

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

Alon Eizenberg[†]

April 2011

First Draft: November 2008

Abstract

This paper investigates the welfare implications of the rapid innovation in Central Processing Units, and, specifically, asks whether it results in inefficient elimination of basic Personal Computer configurations. I analyze a game in which firms make multiple discrete product choices, and tackle challenges such as partial identification and sample selection. The estimated model implies that the demand for PCs is highly segmented. Using this model, I find that Intel's introduction of its Pentium M chip contributed significantly to the growth of the mobile PC segment and to consumer welfare. The lion's share of these consumer benefits were enjoyed by the 20% least price-sensitive consumers. The Pentium M crowded out the Pentium III and Pentium 4 technologies. I find that keeping those products on the shelf could have increased the welfare of price-sensitive consumers, but sizable fixed costs offset these benefits.

Keywords: product variety, personal computers, CPU, innovation, endogenous product choices, discrete product choices, multiple equilibria, partial identification, sample selection

JEL Classification: L10, L11, L13, L15

*This paper is based on the main chapter of my doctoral dissertation. I am indebted to my advisors, Steven Berry and Philip Haile, for their continued advice and encouragement. I am also especially grateful to Donald Andrews for his advice. I have benefited greatly from discussions with Eduardo Faingold, Joshua Lustig, fellow graduate students at Yale, and seminar participants at the Analysis Group, Bates White, Bar-Ilan University, Cornerstone Research, CRA, Federal Reserve Board, Harvard University, Hebrew University, NERA Economic Consulting, Northwestern (Kellogg), Georgetown University (McDonough), Penn State University, Tel Aviv University, The Southern Economic Association meetings, Stanford University, UCLA, UT Austin, Washington University (Olin), University of Pennsylvania (Wharton), and Yale University (SOM). All errors are mine. I am grateful to IDC and to Mr. Steven Clough for making data available. Financial support from the Carl Arvid Anderson Prize Fellowship of the Cowles Foundation (while a graduate student at Yale) and from the Falk Institute is gratefully acknowledged.

[†]Department of Economics, Hebrew University of Jerusalem. e-mail address: aloneiz@mscc.huji.ac.il

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 work 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 or word processing), which require modest computing power.

When basic PC configurations disappear, consumers who would have optimally chosen to purchase them end up buying stronger machines (or choose not to purchase a PC at all). Keeping these basic configurations on the shelf, alongside the more advanced ones, could, therefore, contribute to the welfare of such consumers.² From the point of view of a social planner, however, these benefits to consumers must be weighed against the impact on producer profit, and, specifically, the additional fixed costs that would have been required so as to keep such basic configurations on the shelf (e.g. the cost of technical support, marketing, and inventory management). If these costs are not too large, a social planner may prefer to keep offering these basic configurations, meaning that their elimination by the market was socially inefficient.

Theoretical analyses (Spence (1976), Mankiw and Whinston (1986)) have demonstrated that equilibrium product choices may be characterized by social inefficiency. Two types of externalities drive these failures: on the one hand, a firm that contributes to variety by introducing a differentiated product to the market generally fails to appropriate the full surplus associated with this product introduction, and the implied *positive* externality may lead to insufficient variety. On the other hand, launching an additional product imposes a *negative* externality on rivals, suggesting a potential for excessive product introduction. The lesson from this theoretical literature is that it is necessary to learn the values of market-specific parameters (namely, those governing cost and demand in a given industry) to be able to determine the existence, nature and magnitude of such market failures.

Motivated by this question, I estimate a model of supply and demand in which both the set of PC configurations offered to consumers, and the prices charged for such configurations, are endogenously determined. I then perform counterfactual analyses to determine the impact of innovation on the portfolio of technologies offered to consumers, to determine whether products are inefficiently eliminated, 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.

¹Pakes (2003) reports an average annual attrition rate of 85 percent.

²Adding or removing products from the market affects consumers also via the effect on equilibrium prices. As explained below, my empirical framework treats both product offerings and prices as endogenous, which allows me to capture these effects.

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, 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 outcomes 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, in the taste for portability, and in the degree to which consumer 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.³ 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 innovations to consumer welfare.

I use the estimated model in counterfactual analysis to study the impact of Intel's introduction of its Pentium M chip, which is considered a landmark in mobile computing. I artificially remove this technology from the market, and compute the set of potential equilibria under this "no Pentium M" scenario.⁴ Comparing these outcomes to outcomes in the observed equilibrium (i.e.,

³As 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., in Web applications) is likely to have been the driving force behind this process.

⁴As explained below, the term "potential equilibria" pertains to outcomes which cannot be ruled out as equilibria of the game. The need to work with this concept stems from the partial identification of fixed costs, also to be discussed below.

in the presence of the Pentium M) provides a measure of the Pentium M's effect.

I find that, in the second quarter of 2004, the presence of the Pentium M made a substantial contribution to the growth of the mobile segment of the PC market, and that some of this growth came at the expense of Desktop sales. The presence of the Pentium M increased the total consumer surplus by 3.2%-6.3%. This innovation also led to a significant re-alignment of PC makers' product offerings, and, in particular, crowded out products based on older technologies such as certain Intel Pentium III and Pentium 4 chips.

I also document substantial heterogeneity in the impact of innovation on different consumer types. The 20% least price-sensitive consumers enjoy the bulk of the benefits from innovation, while the impact on other segments of consumer demand appears minimal to nonexistent (at least in the short run). While price-sensitive consumers can be hurt by the elimination of basic technologies, they also benefit from the decrease in prices of surviving technologies prompted by the arrival of cutting-edge innovations. Since both product choices and prices are treated as endogenous, the effects of both forces are taken into account in the welfare analysis.

I use the model to ask whether a social planner could improve welfare by adding back to the market notebook PCs based on the eliminated Pentium III and Pentium 4 technologies, in the sense that the added fixed costs would be outweighed by the contributions to consumer surplus and to variable profit.⁵ If this is the case, we would determine that the elimination of such technologies was socially inefficient. I find that bringing these technologies back to the shelf would have substantially increased the welfare of price-sensitive consumers. On the other hand, the presence of sizable fixed costs offsets these welfare gains to a large extent, so that the scope for inefficient elimination appears to be very limited: the upper bound on the lost welfare is small compared to the benefits from innovation, while the lower bound suggests no welfare loss.

Caveats: long-term benefits and upstream profits. While my estimates capture the process by which consumer utility from a fixed bundle of hardware characteristics falls 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.⁶ My analysis, therefore, does not account for some long-term contributions of CPU innovation to welfare. This motivates future research of dynamic complementarities in innovative activities.⁷

Another aspect not formally addressed by my framework is the quantitative impact of the analyzed downstream game on the profits and decisions of upstream firms such as Intel or Mi-

⁵To be clear, this counterfactual does not "reverse" the effect of innovation, as it keeps the Pentium M technology in the market. It merely adds some eliminated PC products back to the market, alongside the more advanced technologies, to determine if this can be welfare-enhancing.

⁶Gawer 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.

⁷See Rosenberg (1979) for a seminal discussion of this issue.

crosoft. I take the process of CPU innovation (which determines the menu of feasible CPUs) to be exogenous.⁸ This assumption is reasonable since, as discussed below, the pace of innovation in computing power is largely driven by Moore’s Law.⁹ Furthermore, my welfare calculations define “producer profits” as the profits of PC makers, and so I do not formally take into account the impact of, say, adding basic PC configurations to the shelf on the profit of Intel or Microsoft.

I argue that taking these externalities into account would only reinforce my findings: Intel and Microsoft are likely to benefit from the elimination of basic products, as it shifts consumer demand toward higher-margin chips and software products. Since I already find that adding these basic products back to the market does not improve total welfare, taking into account the likely-adverse impact on Intel and Microsoft would only reinforce this finding.

Methodological contribution. To the best of my knowledge, this is the first paper that solves and estimates a model where firms make multiple-discrete product choices (i.e., they choose their product portfolio from a discrete menu of potential products) while providing a full-fledged modeling of cost and demand systems (see the literature review below). This requires me to tackle several methodological challenges.

First, the discrete nature of these product choices implies that the game is not guaranteed to have a unique equilibrium, and, as a consequence, fixed costs are only partially identified. Recent literature (see below) has exploited necessary equilibrium conditions to place bounds on partially-identified parameters, and I extend this approach in this paper to the multiple-discrete product setup described above. I allow for important heterogeneity in fixed costs across firms, and estimate bounds on the mean per-configuration fixed cost for each firm. This approach yields identified intervals for fixed cost parameters which endpoints are defined by means of random variables. As a consequence, I am able to use rather simple techniques to perform inference on the set-identified parameters.

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 predictions, as allowing for n product choices yields 2^n feasible vectors of product offerings. I develop methods which reduce this burden. The intuition behind these methods 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 third, difficult challenge tackled in this paper is sample selection, which arises since firms are explicitly assumed to have selected the set of products observed in the data. I allow for structural

⁸I explain below how I control for Intel’s decisions regarding when to phase out its CPU technologies, which allows me to focus on the downstream decisions regarding which of the remaining chips to adopt.

⁹Another reasonably-exogenous source of innovation is initiative taken by engineers, which is known to have played a key role in the development of the Pentium M.

errors in both the fixed and the variable components of the profit function. As a consequence, potential selection bias could affect both the point estimators of the variable profit parameters, and the set estimators of the fixed cost parameters.

I develop techniques that address both these issues. The selection problem in the estimation of variable profit parameters is handled by imposing a point-identifying assumption, according to which firms commit to product choices before they observe realizations of marginal cost and demand shocks.¹⁰ I then show that a finite-support assumption for fixed costs gives rise to a partial-identification approach that solves the selection bias in the estimation of fixed costs.

Related literature. Song (2007, 2010), Gordon (2009), and Goettler and Gordon (2009) 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. Nosko (2010) also makes the assumption that the CPU is a perfect proxy for the PC, and builds on the methodology developed in this paper to study the product portfolio choices of upstream CPU makers.¹¹

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 (2009)). Ishii (2008) estimates a model in which banks choose an integer number of ATM locations. The discreteness of this choice leads to 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 type. As a consequence, I must consider a vector of product-choice binary indicators for each firm. Moreover, unlike the literature cited above, my framework does not rely on a reduced-form profit function and instead measures profits from an estimated model of cost and demand, while showing how to overcome difficult selection problems that arise with respect to the structural demand and cost errors.

This paper is also related to a long-standing literature (e.g. Trajtenberg (1989), Petrin (2002)) that studies 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 this 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 questions of interest via counterfactual analysis. Concluding remarks are offered in Section 7.

¹⁰Eizenberg (2009) describes an alternative approach (that was not implement in practical estimation) which relaxes this assumption. There, it is shown that the selection mechanism itself can be used to generate moment inequalities which provide partially-identifying information on variable profit parameters.

¹¹This paper predates Nosko's paper.

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.¹² 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.¹³ 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. I define a product as a unique combination of observed characteristics.¹⁴

As discussed below, the demand model employed in this work 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.¹⁵ 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)), I define the size of the market as the number of U.S. households in a quarter, as reported by the U.S. Census Bureau.¹⁶ Since PC makers typically target the Home and Commercial segments with *different product lines*, it is reasonable to study product choices in the Home market separately.¹⁷

The Home PC 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 widespread hardware upgrades 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. The top 5 vendors together accounted for a 60%-70% share of the market. A similar concentration level is reported by Goeree (2008) for the late 1990s.

¹²<http://www.IDC.com/>.

¹³In some cases, slightly less-disaggregated information is available in that sales are split evenly among observations pertaining to the same vendor-quarter cell. This issue is not likely to cause a problem since the implied average prices, computed as explained below, appear very reasonable.

¹⁴These definitions follow 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.

¹⁵Purchases of the latter were studied by Hendel (1999).

¹⁶I interpolate linearly between the 2000 and 2004 household totals to obtain quarter-by-quarter figures.

¹⁷Some overlap exists between these markets, since some "Home" consumers purchase PC products designed for commercial users, and vice versa. As explained below, the estimation of fixed costs focuses on brands that mainly target the Home segment, which should alleviate potential biases resulting from this issue.

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%).¹⁸

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 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.¹⁹ 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”.²⁰ 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, brand and form factor (e.g., “Dell-Inspiron-Portables”). I define *CPU technologies* as combinations 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.²¹ 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 that adopt Celeron processors with CPU speed in the 0.5-0.99 GHz range. These CPUs were utilized in 89% of the product lines in the first quarter, but were rapidly phased out, in parallel to increased adoption of new CPU technologies.

¹⁸After removing Apple products, observations with negligible market shares (defined as selling less than 100 units in the quarter), observations with a dollar value of zero, and observations with missing speed information, I obtain 2,287 observations, each of which is a quarter-product pair.

¹⁹Note, 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 adequately capture CPU performance, and I do so in the empirical application.

²⁰The 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).

²¹The analysis in this paper is restricted to PC makers’ decisions to install Intel’s CPUs, taking offerings based on AMD products as exogenous (the *prices* of all PC products are always endogenous in this paper, though). An analysis of the variety of AMD chips offered would be an interesting extension, but would require careful attention given the asymmetry between the two chip makers.

Table 3b reports such information for portable PC product lines.

Tables 3a and 3b convey significant variation, in that most CPU technologies are only adopted in a subset of product lines at a given point in time. This variation is instrumental in identifying the cost of offering PC configurations. Some of this variation, however, is artificial; first, certain CPUs could not be installed in certain PC product lines due to technical constraints. 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. In such cases, it is likely that Intel has already phased out the relevant technology, and it would therefore be misleading to include it in the menu of CPU technologies that PC makers can feasibly install in their PC product lines. I describe below how I take such issues into account in defining the feasible set of CPU technologies.

3 Model

The primitives of the model are consumer demand for PCs, PC makers' marginal and fixed costs, 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.²² The household makes the discrete choice that maximizes the following indirect utility function, describing the utility derived by household i from PC product j at time t :

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

The following notation is used: x_{jt} is a K -vector of PC product characteristics observed by the econometrician. 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 fixed PC characteristics changes (falls) over time. ξ_{jt} is a quarter-specific demand shock which is unobserved by the econometrician. The product's

²²Gordon (2009) 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.

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 $\theta_2 \equiv (\alpha, \sigma')'$. Conditioning on δ , the utility function can then be written 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:

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

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 observe realizations of shocks to fixed costs that are unobserved by the econometrician; they then simultaneously choose which CPU technologies to offer with each product line, and 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 discussed below, the assumption that firms learn the realizations of the shocks to demand and marginal cost only after committing to product choices makes it possible to overcome the sample selection issue in the estimation of the variable profit parameters. Since I control for brand-specific intercepts (for most brands), these errors should not capture any systematic brand effects that the firms are likely to know prior to committing to their configuration 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. Denote by $L_{dm} \subseteq H$ the set of CPU technologies that firm d chooses to offer with product line m .²³

Stage 1: In the beginning of this stage, firms observe realizations of shocks to fixed costs. I now describe these fixed costs in detail; each potential PC configuration (i.e., CPU option) that firm d could feasibly offer is associated with fixed costs that would be incurred, should the firm choose to offer that configuration. These fixed costs may include physical production costs, inventory management costs that are necessary to ensure that the product configuration is in stock, and administrative, sales and marketing costs.

Let \mathbf{J}^d represent the set of all firm d 's *potential* product configurations, i.e., both those configurations offered, and those which the firm chooses not to offer. This set has $|S_d| \times |H|$ elements. The following specification is chosen for the fixed cost associated with firm d 's product $j \in \mathbf{J}^d$:

$$F_j = F^d + \nu_j, \quad \text{with } E[\nu_j | j \in \mathbf{J}^d] = 0 \quad (4)$$

This specification implies that the fixed costs associated with firm d 's products are given by adding together a mean, F^d , viewed as a parameter to be estimated, and a mean-zero error term. It allows for both heterogeneity in fixed costs across firms (via the mean), and for stochastic fluctuations about that mean.²⁴

In the empirical application, for reasons discussed below, I estimate only the fixed costs associated with the four leading notebook product lines (which account for over 70% of notebook sales). I therefore allow the mean of these costs to be specific not only to the firm (e.g. Dell) but also to the product line.²⁵

Upon observing the shocks ν_j , firms proceed with simultaneous product configuration choices: each firm $d \in D$ determines the sets L_{dm} for each product line $m \in S_d$. Collecting these sets

²³For 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.

²⁴This specification does not allow for economies (or diseconomies) of scope. Such effects could be captured, however, if one would assume instead that the firm's total fixed cost depends non-linearly only on the total number of offered products. Such an extension would require some alterations to the econometric procedures described below.

²⁵There is one exception: I impose that two different notebook product lines by HP are characterized by the same mean (across configurations) of fixed costs.

across all firms yields the set $J = \{L_{dm}\}_{d \in D, m \in S_d}$ of all PC products that would be offered in the quarter. Firms then pay the fixed costs F_j associated with each configuration they offer.

Stage 2. I let the log of marginal costs for a PC product j depend linearly on observed cost shifters, w_j , and on an additive error term ω_j :²⁶

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

In the beginning of Stage 2, firms observe realizations of $e_j = (\xi_j, \omega_j)'$ for each configuration chosen for production in Stage 1 (to re-iterate, these are demand and marginal cost shocks that 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:

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

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 TF_d is firm d 's total fixed cost.

I assume that, given any Stage 1 history (and any parameter values), Stage 2 prices are uniquely determined in a pure-strategy, interior Nash-Bertrand price equilibrium.²⁷ Arranging products in a $|J|$ -dimensional vector, equilibrium prices satisfy a vector of 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 share derivatives depend on the non-linear demand parameters θ_2 .

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. I assume the *existence* of a pure-strategy SPNE for the two-stage game. I do not, however, assume *uniqueness* of the SPNE.

²⁶In 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.

²⁷This is a standard assumption (e.g. Nevo (2001)). The results of Caplin and Nalebuff (1991) guarantee a unique price equilibrium under stronger restrictions than those imposed here.

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 F^d , one such parameter for each firm.

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. These tasks are explained in turn in sections 4.1 and 4.2 below. Both of these tasks involve overcoming sample selection issues.

4.1 Estimating the Variable Profit Parameters $\theta = (\beta', \alpha, \sigma', \gamma)'$

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 offered configurations 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, and these estimates are reported in section 5 below.

The intuition for overcoming a selection bias in the estimation of θ 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”.²⁸ Stating the point-identifying conditions requires a bit more notation. Let us collect all firms’

²⁸See Wooldridge (2002), ch. 17, for a general discussion of the implications of selection mechanisms which depend on variables observed by the econometrician.

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 the potential products, and let F denote 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, $q_j(X, F)$, which is equal to 1 if j was chosen for production, and is 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:

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

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 .²⁹ Notice that we overcome the selection problem by obtaining a moment condition that is defined over observed products only. GMM estimation of θ using the moment conditions (9) follows the BLP method. Additional details regarding this estimation procedure are provided in Appendix A.1.³⁰

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

²⁹Additional regularity conditions are necessary for a formal identification argument.

³⁰Note that this identification strategy for θ 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. Eizenberg (2009) provides the details of an alternative identification strategy (that was not implemented in practical estimation) that relaxes this assumption, and uses the selection mechanism itself to generate bounds on the error terms of "missing" products. This allows the construction of moment inequalities that are defined over the entire set \mathbf{J} of *potential* products.

speed.³¹ 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.

4.2 Estimating the Fixed Cost Parameters F^d

Given the point estimate $\hat{\theta}$ obtained in section 4.1, a set estimate can be obtained for the fixed cost parameter F^d , firm d 's mean fixed cost. 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.³² In particular, it will be shown that an upper bound can be derived on the fixed cost associated with each offered product, and that a lower bound is available for the costs associated with each product the firm does not offer. These bounds, in turn, are then used to construct bounds on the mean fixed costs parameters F_d .

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 offered. Since firm d may have more than one product line, the general 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 corresponding to products offered and not offered, respectively. The set of all the entries in A_d corresponds to \mathbf{J}^d , defined above as the set of all firm d 's potential products.

Upper and lower bounds on F_j . Consider any product configuration j which belongs in the set A_d^1 , i.e., firm d chose to offer this product. A necessary equilibrium condition implies an upper bound on F_j , the fixed costs associated with such a product, at the true parameter values:

$$F_j \leq E_{(e|\theta_0)} \left[VP_d(A_d; e, \theta_0) - VP_d(A_d - \mathbf{1}_d^j; e, \theta_0) \right] \equiv \bar{F}_j(\theta_0), \quad \forall j \in A_d^1 \quad (10)$$

³¹For the purpose of constructing this instrument I compute speed as the middle of the relevant speed range.

³²See cf. Berry and Tamer for a discussion of the use of necessary equilibrium conditions in partially-identified entry models.

where $\mathbf{1}_d^j$ denotes a vector of the same length as A_d which j^{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 the firm's 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.1).

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 on the fixed costs associated with products the firm chose not to offer: a deviation that adds such a product to the firm's portfolio must not be profitable, implying that the added fixed costs must exceed the expected variable profit gains:

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

Using the bounds on F_j to identify F_d . The above arguments demonstrated that one can obtain an upper bound on the fixed cost associated with each product offered by firm d , and a lower bound on the costs associated with each product the firm chose not to offer.³³ Our goal, however, is to estimate F^d , firm d 's mean fixed cost. Recalling that $F_j = F^d + \nu_j$, and applying a conditional expectation to (10) implies:

$$F^d + E[\nu_j | j \in A_d^1] \leq E[\overline{F}_j(\theta_0) | j \in A_d^1]$$

The expectation on the RHS is identified. If we could assert that $E[\nu_j | j \in A_d^1] = 0$, we would have identified an upper bound on the parameter F^d . However, this conditional expectation is not zero: unlike the structural error e_j , the structural error ν_j was known to the firm at the time it committed to its product choices. While its unconditional mean is zero, its mean *conditional* on the product being offered need not be zero. The term $E[\nu_j | j \in A_d^1]$, therefore, represents a selection bias.³⁴

To circumvent this problem, I proceed with a strategy that allows me to obtain bounds on F_j for every potential product $j \in \mathbf{J}^d$, which hold regardless of whether this product is offered or not. It is then possible to obtain inequalities which involve the unconditional mean of ν_j , which does equal zero. To that end, I impose a finite-support condition on firm d 's fixed costs:

³³Note that additional necessary conditions could be exploited, e.g., conditions that involve multi-product deviations from the firm's observed portfolio. While this means that the procedure developed here does not exploit all the information provided by the data and the model, the results section shows that the information used is sufficient to address the questions of interest.

³⁴See Pakes, Porter, Ho and Ishii (2006) for a discussion of this type of selection problems in models involving inequalities.

Assumption 2. $\sup_{j \in \mathbf{J}^d} \{F_j\} = F_d^U < \infty$, $\inf_{j \in \mathbf{J}^d} \{F_j\} = F_d^L > -\infty$

This assumption suggests that fixed costs associated with firm d 's products have a bounded support given by $[F_d^L, F_d^U]$. The assumption that fixed costs are bounded from below is clearly weak as they can be assumed to be nonnegative. The next assumption guarantees that we can identify an interval which contains this support:

Assumption 3. $[F_d^L, F_d^U] \subset \text{supp}(\text{expected change in variable profit due to elimination or addition of a single product by firm } d)$

This assumption states that the support of the fixed costs is contained within the support of the expected change in variable profit resulting from single-product changes to the firm's product portfolio. As indicated in (10) and (11) above, such expected changes in variable profit are identified, and, as a consequence, so is their support, denoted by $[V_d^L(\theta_0), V_d^U(\theta_0)]$.³⁵

For each product $j \in \mathbf{J}^d$, define the following two random variables:

$$L_j(\theta_0) = \begin{cases} V_d^L(\theta_0) & j \in A_d^1 \\ \underline{F}_j(\theta_0) & j \in A_d^0 \end{cases} \quad U_j(\theta_0) = \begin{cases} \overline{F}_j(\theta_0) & j \in A_d^1 \\ V_d^U(\theta_0) & j \in A_d^0 \end{cases}$$

The following bounds on F_j now apply to *any* potential product of firm d (i.e., without conditioning on whether the product is offered):

$$L_j(\theta_0) \leq F_j \leq U_j(\theta_0) \quad \forall j \in \mathbf{J}^d \quad (12)$$

We can now apply an *unconditional expectation* to obtain:

$$EL_j(\theta_0) \leq F^d \leq EU_j(\theta_0) \quad \forall j \in \mathbf{J}^d \quad (13)$$

The inequalities in (13) define the identified set for the firm-specific mean fixed cost parameter F^d , denoted $[\mu^\ell, \mu^u]$ where $\mu^\ell = EL_j(\theta_0)$ and $\mu^u = EU_j(\theta_0)$ for all $j \in \mathbf{J}^d$.

The estimated set is obtained by replacing the true variable-profit parameter vector θ_0 with its consistent estimator $\hat{\theta}$, described in section 4.1 above, and computing the appropriate sample averages.³⁶ That is, the estimated set is given by $[\bar{\ell}_n^d(\hat{\theta}), \bar{u}_n^d(\hat{\theta})]$ where

³⁵Assumption 3 is reasonable since adding or removing a popular product can have a substantial impact on variable profit, while adding or removing a niche product could have a very minimal effect. This suggests that the length of the support $[V_d^L(\theta_0), V_d^U(\theta_0)]$ should be quite large. In contrast, the impact on fixed costs of adding or removing a product primarily involves the added (or saved) per-product inventory management costs, or sales and marketing costs, which support can be assumed to be shorter than the support of the changes to variable profit.

³⁶Plugging in the estimator $\hat{\theta}$ for θ_0 means that the confidence interval described below needs to be corrected to account for the error in estimating θ . Formally this can be done by a bootstrap. I do not perform this adjustment since it would be computationally expensive and, as discussed in the appendix, would not affect the findings of the paper.

$$\bar{\ell}_n^d(\hat{\theta}) = (1/n^d) \sum_{j=1}^{n^d} L_j(\hat{\theta}), \quad \bar{u}_n^d(\hat{\theta}) = (1/n^d) \sum_{j=1}^{n^d} U_j(\hat{\theta})$$

with $n^d = |\mathbf{J}^d|$ denoting the number of firm d 's potential products. Following arguments in Imbens and Manski (2004), we can construct a $(1 - \alpha) \times 100\%$ confidence interval for F^d by constructing appropriate one-sided intervals for the sample averages:

$$\left[\bar{\ell}_n^d(\hat{\theta}) - \frac{S_\ell(\hat{\theta})}{\sqrt{n^d}} z_{1-\alpha}, \bar{u}_n^d(\hat{\theta}) + \frac{S_u(\hat{\theta})}{\sqrt{n^d}} z_{1-\alpha} \right] \quad (14)$$

where $S_\ell(\hat{\theta})$, $S_u(\hat{\theta})$ are estimators of the standard deviation of L_j and U_j , respectively.³⁷

Additional details are available in Appendix A.2 which describes the manner by which the expected variable profits are simulated using the estimated empirical distribution of the error terms, provides details regarding a procedure that adjusts the estimated support bounds $V_d^L(\theta_0)$ and $V_d^U(\theta_0)$ to correct a finite-sample bias, and offers a discussion of sufficient conditions for consistency under possible dependence among the L_j (and the U_j) variables.

5 Estimation Results

5.1 Estimation Results: Variable Profit 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 fixed bundle of characteristics becomes obsolete over time, most likely due to the emergence of advanced software applications which require better hardware.

³⁷Imbens and Manski (2004) discuss an adjustment that accounts for the case where the length of the identified set is "small" such that the estimators of the endpoints may cross due to sample variation. This case is assumed away in the current context, which seems reasonable given the length of the estimated intervals reported in Table 8 below.

Full-model (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-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 logit results, the negative β coefficient on the time trend implies that a fixed bundle of characteristics 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 a given vendor dummy captures the effect of brands of that vendor which were not included, and not an “overall” vendor 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 . This would have made it less reasonable to assume that firms do not observe these errors until after they have committed to their configuration choices.³⁸

Table 6 offers an insight into some important economic implications of the estimated coeffi-

³⁸I do not, however, control for every brand, but rather for a large number of them.

cients. 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. This also suggests that it is important to take into account both CPU and PC characteristics when modeling demand in this market.

An important aspect of PC demand is the pace at which households’ utility from a fixed bundle of characteristics drops over time, as captured by the taste parameters associated with the time trend. Table 6 reports that the average household is “willing to pay” a negative amount of \$(-257) for a passing of one year. This means that, *holding everything else equal*, the average willingness to pay for fixed hardware drops by this amount every year, presumably since new software applications require better hardware over time. A sizable household heterogeneity along this dimension is displayed in Figure 4. Such heterogeneity is to be expected (for example, a gamer’s utility from a fixed PC product may drop much faster than that of a basic user).

To summarize, the estimated demand parameters convey strong heterogeneity among households in terms of their willingness to pay for cutting-edge technology. This heterogeneity affects both PC makers’ incentives to offer vertically-differentiated configurations, and the welfare implications of such choices. Both these issues are investigated in Section 6 below.

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 marginal errors, $\xi_j(\hat{\theta})$ and $\omega_j(\hat{\theta})$.

5.2 Estimation Results: Fixed Cost Parameters F^d

Bounds on fixed cost parameters were estimated as explained in section 4.2 above. As described there, this estimation involves computing the expected changes in variable profit associated with single-product deviations from firms' observed product portfolios.

I focus attention on the three major notebook producers: Dell, HP and Toshiba, and estimate the mean fixed cost F^d associated with notebook products of each of these firms. I include in this procedure notebook product lines that mainly target the Home segment of the market. A total of four such product lines are included (one from Dell, one from Toshiba, and two from HP). These are also the four top-selling notebook brands, and the counterfactual analysis described below focuses on the configurations offered with these four brands. Information from product offerings in quarters 7 through 12 is used in the estimation of the F^d parameters.³⁹

A key step is the definition of the set H , i.e., the set of feasible CPU technologies which vendors could offer with their product lines in the relevant quarter. This set defines the set \mathbf{J}^d , i.e., the set of firm d 's potential products. I include in this set Intel CPU technologies that satisfy the following inclusion criteria: (i) the CPU technology must be offered by at least one of the four leading brands mentioned above (ii) it must sell at least 10,000 units in the quarter, and (iii) it must be offered by at least two firms. These criteria are designed to make sure that "marginal" technologies (e.g. a CPU that Intel no longer offers to PC makers, but a small quantity of it may still be in stock) are not included in this set. The sets H obtained by applying these criteria in the various quarters are reported in Table 7.⁴⁰

I report the estimation results for the mean-fixed cost parameters in Table 8. The table reports estimated sets computed with and without an adjustment to the estimated support bounds of variable profit differences, as explained in Appendix A.2. The adjustment widens the estimated sets, but only in a modest fashion. The estimated sets (in \$ million) for the mean fixed costs of a notebook configuration belonging to the included product lines of Dell and Toshiba appear rather similar: [2.353, 4.559] and [2.555, 4.119], respectively. In contrast, the mean fixed costs for HP's included product lines appear lower, with an estimated set of [1.081, 2.795]. We cannot, however, reject a null hypothesis according to which all these means of fixed costs are equal.

³⁹Not using information from quarters 1 through 6 saves on computation time. It also avoids assuming that the mean fixed cost is stable throughout the entire sample, and instead requires its stability in the second half of the sample only. The focus on the latter part of the sample stems from the fact that the counterfactual analysis reported below, in which fixed cost estimates are utilized, focuses on the last sample quarter, 2004Q2.

⁴⁰A practical issue is that, as explained above, I exclude products that 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.

6 Using the Estimated Model: Counterfactual Analysis

In this section I analyze the impact of Intel's introduction of its Pentium M processor, which is considered a major innovation in mobile computing. Section 6.1 provides background and a description of the questions of interest. Section 6.2 describes the practical details of the counterfactual experiment, and section 6.3 provides the results.

6.1 The Impact of Intel's Pentium M: Background and Questions of Interest

Rather than offering a further increase in clock speed, Intel's Pentium M chip 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 consumption and in longer notebook battery life.⁴¹

Pentium M-based notebooks appear in the sample for the first time in the first quarter of 2003 (see Table 3b). The goal of my analysis is to answer the following questions: (1) what was the impact of the Pentium M's presence on product choices and prices in the notebook segment? (2) what was the impact of this innovation on various consumer types? and (3) did the Pentium M crowd out PC configurations based on older technologies, and, if so, was the elimination of such technologies socially efficient?

The introduction of the Pentium M was accompanied by a gradual exit of older Intel mobile CPUs such as the Pentium III. In the last sample period, i.e., the second quarter of 2004, only 2% of notebooks sold were Pentium III-based.⁴² Among the five top-selling notebook product lines (i.e., notebook brands) in that quarter, only one recorded positive sales of a Pentium III-based configuration.⁴³ In the quarter immediately preceding the Pentium M's introduction, however, Pentium-III based notebooks enjoyed a market share of 14.1%, and were offered by the two top-selling brands. While this could suggest that the Pentium M played a key role in 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 market between 2003Q1 and 2004Q2.

Importantly, the Pentium M's market share in the notebook segment reached 31.8% by 2004Q2. This makes its analysis interesting at that point in time; an earlier analysis, at a point when this chip was making more modest sales, would have been of limited interest.

⁴¹ "Bigger Notebooks Still Using Older Mobile Chips", Tom Krazit, IDG News Service, September 28, 2004.

⁴² Excluding Apple products, PCs with CPUs not made by Intel or AMD, and products with negligible sales.

⁴³ That configuration had very small sales, and it is possible that it recorded positive sales simply because a small remaining stock was cleared.

6.2 Description of the Counterfactual Analysis

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 outcomes in the observed equilibrium provides a measure of the Pentium M’s effect.⁴⁴

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.⁴⁵ This allows me to ask how many Pentium III-based PC configurations would have been offered *in the absence of the Pentium M*. Certain restrictions are imposed on the analysis to overcome the associated computational burden, and they are described below.

Computing “potential 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 the set of counterfactual equilibria, and use this set to place bounds on welfare predictions. What I actually compute, however, is the set of outcomes *that cannot be ruled out* as equilibria of the game. The reason for this approach is the partial identification of the fixed costs, which implies that it is not always possible to unambiguously rule out a particular outcome as an equilibrium.

Recall that A_d was used to denote a vector of binary indicators describing the observed 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 that 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 may not be 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 elements of A^e . Once the set A^{pe} is computed, I can compute welfare measures at each of its elements, and use this information to place bounds on the counterfactual welfare predictions.

⁴⁴As explained below, I compare expected outcomes (i.e. expectations over the distribution of the error terms e) given the counterfactual and observed sets of products.

⁴⁵This 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.

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 has to check for necessarily-profitable deviations from each of the $2^{|A|}$ vectors in \mathbf{A} . This requires evaluating expected variable profits at the vector examined, and at each of the possible deviations, with each such evaluation being expensive: as in the estimation of fixed costs described above, expected variable profits are simulated by drawing from the distribution of the error terms e , computing a price equilibrium and variable profits at each such draw, and averaging over the simulated draws. For this reason I impose several restrictions on the analysis.

First, I focus the analysis on configuration choices by the four top-selling notebook brands in 2004Q2. This means that I only allow product configuration choices that pertain to these notebook brands to vary in the experiment. I refer to these as the “participating” brands. At the same time, *all* PC products (notebooks and desktops) are included in the experiment and their prices are treated as endogenous.

Second, I restrict the set H of potential configurations by requiring that firms offer Pentium 4 configurations in the speed ranges 1.5-1.99 and 2-2.99 GHz. In the observed equilibrium, all the participating brands offer these configurations. Fixing these choices allows me to reduce the computational burden while still treating as endogenous the most interesting choices: those which pertain to offering low-end configurations such as the Pentium III or the Celeron, and to offering Pentium 4 chips with speed range above 3 GHz. The latter technology can be viewed as a competitor of the Pentium M in the high-end segment of the market.⁴⁶

These restrictions imply that the set of CPU technologies over which firms make endogenous choices in this experiment (i.e., over which the choices are not fixed) is:

$$\left\{ \mathbf{P3}_{1.5-1.99}, \mathbf{C}_{1.5-1.99}, \mathbf{C}_{2-2.99}, \mathbf{P4}_{3-3.99} \right\}$$

With these four CPU technologies, and the four participating notebook brands, we have that $|A| = 16$, that is, 16 product choices are treated as endogenous. I reduce this number to 9 by imposing that firms must make a single “Celeron” choice (i.e., they can either offer both of the Celeron configurations, or none), and that HP, that owns two of the four participating brands, must make the same configuration choices for both its brands. This leaves me with the task of evaluating $2^9 = 512$ vectors as candidates for inclusion in the set A^{pe} of potential equilibria. Additional details on computation, including a complete description of the algorithm, are available in Appendix B. Specifically, I explain there how a “strategic substitutes” conjecture helps reduce the computational burden by making it possible to a-priori rule out some product-choice vectors as potential equilibria.

⁴⁶The Pentium M emerged as the winning technology since improving the performance of the Pentium 4 required increasing amounts of power consumption.

6.3 The Impact of the Pentium M: Results

As explained above, the counterfactual experiment evaluates the impact of the presence of Intel’s Pentium M in 2004Q2 by comparing outcomes in the observed equilibrium to counterfactual predictions for a hypothetical “no Pentium M” scenario. Since firms make product choices prior to observing the realizations of the demand and marginal cost errors e , I evaluate the welfare measures in both the counterfactual scenario, and in the observed equilibrium, as simulated expectations over the distribution of e .⁴⁷ I answer, in turn, the three questions stated above: what was the Pentium M’s impact on product choices and prices? What was its impact on various consumer types? and finally, did it prompt inefficient product elimination?

1. The Pentium M’s impact on product offerings and prices. Table 9 reports the impact of the presence of the Pentium M on expected 2004Q2 outcomes. Given the set of products in the observed equilibrium, the expected total notebook sales are 1.707 million. In the absence of the Pentium M, these sales are between 1.379 and 1.614 million.⁴⁸ This suggests that the Pentium M increases the expected total notebook sales by 5.8% to 23.8%. Some of this growth comes at the expense of expected Desktop sales, which are depressed by 0.9% to 3.1%. The expected sales-weighted average notebook price is between \$829 and \$872 in the absence of the Pentium M, compared to \$905 in its presence. These findings suggest that the Pentium M made a significant contribution to the growth of the mobile market segment.

Table 9 continues to report the Pentium M’s impact on the product configurations offered by the four top-selling notebook product lines. In the presence of the Pentium M, none of these brands offered Pentium III configurations with speed in the 1.5-1.99 GHz range. In contrast, in the absence of the Pentium M, between one and four of these brands would have offered such a configuration. The Pentium M also crowds out configurations based on Intel’s Pentium 4 in the 3-3.99 GHz range. The prediction for the Celeron-based products is ambiguous: some potential equilibria imply that they were crowded out, while others imply the opposite.

The bottom panel of Table 9 reports that the presence of the Pentium M reduces the total expected share of the Pentium III in the notebook segment from 15.6%-23.9% to merely 7.7%, suggesting that the Pentium M played a key role in eliminating the Pentium III technology.⁴⁹

2. The Pentium M’s impact on various consumer types. Table 10 reports the impact

⁴⁷To be clear, after computing the set of potential equilibria in the hypothetical scenario as explained in section 6.2 above, I compute welfare measures at each such outcome, and then use these measures to place bounds on the welfare outcomes that would have obtained had the Pentium M been absent from the market. These outcomes are then compared to the welfare outcomes which obtain given the observed sample, i.e., in the presence of the Pentium M. All the compared quantities are computed in terms of expectations over the distribution of the e error terms.

⁴⁸To be clear, these values represent the highest and lowest values recorded over the set of potential equilibria AP^e .

⁴⁹These shares pertain to all Pentium III chips, and not just those at the 1.5-1.99 GHz range.

of Intel's Pentium M on the expected consumer surplus. With the observed set of products, the expected total consumer surplus is \$1.21 billion, whereas in the absence of the Pentium M, it is between \$1.14 billion and \$1.18 billion. The expected consumer surplus is, therefore, boosted by 3.18% to 6.29% as a consequence of the Pentium M's presence.

The table continues to report a breakdown of the expected benefits from the Pentium M's presence that accrue to various consumer segments. That is, the impact on the total expected surplus for *quantiles* of consumer price sensitivity is reported. The benefits from innovation are not evenly distributed among different consumer types, with the vast majority of the benefits garnered by the 20% least price sensitive consumers. The impact on other groups of consumers is much smaller. Innovation may affect price-sensitive consumers via two channels: while they are hurt by the elimination of basic technologies, they are also helped by the innovation's competitive pressure on the prices of surviving technologies.

3. Was the elimination of older technologies efficient? Having established that the Pentium M played a key role in the crowding out of older technologies, we may ask if it was actually efficient for such technologies to leave the market. To investigate whether the absence of the Pentium III from the product lines of the major notebook producers in 2004Q2 reflected a market failure, I consider a hypothetical action by a social planner: adding to the market Pentium III-based configurations (with 1.5-1.99 GHz) of the four top-selling notebook brands. Similarly, I also calculate the impact of adding Pentium 4 configurations in the 3-3.99 GHz speed range (since one of the four brands had such a configuration in the observed sample, such configurations were added to the other three brands).

The results of this analysis are presented in Table 11. Adding the Pentium III-based notebooks to the market increases the total expected consumer surplus by \$9.368 million, or by 0.77%. It also increases total producer variable profit by \$2.802 million. On the other hand, producers (and hence society) would have also incurred additional fixed costs ranging between \$7.071 million and \$14.267 million.⁵⁰

Defining welfare as the sum of consumer and producer surplus, the total expected impact on welfare ranges between a negative effect of \$(-2.097) million, and a positive effect of \$5.099 million. Since the lower bound is negative, we cannot affirmatively conclude that the absence of the Pentium III-based notebooks reflects a market failure. While the positive upper bound does suggest a potential inefficiency, the scope of the lost welfare appears modest.

Table 11 also reports that performing the same analysis for fast Pentium 4 chips yields similar findings: adding such configurations increases the expected totals of consumer surplus and variable profits, but does not necessarily increase the expected total welfare. A key difference

⁵⁰Evaluated using the boundaries of the estimated set for the relevant firms' mean fixed costs.

between adding Pentium III and Pentium 4 based products, however, pertains to the distribution of benefits to consumers: The benefits from adding Pentium 4-based configurations are almost exclusively garnered by the 20% least price-sensitive consumers, whereas adding Pentium III configurations benefits additional consumer segments in a more pronounced fashion.

In sum, adding the eliminated technologies back to the market could have improved consumer welfare. Specifically, adding back the Pentium III technology can be beneficial to consumers who are more price-sensitive than the group that enjoys the short-run benefits of innovation. Defining total welfare as the sum of consumer and producer surplus, however, means that these benefits would be largely offset by the impact on producer profits.

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. The analysis provides identified sets for fixed cost parameters that take a simple form: intervals which endpoints are defined by means of random variables. As a consequence, I am able to employ rather simple techniques to perform inference on the partially-identified parameters.

I provide a rich analysis of PC product variety by allowing for a large product space, and develop techniques that alleviate the burden associated with predicting counterfactual outcomes. A major challenge that arises in this framework is sample selection, which could potentially impact both the point-estimators of variable profit parameters, and the set-estimators of fixed cost parameters. Both issues are addressed to ensure consistent estimation.

I find strong evidence for consumer heterogeneity along key dimensions such as price sensitivity. Using the estimated model in counterfactual analysis, I find that Intel's introduction of its Pentium M chip contributed significantly to the growth of the mobile segment of the PC market, and to total consumer surplus, while crowding out older technologies. The scope for inefficient product elimination appears to be very limited. I also find that the lion's share of the short-run effect of innovation is enjoyed by the 20% least price-sensitive consumers, while other consumer segments are largely unaffected. Importantly, however, I do not account for long-term benefits such as complementary software innovations, and some of these benefits are likely to be enjoyed by price-sensitive consumers.

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

(2009) 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 potential equilibria in a given quarter, and then ask what was special about the equilibrium that was actually played by firms.

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 21st century economy.

References

- [1] Berry, S. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *RAND Journal of Economics*, 25(2): 242-262
- [2] Berry, S., Levinsohn, J., and A. Pakes (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4): 841-90
- [3] 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 University Press
- [4] Caplin, A., and B. Nalebuff (1991): “Aggregation and Imperfect Competition: On the Existence of Equilibrium,” *Econometrica*, 59(1): 25-59
- [5] Ciliberto, F., and E. Tamer (2009): “Market Structure and Multiple Equilibria in Airline Markets,” *Econometrica*, 77(6): 1791-1828
- [6] Eizenberg, A. (2009): “Estimating Multiple-Equilibria Models of Product Variety,” *Yale University Ph.D. Dissertation*.
- [7] Gawer A., and M.A. Cusumano (2002): “Platform Leadership: How Intel, Microsoft, and Cisco Drive Industry Innovation,” Boston, MA: *Harvard Business School Press*
- [8] Goeree, M. S. (2008): “Limited Information and Advertising in the US Personal Computer Industry,” *Econometrica*, 76(5): 1017-1074
- [9] Goettler, R., and B. Gordon (2009): “Does AMD Spur Intel to Innovate more?,” *mimeo*
- [10] Gordon, B. (2009): “A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry,” *Marketing Science*, 28(5): 846-867
- [11] Haile, P.A., and E. Tamer (2003): “Inference with an Incomplete Model of English Auctions,” *Journal of Political Economy*, 111(1): 1-51
- [12] Hall, P., and B.U. Park (2002): “New Methods for Bias Correction at Endpoints and Boundaries,” *The Annals of Statistics*, 30(5): 1460-1479
- [13] Hendel, I. (1999): “Estimating Multiple-Discrete Choice Models: An Application to Computerization Returns,” *Review of Economic Studies*, 66(2): 423-446
- [14] Hoeffding, W. and H. Robbins (1948): “The Central Limit Theorem for Dependent Random Variables,” *Duke Math. J.*, 15: 773-780

- [15] Imbens, G., and C. Manski (2004): “Confidence Intervals for Partially Identified Parameters,” *Econometrica*, 74(6): 1845-1857
- [16] Ishii, J. (2008): “Interconnection Pricing, Compatibility, and Investment in Network Industries: ATM Networks in the Banking Industry,” *mimeo*
- [17] Lehmann, E.L. (1999): “Elements of Large-Sample Theory,” New York: *Springer Texts in Statistics*
- [18] Mankiw, G.N., and M.D. Whinston (1986): “Free Entry and Social Inefficiency,” *RAND Journal of Economics*, 17(1): 48-58
- [19] Manski, C.F. (2003): “Partial Identification of Probability Distributions,” Berlin, Heidelberg, New York: *Springer-Verlag*
- [20] Mazzeo, M. (2002): “Product Choice and Oligopoly Market Structure,” *RAND Journal of Economics*, 33(2): 221-242
- [21] Nevo, A. (2000): “A Practitioner’s Guide to Estimation of Random-Coefficients Logit Models of Demand,” *Journal of Economics and Management Strategy*, 9(4): 513-548
- [22] Nevo, A. (2001): “Measuring Market Power in the Ready-To-Eat Cereal Industry,” *Econometrica*, 69(2): 307-342
- [23] Nosko, C. (2010): “Competition and Quality Choice in the CPU Market,” *mimeo*
- [24] Pakes, A. (2003): “A Reconsideration of Hedonic Price Indexes with an Application to PC’s,” *American Economic Review*, 93(5): 1578-1596
- [25] Pakes, A., J. Porter, K. Ho, and J. Ishii (2006): “Moment Inequalities and their Application,” *mimeo*, Harvard University
- [26] Petrin, A. (2002): “Quantifying the Benefits of New Products: The Case of the Minivan,” *Journal of Political Economy*, 110: 705-729
- [27] Rosenberg, N. (1979): “Technological Interdependence in the American Economy,” *Technology and Culture*, 20: 25-50
- [28] 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
- [29] Song, M. (2010): “A Dynamic Analysis of Cooperative Research in the Semiconductor Industry,” accepted *International Economic Review*

- [30] Spence, M. (1976): “Product Selection, Fixed Costs, and Monopolistic Competition,” *The Review of Economic Studies*, 43(2): 217-235
- [31] Train, K. (2003): “Discrete Choice Methods with Simulation,” Cambridge, MA: *Cambridge University Press*
- [32] Trajtenberg, M. (1989): “The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners,” *Journal of Political Economy* 97: 444-479
- [33] Walters, E.G. (2001): “The Essential Guide to Computing,” *Prentice Hall PTR*
- [34] Wooldridge, J.M. (2002): “Econometric Analysis of Cross Section and Panel Data,” Cambridge, MA: *The MIT Press*

A Estimation Details

A.1 Estimation Details for the Variable Profit Parameters θ

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 draw the v_i household-specific taste shifters for $ns = 3000$ households. To reduce the error induced by simulation, I use antithetic draws.⁵¹ I then obtain the market share predicted by the model for product j (quarter indices suppressed) as follows:

$$s_j(x, p, \delta, P_{ns}; \theta_2) = \frac{1}{ns} \sum_{i=1}^{ns} \frac{\exp(\delta_j + \mu_{ij})}{1 + \sum_{m \in J} \exp(\delta_m + \mu_{im})} \quad (15)$$

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:

$$s(x, p, \delta, P_{ns}; \theta_2) = S \quad (16)$$

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:

$$\delta^{h+1} = \delta^h + \ln(S) - \ln[s(x, p, \delta^h, P_{ns}; \theta_2)] \quad (17)$$

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

$$\xi(\theta^d) = \delta(\theta_2) - x\beta \quad (18)$$

where x is a covariate matrix for the products observed in the sample. Marginal cost unobservables are computed from (7):

$$\omega(\theta) = \log[p - (T * \Delta(\theta_2))^{-1}s] - x\gamma \quad (19)$$

Next, I define the GMM objective function. Recall that $z_j(X)$ is a $1 \times L$ vector, and 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:

⁵¹See Train (2003). Antithetic draws are used in Goeree's (2008) analysis.

$$Q_N(\theta) = \left[\sum_{j=1}^N g_j(\theta) \right]' \Phi^{-1} \left[\sum_{j=1}^N g_j(\theta) \right] \quad (20)$$

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 Estimation Details for the Fixed Cost Parameters F_d

Computing the values $L_j(\hat{\theta})$ and $U_j(\hat{\theta})$ requires estimating the following quantities: $\underline{F}_j(\theta_0)$ for all $j \in A_d^0$ and $\overline{F}_j(\theta_0)$ for all $j \in A_d^1$, as well as the support bounds $V_d^L(\theta_0)$ and $V_d^U(\theta_0)$ of such expected changes in variable profits. The upper bounds $\overline{F}_j(\theta_0)$ for all $j \in A_d^1$ are estimated as follows: the BLP estimate $\hat{\theta}$ implies empirical values $e_j(\hat{\theta})$ for all the products observed in the sample, a total of 2,287 (see Appendix A.1). From this empirical distribution, I draw 1000 vectors of error terms e for all the potential products $j \in \mathbf{J}$ in the relevant quarter. Recalling that $e = (\xi, \omega)$, I am effectively drawing from the *joint distribution* of the shocks to demand and marginal cost, which reflects their positive correlation.

At each such simulated error vector, I compute price equilibria under (A_d) and $(A_d - 1_d^j)$ (i.e., with and without product j), and compute the decrease in variable profit associated with eliminating product j that appears in (10).⁵² Averaging over these simulated decreases in variable profit yields the estimate of $\overline{F}_j(\theta_0)$, denoted $\overline{F}_j(\hat{\theta})$. An analogous procedure yields estimates for $\underline{F}_j(\theta_0)$ for all $j \in A_d^0$, by simulating the expected increase in variable profit associated with adding product j to the firm's portfolio from (11). Such estimates are denoted by $\underline{F}_j(\hat{\theta})$.⁵³

Estimating the support bounds. We can now collect all the estimated $\overline{F}_j(\hat{\theta})$ and $\underline{F}_j(\hat{\theta})$ in a vector V_d , which j^{th} element is given by:

$$V_d(j) = \begin{cases} \overline{F}_j(\hat{\theta}), & j \in A_d^1 \\ \underline{F}_j(\hat{\theta}), & j \in A_d^0 \end{cases}$$

A natural estimator for $V_d^L(\theta_0)$ is given by $\min_{j \in \mathbf{J}^d} \{V_d(j)\}$, and an estimator for $V_d^U(\theta_0)$ is given by $\max_{j \in \mathbf{J}^d} \{V_d(j)\}$. Such estimators, however, are likely to suffer from finite-sample bias.

A simple method of correcting this bias is discussed in Hall and Park (2002). Consider a

⁵²Price equilibria are simulated by iterating on the first-order conditions (7) until convergence, which typically takes a few seconds of computation time.

⁵³A feature of this procedure is that the empirical distribution of the error terms includes some favorable values (e.g., high utility shocks). Since profits are non-linear functions of these error terms, the simulations can overstate products' variable profit potential. I performed several robustness checks (e.g. setting the errors to zero or imposing a finite-support condition on the joint distribution of the mean utility and marginal cost), and found that the qualitative findings of the paper are robust to this issue.

random sample (X_1, \dots, X_n) from the distribution of the random variable X , and order it as $X_{(1)}, \dots, X_{(n)}$ where $X_{(k)} > X_{(k-1)}$ for each $k \in \{2, \dots, n\}$. A downward-biased estimator of the upper bound of the support of X is $X_{(n)}$. The following correction for this estimator is suggested:

$$X_{(n)} + \frac{\sum_{i=1}^m (X_{(n-i+1)} - X_{(n-i)})K(i/m)}{\sum_{i=1}^m K(i/m)}$$

where $K(\cdot)$ is a Kernel function.⁵⁴ Using $m = n^{2/3}$ provides a rate of convergence of $n^{-1/3}$. The results section below reports estimated sets for the fixed cost parameters F^d with and without such adjustment for finite-sample biases in the estimation of support boundaries. Reassuringly, the impact of such adjustments on the estimated intervals is not large.

A second source of bias we may expect has to do with the fact that we do not observe realizations of the actual random variable in which we are interested (that is, variable profit differences), but rather estimates of such quantities that contain sampling error. Estimates of the maximum (minimum) of such quantities suffer from an upward (downward) bias as discussed in Haile and Tamer (2003). However, this is an “outward” bias that does not invalidate the procedure, since all that is required is that the true quantities are contained within the boundaries. This second source of bias does imply, however, that the estimates of the bounds on F^d would be conservative in the finite sample.

Finally, note that the estimation approach for the fixed costs parameters is based on the sample average being a consistent, asymptotically-normal estimator of the mean, which is guaranteed for IID sequences. We can allow for dependence among the L_j variables for $j = 1, \dots, n^d$ given conditions which, heuristically, imply that the dependence decays at a fast-enough rate as observations become further apart from each other. For brevity, I focus the discussion on the L_j variables, but it is understood that the same conditions need to be satisfied by the U_j variables as well.

Following Lehmann (1999) (p. 60), consistency of the set estimator $[\bar{\ell}_n^d(\hat{\theta}), \bar{u}_n^d(\hat{\theta})]$ is retained with dependent variables if the variables L_j have finite second moments and satisfy

$$\sum_{i=1}^{n^d} \sum_{j \neq i}^{n^d} Cov(L_i, L_j) = o((n^d)^2)$$

The asymptotic normality of the sample average $\bar{\ell}_n^d(\hat{\theta})$ (and $\bar{u}_n^d(\hat{\theta})$) is also retained with dependent variables if appropriate conditions are satisfied. In particular, the confidence interval in (14) would be valid with the following: (i) the variables L_j , $j = 1, \dots, n^d$ have finite third moments and are m -dependent.⁵⁵ (ii) $\lim_{p \rightarrow \infty} (1/p) \sum_{h=1}^p A_{i+h} = A$ exists, uniformly for all $i = 0, 1, \dots$

⁵⁴I set $K(u) = (15/16)(1 - u^2)^2 I(|u| \leq 1)$, as in the simulations performed by Hall and Park.

⁵⁵This requires that for some positive integer m , $s - r > m$ implies that the sets (L_1, L_2, \dots, L_r) , $(L_s, L_{s+1}, \dots, L_{n^d})$ are independent.

with $A_i = \text{var}(L_{i+m}) + 2 \sum_{j=1}^m \text{Cov}(L_{i+m-j}, L_{i+m})$. (iii) $S_\ell(\hat{\theta})$ is a consistent estimator of \sqrt{A} .⁵⁶

A HAC covariance matrix estimator can be used to estimate the covariance terms that adjust the confidence interval for dependence. The intervals reported in this paper are not adjusted for dependence in this manner. They also need to be adjusted to account for the variance in the estimation of θ , as well as for the error due to the simulation of expected variable profits. These adjustments can be performed using a computationally-expensive bootstrap approach. As the results sections below indicate, performing this procedure would not change the findings of the paper: it is sufficient to examine the estimated set of fixed costs to notice that one cannot reject the null hypothesis which argues that the elimination of basic product configurations is efficient. Using the confidence interval instead of the estimated set, whether adjusted or not, would not change this.

B Computing Counterfactual Potential Equilibria: The Algorithm

For the purposes of the counterfactual analysis, I set the structural errors ν_j in the specification for fixed costs to zero, which amounts to imposing that the firm's fixed costs are the same for each of its products. This identical per-configuration fixed cost is assumed to lie in the estimated interval for F^d , the firm's mean fixed cost, reported in Table 8 and discussed above. For example, Dell's per-configuration fixed cost is assumed to be between \$2.353 million and \$4.559 million. I denote this estimated interval by \mathbf{C}_d .⁵⁷

Evaluating a product-choice vector A to determine whether it is a potential equilibrium (i.e., a member of A^{pe}) requires computing, for each firm, an interval of its per-configuration fixed cost under which it does not have a profitable deviation from A . Denote this interval, pertaining to firm d , by \mathbf{I}_d^A . If, for each firm d making endogenous product choices, this interval has a non-empty intersection with the estimated interval of its fixed costs, \mathbf{C}_d , the vector A cannot be ruled out as supporting an equilibrium, and is deemed an element of A^{pe} . Also note that \mathbf{I}_d^A itself may be empty, in which case A is clearly ruled out as a potential equilibrium.

Computing the interval \mathbf{I}_d^A tends to be rather expensive, and, in many cases, unnecessary, since one can often quickly verify that a necessarily-profitable deviation exists. For that reason, the actual algorithm used to compute the set of potential equilibria A^{pe} has the following three steps:

1. For each vector $A \in \mathbf{A}$, check whether any firm has a necessarily-profitable single-product deviation from A .

⁵⁶This follows immediately from Theorem 1 in Hoeffding and Robbins (1948).

⁵⁷Another potential approach could be to allow firm d 's per-configuration fixed costs to fluctuate about the mean F_d as suggested by the model, and perform the counterfactual analysis by repeatedly drawing from the distribution of these fixed costs. But since F_d is only partially-identified, it is not clear how to draw from the distribution.

2. For each of the vectors that were not ruled out in step 1, check whether any firm has a necessarily-profitable multi-product deviation.
3. For each vector A that survived step 2, compute \mathbf{I}_d^A for each firm d . If, for every firm d , $\mathbf{I}_d^A \cap \mathbf{C}_d \neq \emptyset$, determine that $A \in A^{pe}$.

Given that steps 1 and 2 already examined each possible deviation and did not find any of them to be necessarily-profitable, step 3 may seem redundant. However, it is possible that firm d 's per-product fixed costs must lie in some interval, denoted $I1$, to prevent one deviation from being necessarily profitable, and must also be inside some other interval, $I2$, to guarantee that *another* deviation is not necessarily profitable. Suppose that both $I1$ and $I2$ have non-empty intersections with \mathbf{C}_d , but that $I1 \cap I2 = \emptyset$. In this case, even though neither of these deviations is necessarily profitable individually, one of them must be profitable. Due to this subtle point, step 3 is necessary. Note that we could perform step 3 only - however, as explained above, performing steps 1 and 2 first saves computation time.

I was able to reduce the computational burden (specifically, the number of vectors that need to be evaluated to determine whether they qualify as potential equilibria) by application of the following conjecture:

Conjecture 1. (*Strategic Substitutes*): *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.⁵⁸ The usefulness of this conjecture is in that, once a certain deviation is found to be necessarily profitable (unprofitable) at some vector A , it can be automatically considered to be profitable (unprofitable) at many other vectors. This made it possible to avoid a direct computation of expected variable profits in about 9% of the 512 vectors in \mathbf{A} . In an earlier draft of this paper, which allowed for a much larger space of potential outcomes (but saved on computation time with a “shortcut” that set the error terms to zero, rather than repeatedly drawing from their estimated distribution as performed in this version), this conjecture allowed me to evaluate 16,384 vectors rather than $2^{24} = 16,777,216$ - an immesne reduction in computation time.

⁵⁸This 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.

C Tables and Figures

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

Year 1		Year 2		Year 3	
Vendor	Share	Vendor	Share	Vendor	Share
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: Adoption Rates of CPU Technologies by Desktop Product Lines

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

Quarter	P3_1-1.49	P4_1-1.49	P4_1.5-1.99	P4_2-2.99	P4_3-3.99
2001Q3	0.67	0.48	0.26	0.00	0.00
2001Q4	0.50	0.65	0.65	0.12	0.00
2002Q1	0.31	0.58	0.73	0.50	0.00
2002Q2	0.13	0.43	0.70	0.65	0.00
2002Q3	0.13	0.26	0.74	0.70	0.00
2002Q4	0.04	0.11	0.37	0.81	0.00
2003Q1	0.00	0.11	0.44	0.81	0.15
2003Q2	0.00	0.11	0.44	0.81	0.15
2003Q3	0.00	0.08	0.28	0.92	0.60
2003Q4	0.00	0.12	0.28	0.92	0.60
2004Q1	0.00	0.08	0.19	0.92	0.65
2004Q2	0.00	0.08	0.17	0.92	0.63

Home market, Intel technologies. Excluding vendors identified as “Others”, Apple products, and products with sales smaller than 100 units in a quarter. C=Celeron, P3=Pentium III, and P4=Pentium 4. P3_0.5-0.99=Pentium III with speed between 0.5-0.99 GHz.

Table 3b: Adoption Rates of CPU Technologies by Portable Product Lines

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.

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 vendors and 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
Dell		12.332	2.603	0.774	0.062
	dimension	-10.426	2.813	-0.915	0.065
	inspiron	-9.908	2.732	-0.838	0.064
	latitude	-7.529	2.137	-0.488	0.071
	optiplex	-13.509	2.819	-0.903	0.064
HP		-0.976	0.334	-0.049	0.021
	evoipaq	-1.651	0.519	-0.174	0.030
	media	7.568	0.870	0.424	0.029
	pavilion	2.625	0.385	-0.015	0.024
	presario	2.593	0.355	0.026	0.020
	cmpq_notebook	1.841	0.653	0.175	0.033
	cmpq_ultprtbl	10.945	2.221	0.741	0.080
Gateway		0.309	0.399	0.068	0.025
	gateway3	-2.619	0.730	-0.408	0.035
	gateway5	1.755	0.865	-0.030	0.048
	gateway7	2.159	0.690	0.077	0.035
	essential	2.124	0.458	-0.098	0.030
	performance	1.751	0.530	0.039	0.034
	media	4.960	0.828	0.365	0.035
	gateway4	-1.320	0.510	-0.139	0.031
	gateway6	4.725	1.077	0.173	0.051
	solo	0.185	0.868	-0.106	0.050
eMachines		0.389	0.602	-0.325	0.050
Toshiba		7.933	1.684	0.479	0.050
	portege	0.593	1.018	0.093	0.059
	port_tablet	2.855	1.303	0.218	0.085
	satellite	-5.405	1.752	-0.517	0.055
	satpro	-2.628	1.141	-0.132	0.053
Sony		5.684	0.821	0.306	0.037
	vaio_ds	-3.909	0.871	-0.265	0.043
	vaio_r	0.500	0.824	0.205	0.052
	vaio_w	3.163	0.979	0.283	0.059
	vaio_505	1.288	0.963	-0.007	0.061
	vaio_fx	0.951	0.734	0.053	0.053
IBM		2.037	1.217	0.208	0.083
	netvista	-3.868	1.307	-0.244	0.087
	thinkCentre	0.419	1.301	0.040	0.095
	thinkpadA	8.348	2.084	0.452	0.097
	thinkpadT	1.253	1.366	-0.016	0.092
	thinkpadR	-3.304	1.291	-0.233	0.085
Acer_veriton		-2.120	0.382	-0.120	0.016
Averatec		1.131	0.688	-0.034	0.048
Fujitsu		-1.090	0.354	-0.018	0.023
MicroElectronics		-1.585	0.236	-0.009	0.017

See notes for Table 5a. Bold-type entries represent a vendor dummy, followed by dummy variables for that vendor's brands. The coefficients on vendors (e.g. Dell) **do not** capture an "overall" vendor effect, but rather the effect of brands of that vendor that were not included.

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: The Sets H of Feasible CPU Technologies

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
P3.0.5-0.99	X	X	X									
P3.1.0-1.49		X	X	X	X	X	X	X	X	X		
C.0.5-0.99	X	X	X									
C.1.0-1.49		X	X	X	X							
C.1.5-1.99					X	X	X	X	X	X	X	X
C.2-2.99						X	X	X	X	X	X	X
P4.1.0-1.49				X	X	X	X					
P4.1.5-1.99			X	X	X	X	X	X	X	X	X	X
P4.2-2.99					X	X	X	X	X	X	X	X
P4.3-3.99										X	X	X
PM.1.0-1.49									X	X	X	X
PM.1.5-1.99								X	X	X	X	X

The table reports, for each of the 12 data quarters, the set H of feasible CPU technologies. See text for the three criteria which determine inclusion in this set.

Table 8: Bounds on Mean Fixed Costs F^d (\$M)

	Estimated Set	Estimated Set (adjusted*)	95% CI**
Dell	[2.427, 4.529]	[2.353, 4.559]	[2.309, 4.914]
HP	[1.094, 2.721]	[1.081, 2.795]	[0.983, 3.009]
Toshiba	[2.622, 4.069]	[2.555, 4.119]	[2.295, 4.453]

This estimation utilizes information on specific notebook product lines of these firms (see text). The number of potential products, n^d , is 41 for Dell and Toshiba and 82 for HP. *With adjustment for support bounds (see Section 4.2). **The confidence interval does not take into account the variance due to the estimation of θ and to the simulation of expected variable profits, see text.

Table 9: Effect of Intel’s Pentium M on Expected 2004Q2 Outcomes

	Observed*	“No Pentium M” Counterfactual	
		Lower bound**	Upper bound**
Total Notebook Sales (M)	1.707	1.379	1.614
Total Desktop Sales (M)	3.831	3.866	3.952
Mean Notebook price*** (\$)	905	829	872
Impact on number of PC configurations (top 4 brands)			
	Observed	“No Pentium M” Counterfactual	
		Lower bound**	Upper bound**
# P3_1.5-1.99	0	1	4
# C_1.5-1.99	3	0	4
# C_2-2.99	3	0	4
# P4_3-3.99	1	2	4
Impact on Pentium III’s share of total Portables sales			
	Observed*	“No Pentium M” Counterfactual	
		Lower bound**	Upper bound**
Share P3	0.077	0.156	0.239

* As explained in the text, these are not outcomes observed in the sample, but rather simulated expected outcomes computed at the set of products observed in the sample, which includes Pentium M based notebooks. This does not pertain to the reported observed number of configurations, which is simply the observed sample quantity. ** The bounds represent the largest and smallest values computed over the set of all potential counterfactual equilibria. *** Sales-weighted average.

Table 10: The Effect of Intel’s Pentium M on Consumers

	Observed*	”No Pentium M” Counterfactual	
		Lower bound**	Upper bound**
Total Expected Consumer Surplus	1213.7	1141.9	1176.3
Expected Surplus for Price Sensitivity Quantiles (see text)			
0-20% Price sensitive	992.7	932.6	955.1
20%-40% Sensitive	130.3	123.1	129.9
40%-60% Sensitive	76.5	73.2	76.8
60%-80% Sensitive	12.6	11.9	12.8
80%-100% Sensitive	1.6	1.1	1.7

All figures in M\$. * As explained in the text, these are not outcomes observed in the sample, but rather simulated expected outcomes at the set of products observed in the sample. ** The bounds represent the largest and smallest values computed over the set of all potential counterfactual equilibria.

Table 11: Expected Welfare Effects of Adding Products Based on Older CPUs

	Added Configurations	
	P3_1.5-1.99	P4_3-3.99
A. Total Welfare Components		
Change to Expected CS:	+9.368	+7.027
Change to PC makers' Expected VP:	+2.802	+1.705
Change to Fixed costs:	[-14.267, -7.071]	[-10.149, -4.515]
Total effect:	[-2.097, +5.099]	[-1.417, +4.216]
B. Effect on Different Consumer Segments (Quantiles of Price Sensitivity)		
0-20% Price sensitive	+6.018	+6.445
20%-40% Sensitive	+2.060	+0.443
40%-60% Sensitive	+0.978	+0.118
60%-80% Sensitive	+0.213	+0.015
80%-100% Sensitive	+0.099	+0.005

Configurations based on Intel's Pentium III chips in the 1.5-1.99 GHz speed range were added to all four participating brands, while configurations based on its Pentium 4 chips with speed above 3 GHz were added to three (see text). All figures in M\$, evaluated at expected 2004Q2 outcomes.




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: WTP for an upgrade from Intel's Celeron to its Pentium M chip (\$)




Figure 4: WTP for "1 year forward" (\$) (see text)