Consumer Demand for New Durable Goods," NBER working paper #W14737.

[30] 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.

[31] 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.

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

[33] Holmes, T., D. Levine, and J. Schmitz, Jr. (2011), "Monopoly and the Incentive to Innovate When Adoption Involves Switchover Disruptions," NBER Working Paper #13864.

[34] Hopenhayn, H. (1992), "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*,60(5), 1127–1150.

[35] Jorgenson, D., M. Ho, and J. Samuels (2010), "Information Technology and U.S. Productivity Growth: Evidence from a Prototype Industry Production Account," in M. Mas and R. Stehrer eds., *Industrial Productivity in Europe: Growth and Crisis*, Edgar Elgar Publishing.

[36] Klette, J. and S. Kortum (2004), "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 112(5), 986–1018.

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

[38] Levin, R., W. Cohen, and D. Mowery (1985), "R&D Appropriability, Opportunity, and Market Structure: New Evidence on Some Schumpeterian Hypotheses," *American Economic Review*, 75(2), 20–24.

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

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

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

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

[43] Nosko, C. (2010), "Competition and Quality Choice in the CPU Market," mimeo, Harvard University.

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

[45] 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.

[46] Pillai, Unnikrishnan (2009), "Technological Progress in the Microprocessor Industry: Explanations of the Acceleration in 1990–2000 and the Slowdown in 2001–2008," mimeo, SUNY Albany.

[47] 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.

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

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

[50] 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.

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

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

[53] Scott, J. (1984), "Firm versus Industry Variability in R&D Intensity," in Z. Griliches, ed., *R&D, Patents, and Productivity*, Chicago: University of Chicago Press for the NBER.

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

[55] 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.

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

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

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

[59] Waldman, M. (2007), "Antitrust Perspectives for Durable-Goods Markets," in J. P. Choi, ed., *Recent Developments in Antitrust: Theory and Evidence*, MIT Press, Cambridge, 1–37.

Table 1: Empirical and Simulated Moments

Moment	Actual	Actual SE	Fitted
Intel Price Equation:			
Average Intel Price	219.7	(5.9)	206.2
$q_{Intel,t} - q_{AMD,t}$	47.4	(17.6)	27.3
$q_{Intel,t} - \bar{\Delta}_t$	94.4	(31.6)	43.0
AMD Price Equation:			
Average AMD Price	100.4	(2.3)	122.9
$q_{Intel,t} - q_{AMD,t}$	-8.7	(11.5)	-22.3
$q_{AMD,t} - \bar{\Delta}_t$	16.6	(15.4)	5.9
Intel Share Equation:			
constant	0.834	(0.007)	0.846
$q_{Intel,t} - q_{AMD,t}$	0.055	(0.013)	0.092
Potential Upgrade Gains:			
Mean $(\bar{q}_t - \bar{\Delta}_t)$	1.146	(0.056)	1.100
Mean Innovation Rates:			
Intel	0.557	(0.047)	0.597
AMD	0.610	(0.079)	0.602
Relative Qualities:			
Mean $q_{Intel,t} - q_{AMD,t}$	1.257	(0.239)	1.352
Mean $\mathcal{I}(q_{Intel,t} \geq q_{AMD,t})$	0.833	(0.054)	0.929
Mean R&D / Revenue:			
Intel	0.114	(0.004)	0.101
AMD	0.203	(0.009)	0.223

Simulated moments, as defined in Section 4.1.1, are averages over 10,000 simulations of 48 quarters of data. Though a constant is in each of the first two regressions, we match each firm's mean price instead. $\mathcal{I}(\cdot)$ is an indicator function.

Table 2: Parameter Estimates

Parameter	Estimate	Std. Error
Price, α	0.0131	0.0017
Quality, γ	0.2764	0.0298
Intel Fixed Effect, ξ_{Intel}	-0.6281	0.0231
AMD Fixed Effect, ξ_{AMD}	-3.1700	0.0790
Intel Innov, $a_{0,Intel}$	0.0010	0.0002
AMD Innov, $a_{0,AMD}$	0.0019	0.0002
Spillover, a_1	3.9373	0.1453
Stage-1 Marginal Cost Equation		
Constant, λ_0	44.5133	1.1113
$\max(0, q_{competitor,t} - q_{own,t})$, λ_1	-19.6669	4.1591

Table 3: Industry Measures under Various Scenarios

	(1) AMD-Intel Duopoly	(2) Symmetric Duopoly	(3) Monopoly	(4) No Spillover Duopoly	(5) Myopic Pricing AMD-Intel	(6) Myopic Pricing Monopoly	(7) Social Planner
Industry Profits ($billions)	408	400	567	382	318	322	-267
Consumer Surplus	2978	3012	2857	3068	2800	2762	4032
CS as Share of Monopoly CS	1.042	1.054	1.000	1.074	0.980	0.967	1.411
Social Surplus	3386	3412	3424	3450	3118	3084	3765
SS as Share of Planner SS	0.929	0.906	0.940	0.916	0.828	0.819	1
Margins	3.434	2.424	5.672	3.478	2.176	2.216	0.000
Price	194.17	146.73	296.98	157.63	140.06	143.16	43.57
Frontier Innovation Rate	0.599	0.501	0.624	0.438	0.447	0.438	0.869
Industry Investment ($millions)	830	652	1672	486	456	787	6672
Mean Quality Upgrade %	261	148	410	187	175	181	97
Intel or Leader Share	0.164	0.135	0.143	0.160	0.203	0.211	0.346
AMD or Laggard Share	0.024	0.125		0.091	0.016		0.014

Reported values are based on 10,000 simulations of 300 periods each. Profits, surplus, and investments are reported in billions of dollars. Profits and surplus are discounted back to period 0. Social surplus is the sum of consumer surplus and industry profits. Symmetric duopoly uses Intel's firm-specific parameters for both firms. Under "myopic pricing" firms choose price ignoring its effect on future demand. The "no spillover" duopoly uses symmetric firms, both with Intel's parameters. The social planner sells two products, but the results are nearly identical for a single-product planner. The monopolist offers one product. Margins are computed as $(p - mc)/mc$. Price and margins are share-weighted averages. In the symmetric duopoly, both firms have Intel's ξ_j and $a_{0,j}$.

Figure 1: CPU Qualities, Prices, Costs, and Shares: 1993 Q1 to 2004 Q4

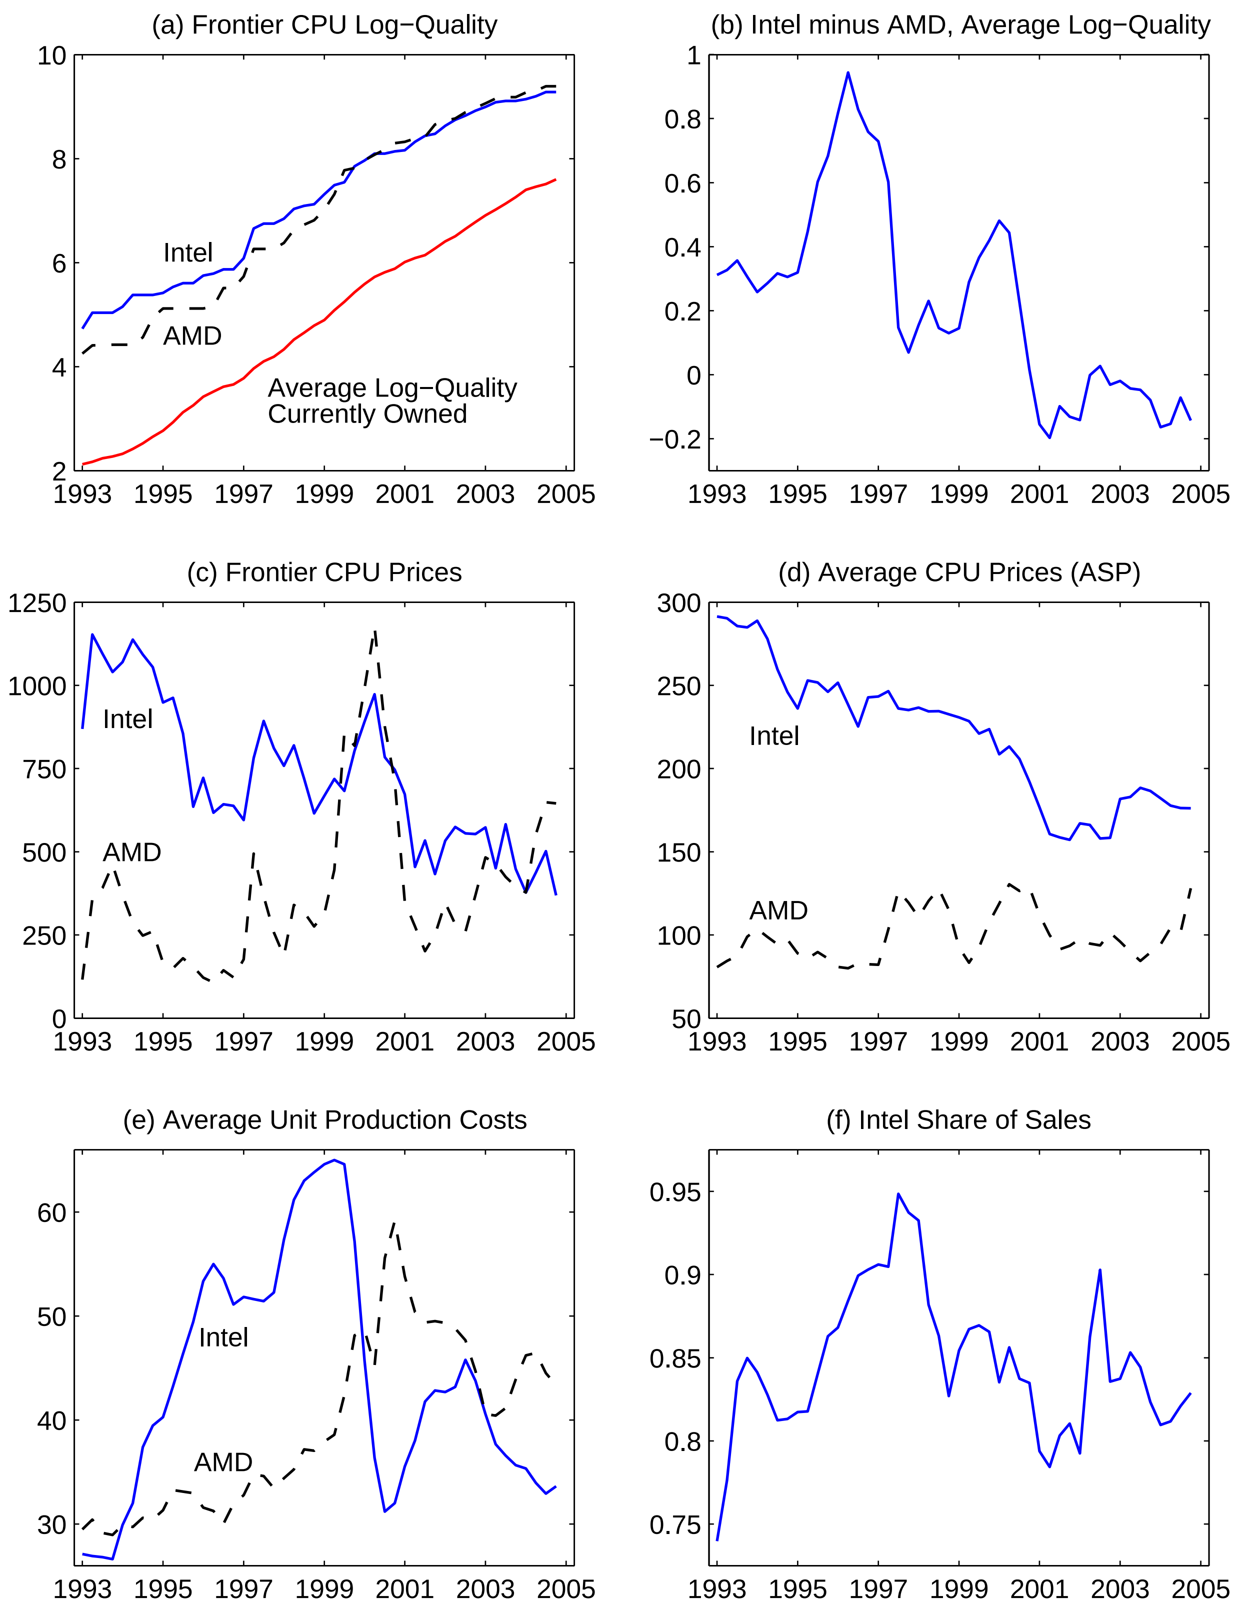

Figure 2: Value and Policy Functions: Monopoly and Symmetric Duopoly: Column 1 corresponds to the monopolist. Columns 2 and 3 correspond to the symmetric duopoly. In columns 1 and 2, the x-axis is the average quality in the ownership distribution Δ. In column 2, values are reported for two scenarios: when the firms are tied and when their qualities differ by 4 δ-steps. In column 3, Δ is fixed at its most frequent value in the duopoly simulations and the x-axis is the quality difference between the leader and laggard, which ranges from 8 to 0 δ-steps. Histograms in the first row provide simulated frequencies of each state.

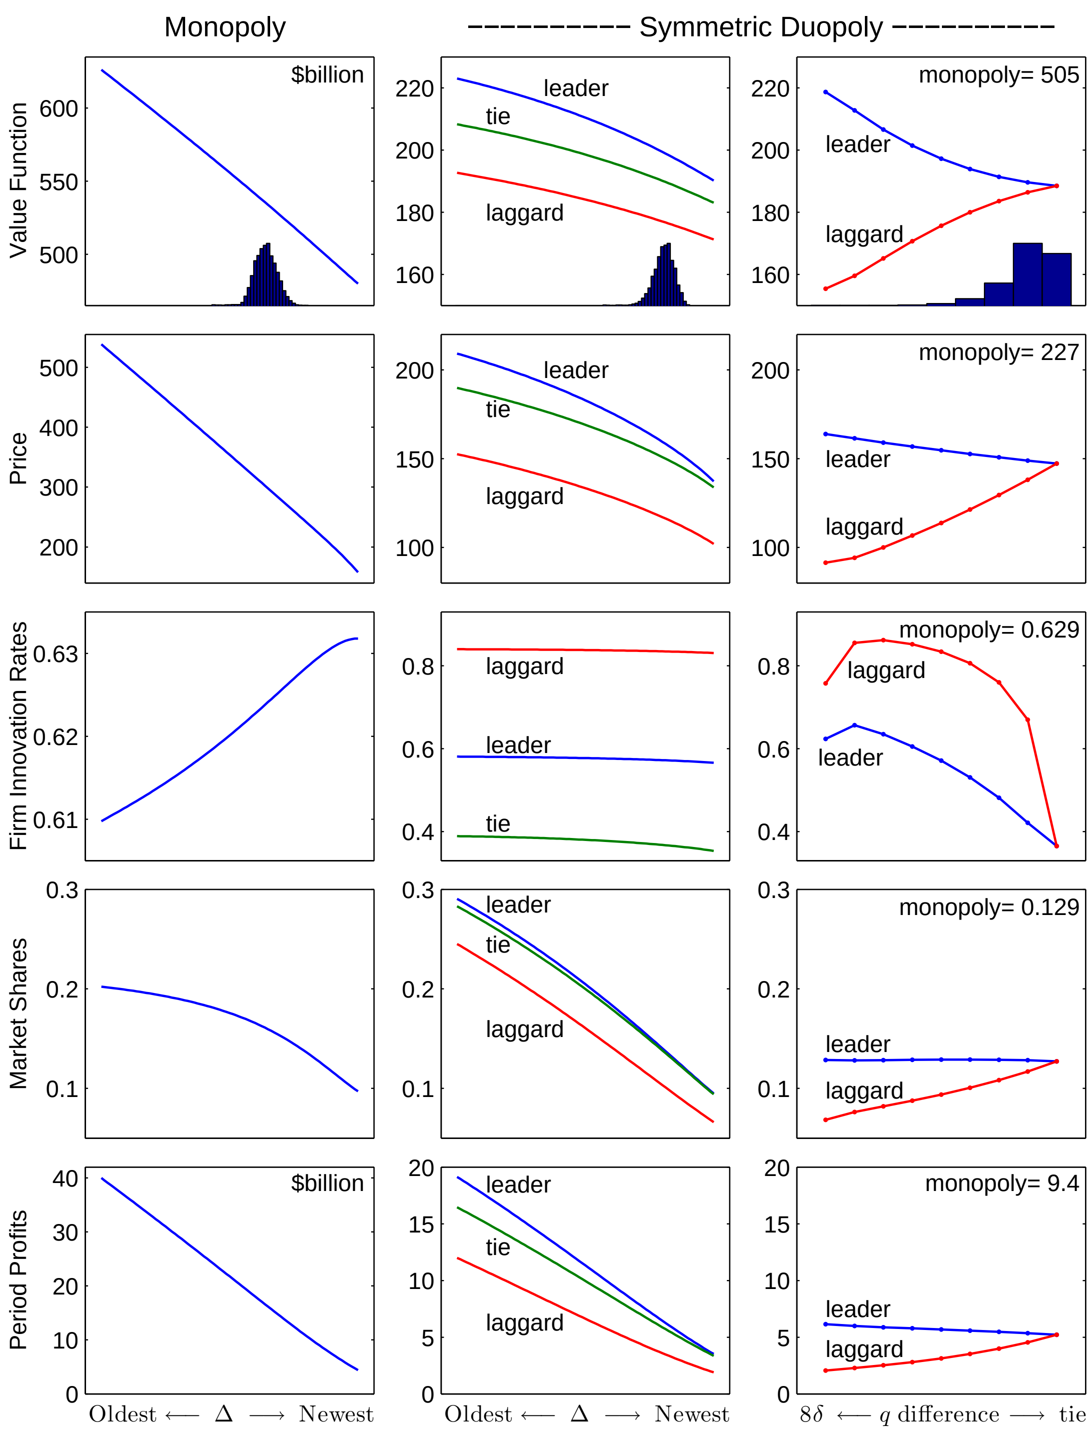

Figure 3: Value and Policy Functions: Monopoly and Symmetric, No Spillover Duopoly: The notes to Figure 2 apply to this figure as well.

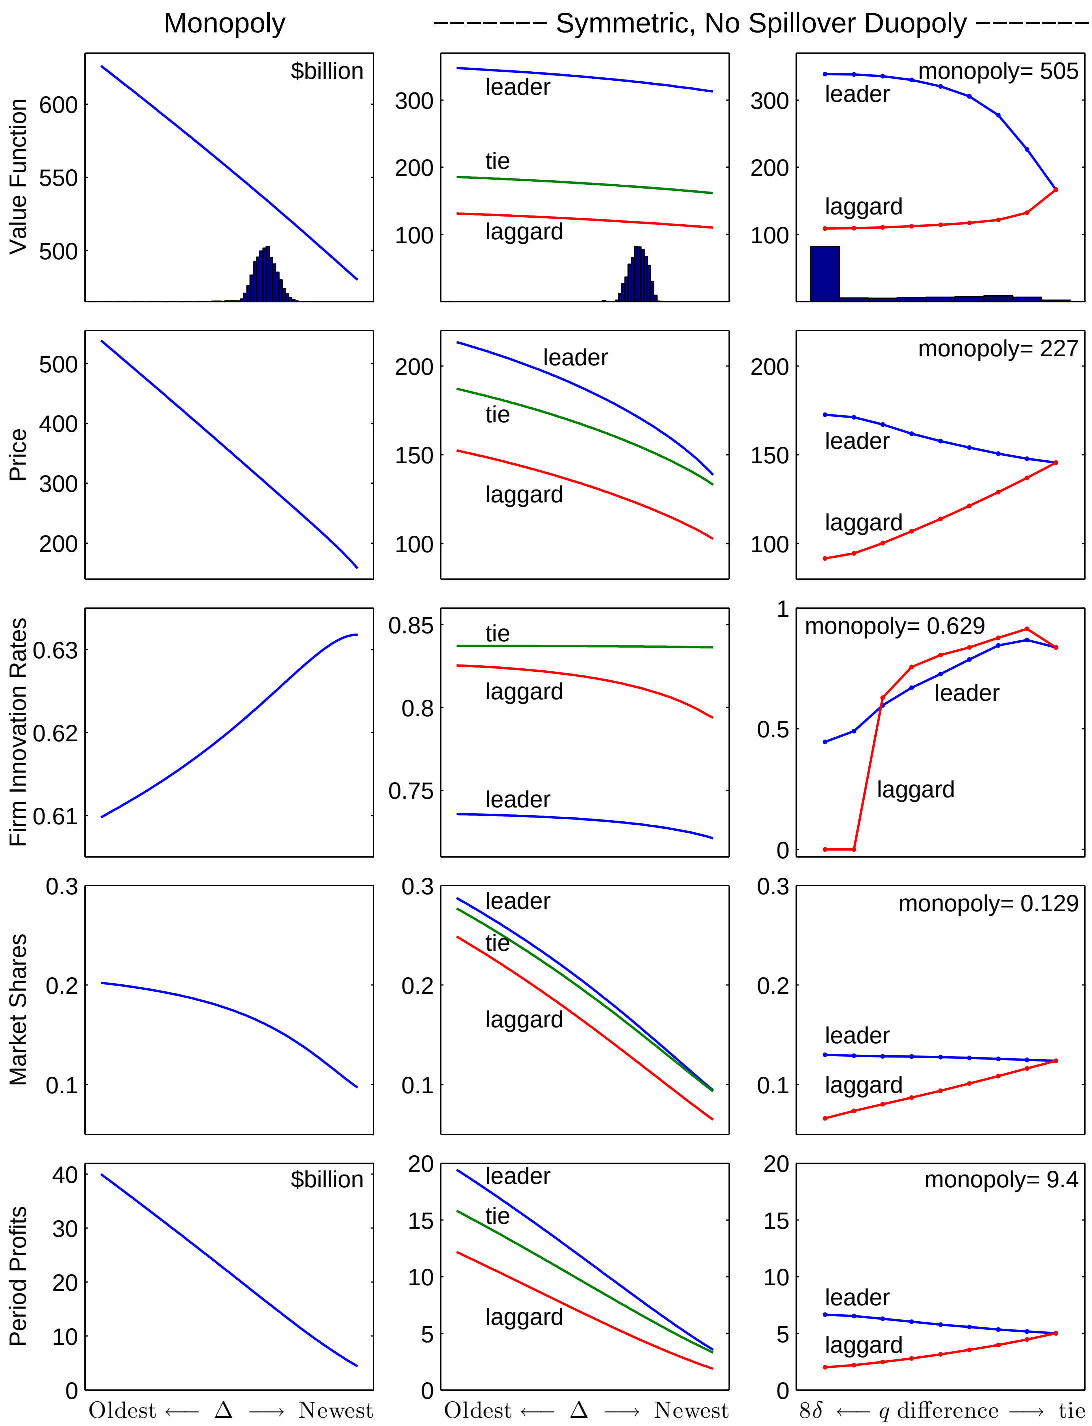

Figure 4: Average Purchase Probabilities by Vintage: AMD-Intel Duopoly

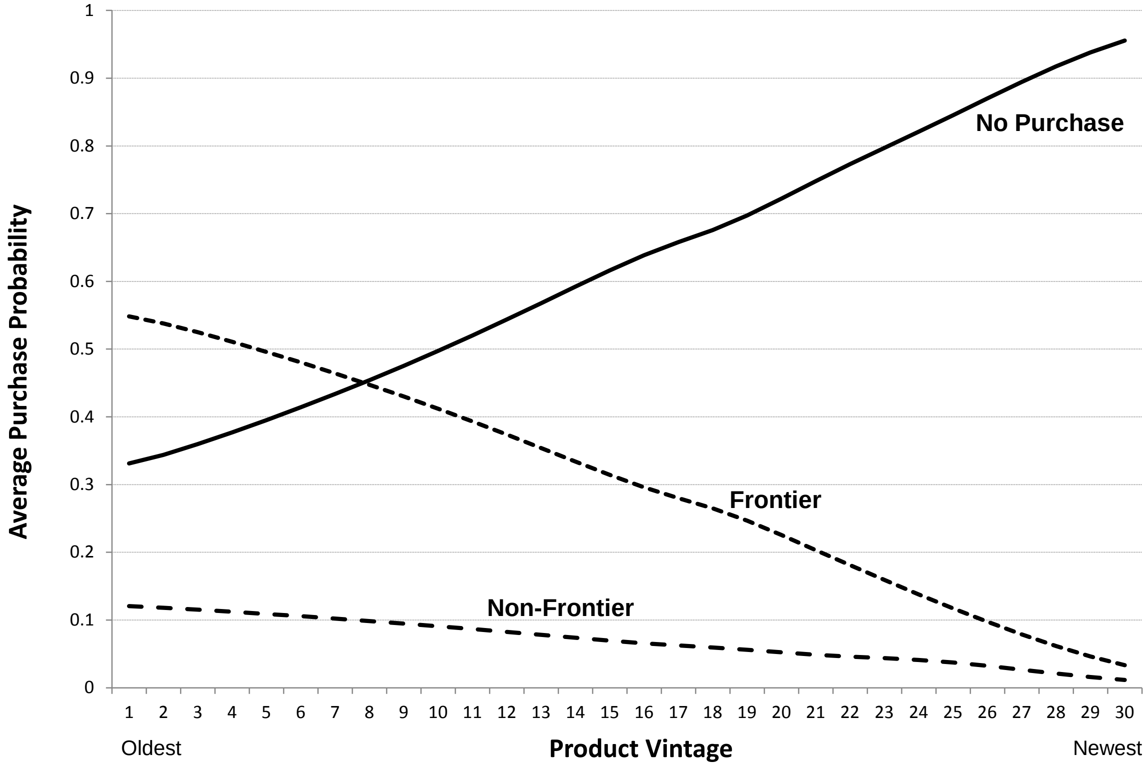

Figure 5: Average Ownership Distributions: AMD-Intel Duopoly and Monopoly

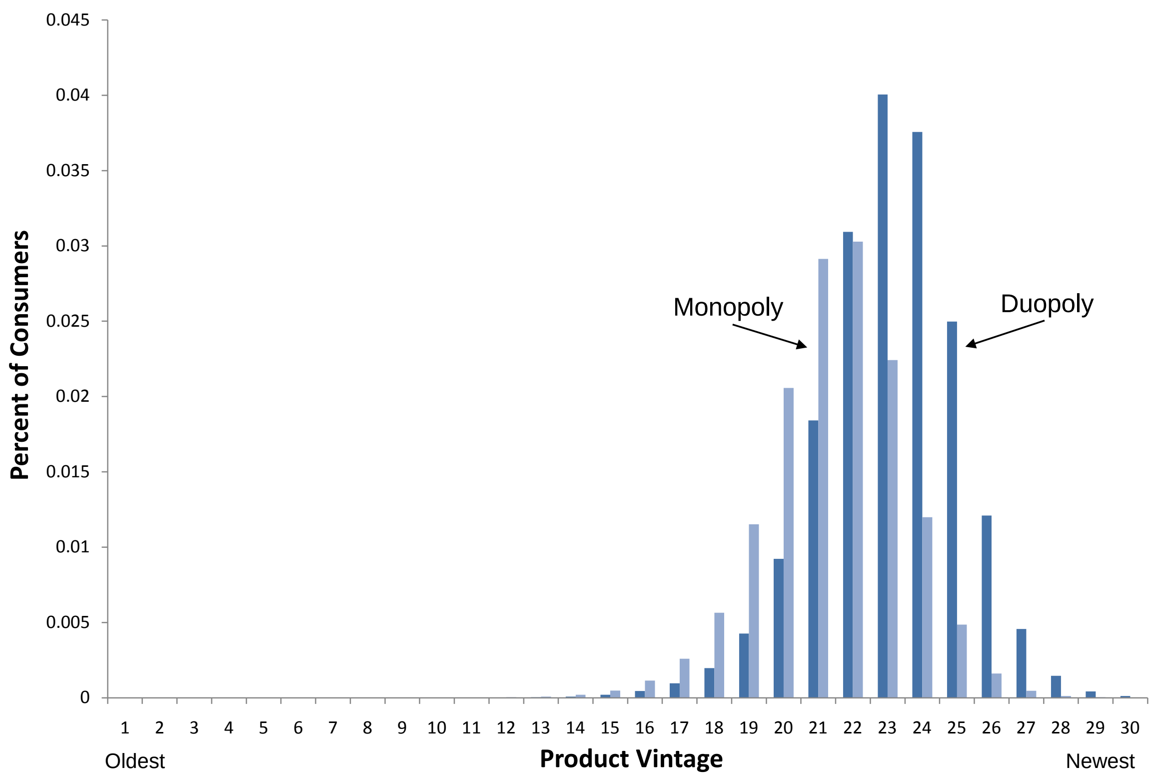

Figure 6: Foreclosing AMD from the Market

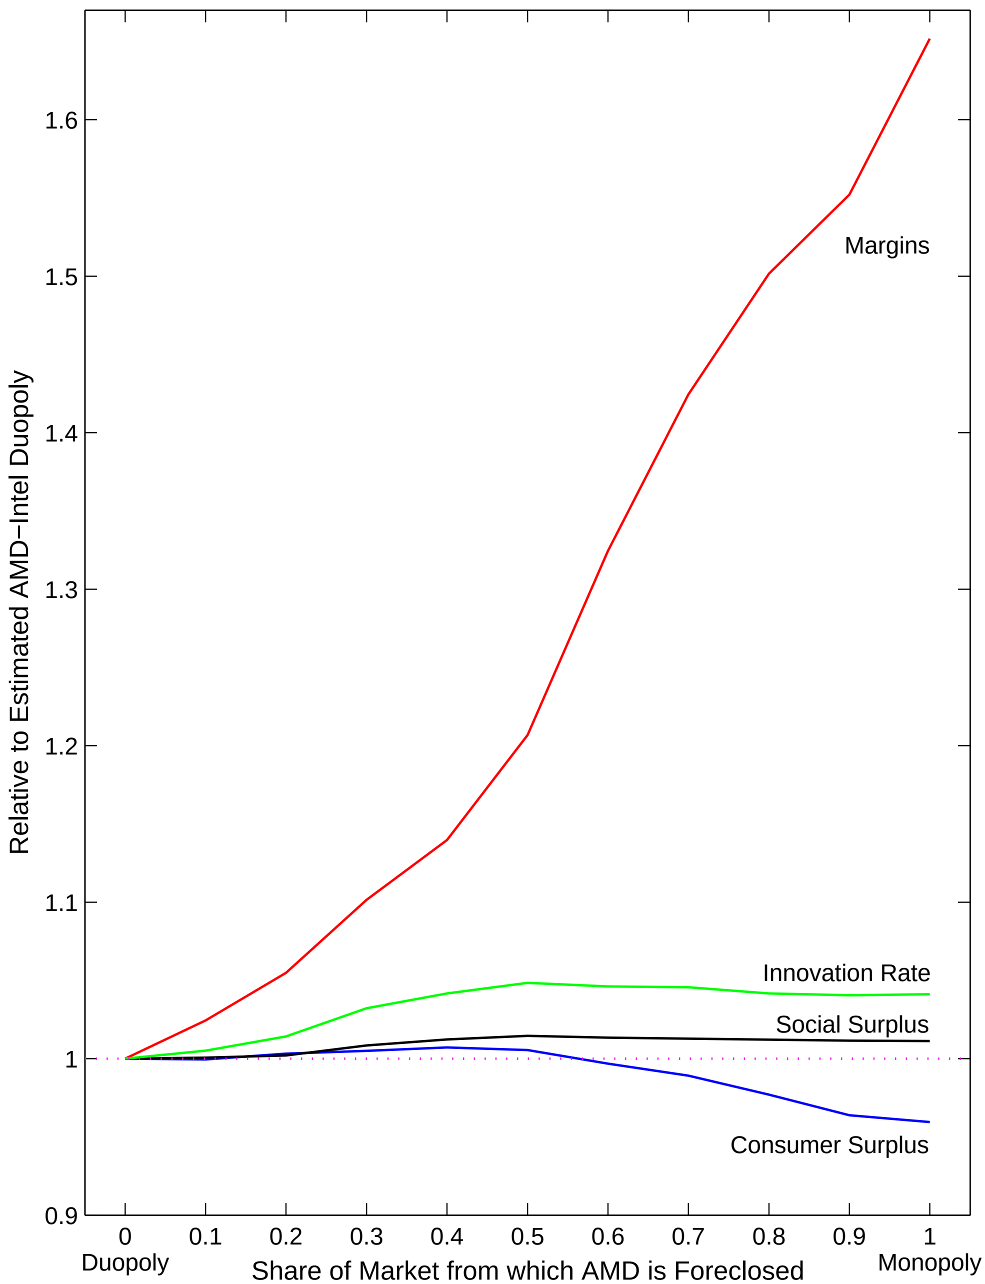

Figure 7: Monopoly versus Duopoly Innovation as (α, γ) Vary

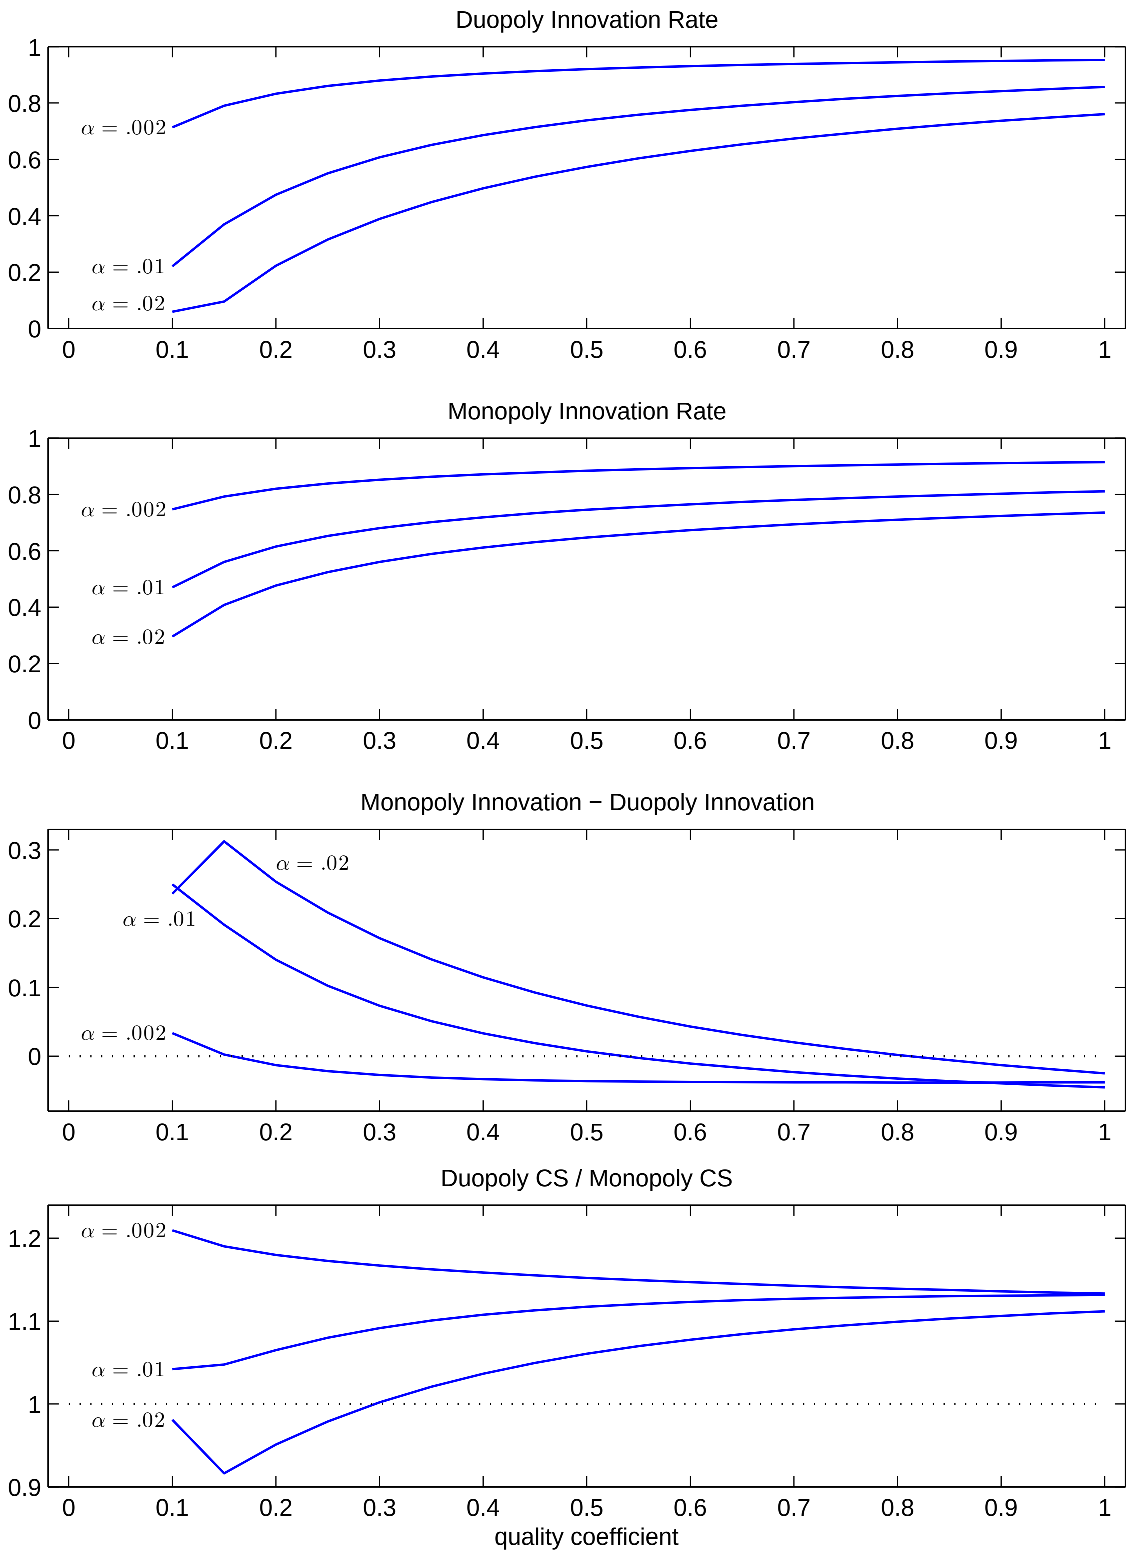

Figure 8: Equilibrium Outcomes in the Monopoly and Symmetric Duopoly as Depreciation Varies: Depreciation is measured as the quarterly probability that $\tilde{q}$ declines one δ-step.

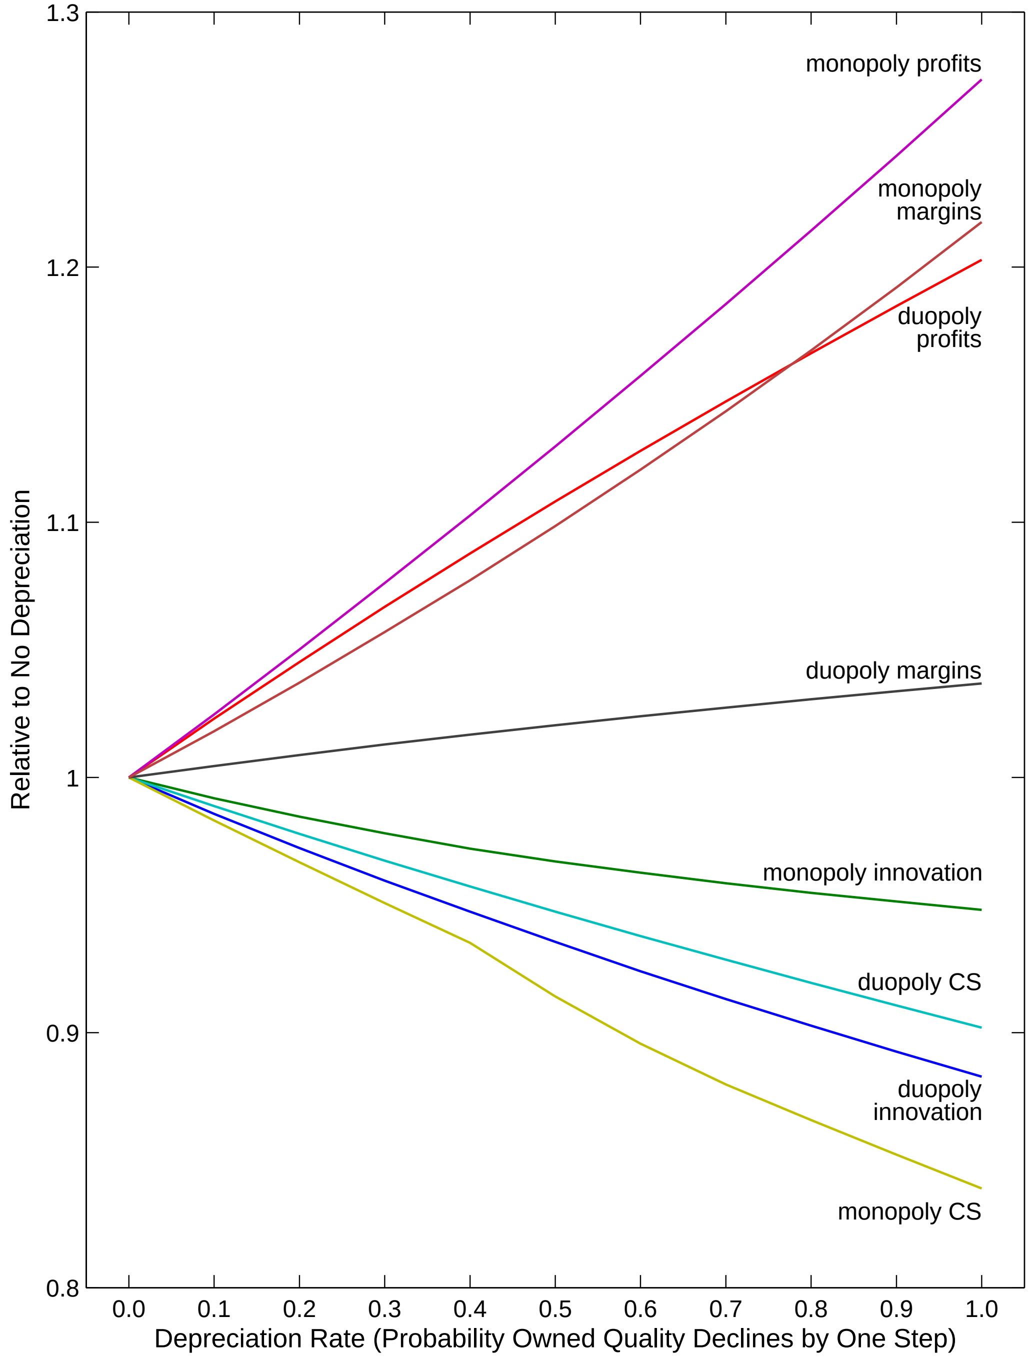

Figure 9: Innovation in the Monopoly and Symmetric Duopoly as Market Growth and Market Size Vary

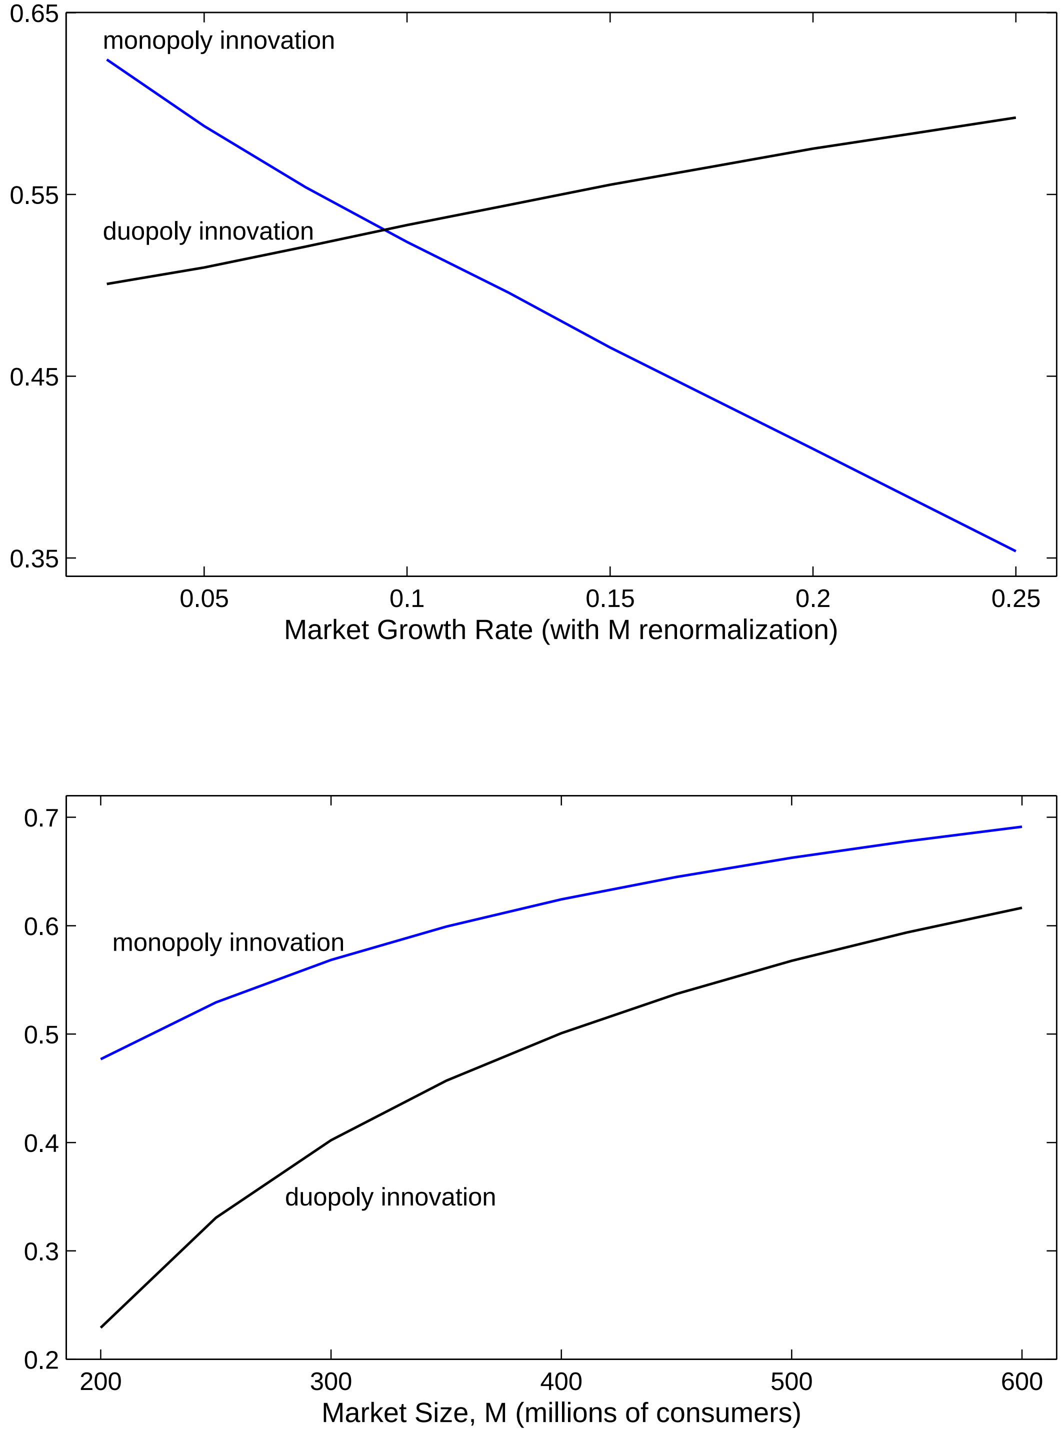

Figure 10: Symmetric Duopoly Innovation as the Spillover Varies: For comparison, the monopolist's values relative to the estimated-spillover symmetric-duopoly are 2.34 for margins, .949 for consumer surplus, 1.004 for social surplus, and 1.246 for innovation.

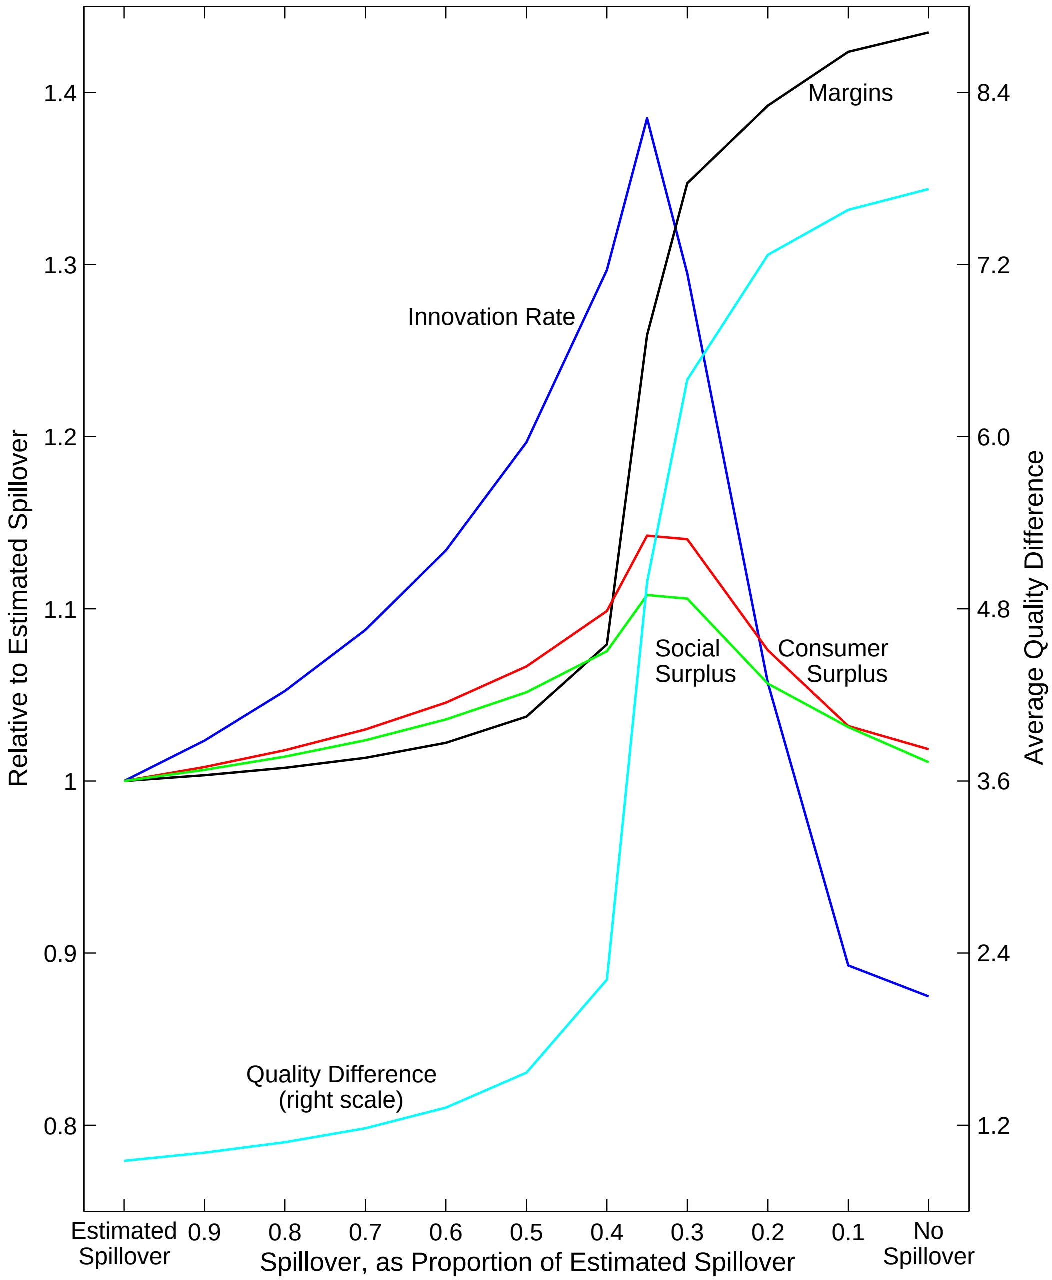

Figure 11: Innovation in the Monopoly and Symmetric Duopoly as Product Substitutability Varies

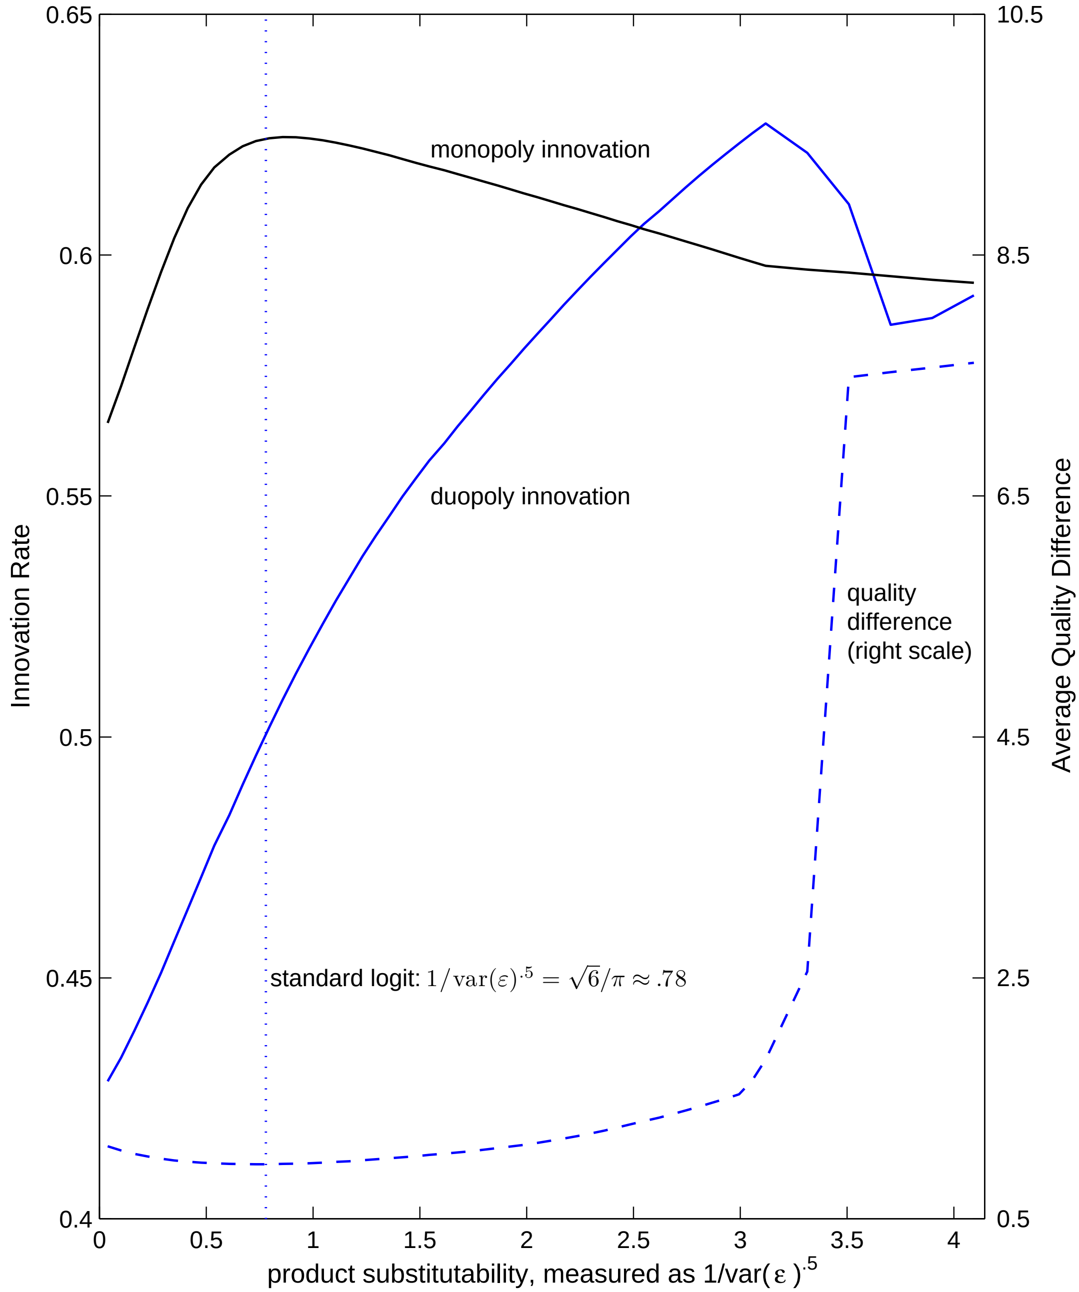

Figure 12: Innovation in the Symmetric Duopoly with Nondurable Goods as Product Substitutability Varies

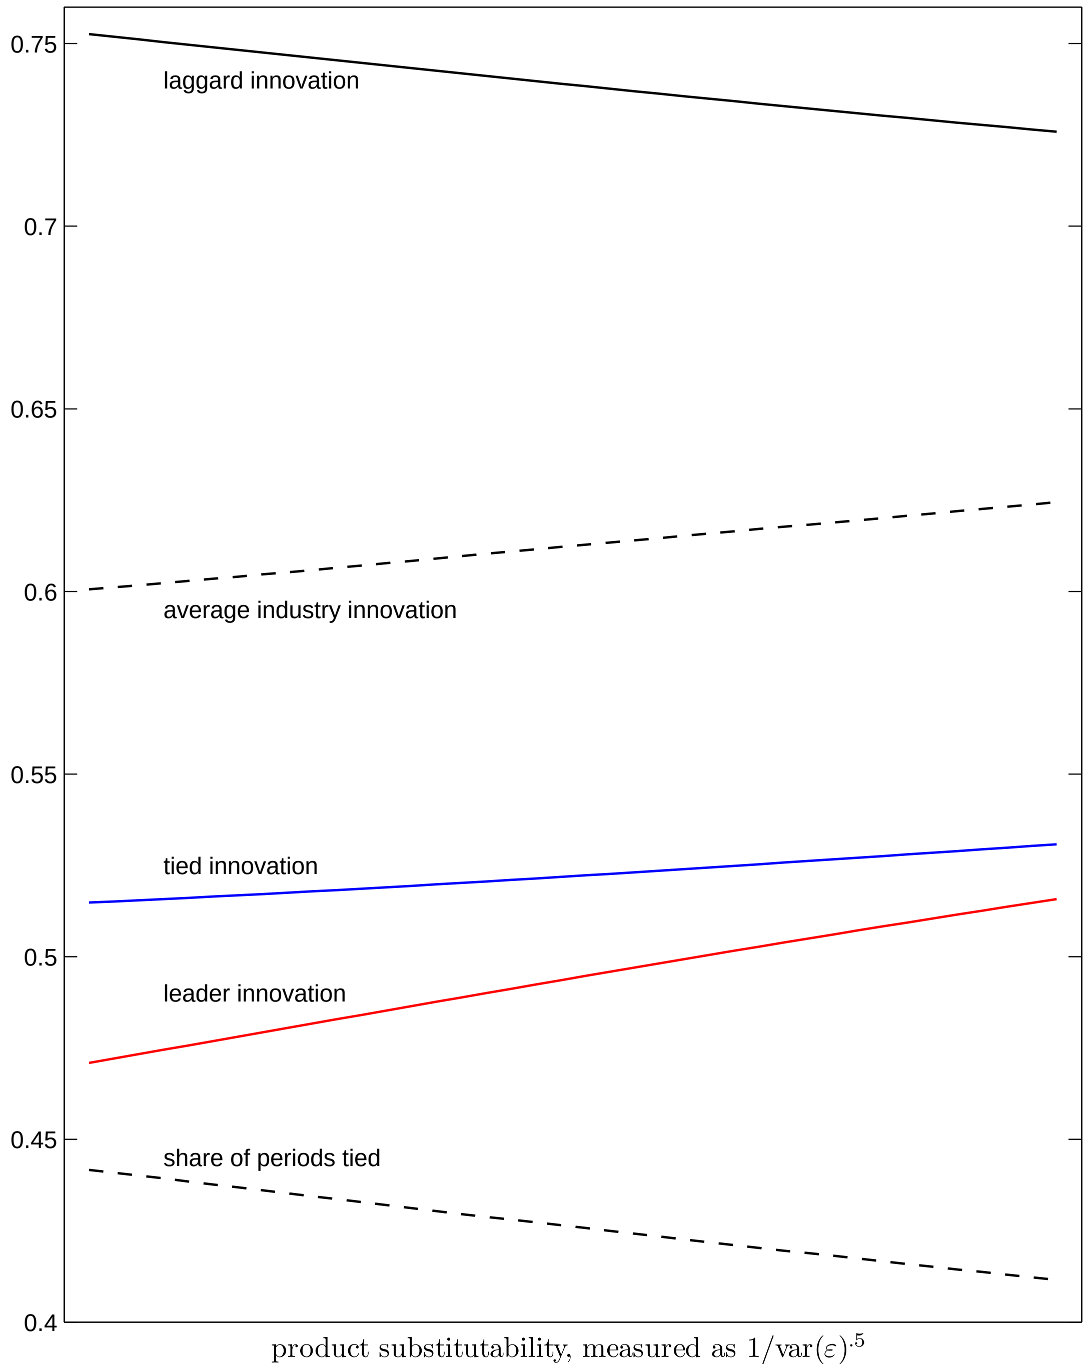

Figure 13: Decomposing Changes in Innovation as γ Increases: Panels 1 and 3 plot the effect of the quality coefficient, γ, on innovation and investment by firms in a symmetric duopoly and monopoly for particular Δ values of 7.25 and 3.5 δ-steps below the frontier, on average across consumers. All panels except panel 5 (bottom-left) use the legend in panel 1 and condition on these same Δ values, which match their respective simulated averages. Panel 2 plots the derivative of firms' innovation probabilities with respect to investment and, for the tied duopolists, the derivative of *either* firm advancing the frontier. Panel 4 is the derivative of investment with respect to γ, and panel 6 is the derivative of innovation with respect to γ. Panel 5 reveals simulated summary measures of which states are visited as γ varies.