

Upstream Innovation and Product Variety in the U.S. Home PC Market*

Alon Eizenberg[†]

November 12, 2008

[JOB MARKET PAPER]

Abstract

This paper asks whether the rapid innovation in Central Processing Units (CPU) results in inefficient elimination of basic Personal Computer (PC) configurations. I estimate a model in which PC makers choose first which CPU options to offer with their products, and then set prices. I contribute to the literature on vertical product choices by relaxing assumptions which guarantee a unique equilibrium outcome, by allowing for a large product space, and by developing techniques which alleviate the burden associated with computing sets of counterfactual equilibria.

My estimates imply that the demand for PCs is highly segmented. Preliminary counterfactual analysis suggests that Intel's Pentium M chip boosted notebook sales by 10.9%-18.9% and increased the average notebook price by \$32 to \$44 in 2004Q2. This innovation also led to a significant re-alignment of PC makers' product offerings, including both the elimination of some PC configurations (e.g. those with Intel's older Pentium III chips), and the addition of other configurations. A traditional model with fixed product offerings does not capture this effect and, as a consequence, significantly understates the impact of the Pentium M on the market share of the Pentium III. The results do not provide strong evidence that this re-alignment has substantially offset consumer welfare gains from the Pentium M's introduction. Finally, I find that the *average* consumer's willingness to pay for a fixed bundle of PC characteristics falls by \$257 every year, consistent with a scenario according to which innovation in software motivates consumers to seek better hardware over time.

*PRELIMINARY RESULTS, COMMENTS WELCOME. I am grateful to my advisors, Steven Berry and Philip Haile, for their continued advice and encouragement. I also thank Donald Andrews, Eduardo Faingold, Josh Lustig, and participants at the Prospectus I.O. Workshop at Yale for many useful comments. I am grateful to IDC and to Mr. Steven Clough for making data available. All errors are mine.

[†]Ph.D. candidate, Department of Economics, Yale University. alon.eizenberg@yale.edu

1 Introduction

Innovation in Personal Computer (PC) technology plays a key role in fostering growth in many economic sectors. A salient feature of this process is a rapid elimination of existing products.¹ The goal of this paper is to ask whether this process results in *inefficient product elimination*. This question is motivated by consumer heterogeneity: while some consumers have a high willingness to pay for the most advanced technology available, others primarily perform basic tasks (e.g. Web browsing) which require modest computing power. This latter group of consumers could be hurt when basic PC configurations exit the market.

To address this question, I estimate a model of supply and demand in which the set of PC configurations offered to consumers is endogenously determined. I then perform counterfactual analysis to determine the extent to which specific innovations crowd out existing products, to evaluate the degree to which the welfare gains from such innovations are offset by product elimination, and to quantify the impact of innovation on various consumer types. The answers to these questions depend on primitives: the distribution of consumer preferences, the variable and fixed costs incurred by PC makers, and the nature of the supply-side game.

I focus on innovation in the Central Processing Unit (CPU), a crucial PC component which is responsible for all calculations. CPU innovation plays a central role in the PC industry: in addition to directly improving PC performance, faster chips also increase the marginal value of complementary innovations in both software and hardware. The CPU market is controlled by two main vendors: Intel, and its smaller competitor Advanced Micro Devices (AMD). Downstream PC makers (e.g. Dell, Hewlett-Packard (HP), Gateway) purchase these chips and install them in their various PC products.

I model a two-stage game played by PC makers: in the first stage, they face a discrete menu of vertically differentiated CPUs, and simultaneously choose which CPU options to offer with their PC products. While consumer heterogeneity provides incentives to offer vertically differentiated PC configurations, offering each such configuration results in fixed costs. In the second stage,

¹Pakes [2003] cites an average annual attrition rate of 85 percent.

the chosen configurations are sold to consumers in a simultaneous price-setting game. CPU innovation expands the menu of CPU options, and I use the model to predict the impact of this expansion on both product choices and prices in the PC market.

I use data on PC prices, characteristics and sales to estimate demand and marginal costs for PC products. These estimates reveal producers' variable-profit benefits from offering PC configurations. I also use the observed variation in product offerings to make inference on fixed cost parameters. For example, an observed decision to offer a certain PC configuration implies an upper bound on the fixed costs associated with it. Having estimated both the benefits and the costs which accrue to PC makers from offering PC configurations, I simulate equilibria of the two-stage game to study the impact of innovation.

My estimates imply that the demand for PCs is highly segmented. In particular, strong consumer heterogeneity is detected in price sensitivity, as well as in the degree to which a consumer's utility from any fixed bundle of PC characteristics falls over time. I find that the *average* willingness to pay for a fixed product falls by \$257 every year. I interpret this as evidence that innovation in software drives the average consumer toward being more of an "advanced PC user" over time.² Consumers also display a considerable willingness to pay for PC brands, suggesting that product choices by some PC makers can have an important impact on the map from upstream CPU innovation to consumer welfare.

I use the estimated model to study the impact of Intel's introduction of its Pentium M chip, which is considered a landmark in mobile computing. Preliminary counterfactual analysis suggests that, in the second quarter of 2004, the presence of the Pentium M boosted total notebook sales by 10.9%-18.9% and increased the average notebook price by \$32 to \$44.

The introduction of the Pentium M also led to a significant re-alignment of PC makers' product offerings; while PC configurations based on Intel's Pentium III (and some very fast Pentium 4 chips) were crowded out, other configurations (mostly based on Intel's Celeron and slow Pentium 4 chips) were added. The presence of the Pentium M decreased the market share of the Pentium

²As discussed below, my sample period was not characterized by significant hardware upgrades driven by a new operating system from Microsoft, so other innovation (e.g. Web applications) is likely to have been the driving force behind this process.

III in the Portables segment from 13.6%-17.6% to merely 2%. Since the bulk of this decrease was due to product elimination, a restricted model in which PC makers' product choices are fixed significantly understates this effect.

I find that 20% of consumers, characterized as being the least price sensitive, garnered as much as 91%-93% of the consumer welfare gains associated with the Pentium M's presence in the second quarter of 2004. At the same time, even though some basic PC configurations were crowded out by this innovation, price sensitive consumers do not appear to be hurt. Moreover, I do not find strong evidence that the re-alignment of PC product offerings prompted by the Pentium M has significantly offset consumer welfare benefits from this innovation.

An important caveat: complementary innovation. While my estimates capture the process by which consumers seek better hardware over time, my framework does not account for the crucial role played by CPU innovation in fostering complementary innovation in software, which fuels this shift in consumer preferences.³

My analysis, therefore, does not account for some long-term contributions of CPU innovation to welfare. For example, some basic users may not benefit from the introduction of an advanced chip in 2004. If, however, this innovation facilitates the emergence of new software applications, these basic users may become more advanced users, and benefit substantially from that CPU innovation by, say, 2006. This motivates future quantitative research of dynamic complementarities in innovative activities.⁴

Multiple equilibria, partial identification, and sample selection. The paper offers a couple of methodological contributions. First, in contrast to previous work with vertical differentiation (e.g. Mazzeo [2002]), I relax assumptions which guarantee a unique equilibrium outcome. This results in partial identification of fixed costs. Following recent literature, I exploit necessary equilibrium conditions to estimate bounds on fixed cost parameters.

³Gawer and Cusumano [2002] describe the manner by which Intel acts to coordinate standards used by hardware and software developers in order to foster complementary innovation, which, in turn, increases the demand for new chips.

⁴See Rosenberg [1979] for a seminal discussion of this issue.

Second, I allow for a large, discrete product space, which provides a detailed picture of PC product variety. This exacerbates the computational burden associated with simulating sets of counterfactual equilibria, as allowing for n product choices yields 2^n feasible vectors of product offerings.⁵ I develop techniques which alleviate this burden. The intuition behind these techniques is that, if a firm can profitably deviate by offering an additional product at a given situation, it would have the same profitable deviation when facing fewer competing products.

A difficult challenge tackled in this paper is sample selection, which arises since firms are explicitly assumed to have chosen the set of products observed in the data. This may bias familiar estimators of parameters governing variable profits. I impose a point-identifying assumption, according to which firms commit to product choices before they observe realizations of cost and demand shocks. In an appendix, I consider relaxing this assumption, and show that the selection mechanism itself can be used to generate moment inequalities which provide partially-identifying information on variable profit parameters. Since I have not yet implemented this alternative approach in practice, its discussion should be viewed as preliminary.

Related literature. Spence [1976] argues that fixed costs restrict the number and variety of products offered in equilibrium, and that the set of products offered by firms may fail to be socially optimal. The potential for such market failures depends on market-specific parameters, motivating empirical research on the determinants of product variety in specific industries.

Song [2006, 2007], Gordon [2008], and Goettler and Gordon [2008] study the upstream CPU market. These papers assume that the CPU serves as a perfect proxy for the PC. The current paper addresses a different set of questions (i.e., PC product variety), and, as a consequence, develops a very different framework.

A vast industrial organization literature considers estimation of partially-identified models (e.g. Haile and Tamer [2003], Pakes, Porter, Ho and Ishii [2006], Berry and Tamer [2006], Ciliberto and Tamer [2007]). Ishii [2006] estimates a model in which banks choose an integer number of ATM

⁵As explained in Section 6, I compute a set of outcomes that is potentially larger than the actual set of equilibria outcomes due to the partial identification.

locations. The discreteness of this choice leads to multiple equilibria and partial identification, similarly as in my framework. My focus on product variety, however, implies that I am interested not only in the total number of PC configurations offered by a firm, but also in their variety. As a consequence, I solve for a vector of binary indicators for each firm.

Trajtenberg [1989] and Petrin [2002] study the welfare benefits associated with new goods. My work adds to this literature by explicitly modeling the impact of innovation on the entire portfolio of products offered, thus taking into account the lost welfare from eliminated technologies.

The rest of the paper is organized as follows: Section 2 describes the industry and the data used. Section 3 presents the model, and Section 4 discusses identification and estimation. Section 5 reports structural estimation results, while Section 6 addresses the economic question of interest via counterfactual analysis. Concluding remarks are offered in Section 7.

2 Data and Industry

The data used in this research come from a number of sources. PC market data is from IDC's Quarterly PC Tracker database. I observe three years of quarterly data (2001Q3-2004Q2) from the U.S. market, including the number of units sold and total dollar value by quarter (e.g. 2002Q3), segment (e.g. Home), vendor (e.g. Dell), brand (e.g. Inspiron), form factor (e.g. Portables), CPU vendor (e.g. Intel), CPU brand (e.g. Pentium 4) and CPU speed range (e.g., 1.0-1.49 GHz) combinations.

As discussed below, I rely on a demand model which assumes that a consumer buys at most one unit of some PC product in a quarter. This is a reasonable assumption for households, but not for commercial PC consumers.⁶ I therefore use only the portion of the data which pertains to the Home segment of the market, and, following previous work (e.g. Goeree [2008]), define the size of the market as the number of U.S. households in a quarter, as reported by the U.S. Census Bureau.⁷ Since PC makers typically target the Home and Commercial segments with

⁶Purchases of the latter were studied by Hendel [1999].

⁷I interpolate linearly between the 2000 and 2004 household totals to obtain quarter-by-quarter figures.

different product lines, it is reasonable to study product choices in the Home market separately.⁸

For each observation, I compute the average price by dividing total value by total sales. I convert values to constant dollars using the Consumer Price Index (CPI), reported by the Bureau of Labor Statistics (BLS). I define a product as a unique combination of observed characteristics.⁹ After removing observations with negligible market shares (defined as selling less than 100 units in the quarter), I obtain 2,287 observations, each of which is a quarter-product pair.

The Home Personal Computer market. The sample period corresponds to the early years of Microsoft’s Windows XP operating system. Due to modest system requirements, the launch of Windows XP did not prompt a widespread hardware upgrade by consumers. This makes the sample period appropriate for the estimation of a model in which the distribution of consumers’ willingness to pay for computing power plays an important role.

Sales in the Home segment accounted for about 38% of total U.S. PC sales during the studied period. While many firms operate in this competitive market, some vendors (most notably Dell and HP) enjoy sizable market shares, as reported in Table 1 (see appendix C for all tables and figures). The top 5 vendors together accounted for a 60%-70% share of the market. Similar concentration is reported by Goeree [2008] for the late 1990s.

The upstream market for CPUs is, by contrast, significantly more concentrated. Table 2 shows that more than 70% of the PCs sold in the Home market had an Intel CPU installed, while slightly over 20% had a CPU from AMD. IBM had a small market share by virtue of making the CPUs used in Apple’s computers. I exclude Apple products from the empirical analysis since I do not have processor speed information for them (Apple’s market share during the sample period hovered about 3%).

Evidence for the rapid innovation in CPU technology is offered in Figure 1, which depicts the share of various CPU clock speed ranges in the three years of the sample. The market share of

⁸Some overlap exists between these markets, since some “Home” consumers purchase PC products designed for commercial users. I discuss below the steps I take to insulate the analysis of the Home market from such spillover effects.

⁹These definitions follows Goeree [2008]. The data used in that paper has a somewhat similar structure to that used in this paper, in that it also consists of 12 quarters, and has similar observed product characteristics.

CPUs with clock speeds in the 2-2.99 GHz range jumped from merely 5% in the first year of the sample to almost 60% by the second year. In parallel, the share of slower CPUs fell sharply over time. It is important to note, however, that clock speed alone is a poor indicator of CPU performance. CPUs of advanced generations (e.g. Intel’s Pentium 4) are differentiated from their predecessors along dimensions other than raw clock speed: they may have more cache memory on board the chip, have better designs, or use more sophisticated algorithms. It is, therefore, important to control for both CPU brand and clock speed to capture CPU performance.

A fundamental force behind CPU innovation has been the ability of manufacturers to double the number of transistors on an integrated circuit every 18-24 months, a regularity known as “Moore’s law”.¹⁰ As a consequence, chips become smaller, faster, less power-consuming, and cheaper to produce. Lower levels of power consumption played a key role in the growth of the mobile PC segment, while lower CPU production costs contributed (among other forces) to a rapid decline in average PC prices. Both these PC market trends are underscored in Figure 2.

PC product lines and CPU technologies. This paper is interested in the portfolio of CPU options offered with PC product lines. I define PC *product lines* as combinations of PC vendor-PC Brand-Form factor (e.g. “Dell-Inspiron-Portables”). I define a *CPU technology* as a combination of CPU brand and speed range (e.g., Intel’s Pentium 4 1.5-1.99 GHz). Typically, multiple configurations of each product line are observed in the data, each with a different CPU technology installed.

Table 3a reports the rate of adoption of Intel’s CPU technologies in Desktop PC product lines.¹¹ The columns of the table correspond to CPU technologies, and the entries report the fraction of PC product lines in which these technologies were offered. The first column, for example, reports the fraction of product lines to adopt Celeron processors with CPU speed in the 0.5-0.99 GHz range. These CPUs were utilized in 89% of product lines in the first quarter,

¹⁰The original prediction by Intel’s Gordon Moore was that the number of transistors on a chip would double and costs would fall by 50% every 18 months (Walters [2001], p.22).

¹¹The analysis in this paper is restricted to PC makers’ decisions to install Intel’s CPUs. An analysis of the variety of AMD chips offered in PCs would be an interesting extension, but would require some careful attention given the asymmetry between the two chip makers.

but were rapidly phased out, in parallel to increased adoption of new CPU technologies. Table 3b reports such information for *portable* product lines.

These tables convey significant variation, in that most CPU technologies are adopted in a subset of product lines only. A substantial amount of this variation, however, stems from technical issues; first, certain CPUs could not be installed in certain PC product lines due to technical constraints (e.g. CPUs with high power consumption could not be installed in some “thin and light” notebooks). Second, some PCs with obsolete CPU technologies may be sold in a given quarter, in small amounts, simply because some firm still has them in stock. I describe below how I take such issues into account when defining the feasible set of CPU technologies.

Importantly, the mobile segment of the market was characterized by both a wider array of Intel CPU technologies, and by a greater degree of variation in installation decisions. This makes this segment more suitable for the analysis of this paper, as this variation is key to the identification of the model. As explained below, while I estimate demand and marginal costs using information on both portable and Desktop machines, I estimate fixed costs and perform counterfactual analysis by focusing on portables only (the interpretation of the estimated fixed cost parameters is, then, that they pertain to the fixed costs of offering notebook product configurations).

3 Model

The primitives of the model are consumer demand for PCs, PC makers’ marginal costs, the fixed costs associated with offering PC product configurations, and the Subgame Perfect Nash Equilibrium (SPNE) concept of a game played by the oligopoly of PC makers. I now describe the model in detail.

3.1 Household Demand

Following Berry, Levinsohn, and Pakes [1995] (BLP), and Goeree [2008], the demand for PCs is modeled by a random-coefficient-logit specification. A set J_t of PC products is available for purchase in quarter t . Each household chooses at most one of the products in J_t , or chooses the

outside option of not purchasing any of the PCs offered. The latter option may include buying a used PC, or buying an Apple computer.¹² The household makes the discrete choice which maximizes the following indirect utility function, describing the utility derived by household i from PC product j at time t :

$$u_{ijt}(\zeta_{it}, x_j, p_{jt}, \xi_{jt}; \theta^d) = \underbrace{x_j \beta + \xi_{jt}}_{\delta_{jt}} + \underbrace{[-\alpha_i \times p_{jt}] + \sum_{k=1}^K \sigma^k x_j^k v_i^k}_{\mu_{ijt}} + \epsilon_{ijt} \quad (1)$$

The following notation is used: x_j is a K -vector of PC product characteristics observed by the econometrician. In the empirical application, these include a constant term, a laptop dummy variable, and dummy variables for PC brands, CPU brands, and CPU speed ranges. I also include a time trend, which captures the degree to which the utility from a fixed bundle of characteristics changes (falls) over time. ξ_{jt} is a quarter-specific demand shock which is unobserved by the econometrician. The product's price is p_{jt} , and $\zeta_{it} \equiv (v_i, \{\epsilon_{ijt}\}_{j \in J_t})$ are household-specific variables: v_i is a $(K+1)$ -vector of standard-normal variables (assumed IID across households, as well as across the $(K+1)$ product characteristics, one of which is price), and ϵ_{ijt} are IID (across households and products) Type-I Extreme Value taste shifters.

I define $\alpha_i \equiv \exp(\alpha + \sigma^p v_i^p)$, so that the price sensitivity is log-normal with parameters (α, σ^p) . The demand parameters are $\theta^d = (\beta', \alpha, \sigma')'$. Note that utility is separated into a mean-utility component δ_{jt} , and a household-specific deviation $\mu_{ijt} + \epsilon_{ijt}$. I further define $\theta_2 \equiv (\alpha, \sigma')'$, and, conditioning on δ , I can write the utility function as $u_{ijt}(\zeta_{it}, x_j, p_{jt}, \delta_{jt}; \theta_2)$.

This specification allows households' taste toward a characteristic $k \in \{1, 2, \dots, K\}$ to shift about its mean, β^k , with the heterogeneous term $\sigma^k v_i^k$. For computational reasons, I restrict many of the σ^k to equal zero in the empirical application. I do allow for heterogeneity in price sensitivity, in the taste for portability, in the taste for the outside option, and in the degree to which that taste changes over time. Heterogeneity along these dimensions governs firms' incentives to provide product variety. I define the utility from the outside option by:

¹²Gordon [2008] models the consumer replacement cycle with respect to CPU products. In order to keep the analysis of product variety tractable, my framework abstracts from durable good aspects of the PC. Incorporating such aspects is an important extension.

$$u_{i0t} = \epsilon_{i0t} \tag{2}$$

The model-predicted market share of product $j \in J_t$ is given by:

$$s_{jt}(x, p, \delta, v; \theta_2) = \int \frac{\exp[\delta_{jt} + \mu_{ijt}(x_j, p_{jt}, v_i; \theta_2)]}{1 + \sum_{m \in J_t} \exp[\delta_{mt} + \mu_{imt}(x_m, p_{mt}, v_i; \theta_2)]} dP_v(v_i) \tag{3}$$

Where $P_v(\cdot)$ is the joint distribution of the taste shifters v_i .

3.2 Supply

I assume that, in each quarter, each PC maker is endowed with a pre-determined set of PC product lines. This assumption is justified by the fact that product lines (e.g. “Dell Inspiron Notebook”) are typically well-established brands that do not frequently enter or exit the market. PC makers also face a menu of CPU technologies which they can offer with their various product lines. The timeline for a two-stage game, played by PC makers in each quarter, is:

1. PC makers simultaneously choose which CPU technologies to offer with each product line; they incur fixed costs for each such offered configuration.
2. For each PC configuration chosen in Stage 1, PC makers observe realizations of demand and marginal cost shocks that are unobserved by the econometrician; they then simultaneously set PC prices for these configurations.

As explained in Section 4 below, the assumption that firms learn the realizations of the error terms only after committing to product choices is key to overcoming a sample selection problem. Since I control for brand-specific intercepts (for most brands), these error terms should not capture any systematic brand effects which the firms are likely to know prior to committing to product choices.

I now turn to a formal description of the game, beginning with some notation. Denote by D the set of active PC vendors (quarter indices suppressed), and define S_d as the set of product

lines for firm $d \in D$. Let H represent the menu of feasible CPU technologies (defined in Section 2 above). Denote by $L_{dm} \subseteq H$ the set of CPU technologies that firm d chooses to offer with product line m .¹³

Stage 1: In this stage, each firm $d \in D$ determines the sets L_{dm} for each product line $m \in S_d$. These decisions are made simultaneously. Collecting all these sets yields the set $J = \{L_{dm}\}_{d \in D, m \in S_d}$ of all PC products that would be offered to consumers in the quarter.

Firm d incurs a fixed cost for each configuration it offers. These costs may include the engineering costs associated with developing and testing product configurations, the inventory management costs necessary to ensure the product is in stock, as well as administrative, sales and marketing costs. I assume that the total fixed costs incurred by firm d are given by:

$$F_d = V_d \lambda \times \sum_{m \in S_d} |L_{dm}| \quad (4)$$

This assumption implies that firm d incurs a constant fixed cost of magnitude $V_d \lambda$ for each configuration, where λ is a parameter vector to be estimated.¹⁴ This implies that total fixed costs are proportional to the total number of configurations offered. This assumption could be relaxed to capture economies (or diseconomies) of scope. V may include a constant and PC maker dummies (which capture systematic firm heterogeneity in fixed costs). The results reported in this paper are based on a simple specification for V which includes a constant, and a dummy variable which receives the value 1 for major PC producers.

Stage 2. I let the log of marginal costs for a PC product $j \in J$ depend linearly on observed cost shifters, w_j , and on an additive error term ω_j :¹⁵

¹³For instance, if $d = \text{"Dell"}$, $m \in S_d$ is Dell's "Inspiron" notebook product line, and $L_{dm} = \{\text{Pentium 4 1-1.49 GHz, Pentium 4 1.5-1.99 GHz}\}$, then Dell has chosen to sell two Inspiron configurations, based on Intel's Pentium 4 CPUs with the specified speed ranges.

¹⁴Ideally, one would like to allow for a structural error term in the per-configuration cost. That, however, leads to selection bias (cf. Pakes et al.). Devising a strategy that would allow for such an error term while overcoming the selection problem in the current context is a work-in-progress.

¹⁵In the empirical application I set $w_j = x_j$, i.e., I let the same observed characteristics shift both utility and marginal cost. Note that the CPU price, charged by Intel or AMD, is a component of PC marginal costs. As a consequence, the γ coefficients on CPU brand and speed provide reduced-form evidence with respect to the manner in which CPU prices vary with such attributes.

$$\log(mc_j) = w_j\gamma + \omega_j \quad (5)$$

In the beginning of Stage 2, firms observe realizations of $e_j = (\xi_j, \omega_j)'$ for each $j \in J$, i.e., for each configuration to which they committed in Stage 1. To re-iterate, these are demand and marginal cost shocks which are unobserved by the econometrician, and appear in (1) and (5) above.

After observing these shocks, firms simultaneously set prices for products $j \in J$ to maximize profits. Firm d 's profits are given by:

$$\pi_d = \sum_{m \in S_d} \sum_{\ell \in L_{dm}} [p_{m\ell} - mc_{m\ell}] s_{m\ell}(p) \times M - F_d \quad (6)$$

where $p_{m\ell}$, $s_{m\ell}$, and $mc_{m\ell}$ are the price, market share and the (assumed constant) marginal cost associated with configuration ℓ of product line $m \in S_d$. M is market size, p is a $|J|$ -vector of prices, and F_d is firm d 's total fixed cost, specified in (4) above.

I assume that, given any Stage 1 history (and any parameter values), Stage 2 prices are determined in a pure-strategy, interior Nash-Bertrand price equilibrium.¹⁶ Arranging products in a $|J|$ -dimensional vector, equilibrium prices satisfy the following (vector) first-order conditions:

$$p - mc = (T * \Delta(p; \theta_2))^{-1} s(p) \quad (7)$$

where T is a $|J| \times |J|$ PC product ownership matrix (i.e., $T_{i,j}=1$ if i, j are produced by the same PC vendor, and is equal to zero otherwise), $\Delta_{i,j}$ is the derivative of the market share of product j with respect to the price of product i , and $*$ represents element-by-element multiplication. It is easy to show that the market share derivatives depend on the non-linear demand parameters θ_2 .

¹⁶This is a familiar assumption (e.g. Nevo [2001]). The results of Caplin and Nalebuff [1991] guarantee a unique price equilibrium under stronger restrictions than those imposed here.

Solution Concept and Multiple Equilibria. A Subgame-Perfect Nash Equilibrium consists of product choices and prices $(J, p(J))$ which constitute a Nash equilibrium in every subgame. As explained above, I assume that Stage 2 prices $p(J)$ are set in a unique Nash-Bertrand equilibrium. In addition, I assume the *existence* of an SPNE for the two-stage game. I do not, however, assume *uniqueness* of the SPNE.

To gain intuition regarding the potential for multiple equilibria, consider the following simple example: suppose we have only two heterogeneous PC makers, each with a single product line. We may have one equilibrium in which only firm A caters to the value segment of the market by offering a PC configuration with a slow CPU installed, and a second equilibrium, in which only firm B chooses to do so.

Finally, recall that even though period indices were suppressed for convenience, the two-stage game is assumed to be played in every quarter. This frequency is justified by the rapid entry and exit of products in the PC market.

4 Identification and Estimation

The parameters to be estimated are the demand parameters $\theta^d = (\beta', \alpha, \sigma)'$, the marginal cost parameters γ , and the fixed cost parameters λ .

Let $\theta = (\theta'_d, \gamma)'$. The estimation strategy employed obtains an estimate of θ first, revealing information on variable profits associated with product configurations. Given the estimate $\hat{\theta}$, necessary equilibrium conditions are used to estimate bounds on the fixed cost parameters λ . These tasks are taken in turn in sub-sections 4.1 and 4.2 below.

4.1 Identification and Estimation of $\theta = (\beta', \alpha, \sigma, \gamma)'$

Intuitively, the demand parameters are identified from the joint distribution of prices, sales, and observed PC characteristics. Marginal cost parameters γ are identified as follows: the pricing FOCs in (7) identify markups, allowing us to identify marginal costs as the difference between observed prices and these markups. The co-movement of these identified marginal costs with PC

characteristics identifies γ .

Identification of θ is jeopardized, however, by sample selection, as the set J of product configurations offered to consumers was selected by firms. The econometrician, therefore, does not observe a random sample from the underlying distribution of product characteristics. In this section, I describe a standard approach which allows point-identification of θ . It also allows me to consistently estimate θ following the BLP method. These estimates are reported in Section 5 below. In appendix B, I describe an alternative approach, in which point-identifying assumptions are relaxed.

The intuition for the point-identification approach is that, under the assumption that firms do not observe the error terms $e_j = (\xi_j, \omega_j)'$ until after they have selected their products, the selection does not depend on unobservables, and is therefore “ignorable”.¹⁷ Stating the point-identifying conditions requires a bit more notation. Let us collect all firms’ product lines in the set $P = \{S_d\}_{d \in D}$. Denote by \mathbf{J} the set of all $|H| \times |P|$ *potential* product configurations. It is from this set that firms pick, in Stage 1, the subset $J \subseteq \mathbf{J}$ actually offered to consumers. Let X denote a $|\mathbf{J}| \times K$ matrix of product characteristics for all potential products.

Denote by F the fixed costs of all PC makers. I make the following assumption:

Assumption 1. $E[e_j|X, F] = 0$ for each $j \in \mathbf{J}$

Assumption 1 is very similar to the mean-independence assumption made by BLP, except that the relevant population here is that of all potential PC configurations, rather than the sub-population of products actually offered to consumers.

For each potential product configuration $j \in \mathbf{J}$, I define a selection indicator, $q_j(X, F)$, equal to 1 if j was chosen for production, and equal to zero otherwise. This indicator does not depend on the error terms e_j because firms do not know these values when making their Stage 1 product choices. This allows for a standard identification approach: let $z_j(X)$ be a $1 \times L$ vector of instrument functions pertaining to product j , where $L \geq \dim(\theta)$. By application of the Law of Iterated Expectations, and using Assumption 1, we obtain:

¹⁷See Wooldridge [2000], ch. 17. for a general discussion of the implications of selection mechanisms which depend on variables observed by the econometrician.

$$E[q_j(X, F)e_j z_{j\ell}(X)] = 0 \text{ for } \ell = 1, \dots, L \quad (8)$$

BLP show that a generic value for the parameter θ implies a unique solution $e_j(\theta)$ for each observed product $j \in J$. As a consequence, as long as $Pr[q_j = 1] > 0$, condition (8) implies:

$$E[e_j(\theta_0)z_{j\ell}(X)|q_j = 1] = 0 \text{ for } \ell = 1, \dots, L \quad (9)$$

where θ_0 is the true parameter value. Equation (9) defines L moment conditions that provide point identification of θ .¹⁸ Notice that we overcome the selection problem by obtaining a moment condition that is defined over observed products only.

Since firms observe the errors e before setting prices, it is necessary to account for price endogeneity. In choosing the instruments $z_j(X)$, I follow Berry [1994] and BLP by using variables that should be correlated with markups, and, therefore, with prices. In addition to the x_j vector of PC characteristics, I use the number of product lines for both the vendor and competitors in various data cells (e.g., formfactor-speed cells), the number of competitors' Celeron-based configurations, the squared time trend, and the ratio of average rivals' speed to vendor's average speed.¹⁹ I also use interactions of observed PC characteristics (laptop, Pentium and Celeron dummy variables) with a time trend to obtain additional instruments. These terms can be viewed as cost shifters excluded from the demand side, since they capture the decrease in the marginal costs of providing these PC characteristics over time.

GMM estimation of θ using the moment conditions (9) follows the BLP method. Additional details regarding this estimation procedure are provided in Appendix A.1.

Finally, note that the identification strategy for θ outlined here relies heavily on the assumption that firms observe the errors e_j only after committing to product choices. In the absence of this assumption, the selection indicator $q_j(\cdot)$ would depend on these errors, and, as a consequence, condition (8) could fail. Appendix B offers the details of an alternative identification

¹⁸Additional regularity conditions are necessary for a formal identification argument.

¹⁹For the purpose of constructing this instrument I compute speed as the middle of the relevant speed range.

strategy which relaxes this assumption. That strategy shows that, given a parameter value θ , bounds can be placed on the e_j error terms associated with products that firms chose *not to introduce*. This allows me to construct moment inequalities that are defined over the entire set \mathbf{J} of *potential* products. As explained in the introduction, that alternative strategy should be viewed as preliminary and has not been implemented yet in practice.

4.2 Identification and Estimation of Fixed Cost Parameters λ

Given the point estimate $\hat{\theta}$ obtained in section 4.1, a set estimate can be obtained for λ . I assume that the product choices and prices observed in the data constitute an SPNE of the two-stage game. A necessary equilibrium condition is, then, that no firm could increase its expected profit by unilaterally altering its first-stage product choices, taking into account the impact of that deviation on second-stage prices (the simultaneous-move nature of the first stage implies that the firm need not consider an impact of its deviation on rivals' *product choices*). Such conditions imply bounds on expressions involving fixed cost parameters.²⁰

I let the vector A_d denote firm d 's observed product choices. Each entry in this vector is a binary variable, which takes the value 1 if the relevant product configuration is chosen for production. Since firm d may have more than one product line (i.e. the set S_d may not be a singleton), the typical form of this vector is:

$$A_d = \left\{ \underbrace{0\ 1\ 1\ 0\ 1}_{\text{Product Line 1}}, \quad \underbrace{1\ 1\ 0\ 1\ 1}_{\text{Product Line 2}}, \dots \right\}$$

I define the sets $A_d^1 = \{k : A_d(k) = 1\}$ and $A_d^0 = \{k : A_d(k) = 0\}$, which collect the indices in A_d corresponding to products offered and not offered in the observed sample, respectively.

Upper and lower bounds on $V_d\lambda$. Recalling that firm d 's per-configuration fixed costs are given by $V_d\lambda$, upper bounds can be placed on this quantity at the true parameter values:

²⁰See cf. Berry and Tamer for a discussion of the use of necessary equilibrium conditions in the context of partially-identified entry models.

$$V_d \lambda_0 \leq E_{(e|\theta_0)} \left[VP_d(A_d; e, \theta_0) - VP_d(A_d - \mathbf{1}_d^k; e, \theta_0) \right] \equiv U_{d,k}(\theta_0), \quad \forall k \in A_d^1 \quad (10)$$

where $\mathbf{1}_d^k$ denotes a vector of the same length as A_d which k^{th} entry is equal to 1, and all its other entries are equal to zero. $VP_d(\cdot)$ denotes the variable profit firm d garners as a consequence of choosing various product portfolios (taking into account the impact of such portfolios on second-stage prices). $E_{(e|\theta_0)}$ denotes expectation over the true joint distribution of the error terms associated with all products. This notation reflects the fact that this distribution is indexed by the parameter θ (see Appendix A.2).

In words, condition (10) states that a deviation by firm d which eliminates one of its observed products must not be profitable. To ensure that, firm d 's savings in fixed costs cannot exceed the expected drop in its variable profit. An analogous argument generates lower bounds as follows:

$$V_d \lambda_0 \geq E_{(e|\theta_0)} \left[VP_d(A_d + \mathbf{1}_d^j; e, \theta_0) - VP_d(A_d; e, \theta_0) \right] \equiv L_{d,j}(\theta_0), \quad \forall j \in A_d^0 \quad (11)$$

Using (10) and (11), the identified set for $V_d \lambda$ is:

$$\left[\max_{j \in A_d^0} \{L_{d,j}(\theta_0)\}, \min_{k \in A_d^1} \{U_{d,k}(\theta_0)\} \right] \quad (12)$$

where θ_0 is a parameter that was already identified in sub-section 4.1 above.

One possible estimation approach is to treat firm d 's per-configuration fixed cost $V_d \lambda$ as a firm-specific parameter to be estimated, and place bounds on this parameter using (12). To do that, it is necessary to replace the bounds $L_{d,j}(\theta_0)$, $U_{d,k}(\theta_0)$ which appear in (12) with estimates $\hat{L}_{d,j}(\hat{\theta})$ and $\hat{U}_{d,k}(\hat{\theta})$. The computational details associated with estimating these quantities are provided in Appendix A.2 (this estimation requires simulating the expectations which appear in (10) and (11) by drawing from an empirical distribution of e given $\hat{\theta}$, and, at each such draw, computing price equilibria which would prevail under the observed action, and under the deviation).

A practical problem which arises in following this approach is that the estimated endpoints of the interval given in (12) tend to cross. As pointed out by Haile and Tamer [2003], this may

occur even if the model is correctly specified due to finite-sample bias. Intuitively, by focusing on $\max_{j \in A_d^0} \{L_{d,j}(\theta_0)\}$ we choose a lower bound which tends to contain an upward bias (and, analogously, we choose an upper bound that tends to be downward-biased).

As an alternative approach, I treat the parameters λ , rather than the firm-specific index $V_d\lambda$, as the object to be estimated. Recent econometric literature has developed techniques for making inference on partially-identified parameters (e.g. Andrews and Guggenberger [2007], Chernozhukov, Hong and Tamer [2007]).

The estimation results for λ reported in Section 5 below, however, were obtained by a simple heuristic approach: I searched for the values of λ which made the condition $V_d\lambda \leq \hat{U}_{d,k}(\hat{\theta})$ hold on average across all $d \in D$, $k \in A_d^1$, and at the same time made the condition $V_d\lambda \geq \hat{L}_{d,j}(\hat{\theta})$ hold on average across all $d \in D$, $j \in A_d^0$. In words, the estimated set for λ consists of the parameter values which make $V_d\lambda$ respect both the upper bounds and the lower bounds *on average*. I plan to improve on this in future versions by following the literature cited above.

Additional restrictions implied by the model. Necessary equilibrium conditions could be used to obtain additional information on the model's parameters. The results reported in Section 5 below *do not* make use of such additional information. First, one could provide additional bounds on $V_d\lambda$ by means of considering more complex deviations: for example, a deviation in which two observed products are eliminated, and one unobserved product is introduced would provide an upper bound on $V_d\lambda$.

Second, we can consider a deviation in which the firm eliminates an observed product located at k and launches an unobserved product located at j instead. This deviation does not alter total fixed costs, since these depend only on the number of configurations offered (see (4) above). Requiring that such deviations are not profitable would yield the following conditions:

$$N_{d,j,k}(\theta_0) \equiv E_{(e|\theta_0)} \left[VP_d(A_d; e, \theta_0) - VP_d(A_d + \mathbf{1}_d^j - \mathbf{1}_d^k; e, \theta_0) \right] \geq 0, \quad \forall d \in D, j \in A_d^0, k \in A_d^1 \quad (13)$$

This condition could, in principle, be used to provide additional information on θ (on top of the point-identifying information described in sub-section 4.1 above)²¹.

5 Estimation Results

Section 5.1 below reports estimation results for θ , beginning with descriptive results based on the simple logit demand model, and continuing with results from the full model described in Section 3. Section 5.2 provides estimated bounds on fixed cost parameters.²²

5.1 Estimation Results: Demand and Marginal Cost Parameters, θ

It is instructive to begin with a simple, descriptive outlook on the demand system. Table 4 reports demand estimation results based on the simple logit model, which is obtained from the demand model described in Section 3.1 by setting all the σ coefficients to zero, so that consumer heterogeneity is only allowed via the additive IID ϵ_{ijt} term. Estimation is performed via linear regressions following Berry [1994]. The first column provides OLS estimates of the mean utility parameters β , while the second column employs 2SLS to account for the endogeneity of price using the instruments described in Section 4.1 above.

These results demonstrate the importance of correcting for price endogeneity. While demand is downward-sloping in both specifications, the price sensitivity coefficient is much larger (in absolute value) in the IV case. The results suggest that households value CPU speed as well as high-end CPU brands (the omitted CPU brand is Intel's Celeron). The taste for portability appears negative and insignificant, a point to which I return below. The negative sign on the time trend reflects the fact that a given bundle of characteristics becomes obsolete over time, most likely due to the emergence of advanced software applications which require better hardware.

²¹In Appendix B I show that a similar condition to (13) facilitates partial identification of θ even if we relax the assumption that firms observe the shocks e only after committing to product choices.

²²Some robustness checks are still needed with respect to the results reported. First, somewhat (but not dramatically) different estimates obtain for different starting values. Second, potential measurement error could stem from the fact that, in some cases, observations pertaining to the same PC vendor and quarter report identical unit sales.

BLP estimation results for θ . By contrast to the simple logit model, the random-coefficient demand model described in Section 3 allows for more realistic substitution patterns (see the discussion in BLP), and captures consumer heterogeneity along important dimensions. Tables 5a and 5b provide estimation results for θ obtained by following the BLP estimation procedure. Table 5a reports the estimated coefficients on main PC characteristics, while Table 5b reports estimated coefficients on a large number of dummy variables for PC vendors and brands. Economic implications of these estimates are offered in Table 6. The estimated parameters include mean utility parameters (β), parameters which capture heterogeneity in household tastes (σ), marginal cost parameters (γ), and the parameters of the distribution of price sensitivity.

The results in Table 5a reveal precise estimates of both the mean (α) and the variance (σ^p) parameters of the log-normal price sensitivity. As in the simple logit results, households value CPU speed, as well as CPU brands, and these effects are very precisely estimated. The mean taste for laptop products is negative and imprecisely estimated, but significant heterogeneity in this taste is captured by the precisely-estimated σ coefficient on the laptop dummy. Heterogeneity along this dimension is to be expected.

As in the simple logit results, the negative β coefficient on the time trend implies that a fixed bundle of characteristic is becoming obsolete over time. The random-coefficient model allows me to precisely estimate, in addition, the degree of household heterogeneity in this important effect. I return to this issue below in the discussion of the quantitative economic implications of the estimated coefficients.

The marginal cost coefficients γ are all very precisely estimated and economically reasonable. Producing a laptop is found to be 31.2% more expensive than producing a Desktop. Installing an Intel Pentium 4 instead of a Celeron CPU drives PC marginal costs up by a similar magnitude of 30.5%. The negative coefficient on the time trend implies that PC marginal costs fell at a rate of 9% per quarter. This is consistent with the sharp decline in PC prices depicted in Figure 2.

Table 5b reports a large number of estimated coefficients on dummy variables for PC vendors (e.g. Dell) and their various brands (e.g. Inspiron). Importantly, the coefficient on “Dell”

captures the effect of Dell brands not included, and not an “overall” Dell effect. Most of the effects are very precisely estimated. Controlling for brand and vendor information is useful, as these should be strongly correlated with unobserved quality. Moreover, had I not controlled for these brand effects, they would have showed up in the error terms e_j , making it less reasonable to assume that firms do not observe these errors until after committing to their first-stage configuration choices²³.

Table 6 offers an insight into some important economic implications of the results presented above. Panel A of this table reports the willingness of the average household to pay for various product characteristics. The average household is willing to pay up to \$150.1 to upgrade from CPU speed in the 2-2.99 GHz range to the next speed range, 3-3.99 GHz. It is also willing to pay up to \$171.5 for an upgrade from the Intel Celeron to the Intel Pentium 4 brand, and up to \$447.3 for an upgrade to Intel’s Pentium M.

These are considerable amounts, suggesting that CPU characteristics are important to the average PC consumer. Recall also that an entire distribution of these figures was actually estimated. One would expect some consumers (e.g. gamers, engineers) to be willing to pay much more than the average consumer for a better CPU. Figure 3 plots the estimated distribution of households’ willingness to pay for an upgrade from Intel’s Celeron to its Pentium M brand and reveals significant heterogeneity along this dimension.

Households are also willing to pay considerable amounts for a familiar PC brand name. The average household is willing to pay \$107.8 to upgrade from a non-branded notebook computer to Dell’s Inspiron brand, and \$462.1 for IBM’s ThinkPad A series. These results indicate that downstream PC makers possess powerful brand names, suggesting that their product choices may have an important impact on welfare.

A crucial parameter in the analysis is the pace at which households’ utility from a given bundle of PC characteristics drops over time. Using the estimated coefficients on the time trend, Table 6 reports that the average consumer is “willing to pay” a negative amount of \$(-257) for a passing

²³I do not, however, control for every brand, but rather for a large number of them.

of one year. My interpretation of this finding is that, *holding everything else equal*, the average household’s willingness to pay drops by \$257 every year. This finding is consistent with a scenario according to which new software applications make a fixed hardware system less attractive over time. One may also expect households to be strongly heterogeneous in this respect, and, indeed, such heterogeneity is displayed in Figure 4.

To summarize, a key finding stemming from the estimated demand parameters is that households display strong heterogeneity in price sensitivity, as well as in the rate at which products become obsolete from their point of view. This heterogeneity affects both PC makers’ incentives to offer vertically-differentiated product configurations, and the welfare implications of such product choices.

Panel B provides some additional economic implications of the BLP estimates for θ . The median markup for a PC manufacturer is \$76.4, and the median price-cost margin (markup as a percentage of price) is 7.8%. As expected, markups are positively and strongly correlated with prices. Another intuitive finding is the positive correlation between the estimated demand and cost-side errors, $\xi_j(\hat{\theta})$ and $\omega_j(\hat{\theta})$.

5.2 Estimation Results: Fixed Cost Parameters λ

Bounds on fixed costs were estimated as explained in section 4.2 above. The results below are based on constructing bounds on fixed costs associated with offering product configurations of portable product lines in the last quarter of my sample (2004Q2).²⁴

As a first step, it is necessary to determine the set H of CPU technologies available in 2004Q2. As reported in Table 3b, many CPU technologies were installed in PCs sold in this quarter. Some of these, however, were very obsolete technologies and sold in rather small quantities. I conjecture that such technologies were on the market due to dynamic inventory issues (e.g., some retailer clearing a stock of rather obsolete PCs). While data limitations prohibit me from directly addressing such issues, I do not want my estimates of fixed costs to be biased by such scenarios,

²⁴I restrict attention to this quarter, at this point, due to the associated computational burden.

and I therefore include in the set H only CPU technologies that appear to have been installed in a significant number of PC products, and have generated significant sales:²⁵

$$H = \left\{ \mathbf{C}_{1.5-1.99}, \mathbf{C}_{2-2.99}, \mathbf{P4}_{1.5-1.99}, \mathbf{P4}_{2-2.99}, \mathbf{P4}_{3-3.99}, \mathbf{Pm}_{1-1.49}, \mathbf{Pm}_{15-1.99} \right\}$$

where \mathbf{C} , $\mathbf{P4}$ and \mathbf{P}_m stand for Intel’s Celeron, Pentium 4 and Pentium M brands, respectively, and the number ranges pertain to clock speed (e.g., $\mathbf{C}_{1.5-1.99}$ are Celeron chips with clock speed between 1.5-1.99 GHz). I also applied some refinements to the set of portable product lines considered for generating bounds on fixed costs. I excluded some product lines that either primarily targeted the commercial PC market, or could not install certain CPU technologies due to technical constraints, in cases where I was aware of such issues.²⁶

I report the estimation results for λ in Table 7. Specification (a) includes only a constant term, and specification (b) allows for both a constant, and a dummy variable getting the value 1 for manufacturers which produce a large volume of notebooks. Including only a constant in V amounts to assuming that the per-configuration fixed costs are the same for all firms, and the estimated set for these costs is the interval [1,165,999 1,506,040] (\$). Specification (b) implies that this cost is in the [980,000 1,500,000] interval for small firms and in the [1,170,000 2,270,000] interval for large firms.

6 Using the Estimated Model: Counterfactual Analysis

I analyze the impact of Intel’s introduction of its Pentium M processor, which is considered a major innovation in mobile computing. Rather than offering a further increase in clock speed, this innovation introduced major improvements in chip design that allowed chips to achieve top performance at modest clock speeds. This resulted in a substantial reduction in power

²⁵Robustness checks and further refinement of these judgment calls are necessary.

²⁶The judgment calls I made in this respect require some additional refinement. An additional issue is that, as explained above, I exclude products which sold less than 100 units in a quarter from the sample due to computational reasons, and I also consider such a product as “not offered” for the purpose of constructing bounds.

consumption and in longer notebook battery life.²⁷ Pentium M-based notebooks appear in the sample for the first time in the first quarter of 2003 (see Table 3b).

In parallel to the introduction of the Pentium M, older Intel mobile CPUs such as the Pentium III gradually left the market. In the last sample period, i.e., the second quarter of 2004, only 2% of notebooks sold were Pentium III-based.²⁸ Among the five top-selling notebook brands in that quarter, only one recorded positive sales of a Pentium III-based configuration.²⁹

In the quarter immediately preceding the Pentium M's introduction, however, Pentium-III based configurations were offered by three of these five brands (including the two with the highest sales), and had a market share of 14.1%. While this could suggest that the Pentium M was responsible for the elimination of the Pentium III, a more careful analysis is required in order to isolate the effect of the Pentium M's presence from the many other forces that operated in the PC market between 2003Q1 and 2004Q2.

To identify the effect of the Pentium M on product offerings and prices in the PC market, I perform the following counterfactual analysis for the 2004Q2 period: I remove the Pentium M chips from the set H of CPU technologies available for installation. Then, I use the estimated model to compute the set of PC configurations, and PC prices, that would have prevailed in the market in the absence of the Pentium M. Comparing these predictions to the observed outcomes provides a measure of the Pentium M's effect. Since I am especially interested in the effect of the Pentium M on the Pentium III, I include in the set H a Pentium III option with speed in the 1.5-1.99 GHz range.³⁰ This allows me to ask how many Pentium III-based PC configurations would have been offered *in the absence of the Pentium M*.

I focus my analysis on configuration choices by the five top-selling notebook brands in 2004Q2. By 2004Q2, the Pentium M's market share in the notebook segment of the market reached 31.8%. This makes its analysis interesting at that point in time; an earlier analysis, at a point when this

²⁷ "Bigger Notebooks Still Using Older Mobile Chips", Tom Krazit, IDG News Service September 28, 2004.

²⁸ Excluding Apple products, PCs with CPUs not made by Intel or AMD, and products with negligible sales.

²⁹ That configuration had very small sales, and it is possible that it recorded positive sales simply because a small remaining stock was cleared.

³⁰ This is the fastest Pentium III chip observed in a mobile PC in the sample. It was actually offered in a handful of PC product lines only. To clarify: the set H in this experiment is obtained from that described in Section 5.2 above by removing the Pentium M technologies and adding the Pentium III 1.5-1.99 GHz technology.

chip was making more modest sales, would have been of limited interest.

Computing the set of counterfactual equilibria. We are interested in the set of SPNE outcomes of the two-stage game under the “no Pentium M” scenario. No equilibrium selection mechanism is imposed. Instead, I would like to compute welfare measures at each of the multiple equilibria, and place bounds on welfare predictions. As explained next, however, partial identification makes it necessary to perform this for a set of outcomes that is potentially larger than the actual set of equilibria.

Recall that A_d was used to denote a vector of binary indicators describing the product choices of firm $d \in D$. I will now use this notation more generally to describe product choices by firm d (not necessarily the observed ones). Let $A = \{A_d\}_{d \in D}$ be a long vector which describes product choices by all firms, and let \mathbf{A} be the set of all such vectors. The set \mathbf{A} has $2^{|A|}$ elements. I define the subset $A^e \subseteq \mathbf{A}$ as the collection of product choice vectors which can be supported in an SPNE of the two-stage game.

In order for a vector A to be an element of A^e , it must be the case that no firm has a unilateral, profitable deviation from A . Fixed costs, however, are only partially-identified, and so is the profitability of deviations. As a consequence, it is not always possible to unambiguously determine whether $A \in A^e$.

To deal with this issue, I define a set $A^{pe} \supseteq A^e$ which contains all elements $A \in \mathbf{A}$ that cannot be unambiguously ruled out as members of A^e . Recall that I assume that A^e is not empty (i.e., at least one equilibrium outcome exists). Under this assumption, the equilibrium played must be included in A^{pe} . Computation of the set A^{pe} , which I refer to as the set of “potential equilibria,” is a very difficult computational task: in principle, one needs to check for profitable deviations from each of the $2^{|A|}$ vectors in \mathbf{A} .

I allow for six CPU options and five PC brands, and so $|A| = 30$. I reduce this number to 24 by requiring that a firm which owns two of the five brands makes the same configuration portfolio choice on both. This leaves me with the task of evaluating 2^{24} vectors (with each such

evaluation requiring computation of the price equilibria that prevail under the various product-choice deviations). I was able to significantly reduce the computational burden by application of the following conjecture:

Conjecture 1. (*Strategic Complements*): *The increase in firm d 's variable profit from adding a product configuration at $A = (A_d, A_{-d})$ is at least as large as at (A_d, A_{-d}^*) where $A_{-d}^* \geq A_{-d}$*

where A_{-d} denotes product choices by firm d 's competitors, and $A_{-d}^* \geq A_{-d}$ implies element-by-element inequality. Conjecture 1 is very intuitive: it suggests that the benefit from adding a product configuration is lower when the firm faces more competing products.³¹ The usefulness of this conjecture is in that, once a certain element of \mathbf{A} is ruled out as a “potential equilibrium”, many other vectors can be automatically ruled out as well. This allowed me, in practice, to evaluate 16,384 vectors rather than $2^{24} = 16,777,216$, resulting in an immense reduction in computation time.³²

The results reported below were obtained by setting the shocks to mean utilities and marginal costs e_j at their mean of zero. In fact, to be consistent with the model, I should have simulated firms' expected variable profits by drawing from the distribution of the error terms, as I did in the estimation of the fixed costs. The results reported below should, therefore, be viewed as preliminary, and I plan to improve on this in future versions.

Counterfactual results. Table 8 reports the results obtained from the counterfactual experiment which removes the Pentium M CPU technology from the market in 2004Q2. In addition to the results from the full model, in which manufacturers endogenously choose their PC configuration portfolio under the “no Pentium M” scenario, I also report the results from a restricted model, which does not allow for such endogenous product choices. This restricted analysis simply removes Pentium M-based PCs from the sample, and calculates a counterfactual price equilibrium. The two left-hand columns report upper-bound and lower-bound results from the full

³¹This conjecture is difficult to prove. I did, however, test it directly in more than 20,000 simulations, and found that it was validated in each of them.

³²Total run time for the Matlab code was about 13 hours on a Desktop PC with an Intel Quad Core processor.

model, and the most-right hand column reports results from the restricted model.

The full model finds that Intel's Pentium M "crowded out" between 2 and 4 PC configurations based on Intel's Pentium III with speed in the 1.5-1.99 GHz range, and between 1 and 3 configurations based on Intel's Pentium 4 in the 3-3.99 GHz range. The latter technology was a rather direct competitor for the Pentium M in the high-end notebook market. On the other hand, this product elimination was accompanied by more intense offerings of PC configurations based on Intel's Celeron and slower Pentium 4 chips.

The counterfactual market share of Pentium III-based notebooks is between 13.6% and 17.6% under the full model, and only 3.3% under the restricted model. The observed share was only 2%. The table therefore reports that the Pentium M's impact on this share was between (-11.6) and (-15.6) percentage points under the full model, but only (-1.2) points under the restricted model. This demonstrates the importance of treating product choices as endogenous: the restricted model does not capture product elimination, and, as a result, significantly understates the Pentium M's impact on the market share of the Pentium III.

The Pentium M boosted total notebook sales by 10.9%-18.9% according to the full model. Some of this growth came at the expense of Desktop sales, which were depressed by 1.6% to 2.6%. This high-quality chip also increased the sales-weighted average notebook price by \$32 to \$44.

The bottom part of the table reports the innovation's impact on consumer surplus. Both the full model and the restricted model suggest that the vast majority (91%-93%) of consumer benefits from the innovation were garnered by the 20% least price sensitive consumers. The full model accounts for product elimination of basic PC configurations, and, therefore, can be used to ask if some household types are actually *hurt* by innovation. The table reveals that no such effect is observed: consumers who are strongly price-sensitive are largely unaffected by this innovation, but do not appear to be hurt. Moreover, recall from the introduction that my framework does not account for a long-run effect of CPU innovation, i.e., the fact that it fosters complementary innovation that is likely to benefit even price-sensitive consumers.

The impact of product re-alignment on welfare. The results discussed above reveal that the presence of the Pentium M led to a re-alignment of PC makers' product offerings: in particular, PC configurations with Pentium III and very fast Pentium 4 chips were crowded out, while some PC configurations with other chips (the Celeron and slow Pentium 4) were actually added. One way to assess the welfare implications of this re-alignment is to compare the welfare predictions of the restricted model, which does not take this effect into account, to those of the full model.

The restricted model finds that the Pentium M contributed \$42.1 million to total consumer surplus in 2004Q2, corresponding to a 4.9% increase. The full model, which takes the product re-alignment effect into account, finds this contribution to be in the [27.92M, 41.87M] interval. Since the upper bound of this interval is close to the figure implied by the restricted model, it is not possible to determine that the welfare gains from the Pentium M were substantially offset by PC makers' re-alignment of their product offerings. On the other hand, the lower bound figure of \$27.92 million still leaves as open the possibility that this re-alignment offset the benefits from the innovation by about a third.³³

Finally, recall that the results reported here are preliminary at this point. An important robustness check would involve altering counterfactual marginal costs; in the absence of the Pentium M, Intel might have charged a higher price for its other chips, which would have increased the marginal costs associated with some PC configurations. Adjusting the analysis to account for that is a work-in-progress. The reduced-form evidence on systematic differences in CPU prices available from the estimated marginal costs of PC makers provide a useful source of information in this context.

³³This is merely a preliminary analysis of the efficiency properties of post-innovation product re-alignment. I plan to perform a more complete analysis which would consider a social planner's problem of choosing the optimal product portfolio. That analysis would take into account the profits of PC makers.

7 Concluding Remarks

This paper asks whether CPU innovation leads to an inefficient elimination of existing PC products. To address this question, I estimate a model in which PC makers endogenously choose which CPU options to offer with their PC product lines. I relax strong assumptions which guarantee a unique equilibrium outcome, and exploit necessary equilibrium conditions to tackle the resulting partial identification of fixed costs.

I provide a rich analysis of PC product variety by allowing for a large product space, which requires me to develop computational techniques which alleviate the burden associated with computing sets of counterfactual equilibria. I overcome a sample selection problem by imposing a point-identifying assumption, and provide the details of an alternative approach, which allows one to relax this assumption and obtain partial identification of variable-profit parameters.

I find that the demand for PCs is highly segmented, such that households differ considerably in their price sensitivity, and in the pace at which PC products become obsolete from their point of view. I also find that the average household's willingness to pay for a fixed bundle of PC characteristics falls substantially over time, consistent with a scenario according to which software innovations create incentives for hardware upgrades.

Preliminary counterfactual results reveal that the introduction of Intel's Pentium M has led to a significant re-alignment of PC maker's product offerings: PC configurations based on the Pentium III and on very fast Pentium 4 chips were crowded out, while configurations based on the Celeron and slow Pentium 4 chips were added. A traditional model with fixed product choices fails to account for this effect, and, as a consequence, substantially understates the effect of the Pentium M on the market share of the Pentium III.

The results reported indicate that the vast majority of the benefits from innovation accrue to the least price sensitive consumers. At the same time, even though some Pentium III-based PC configurations were eliminated, price-sensitive consumers do not appear to be hurt by innovation (which, most likely, stems from the more intense offerings of Celeron-based PCs). Moreover, I do not find strong evidence that consumers' welfare gains from the Pentium M's introduction were

substantially offset by the re-alignment of PC product offerings prompted by this innovation.

A couple of interesting issues are left for future research. While I do not impose an equilibrium selection mechanism, my framework could be used to investigate it. Ciliberto and Tamer [2007] test (and reject) the hypothesis that firms coordinate on the equilibrium outcome which maximizes total industry profits in their study of the airline industry. An interesting exercise in the current framework could be to compute the set of equilibria for a given quarter, and then ask what was special about the equilibrium that was actually played in the observed sample.

An important aspect of CPU innovation is that it fosters complementary innovation in software and hardware. Such complementary innovation prompts households to use more advanced applications, which, in turn, increases the demand for advanced CPUs. A quantitative, dynamic analysis of this “positive feedback loop” is likely to improve our understanding of the singular contribution of CPU innovations to growth in the 21th century economy.

References

- Andrews, D. W. K., and P. Guggenberger (2009): "Validity of Subsampling and 'Plug-in Asymptotic' Inference for Parameters Defined by Moment Inequalities," forthcoming *Econometric Theory*
- Berry, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25(2): 242-262
- Berry, S., Levinsohn, J., and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4): 841-90
- Berry, S., and E. Tamer (2006): "Identification in Models of Oligopoly Entry," in *Advances in Economics and Econometrics: Theory and Applications*, Ninth World Congress, vol. 2, R. Blundell, W.K. Newey and T. Persson, eds., Cambridge Univ. Press
- Caplin, A., and B. Nalebuff (1991): "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica*, 59(1): 25-59
- Chernozhukov, V., Hong, H., and E. Tamer (2007): "Estimation and Confidence Regions for Parameter Sets in Econometric Models," *Econometrica*, 75(5): 1243 - 1284
- Ciliberto, F., and E. Tamer (2007): "Market Structure and Multiple Equilibria in Airline Markets," *mimeo*, Northwestern University
- Gawer A., and M.A. Cusumano (2002): "Platform Leadership: How Intel, Microsoft, and Cisco Drive Industry Innovation," *Harvard Business School Press*
- Goeree, M. S. (2008): "Limited Information and Advertising in the US Personal Computer Industry," forthcoming *Econometrica*
- Goettler, R., and B. Gordon (2008): "Durable Goods Oligopoly with Innovation: Theory and Empirics"
- Gordon, B. (2008): "A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry," Forthcoming *Marketing Science*

- Haile, P.A., and E. Tamer (2003): "Inference with an Incomplete Model of English Auctions," *Journal of Political Economy*, University of Chicago Press, 111(1): 1-51
- Heckman, J. (1976): "The common structure of statistical models of truncation, sample selection, and limited dependent variables and a simple estimator for such models," *Annals of Economic and Social Measurement* 5, 475-592
- Hendel, I. (1999): "Estimating Multiple-Discrete Choice Models: An Application to Computerization Returns," *Review of Economic Studies*, 66(2): 423-446
- Ishii, J. (2006): "Interconnection Pricing, Compatibility, and Investment in Network Industries: ATM Networks in the Banking Industry," *mimeo* Stanford University
- Mazzeo, M.J. (2002): "Product Choice and Oligopoly Market Structure," *RAND Journal of Economics*, 33(2): 221-242
- Nevo, A. (2000): "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand," *Journal of Economics and Management Strategy*, 9: 513-548
- Nevo, A. (2001): "Measuring Market Power in the Ready-To-Eat Cereal Industry," *Econometrica*, 69(2): 307-342
- Pakes, A. (2003): "A Reconsideration of Hedonic Price Indexes with an Application to PC's," *American Economic Review*, 93(5): 1578-1596
- Pakes, A., J. Porter, K. Ho, and J. Ishii (2006): "Moment Inequalities and their Application," *mimeo*, Harvard University
- Petrin, A. (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110: 705-729
- Rosenberg, N. (1979): "Technological Interdependence in the American Economy," *Technology and Culture*, 20: 25-50
- Song, M. (2006): "A Dynamic Analysis of Cooperative Research in the Semiconductor Industry"

- Song, M. (2007): “Measuring Consumer Welfare in the CPU Market: An Application of the Pure Characteristics Demand Model,” *RAND Journal of Economics*, 38: 429-446
- Spence, M. (1976): “Product Selection, Fixed Costs, and Monopolistic Competition,” *The Review of Economic Studies*, 43(2): 217-235
- Train, K. (2003): “Discrete Choice Methods with Simulation”, *Cambridge University Press*
- Trajtenberg, M. (1989): “The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners,” *Journal of Political Economy* 97: 444-479
- Walters, G.E. (2001): “The Essential Guide to Computing,” *Prentice Hall PTR*
- Wooldridge, J.M. (2002): “Econometric Analysis of Cross Section and Panel Data,” *The MIT Press*, Cambridge, Massachusetts

A Details Concerning the Estimation Procedure

A.1 Estimation of θ : Computational Details

Estimating θ following the BLP method requires one to compute the errors $e_j(\theta) = (\xi_j(\theta), \omega_j(\theta))'$ for any generic value of the parameter θ . The integral in (3) is approximated via simulation; I simulate the v_i household-specific taste shifters for $ns = 3000$ households. To reduce the error induced by simulation, I use antithetic draws.³⁴ I then obtain the market share predicted by the model for product j (quarter indices suppressed) as follows:

$$s_j(x, p, \delta, P_{ns}; \theta_2) = \frac{1}{ns} \sum_{i=1}^{ns} \frac{\exp(\delta_j + \mu_{ij})}{1 + \sum_{m \in J} \exp(\delta_m + \mu_{im})} \quad (14)$$

where P_{ns} is the distribution of the simulation draws. The market share equation, which should hold exactly at θ_0 , is given in vector form:

$$s(x, p, \delta, P_{ns}; \theta_2) = S \quad (15)$$

where S denotes *observed* market shares. Given a fixed value for θ_2 , we invert this equation to retrieve a vector of mean utility levels, $\delta(\theta_2)$, using the BLP contraction mapping:

$$\delta^{h+1} = \delta^h + \ln(S) - \ln[s(x, p, \delta^h, P_{ns}; \theta_2)] \quad (16)$$

The vector of demand-side unobservables ξ can now be computed by:

$$\xi(\theta^d) = \delta(\theta_2) - x\beta \quad (17)$$

Marginal cost unobservables are computed from (7):

$$\omega(\theta) = \log[p - (T * \Delta(\theta_2))^{-1}s] - x\gamma \quad (18)$$

Next, I define the GMM objective function. Let z_j denote a $1 \times L$ vector of functions of X . Define:

$$Z_j = \begin{bmatrix} z_j & 0 \\ 0 & z_j \end{bmatrix}_{2 \times 2L}, \quad g_j(\theta) = Z_j' e_j(\theta)$$

Letting N denote the total number of products in the sample, the objective function is given by:

³⁴See Train [2003]. Antithetic draws are used in Goeree's [2008] analysis.

$$Q_N(\theta) = \left[\sum_{j=1}^N g_j(\theta) \right]' \Phi^{-1} \left[\sum_{j=1}^N g_j(\theta) \right] \quad (19)$$

where Φ^{-1} is a $2L \times 2L$ PD weight matrix. The initial choice for this matrix is $[\sum_{j=1}^N Z_j' Z_j]^{-1}$. With an initial estimate for θ at hand, denoted $\hat{\theta}^1$, I estimate the *optimal weight matrix* by $[\sum_{j=1}^N g_j(\hat{\theta}^1) g_j(\hat{\theta}^1)']^{-1}$. Re-estimating θ using the updated matrix yields the estimates reported in Tables 5a-5b.

A.2 Set Estimation of Fixed Cost Parameters λ : Computational Details

Estimation of the quantities $U_{d,k}(\theta_0)$ and $L_{d,j}(\theta_0)$ which appear on the RHS of (10) and (11), respectively, is performed as follows. The BLP estimate $\hat{\theta}$ implies empirical values $e_j(\hat{\theta})$ for all products observed in the sample (a total of 2,287) using equations (17) and (18) above. From this empirical distribution, I draw 500 vectors of error terms e for all products that need to be considered (in the current case, observed and potential products in 2004Q2). Note that we draw from the *joint distribution* of (ξ, ω) .

At each such error vector, I simulate price equilibria under (A_d) and $(A_d - 1_d^k)$, and compute the variable profit figures which appear in (10). Averaging over the variable profit differences yields the estimate of $U_{d,k}(\theta_0)$. An analogous procedure yields estimates for $L_{d,j}(\theta_0)$. Price equilibria are simulated by iterating on the first-order conditions (7) until convergence, which typically takes a few seconds of computation time.

B Relaxing Point Identifying Assumptions for Variable Profit Parameters θ

Relaxing the assumption that firms only observe the error terms e after making their product choices implies that the selection indicator would now depend on unobservables, i.e., it would be written as $q_j(X, F, e)$. Correcting the resulting selection bias with traditional, point-identifying approaches (such as those in Heckman [1976]) is not possible, since that would require specifying a simple model for $q_j(\cdot)$, whereas this object actually depends in a very complex fashion on the observed and unobserved features of *all products*, and is not even uniquely determined in equilibrium.

I describe here an alternative approach which seeks *partial identification* of θ . I have not pursued actual estimation of θ following this strategy, and it is likely to involve a substantial computational burden.

Under the assumptions specified below, I show that given a generic value for θ , necessary

equilibrium conditions can be used to obtain upper (lower) bounds on the ξ (ω) error terms associated with products that are *unobserved by the econometrician*. In addition, equations (17) and (18) above show how to compute exact values for these errors for observed products. The resulting partial information on the distribution of these error terms is translated into a partial identification argument for θ via moment inequalities. These inequalities would reduce to BLP's moment equalities if all products were observed.

I now provide the details of the identification argument. A formal proof of identification is outside the scope of this paper. I maintain Assumption 1, which stated the mean-independence of e from X in the population of all *potential products* \mathbf{J} (see Section 4.1). I also add the following assumption:

Assumption 2. *The support of the marginal distribution of ω has an upper bound, ω^U . The support of the marginal distribution of ξ has a lower bound, ξ^L .*

I now show that Assumptions 1 and 2 yield moment inequalities involving θ if the support bounds ω^U , ξ^L are known. Since they are not likely to be known to the econometrician, I return below to the issue of identifying these support bounds. This latter task would require additional assumptions.

For observed products $j \in J$, equations (17)-(18) imply an exact value $e_j(\theta) = (\xi_j(\theta), \omega_j(\theta))'$. For *unobserved* products $j \in \mathbf{J} \setminus J$, I compute, for each guess of θ , an upper bound $\bar{\xi}_j(\theta)$ and a lower bound $\underline{\omega}_j(\theta)$.

Without loss of generality, assume that the unobserved product $j \in \mathbf{J} \setminus J$ belongs to firm d , and, with a slight abuse of notation, let $j \in A_d^0$. Consider any $k \in A_d^1$, i.e., any observed product offered by this firm (I assume that this set is not empty, i.e., the firm offers at least one product). Using necessary equilibrium conditions, the bounds $\bar{\xi}_j(\theta)$ and $\underline{\omega}_j(\theta)$ are defined as the implicit solutions to the following equations:

$$VP_d(A_d + \mathbf{1}_d^j - \mathbf{1}_d^k; \theta, \omega_j = \omega^U, \xi_j = \bar{\xi}_j) - VP_d(A_d; \theta) = 0 \quad (20)$$

$$VP_d(A_d + \mathbf{1}_d^j - \mathbf{1}_d^k; \theta, \xi_j = \xi^L, \omega_j = \underline{\omega}_j) - VP_d(A_d; \theta) = 0 \quad (21)$$

In words, we consider a deviation in which the firm eliminates the observed product located at k and introduces product j instead. This deviation does not alter total fixed costs, and so its profitability hinges entirely on its variable profit implications. The idea underlying (20) is that, all else equal, the higher is the demand shifter ξ_j , the more profitable is the deviation. An upper bound on ξ_j , therefore, is the value at which the firm is indifferent about performing this deviation. The profitability of the deviation also depends on the unknown ω_j value, but fixing

it at the upper bound of its support yields a valid upper bound on ξ_j . An analogous argument, specified in (21), places a lower bound on the ω_j value.

I now set up moment inequalities as follows: define $|\mathbf{J}|$ -vectors $\tilde{\xi}$ and $\hat{\xi}$, which j^{th} elements are given by:

$$\tilde{\xi}_j(\theta) = \begin{cases} \xi_j(\theta) & j \in J \\ \bar{\xi}_j(\theta) & j \in \mathbf{J} \setminus J \end{cases} \quad \hat{\xi}_j(\theta) = \begin{cases} \xi_j(\theta) & j \in J \\ \xi^L & j \in \mathbf{J} \setminus J \end{cases}$$

At the true parameter value θ_0 , we have $\hat{\xi}_j(\theta_0) \leq \xi_j \leq \tilde{\xi}_j(\theta_0)$ for each $j \in \mathbf{J}$. For each product j , let $z_j : R^{|\mathbf{J}| \times K} \rightarrow R_+^L$ be a nonnegative vector-valued function of instruments. We get:

$$E[\hat{\xi}_j(\theta_0)z_j(X)|j \in \mathbf{J}] \leq E[\xi_j z_j(X)|j \in \mathbf{J}] \leq E[\tilde{\xi}_j(\theta_0)z_j(X)|j \in \mathbf{J}] \quad (22)$$

By Assumption 1 and the Law of iterated expectations, $E[\xi_j z_j(X)|j \in \mathbf{J}] = 0$, implying a set of moment inequality conditions involving θ :

$$E[\hat{\xi}_j(\theta_0)z_j(X)|j \in \mathbf{J}] \leq 0 \leq E[\tilde{\xi}_j(\theta_0)z_j(X)|j \in \mathbf{J}] \quad (23)$$

Note that, if all products are observed, i.e., $J = \mathbf{J}$, the selection problem disappears, and (23) is reduced to the BLP moment equalities. Intuitively, the more severe is the selection problem, the further away from point identification we get. Sets of “supply-side” moment inequality conditions can be analogously obtained. Define the $|\mathbf{J}|$ -vectors $\tilde{\omega}$, $\hat{\omega}$ by their j^{th} elements:

$$\tilde{\omega}_j(\theta) = \begin{cases} \omega_j(\theta) & j \in J \\ \underline{\omega}_j(\theta) & j \in \mathbf{J} \setminus J \end{cases} \quad \hat{\omega}_j(\theta) = \begin{cases} \omega_j(\theta) & j \in J \\ \omega^U & j \in \mathbf{J} \setminus J \end{cases}$$

Which yields:

$$E[\tilde{\omega}_j(\theta_0)z_j(X)|j \in \mathbf{J}] \leq 0 \leq E[\hat{\omega}_j(\theta_0)z_j(X)|j \in \mathbf{J}] \quad (24)$$

Identifying the support bounds. Throughout the discussion above I made reference to the support bounds ξ^L and ω^U without specifying how one may learn their values. One possible avenue to do that is to define a subset of products which introduction is assumed to be predetermined to the game. Observe in Table 3a that, at each point in time, at least one CPU option was offered in the vast majority of product lines (e.g. **P3**_0.5-0.99 in 2001Q3). Such a CPU option may be viewed as an “industry standard” option, offered by most firms (and virtually all major firms). This motivates the following assumption:

Assumption 3. A subset $J^P \subset \mathbf{J}$ of the potential products is offered in a manner that is pre-determined to the two-stage game. In addition, $E[e_j|X, j \in J^P] = 0$ holds, and the distributions of ξ_j and ω_j conditional on $j \in J^P$ have the same support bounds as those of the unconditional distributions.

The idea behind this assumption is that, since the products in J^P are not selected on account of unobservables, the mean-independence condition should not be violated within this sub-population. Assumption 3 yields *moment equalities*: since $J^P \subset J$, i.e., all the pre-determined products are observed, a generic value for θ implies exact values $e_j(\theta)$ for each $j \in J^P$ using (17) and (18) above, and we get:

$$E[e_j(\theta_0)z_j(X)|j \in J^P] = 0 \tag{25}$$

These moment conditions point-identify a subset of the elements of θ , denoted θ^P , while providing no information on other elements. In addition, exact values $e_j(\theta^P)$ are implied for each $j \in J^P$ by using (17)-(18) once again. The support bounds are then identified by $\xi^L = \inf_{j \in J^P} \xi_j(\theta^P)$, $\omega^U = \sup_{j \in J^P} \omega_j(\theta^P)$.

Discussion. An unattractive feature of the identification strategy above is that it relied on the assumption of a pre-determined set of products. This may be more reasonable in some applications than in others. Coming up with an alternative approach for pinning down the support bounds, therefore, would be desirable.

Finally, actual estimation of θ following this approach would require constructing the sample analogs of the moments in (23)-(25). While efficient estimation would use all these conditions simultaneously, the computational burden would be significantly alleviated if one uses (25) first to obtain a point estimate of $\theta^P \subset \theta$, and then obtains set estimates of the remaining parameters using (23)-(24), holding the point estimates fixed. This is still likely to be computationally expensive, so that a parsimonious specification for utility and marginal cost is likely to be necessary.

C Tables

Table 1: Top Vendors' Market Shares, US Home PC Market

Year 1		Year 2		Year 3	
Vendor	Market Shares	Vendor	Market Shares	Vendor	Market Shares
Dell	0.190	Dell	0.263	Dell	0.279
HP*	0.185	HP	0.234	HP	0.258
Compaq*	0.092	eMachines	0.076	eMachines*	0.070
Gateway	0.091	Gateway	0.070	Gateway*	0.053
eMachines	0.060	Toshiba	0.042	Toshiba	0.043
Top 5 vendors	0.618	Top 5 vendors	0.685	Top 5 vendors	0.704

Years: 01Q3-02Q2, 02Q3-03Q2, 03Q3-04Q2. *Compaq and HP merge in Year 1, eMachines and Gateway merge in Year 3.

Table 2: CPU Vendor Shares

Vendor	Market Shares		
	Year 1	Year 2	Year 3
Intel	0.71843	0.72246	0.74496
AMD	0.24429	0.23643	0.22032
IBM	0.03230	0.03450	0.03048
Others	0.00477	0.00524	0.00323
Transmeta	0.00022	0.00135	0.00097
Via	0.00000	0.00002	0.00005

Years: 01Q3-02Q2, 02Q3-03Q2, 03Q3-04Q2, U.S. Home market

Table 3a: Fraction of Desktop Product Lines to Install Intel's CPUs

Quarter	C_0.5-0.99	C_1-1.49	C_1.5-1.99	C_2-2.99	P3_0.5-0.99	P3_1-1.49
2001Q3	0.89	0.00	0.00	0.00	0.93	0.67
2001Q4	0.46	0.42	0.00	0.00	0.46	0.50
2002Q1	0.35	0.58	0.00	0.00	0.31	0.31
2002Q2	0.13	0.57	0.00	0.00	0.17	0.13
2002Q3	0.09	0.39	0.48	0.13	0.13	0.13
2002Q4	0.07	0.04	0.44	0.41	0.11	0.04
2003Q1	0.04	0.04	0.41	0.41	0.04	0.00
2003Q2	0.04	0.04	0.37	0.41	0.04	0.00
2003Q3	0.04	0.04	0.24	0.48	0.04	0.00
2003Q4	0.04	0.04	0.20	0.52	0.04	0.00
2004Q1	0.00	0.04	0.15	0.54	0.00	0.00
2004Q2	0.00	0.00	0.17	0.54	0.00	0.00

Quarter	P4_1-1.49	P4_1.5-1.99	P4_2-2.99	P4_3-3.99
2001Q3	0.48	0.26	0.00	0.00
2001Q4	0.65	0.65	0.12	0.00
2002Q1	0.58	0.73	0.50	0.00
2002Q2	0.43	0.70	0.65	0.00
2002Q3	0.26	0.74	0.70	0.00
2002Q4	0.11	0.37	0.81	0.00
2003Q1	0.11	0.44	0.81	0.15
2003Q2	0.11	0.44	0.81	0.15
2003Q3	0.08	0.28	0.92	0.60
2003Q4	0.12	0.28	0.92	0.60
2004Q1	0.08	0.19	0.92	0.65
2004Q2	0.08	0.17	0.92	0.63

Calculations pertain to the Home market and exclude vendors identified as "Others", Apple products, and PC products which sold under 100 units in a quarter. C,P3, and P4 stand for Intel's Celeron, Pentium III and the Pentium 4 brands, respectively. **P3**_0.5-0.99 implies Pentium III with clock speed range between 0.5 and 0.99 GHz

Table 3b: Portable Product Lines Installations

Quarter	C_0.5-0.99	C_1-1.49	C_1.5-1.99	C_2-2.99	P3_0.5-0.99	P3_1-1.49
2001Q3	0.81	0.00	0.00	0.00	1.00	0.15
2001Q4	0.59	0.21	0.00	0.00	0.79	0.72
2002Q1	0.36	0.25	0.00	0.00	0.64	0.86
2002Q2	0.12	0.31	0.00	0.00	0.54	0.62
2002Q3	0.11	0.21	0.07	0.00	0.18	0.64
2002Q4	0.10	0.03	0.23	0.10	0.16	0.42
2003Q1	0.03	0.06	0.26	0.13	0.19	0.39
2003Q2	0.03	0.03	0.21	0.12	0.15	0.42
2003Q3	0.03	0.00	0.25	0.13	0.16	0.34
2003Q4	0.03	0.03	0.22	0.13	0.13	0.28
2004Q1	0.00	0.03	0.18	0.18	0.03	0.18
2004Q2	0.00	0.03	0.19	0.19	0.03	0.19

Quarter	P4_1-1.49	P4_1.5-1.99	P4_2-2.99	P4_3-3.99	Pm_1-1.49	Pm_1.5-1.99
2001Q3	0.00	0.00	0.00	0.00	0.00	0.00
2001Q4	0.00	0.00	0.00	0.00	0.00	0.00
2002Q1	0.07	0.18	0.00	0.00	0.00	0.00
2002Q2	0.19	0.38	0.00	0.00	0.00	0.00
2002Q3	0.14	0.46	0.32	0.00	0.00	0.00
2002Q4	0.10	0.52	0.58	0.00	0.00	0.00
2003Q1	0.13	0.52	0.58	0.00	0.10	0.06
2003Q2	0.12	0.48	0.55	0.00	0.09	0.09
2003Q3	0.09	0.53	0.59	0.06	0.22	0.19
2003Q4	0.13	0.50	0.50	0.09	0.31	0.25
2004Q1	0.12	0.42	0.52	0.12	0.21	0.48
2004Q2	0.06	0.44	0.50	0.13	0.28	0.56

See notes for Table 3a. Pm stands for Intel's Pentium M brand. CPU technologies with very small installation rates excluded.

Table 4: Descriptive Results, logit Demand

β	Logit_OLS	Logit_IV
Price (00\$)	-0.0395*** (0.0135)	-0.157** (0.0649)
Laptop dummy	-0.616*** (0.0999)	-0.298 (0.199)
Trend	-0.0398** (0.0171)	-0.138** (0.0567)
CPU Speed Range Dummies		
1-1.49 GHz	0.200* (0.107)	0.385** (0.152)
1.5-1.99 GHz	0.383*** (0.138)	0.660*** (0.208)
2-2.99 GHz	0.752*** (0.156)	1.223*** (0.303)
3-3.99 GHz	0.779*** (0.253)	1.586*** (0.508)
CPU Brand Dummies		
AMD Duron	0.694*** (0.208)	0.544** (0.254)
AMD Athlon	0.691*** (0.115)	0.695*** (0.133)
Intel Pentium III	0.227** (0.116)	0.507*** (0.189)
Intel Pentium 4	0.359*** (0.103)	0.629*** (0.176)
Intel Pentium M	0.724*** (0.215)	1.554*** (0.489)
Constant	-10.66*** (0.183)	-9.441*** (0.699)
Observations	2287	2287
R-squared	0.491	0.473

Standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Dummy variables for PC vendors and brands included, not reported

Table 5a: BLP Estimates for θ , Main PC Characteristics

	β	SE	σ	SE	γ	SE
Constant	4.479	3.108	1.546	1.933	6.759	0.020
Laptop Dummy	-0.690	1.158	3.785	0.518	0.312	0.013
Trend	-1.444	0.263	0.430	0.081	-0.090	0.002
CPU Speed Range Dummies						
1-1.49 GHz	2.390	0.386			0.156	0.013
1.5-1.99 GHz	3.621	0.521			0.232	0.016
2-2.99 GHz	6.212	0.809			0.412	0.017
3-3.99 GHz	9.584	1.374			0.709	0.030
CPU Brand Dummies						
AMD Duron	-0.915	0.443			-0.120	0.023
AMD Athlon	0.912	0.217			0.031	0.013
Intel Pentium III	3.517	0.484			0.272	0.014
Intel Pentium 4	3.855	0.487			0.305	0.010
Intel Pentium M	10.051	1.361			0.741	0.032
Price sensitivity	α	SE	σ^p	SE		
	0.810	0.179	0.301	0.060		

Obs: 2287. Dummies for PC vendor, brands included, reported in 5b. Standard errors do not take into account simulation error, which is mitigated via antithetic draws.

Table 5b: BLP Estimates for θ , PC Vendor & Brand Dummies

	β	SE	γ	SE		β	SE	γ	SE
Dell	12.332	2.603	0.774	0.062	Toshiba	7.933	1.684	0.479	0.050
dimension	-10.426	2.813	-0.915	0.065	portege	0.593	1.018	0.093	0.059
inspiron	-9.908	2.732	-0.838	0.064	port_tablet	2.855	1.303	0.218	0.085
Latitude	-7.529	2.137	-0.488	0.071	satellite	-5.405	1.752	-0.517	0.055
OptiPlex	-13.509	2.819	-0.903	0.064	satpro	-2.628	1.141	-0.132	0.053
HP	-0.976	0.334	-0.049	0.021	Sony	5.684	0.821	0.306	0.037
evoipaq	-1.651	0.519	-0.174	0.030	VAIO_DS	-3.909	0.871	-0.265	0.043
media	7.568	0.870	0.424	0.029	VAIO_R	0.500	0.824	0.205	0.052
pavilion	2.625	0.385	-0.015	0.024	VAIO_W	3.163	0.979	0.283	0.059
presario	2.593	0.355	0.026	0.020	VAIO_505	1.288	0.963	-0.007	0.061
cmpq_ntbk	1.841	0.653	0.175	0.033	VAIO_FX	0.951	0.734	0.053	0.053
cmpq_ultprtbl	10.945	2.221	0.741	0.080	IBM	2.037	1.217	0.208	0.083
Gateway	0.309	0.399	0.068	0.025	IBM_netvista	-3.868	1.307	-0.244	0.087
Gateway3	-2.619	0.730	-0.408	0.035	IBM_thinkCentre	0.419	1.301	0.040	0.095
Gateway5	1.755	0.865	-0.030	0.048	IBM_thinkpadA	8.348	2.084	0.452	0.097
Gateway7	2.159	0.690	0.077	0.035	IBM_thinkpadT	1.253	1.366	-0.016	0.092
essential	2.124	0.458	-0.098	0.030	IBM_thinkpadR	-3.304	1.291	-0.233	0.085
performance	1.751	0.530	0.039	0.034	Acer_veriton	-2.120	0.382	-0.120	0.016
media	4.960	0.828	0.365	0.035	Averatec	1.131	0.688	-0.034	0.048
Gateway4	-1.320	0.510	-0.139	0.031	Fujitsu	-1.090	0.354	-0.018	0.023
Gateway6	4.725	1.077	0.173	0.051	MicroElectronics	-1.585	0.236	-0.009	0.017
solo	0.185	0.868	-0.106	0.050					
eMachines	0.389	0.602	-0.325	0.050					

See notes for Table 5a. The coefficients on vendors (e.g. Dell) **do not** capture an “overall” vendor effect, but rather the effect of omitted brands of that vendor (see text).

Table 6: Economic Implications of BLP Estimates

A. Willingness to pay	
	Average Consumer WTP (\$)
1-1.49 GHz → 1.5-1.99 GHz	54.8
1.5-1.99 GHz → 2-2.99 GHz	115.3
2-2.99 GHz → 3-3.99 GHz	150.1
Celeron → Pentium III	156.5
Celeron → Pentium 4	171.5
Celeron → Pentium M	447.3
HP (Compaq) Presario	71.9
Dell Inspiron	107.8
Sony VAIO R	275.2
IBM Thinkpad A	462.1
1 year forward*	-257.0

B. Additional Information	
Median Markup (\$)	76.4
Median (p-mc)/p	0.078
Corr(markup, price)	0.912
Corr(ξ, ω)	0.820

*Change in willingness to pay over one year, see text

Table 7: Fixed Costs Parameters λ

Parameter	(a)		(b)	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound
Constant	1,165,999	1,506,040	980,000	1,500,000
Big Notebook Producer			0	1,290,000
Per-configuration fixed cost ($V_d\lambda$) \$				
Big Producer			1,170,000	2,270,000
Small Producer			980,000	1,500,000

Confidence intervals not yet provided.

Table 8: The Effect of Intel's Pentium M on 2004Q2 Outcomes

	Full Model		Restricted model (fixed products)
	Lower bound	Upper bound	
Impact on number of PC configurations			
# P3.1.5-1.99	-4	-2	-
# C.1.5-1.99	+1	+3	-
# C.2-2.99	+1	+1	-
# P4.1.5-1.99	+2	+2	-
# P4.2-2.99	-1	+2	-
# P4.3-3.99	-3	-1	-
Impact on Pentium III's share of total Portables sales			
Share P3	-0.156	-0.116	-0.012
Impact on PC sales and prices			
Total Notebook Sales	+10.9%	+18.9%	+17.6%
Total Desktop Sales	-2.6%	-1.6%	-2.5%
Mean Notebook price* (\$)	+32	+44	+55
Impact on consumer surplus			
0-20% Price sensitive (\$M)	+25.90	+38.19	+39.31
20%-40% Sensitive (\$M)	+1.24	+2.52	+1.93
40%-60% Sensitive (\$M)	+0.53	+1.04	+0.75
60%-80% Sensitive (\$M)	+0.06	+0.09	+0.08
80%-100% Sensitive (\$M)	+0.01	+0.03	+0.03
Total Consumer Surplus (\$M)	+27.92	+41.87	+42.10

The full model allows for endogenous product choices and endogenous prices, whereas the restricted model simply removes Pentium-M based PCs from the sample and computes counterfactual prices. Entries with a positive (negative) sign imply that the presence of the Pentium M has increased (decreased) the relevant quantity by the reported amount.

For instance, the restricted model finds that the Pentium M increased total consumer surplus by \$42.10 million, (corresponding to a 4.9% increase).

* Sales-weighted average




Figure 1: CPU Speed Range Shares, U.S. Home Market, over the three sample years




Figure 2: Sales-Weighted Average Prices (\$1000's) and Share Portables, U.S. Home PC Market




Figure 3: Willingness-to-pay for an upgrade from Intel's Celeron to its Pentium M processor (\$)




Figure 4: Willingness-to-pay for "1 year forward" (\$) (see text)