A Hybrid Discrete Choice Model of Differentiated Product Demand with An Application to Personal Computers

Minjae Song*

Simon Graduate School of Business, University of Rochester

January 2014

Abstract

In this paper I consider a new discrete choice model of differentiated product demand that distinguishes a brand-level differentiation from a product-level differentiation. The model is a hybrid of the random coefficient logit model of Berry, Levinsohn, and Pakes (1995) and the pure characteristics model of Berry and Pakes (2007), and describes markets where firms offer multiple products of different qualities under the same brand name. I compare the hybrid model with existing models using data on personal computers. Using the estimates of the hybrid model, I also provide empirical evidence that firms reposition their brands in a post-merger market.

JEL classification numbers: C51, D43, L13, L63

Keywords: Discrete choice models, product differentiation, personal computer market, post-merger analysis

*Please send your comments to minjae.song@simon.rochester.edu. I am grateful to Ana Aizcorbe, Yuk-Fai Fong, Maurizio Iacopetta, Sarit Markovich, Aviv Nevo, Robert Porter, Holger Sieg, anonymous referees, and participants at various seminars and conferences for their suggestions and comments. I especially thank Bureau of Economic Analysis for providing me with data used in this paper. Yang Lin and Mark Yuan provided excellent research assistance. All errors are mine.

1 Introduction

Consider a market where brands compete with multiple products, and products of a given brand can be ranked with respect to product characteristics. The auto market is a good example. There are different brands of cars, and each brand has a low-end economy car, a mid-range car, and a high-end luxury car. The personal computer (PC) market is another example. Products of each brand range from computers of a low CPU speed with low memory and hard disk capacities to computers of a high CPU speed with large memory and hard disk capacities. In such markets each brand targets "high-end" consumers with high quality and high price products and "low-end" consumers with low quality and low price products.

In this paper I consider a new discrete-choice model of differentiated product demand in which product groups are "horizontally" differentiated while products of a given group are "vertically" differentiated with respect to observed characteristics. I call a product group that contains "vertically" differentiated products a brand. This is consistent with the role of brands found in many markets. In my model products of the same brand are not necessarily ranked on the single dimension. In the car market case, for example, it is not straightforward to compare SUVs with sports cars on the single quality dimension. My model can accommodate this case by allowing consumers to differ in multiple dimensions. In the car example, consumers may value the car size differently but agree on the quality ranking conditional on the size.

I distinguish these two types of differentiation by assuming that consumers receive an idiosyncratic shock only at a brand level so that all products within a brand share the same shock. This shock represents each consumer's idiosyncratic "taste" for a brand that is not captured by observed and unobserved product characteristics, and I assume that it is independent and identically distributed with Type I extreme value distribution. As a result, consumers' product choice within a brand is solely determined by their preference for product characteristics, and their brand choice is determined by the realization of the brand-level shock in addition to the preference. The model has at least one random coefficient that describes heterogeneous consumer preferences. With one random coefficient, for example, consumers endowed with different degrees of willingness to pay for overall product quality know for sure which product to choose in each brand, and choose among those products after the taste shock for brands is realized.

I call this model a hybrid model because it resembles the random coefficient logit demand model in describing a brand choice (Berry, Levinsohn, and Pakes 1995; BLP hereafter) and the pure characteristics demand model in describing a within-brand product choice (Berry and Pakes 2007; the PCM hereafter). For the markets described above the hybrid model provides a more realistic substitution pattern than the existing demand models. Since the within-brand product choice is not affected by the idiosyncratic shock, only a subset of products are substitutes for one another within a brand and these substitutes are similar in their characteristics. Thus, the most high-end PC, for example, is the closet substitute for the second most high-end PC and not necessarily a substitute for some low-end PCs. Yet, all brands are substitutable because each brand has a product that consumers consider buying given their types.[1] To a high-end type consumer, for example, all brands are substitutable through their high-end PCs. This is an important feature for the aforementioned markets where firms compete to attract each type of consumer with products of different qualities.

Another important feature is that the hybrid model distinguishes brand entry from product entry. New products introduced by existing brands do not add idiosyncratic shock, while new products by a new brand add one brand-level shock. As a result, consumer welfare does not necessarily grow by adding more products. This is in contrast to models with product-level idiosyncratic shock such as BLP where consumer welfare grows without limit as the number of products increases.

One should note that my hybrid model is different from the nested logit model. The two models are similar in that both divide products into groups and distinguish within-group substitution from cross-group substitution. However, in the nested logit model all products belonging to the same group (brand) are substitutes and they are closer substitutes for one another than for any other brands' products regardless of their characteristics. This means that in the PC example a high-end PC and a low-end PC can be very close substitutes as long as they belong to the same brand and this high-end PC cannot be a better substitute for another brand's high-end PC than

[1] All products are substitutes for one another in BLP and all brands are not necessarily substitutable in the PCM.

for the same brand's low-end PC. This feature arises because the nested logit model does not allow substitution patterns to be determined by product characteristics through the random coefficient. In the hybrid model, on the other hand, another brand's high-end PC could be a better substitute than the same brand's low-end PC and products of the same brand are not necessarily substitutes if their characteristics are substantially different. As shown later, this difference leads to drastically different markup estimates. In the nested logit model firms have substantial market power over their brands because of the nesting structure and this can result in unrealistically high markup estimates. This does not happen in the hybrid model because only a small subset of products are substitutes within brands.

One should also note that the hybrid model is different from the PCM with random coefficients on the brand dummy variables. In the hybrid model the brand-level shock is idiosyncratic, so all consumers have a positive probability of choosing any brand, which makes all brands substitutable for one another. On the other hand, a random coefficient on the brand dummy in the PCM represents (non-idiosyncratic) consumer types, so only consumers endowed with high values of the random coefficient choose that brand, which makes brands less substitutable.[2]

I show the advantages of using the hybrid model in a real world application by estimating demand for desktop PCs.[3] I use monthly product-level data over three years starting from October 2001. The data set covers about 65 percent of PCs sold at U.S. retail stores. Data on prices, sales in unit, and ten product characteristics are available at a UPC bar code level. So I can define the product at the same level as products that consumers actually buy in the market. There are six brands included in the data: Sony, Hewlett-Packard (HP hereafter), Gateway, Compaq, eMachines and Apple.

Previous studies on PC demand use BLP-type models with much more aggregate data. For example, Chu, Chintagunta and Vilcassim (2007) define a desktop PC as a brand-model-CPU

[2] It may not be computationally feasible to include the random coefficients for all brands in the PCM. Song (2007) estimates the PCM with two random coefficients and I do not know any studies that estimate the PCM with more than two random coefficients.

[3] Other researchers have adopted my model for various applications. Byzalov (2010) applies the hybrid model to the cable TV market and finds that it fits data better than the PCM and gives more reasonable predictions in counterfactual exercises than BLP. Ghose, Ipeirotis, and Li (2012) apply the hybrid model to the hotel industry and find that it outperforms the nested logit and BLP in out-of-sample predictions.

type combination. Goeree (2008) defines it as a brand-model-CPU type-CPU speed combination. These data formats eliminate the quality ranking within brands such that products are presented as much more horizontally differentiated than they actually are, and their demand models ignore consumer preferences for non-CPU characteristics. An exception is Bajari and Benkard (2005) who use the same product-level data I use here. Their model is similar to the PCM in that it does not have the idiosyncratic taste shock, but it is more flexible in a sense that they do not make any distributional assumption on the random coefficient for product characteristics. However, they can only set-identify preference, and their demand estimates cannot be used to learn about supply-side parameters such as the marginal cost and markup.

Estimation results show that in the hybrid model almost all coefficients on product characteristics are statistically significant and have expected signs and reasonable magnitudes. The random coefficient on the price variable is also statistically significant, which is especially encouraging because no individual-level data such as the current population survey or the consumer expenditure survey is used in estimation. On the other hand, in BLP no coefficient other than those on the CPU speed and the DVD writer is statistically significant.[4] Although the demand estimates in the PCM are statistically significant and have expected signs, the consumer's willingness to pay for quality improvement is unrealistically low. For example, the average consumer is willing to pay about \$30 to increase CPU speed by 1 GHz. This is much less than the money needed to put in a 1 GHz faster CPU. This low magnitude is a result of ignoring the brand-level differentiation and ranking all products on the single quality dimension. In the hybrid model the average consumer's willingness to pay for the same CPU speed difference is about \$200.

Estimation results also show that the markup level is reasonable (lower than 30 percent) in the hybrid model when the market is defined to include potential consumers, while it is consistently unrealistically high in BLP (higher than 50 percent) and unrealistically low in the PCM (lower than 1 percent). As expected, consumer welfare is less sensitive to new product introduction than in BLP, but somewhat surprisingly it is less sensitive than in the PCM as well. This is because in the

[4] Among the models with the product-level logit error term, demand estimates are most reasonable in the nested logit model. However, its markup estimates are unrealistically high, as mentioned above, with the average percentage markup higher than 1 for all periods.

PCM consumers who buy new products have much higher willingness to pay as the new products are positioned on the far left end of the random coefficient distribution.

Lastly, I show how demand estimates can be used to analyze product repositioning in the post-merger market. In particular, I analyze the PC market after the merger between HP and Compaq, which took place in February 2002. The price gap between HP and Compaq was less than $100 right before the merger but reached about $300 by the end of the sample period mainly due to Compaq's price decline. This price trend is not consistent with the usual prediction that the merged firm will increase prices, at least relative to those of non-merged firms, by exploiting its market power.

To see if this post-merger market trend is associated with product repositioning, I construct a quality index for each brand using the demand estimates in the hybrid model. This index is a linear function of observed product characteristics with each characteristic weighted by the estimated marginal utility. Having statistically significant estimates with realistic magnitudes is important in correctly projecting product positioning on characteristics space. The trend of the quality index shows that Compaq's quality improves more slowly than that of HP such that their difference becomes increasingly larger over time. Moreover, Compaq's quality becomes almost identical to that of eMachines from July 2003, suggesting that Compaq now competes with eMachines at the low end of the market. This provides the evidence that the merged firm repositioned Compaq and HP by including more low-end products in Compaq's product line.

Although I focus in this paper on consumer demand for products that are differentiated both vertically and horizontally, my hybrid model can also be used to describe economic agents' (discrete) choices in other settings. For example, it can provide an alternative way of modeling housing/community choices (Epple and Sieg, 1999; Bayer, Ferreira, and McMillan, 2007). In Epple and Sieg (1999) communities are sorted out in equilibrium by households' income and unobserved types. Their model is similar to the PCM in a sense that agent types determine their choices and equilibrium outcomes are characterized by a joint distribution of agent types. Bayer, Ferreira, and McMillan (2007), on the other hand, take the BLP approach and allow for the idiosyncratic taste shock at each location. The hybrid model can be used to combine these two approaches.

Consider modeling households' neighborhood choices in a region that consists of multiple towns. One can treat towns in the same way as I treat brands. As brands include multiple products that are targeted to different types of consumers, towns include neighborhoods that are best matched to different types of households. In the hybrid modeling approach households receive an idiosyncratic shock for each town and choose a neighborhood to live in by comparing the neighborhoods that they prefer most in each town.

The rest of the paper is organized as follows. Section 2 describes the hybrid demand model, followed by an estimation procedure in section 3. Section 4 compares the hybrid model with BLP and the PCM using product-level data on PCs. Section 5 analyzes the post-merger PC market. Section 6 concludes.

2 Hybrid Demand Model

2.1 Utility Function

Suppose there are $t = 1, 2, ..., T$ markets with K brands. Each brand has J_k products that are ranked on a single quality dimension, i.e., vertically differentiated. Products of the same brand do not have to be vertically differentiated. The model I present below becomes more general by adding more random coefficients, but I will confine my exposition to the simplest case. Each market has $i = 1, 2, ..., I_t$ consumers. Given a market, the indirect utility of consumer i from product j of brand k is

$$u_{ij_k} = \delta_{j_k} - \alpha_i p_{j_k} + \varepsilon_{ik}, \quad \text{for } 1 \leq j_k \leq J_k \text{ and } 1 \leq k \leq K \tag{1}$$

where δ_{j_k} and p_{j_k} are quality and price of product j of brand k respectively, α_i is the individual-specific price coefficient with $\alpha_i \sim F(\theta)$, and ε_{ik} is consumer $i's$ idiosyncratic taste for brand k which is assumed to be a Type I extreme-value random variable. Note that ε_{ik} is independent and identically distributed across consumers and brands, but not across products of the same brand. This means that consumers receive an idiosyncratic taste shock when they consider which brand to choose, but have the same shock for products of the same brand.

I assume that product quality δ_{j_k} is a linear function of product characteristics so that

$$\delta_{j_k} = \mathbf{x}_{j_k}\boldsymbol{\beta} + \xi_{j_k}$$

where $\mathbf{x}_{j_k}$ is a vector of observable characteristics of product j of brand k, $\boldsymbol{\beta}$ represents a marginal utility that a consumer derives from product characteristics, and ξ_{j_k} is the mean quality of characteristics that the consumer observes but the econometrician does not. In a more general version of the hybrid model a subset of $\boldsymbol{\beta}$ can be random variables to describe consumer heterogeneity in multi-dimensions.

2.2 Market Share

Let products in each brand be sorted in the order of ascending price such that $p_{j-1_k} < p_{j_k}$ for $j = 2, ..., J$. A rational consumer with a value of α_i chooses product j of brand k over other products if and only if

$$\delta_{j_k} - \alpha_i p_{j_k} + \varepsilon_{ik} > \delta_{l_r} - \alpha_i p_{l_r} + \varepsilon_{ir}, \quad \forall l_r \neq j_k$$

and her probability of choosing product j of brand k is

$$\frac{\exp\left(\delta_{j_k} - \alpha_i p_{j_k}\right)}{\sum_{r=1}^{K} \exp\left(\max_{l_r \in r}\left(\delta_{l_r} - \alpha_i p_{l_r}\right)\right)} \tag{2}$$

if $\frac{\delta_{j+1_k} - \delta_{j_k}}{p_{j+1_k} - p_{j_k}} \leq \alpha_i < \frac{\delta_{j_k} - \delta_{j-1_k}}{p_{j_k} - p_{j-1_k}}$ and 0 otherwise. Appendix I provides a derivation of this probability using the random utility maximization.

Note that the consumer's value of α_i must fall in the interval given above to have a positive probability of choosing product j_k. That is because there exists only one product in brand k that maximizes her utility such that

$$\delta_{j_k} - \alpha_i p_{j_k} > \delta_{h_k} - a_i p_{h_k}, \quad \forall h \neq j$$

after canceling out ε_{ik} from both sides. This implies that the probability of product j_k being chosen conditional on brand k being chosen is

$$\Pr\left(j_k \text{ being chosen } | k \text{ being chosen}\right) = F\left(\overline{\triangle}_{j_k}\left(\delta_k, p_k\right)|\boldsymbol{\theta}\right) - F\left(\underline{\triangle}_{j_k}\left(\delta_k, p_k\right)|\boldsymbol{\theta}\right), \tag{3}$$

where $F(.)$ is the *cdf* of α_i, $\overline{\triangle}_j\left(\delta_k, p_k\right) = \frac{\delta_{j_k} - \delta_{j-1_k}}{p_{j_k} - p_{j-1_k}}$, and $\underline{\triangle}_{j_k}\left(\delta_k, p_k\right) = \frac{\delta_{j+1_k} - \delta_{j_k}}{p_{j+1_k} - p_{j_k}}$. Note also that equation (2) depends on the consumer's product choice in other brands through a max function in the denominator.

The market share of product j of brand k can be obtained by integrating equation (2) over consumers who choose product j_k among products that belong to brand k such that

$$s_{j_k} = \int_{\alpha_i \in j_k} \frac{\exp\left(\delta_{j_k} - \alpha_i p_{j_k}\right)}{\sum_{r=1}^{K} \exp\left(\max_{l_r \in r}\left(\delta_{l_r} - \alpha_i p_{l_r}\right)\right)} f(\alpha)\, d\alpha$$

where $\alpha_i \in j_k$ indicates those consumers who choose product j in brand k. This share function can be rewritten as

$$s_{j_k} = \left[F\left(\overline{\triangle}_{j_k}\left(\delta_k, p_k\right)|\boldsymbol{\theta}\right) - F\left(\underline{\triangle}_{j_k}\left(\delta_k, p_k\right)|\boldsymbol{\theta}\right)\right] \int_{\alpha_i \in j_k} \frac{\exp\left(\delta_{j_k} - \alpha_i p_{j_k}\right)}{\sum_{r=1}^{K} \exp\left(\max_{l_r \in r}\left(\delta_{l_r} - \alpha_i p_{l_r}\right)\right)} g(\alpha)\, d\alpha \tag{4}$$

where

$$g(\alpha) = \frac{f(\alpha)}{\left[F\left(\overline{\triangle}_{j_k}\left(\delta_k, p_k\right)|\boldsymbol{\theta}\right) - F\left(\underline{\triangle}_{j_k}\left(\delta_k, p_k\right)|\boldsymbol{\theta}\right)\right]},$$

and its numerical approximation is

$$\widehat{s}_{j_k} = \left[F\left(\overline{\triangle}_{j_k}\left(\delta_k, p_k\right)|\boldsymbol{\theta}\right) - F\left(\underline{\triangle}_{j_k}\left(\delta_k, p_k\right)|\boldsymbol{\theta}\right)\right] \frac{1}{ns} \sum_{i \in j_k} \frac{\exp\left(\delta_{j_k} - \alpha_i p_{j_k}\right)}{\sum_{r=1}^{K} \exp\left(\max_{l_r \in r}\left(\delta_{l_r} - \alpha_i p_{l_r}\right)\right)},$$

where ns is the number of simulated consumers whose $\alpha \in \left[\overline{\triangle}_{j_k}\left(\delta_k, p_k\right), \underline{\triangle}_{j_k}\left(\delta_k, p_k\right)\right)$.

Notice that the first part of equation (4) is the same as equation (3) and this is identical to the choice probability of the PCM when choices are confined to products of brand k. The second part of equation (4) is the probability of brand k being chosen by those with $\frac{\delta_{j+1_k} - \delta_{j_k}}{p_{j+1_k} - p_{j_k}} \leq \alpha_i < \frac{\delta_{j_k} - \delta_{j-1_k}}{p_{j_k} - p_{j-1_k}}$

and this resembles the choice probability of BLP.[5] The market share function of equation (4) shows why the model is a hybrid model. By limiting the idiosyncratic taste shock at a brand level, the model combines the choice probabilities of the PCM and BLP. Yet, because of the Type I extreme value distribution assumption the hybrid model still belongs to the probability choice system that is consistent with random utility maximization models.

The random coefficient captures heterogeneous consumer preferences in product characteristics. In this particular model, a value of α_i determines which product to choose within each brand. A consumer with a low value of α_i chooses a high quality product (almost equivalently, an expensive product) within brands and a consumer with a high value chooses a low quality product. In equation (4) consumers endowed with $\alpha \in \left[\overline{\triangle}_{j_k}\left(\delta_k, p_k\right), \underline{\triangle}_{j_k}\left(\delta_k, p_k\right)\right)$ choose product j of brand k for sure, given that they choose brand k.

However, their brand choice is stochastic because of the brand-level idiosyncratic shock, and is described by the integral part in equation (4) . There is a product in each brand that maximizes the (non-stochastic part of) utility of a consumer endowed with α_i and the brand-choice probability for this consumer is a function of these products. This means that even consumers who choose the same product of a certain brand can have different brand-choice probabilities as the utility-maximizing products of other brands can differ. In equation (4) this is determined by the max function in the denominator of the integral part. This also means consumers who choose the same product of a certain brand may switch to different products of another brand depending on their values of α.

Because consumer types determine which products are substitutes, they are similar in their product characteristics. This is a realistic description of many markets, including the PC market where consumers rarely switch to very low-end products when the price of a very high-end product goes up. Also, because the idiosyncratic shock is at a brand level, a product has fewer substitutes than in BLP but more substitutes than in the PCM.

Equation (4) also shows how the hybrid model distinguishes the within-brand quality (vertical) differentiation from the brand-level differentiation. With one random coefficient, as in this

[5] The probability that brand k is chosen by all consumers is $\int \frac{\exp\left(\max_{j_k \in k}\left(\delta_{j_k} - \alpha_i p_{j_k}\right)\right)}{\sum_{r=1}^{K} \exp\left(\max_{j_r \in r}\left(\delta_{j_r} - \alpha_i p_{j_r}\right)\right)} f\left(\alpha\right) d\alpha$

specification, prices determine the quality ranking within a brand, meaning that $\delta_{j_k} > \delta_{h_k}$ if $p_{j_k} > p_{h_k}$. The brand-level differentiation, on the other hand, is horizontal in the sense that all brands are substitutes for one another, but can still reflect brand-level quality differences. In other words, if one brand is "superior" to another, the former brand's low-end product can be a substitute for the latter brand's high-end product.

The hybrid model accommodates different types of the outside option. The "default" outside option is an outside option within each brand, which I call the within-brand outside option. This includes products that belong to brands covered by a data set but which are not included in the data set. For example, the PC data I use for an empirical application include six brands and cover about 65 percent of PCs sold to US households through retail stores. Thus, the within-brand outside option includes PCs sold at retail stores that belong to one of the six brands but are not included in the data set. The utility of this option is set to zero and, thus, consumers whose α_i is very high would choose this option.[6]

Another outside option includes brands that a data set does not cover. Dell products, for example, are sold through the internet, so they are not included in my data set. This outside option should be treated differently from the within-brand outside option because of the brand-level idiosyncratic shock. The value of this option, say δ_{Dell}, can be put into the model by adding $\exp(\delta_{Dell})$ to the denominator inside the integral in equation (4) such that

$$\exp(\delta_{Dell}) + \sum_{r=1}^{K} \exp\left(\max_{l_r \in r} (\delta_{l_r} - \alpha_i p_{l_r})\right)$$

This value can be estimated with data on Dell's overall market share which is readily available in industry reports.

Lastly, there are also consumers who consider buying PCs but buy none. I assume that these consumers also receive idiosyncratic shocks. Thus, the average value of buying no PC, say δ_0, can be put into the model by adding $\exp(\delta_0)$ to the above denominator as with the second outside option. The share of these consumers is determined by how the market size is defined. Details on

[6] The zero value assumption does not mean that there is no systematic quality difference across brands. See Appendix II for details.

how to estimate these values are provided in section 4.2.

2.3 Price Elasticity

The own-price elasticity of demand is

$$
\begin{aligned}
\frac{\partial s_{j_k}}{\partial p_{j_k}}\frac{p_{j_k}}{s_{j_k}} &= \frac{p_{j_k}}{s_{j_k}}\frac{\partial}{\partial p_{j_k}}\left(\int_{\alpha_i\in j_k}\frac{\exp\left(\delta_{j_k}-\alpha_i p_{j_k}\right)}{\sum_{r=1}^{K}\exp\left(\max_{l_r\in r}\left(\delta_{l_r}-\alpha_i p_{l_r}\right)\right)}f\left(\alpha\right)d\alpha\right)\\
&= \frac{p_{j_k}}{s_{j_k}}\left(\Psi\left(\overline{\triangle}_{j_k}\right)f\left(\overline{\triangle}_{j_k}|\boldsymbol{\theta}\right)\frac{\partial\overline{\triangle}_{j_k}}{\partial p_{j_k}}-\Psi\left(\underline{\triangle}_{j_k}\right)f\left(\underline{\triangle}_{j_k}|\boldsymbol{\theta}\right)\frac{\partial\underline{\triangle}_{j_k}}{\partial p_{j_k}}\right)\\
&\quad +\frac{p_{j_k}}{s_{j_k}}\left(\int_{\alpha_i\in j_k}-\alpha_i\Psi\left(\alpha_i\right)\left(1-\Psi\left(\alpha_i\right)\right)f\left(\alpha\right)d\alpha\right)
\end{aligned}
\tag{5}
$$

where

$$
\Psi\left(\alpha_i\right)=\frac{\exp\left(\delta_{j_k}-\alpha_i p_{j_k}\right)}{\sum_{r=1}^{K}\exp\left(\max_{l_r\in r}\left(\delta_{l_r}-\alpha_i p_{l_r}\right)\right)}.
$$

Equation (5) shows that the own-price elasticity equation consists of two parts. The term in the first bracket captures changes in market share due to substitution among products of the same brand, and the term in the second bracket captures changes due to substitution across brands. The former resembles the own elasticity in the PCM, but is a function of products of the same brand and is weighted by $\Psi\left(\overline{\triangle}_{j_k}\right)$ and $\Psi\left(\underline{\triangle}_{j_k}\right)$. The latter resembles the own elasticity in BLP, but only a subset of consumers is considered such that $-\alpha_i\Psi\left(\alpha_i\right)\left(1-\Psi\left(\alpha_i\right)\right)$ is integrated over consumers who buy product j of brand k.

There are two types of cross-price elasticity: within-brand cross elasticity and cross-brand elasticity. Within-brand cross elasticity measures the impact of price changes on products of the same brand, while cross-brand elasticity measures the impact of price changes on products of different brands. The former is defined as

$$
\begin{aligned}
\frac{\partial s_{j_k}}{\partial p_{h_k}}\frac{p_{h_k}}{s_{j_k}} &= \frac{p_{h_k}}{s_{j_k}}\frac{\partial}{\partial p_{h_k}}\left(\int_{\alpha_i\in j_k}\frac{\exp\left(\delta_{j_k}-\alpha_i p_{j_k}\right)}{\sum_{r=1}^{K}\exp\left(\max_{l_r\in r}\left(\delta_{l_r}-\alpha_i p_{l_r}\right)\right)}f\left(\alpha\right)d\alpha\right)\\
&= \frac{p_{h_k}}{s_{j_k}}\left(\Psi\left(\overline{\triangle}_{j_k}\right)f\left(\overline{\triangle}_{j_k}|\boldsymbol{\theta}\right)\frac{\partial\overline{\triangle}_{j_k}}{\partial p_{h_k}}-\Psi\left(\underline{\triangle}_{j_k}\right)f\left(\underline{\triangle}_{j_k}|\boldsymbol{\theta}\right)\frac{\partial\underline{\triangle}_{j_k}}{\partial p_{h_k}}\right),
\end{aligned}
\tag{6}
$$

where $\partial\overline{\triangle}_{j_k}/\partial p_{h_k}$ $\left(\text{or } \partial\underline{\triangle}_{j_k}/\partial p_{h_k}\right)$ is zero if $\overline{\triangle}_{j_k}$ $\left(\text{or } \underline{\triangle}_{j_k}\right)$ is not a function of p_{h_k}. The latter cross elasticity is defined as

$$\frac{\partial s_{j_k}}{\partial p_{h_m}}\frac{p_{h_m}}{s_{j_k}} = \frac{p_{h_m}}{s_{j_k}}\left(\int_{\alpha_i \in j_k} \frac{\alpha_i \exp\left(\delta_{j_k} - \alpha_i p_{j_k}\right) \exp\left(\delta_{h_m} - \alpha_i p_{h_m}\right)}{\left[\sum_{r=1}^{K} \exp\left(\max_{l_r \in r}\left(\delta_{l_r} - \alpha_i p_{l_r}\right)\right)\right]^2} f\left(\alpha\right) d\alpha\right) \tag{7}$$

if $\left(\delta_{h_m} - \alpha_i p_{h_m}\right) = \arg\max_{j_m}\left(\delta_{j_m} - \alpha_i p_{j_m}\right)$ for $\alpha_i \in j_k$. If not, $\partial s_{j_k}/\partial p_{h_m} = 0$.

Equations (4) $\sim$ (7) show that when the price of a given product changes it affects products of the same brand in its adjacent neighborhood as well as the products of different brands of a similar quality. This is an important difference from other discrete choice models. For example, in the PCM with one random coefficient a price change only affects products in the adjacent neighborhood of a given product. In any sort of logit demand model, it affects all other products in the market. The functional form of these equations also shows that the degree of the within-brand substitution is likely larger than that of the cross-brand substitution. This means that products of the same brand are likely much closer substitutes than products of different brands.

As will be shown in section 4 and Appendix III, despite these differences in the substitution pattern, estimated markup in the hybrid model is not much lower than that of BLP because the degree of the cross-brand substitution is not drastically different between the two models. However, the markup is more responsive to the market size in the hybrid model such that the markup level goes down to a realistic level when the market is defined to include potential consumers. A larger market size does not have much impact on the second part of equation (5) but increases the magnitude of the first part through a smaller s_{j_k}, and thus increases the own-price elasticity. An intuitive explanation for the larger impact on the first part is as follows. When the market size becomes larger by including more potential consumers, it does not affect how many of the existing consumers switch to the same-brand products from a price increase, but this same change is now larger relative to s_{j_k}, resulting in a larger percentage change.

2.4 Consumer Welfare

Consumer welfare of a consumer with a value of α_i, converted into a monetary unit, is given as

$$\frac{\ln\left(\sum_{r=1}^{K}\exp\left(\max_{l_r\in r}\left(\delta_{l_r}-\alpha_i p_{l_r}\right)\right)\right)}{\alpha_i}. \tag{8}$$

The summation in the numerator is over brands, not over all products in the market, and there is a maximum function inside the exponential function. The model links the consumer to a product in each brand that maximizes her utility, and computes her welfare over products of her choice in each brand. This implies that when firms introduce new products that are better than the existing products, their welfare contribution is limited to a subset of consumers whose willingness to pay for quality improvement is high. The equation also shows that the overall welfare grows more slowly than that of models with the product-level taste term. The average consumer welfare is given as

$$\int \frac{\ln\left(\sum_{r=1}^{K}\exp\left(\max_{l_r\in r}\left(\delta_{l_r}-\alpha_i p_{l_r}\right)\right)\right)}{\alpha_i} dF\left(\alpha\right) \tag{9}$$

Note that, as the value of α_i changes, a product choice within brands also changes by the maximum function.

3 Identification and Estimation

An identification strategy is similar to the other demand models (BLP and the PCM) in the sense that the preferences for characteristics are identified from changes in market shares associated with changes in product characteristics. Roughly speaking, when product j's characteristic (exogenously) changes, how much market share it gains or loses and what products it gains market shares from or loses to identify the preference for that characteristic. Note, however, that there is an important difference between the PCM and BLP in the way that they identify preferences. When product j's desirable characteristic improves in the PCM, it steals market shares only from products with similar characteristics so the preference is identified by which products it steals market shares from and how much market shares it steals from them. In BLP, on the other hand, product j draws market shares from all products so only how much market shares it draws from each product is used in identifying the preference.

The hybrid model is similar to the PCM when it exploits within-brand substitution patterns. When product j's characteristic improves, which products of the same brand product j draws market shares from and how much market shares it draws from them identify the preference. Because consumers are not exposed to the idiosyncratic shock in choosing a product within brands, their within-brand product choices are solely driven by their intrinsic types, or willingness to pay.

Unlike the PCM, however, substitution patterns across brands are also important in identifying preferences. Because of the idiosyncratic shock at a brand level, product j draws market shares from all brands but only from a subset of products in each brand. Thus, which products in each brand it draws markets shares from (as well as how much) identifies preferences, especially the preference for characteristics relative to the preference for brands. Suppose Compaq's high-end PC improves its CPU speed. If it draws market shares from the low end of Sony's product line as opposed to the high end, it would imply that consumers prefer Sony's PCs to Compaq's for the same CPU speed or that consumers care less about the CPU speed than brand-level unobserved quality. This would be translated into a larger brand dummy estimate for Sony or a smaller estimate for the CPU speed; the more consumers care about the CPU speed, the less they care about which brand provides it.

Note that the distribution of consumer types or the random coefficient is identified only relative to the distribution of the brand-level shock. For example, if product j draws more market shares from other brands, it implies that brands are more easily substitutable and the brand-level shock is less important than consumer types in explaining choices. Thus, the degree of heterogeneity in consumer types would be estimated to be larger.

As in the other demand models, price is treated as endogenous and instruments are used to consistently estimate the model.[7] I use a non-linear GMM estimation procedure because the parameters for the distribution of α enter the model in a nonlinear fashion. Given values of the

[7] See Section 4.3 for choices of instruments and identifying assumptions for the PC market.

nonlinear parameters, I search for a vector of product quality, $\boldsymbol{\delta}$, such that the model-predicted market shares equate observed market shares. Appendix II provides computational details of finding $\boldsymbol{\delta}$ and shows that there exists unique $\boldsymbol{\delta}$ in the one random coefficient hybrid model. Given a vector of product quality I compute a GMM objective function using the moment condition that the mean of unobserved demand components is uncorrelated with instrumental variables. By repeating this for different sets of parameter values, I find a set of parameter values that minimizes the GMM objective function. The coefficients of product characteristics are estimated as linear parameters as part of this estimation procedure.

The GMM estimator in the one random coefficient hybrid model belongs to the class of estimators that Berry, Linton, and Pakes (2004) consider. The share function in equation (4) satisfies their regularity conditions, and the share function simulator is the same as theirs. They show that the GMM estimates of BLP and the PCM converge to true values as the number of products increases, provided that the number of simulation draws grows large at a fast enough rate, and that the PCM requires much fewer simulation draws than BLP because the substitution pattern in the former model is more local. By the same logic the hybrid model should not require more simulation draws than BLP. Note that in the hybrid model the number of products can increase either by adding products per brand or by adding more brands. The number of products should increase in both ways because of the way that preferences are identified.

The number of products can also increase by adding more independent markets, either by having more regional markets or by having data on more time periods, and this also makes estimates converge to true values in the same way as in the other demand models. In practice, having a larger number of markets is especially helpful in identification because it provides more variations in product characteristics offered by each brand. In some markets brands offer different sets of products in different regions, and when a market is defined by a time period as in the PC market, researchers can observe a wide range of product characteristics such as a wide range of CPU speed by having a longer time period data set.

In Appendix III I use Monte Carlo simulations to show how the hybrid model is different from BLP and the PCM.[8] Rather than choosing the hybrid model as a true model, I use all of the three models as a true model one at a time and estimate the other two models. All models have the common feature that multiple brands (firms) compete with multiple products and prices are set according to their within-brand quality ranking. Depending on which model is chosen as a true model, the idiosyncratic error structure differs and, hence, simulated market shares differ as well.

Simulation results show that the hybrid model is more robust to the (mis-)selection of models. Comparing the hybrid model and BLP, the mean square errors of the hybrid estimates (when BLP is treated as a true model) are about one third of those of the BLP estimates (when the hybrid model is a true model). The BLP estimates are close to true values on average but their variances are much larger, resulting in larger MSEs. However, both the hybrid model and BLP perform equally poorly when the PCM is a true model and the PCM performs poorly when either of the two models is a true model.

The results on the markup and consumer welfare show that the model structure determines their magnitudes while the choice of a true model only has marginal impact. For example, the PCM's markup is drastically lower than those of the other two no matter which model is chosen as a true model. They also show that, despite much higher own-price elasticity, the hybrid model's markup is not substantially different from that of BLP. This is because (1) closest substitutes belong to the same brand and (2) the brand-level taste renders the degree of the cross-brand substitution not much different from that of BLP. Yet, the welfare effects of new products introduced by existing brands are similar to those of the PCM because the welfare contribution of the taste term is limited to the brand level.

4 Demand for Personal Computers

4.1 Data and Industry

Scanner (barcode-level) data on desktop PCs are used to estimate household demand for desktop PCs. The data were collected from US retail stores by NPD Techworld, a private consulting

[8] Among the existing demand models, only BLP and the PCM allow substitution patterns to be determined by product characteristics.

firm specializing in information technology. Thus, my definition of product refers to an actual PC unit consumers buy at retail stores. I exclude laptops from the data set and estimate the one random coefficient hybrid model.[9] The data provide information on revenue, quantity sold, and characteristics at a product level. I calculate the price by dividing revenue by quantity sold. Observations are on a monthly basis and the sample period starts in October 2001 and ends in September 2004. Product characteristics include CPU speed, CPU type (Intel, AMD, or Apple), memory capacity, hard drive capacity, the size of the level 2 cache, the screen size, whether the screen is LCD, whether the DVD reader is included, and whether the DVD writer is included. I select observations with more than 1,000 units of sale, which cover about 90% of the total sale. My sample consists of 534 unique products and 1,875 observations for 36 months.[10]

There are six brands in the sample: Sony, HP, Gateway, Compaq, eMachines, and Apple. Dell is not included as it sells only through the internet. However, I include Dell's presence in the hybrid model and estimate the mean utility of buying Dell PCs using its aggregate market share. Details are provided in the next section. In the sample HP has the highest number of products, 20.3 products per month on average, followed by Compaq with 13.1 products and eMachines with 9.3 products per month on average. Sony has 6.7 products and Apple has 2.5 products per month on average. Gateway products appear after February 2004 and only 1 product has more than 1,000 units of sale per month. The number of products is closely correlated with the market share. HP has the largest market share (39 percent on average), followed by eMachines (25 percent) and Compaq (24 percent). Sony and Apple have much smaller market shares, 10 percent and just over 1 percent respectively. Apple computers are the most expensive ones: $1,400 on average. Sony also sells expensive products with the average price around $1,200. eMachines sells the cheapest computers with the average price less than $600. Compaq computers are about $150 more expensive than eMachines, while HP computers are about $300 more.

Summary statistics are reported in Tables 1, IV-1, and IV-2. For presentation I first averaged all variables using sales as weights for each month and then simply averaged them at a quarterly level. The tables show the well known trend in the PC industry that product attributes

[9] If laptops were included, I could put another random coefficient on the dummy variable for portability.

[10] See Appendix IV for details of the data.

improve over time without a significant price increase. During the sample period the average CPU speed increased from 1.2 GHz to 2.48 GHz, the average memory capacity from 232.78 MB to 445.36 MB, the average hard drive capacity from 42.29 GB to 113.17 GB, the proportion of computers with a DVD reader from 0.51 to 0.93, and the proportion of computers with a DVD writer from 0.02 to 0.44. Despite these quality improvements, the average price decreased from $826 to $691. The average and total sales follow a seasonal cycle with the fourth quarter of each year showing higher demand than the other three quarters.

Figure 1 shows a brand-level market share during the sample period.[11] I do not have product-level data on Dell, but used an industry report to recover its brand-level share. The figure shows that HP's market share, after accounting for Dell, decreased from 27 percent to 16 percent over time, while Compaq's share increased from 12 percent to 16 percent. Note that HP and Compaq merged in February of 2001. The merged firm's total share sharply decreased from 40 percent to 30 percent in six months after the merger but stabilized at around 30 percent afterwards. Dell and eMachines both gained market share during the sample period, while Sony lost its share. Apple's market share rarely exceeded one percent for the entire sample period. At the end of the sample period, HP, Compaq and eMachines had very similar market shares. All brands' total sales did not change much after the first quarter of 2002, which suggests that most of the consumers that HP lost switched to other brands like Compaq, eMachines and Dell.

Figure 2 shows a brand-level average price during the sample period. I take the simple average across products instead of using sales as a weight. All brands do not follow the same price trend. HP's average price fluctuated between $800 and $1,000 and Apple's average price fluctuated around $1,500. eMachines' average price did not change significantly. But Compaq's average price went down from $900 to $600 and Sony's average price went up from below $1,200 to $1,400. Most interestingly, the price gap between HP and Compaq became wider over time. It was less than $100 right before the merger, but reached about $300 by the end of the sample period.

These figures suggest that the merged firm did not necessarily benefit from the merger; it did not gain market share, nor raised prices. Compaq's increased market share seems to be a result

[11] The brand-level market share here does not include the within-brand outside option. I simply divided the total sales of products of each brand by the total sales of all products in the sample.

of the decreasing price, but the impact of the merger on the firm's profit is unclear without having markup estimates. The source of Compaq's price decline is unclear as well. It could be due to a decreasing marginal cost or product repositioning. More discussion on the post-merger PC market is provided in section 5.

4.2 Market Size and Outside Options

I consider two types of the "market" in my estimation. One is the US household desktop computer market that only includes active desktop PC consumers and the other is a potential household computer market that includes all potential consumers. I set the size of the former market based on the number of computers sold to household consumers. According to IDC, a research company specializing in information technology, 17.1 million computers were sold to household consumers in 2002 and 20 million in 2003. Based on these numbers I set the monthly market size, accounting for seasonality. The mean size per month is about 1.4 million. For the latter type of the market I choose three sizes: three million, five million and ten million.

I consider three outside options based on these two market definitions. The first is the within-brand outside option which is to buy products that belong to the six brands but do not appear in the sample. I assume that consumers receive the lowest utility from these products and set their values to zero.[12] Thus, consumers with high values of α are likely to choose these products. As shown in Appendix II, one does not need data on the market share of this outside option to estimate the model. The second outside option is to buy Dell products and its share is calculated from the monthly sales of Dell computers reported by industry reports. The third option is to make no purchase and its share is calculated from the difference between the size of the potential household computer market and the size of the household desktop computer market.

With these three outside options the denominator inside the integral in equation (4) becomes

[12] It is possible that this outside option can include high-end products that are dropped from the sample either because their market shares are tiny or because they are dropped in a data collection process by mistake.

$$\exp\left(\delta_{Dell}\right) + \exp\left(\delta_0\right) + \sum_{r=1}^{K} \exp\left(\max_{l_r \in r}\left(\delta_{l_r} - \alpha_i p_{l_r}\right)\right), \tag{10}$$

where δ_{Dell} and δ_0 are the average value of choosing Dell products and the average value of making no purchase respectively and are estimated.[13] Note that both BLP and the PCM have only one outside option. In BLP the outside option enters in the denominator of its share function and its value is set to zero. In the PCM the outside option is similar to the within-brand outside option described above and its value is also set to zero.

4.3 Demand Estimates

Regardless of which model I estimate, I assume that product characteristics except prices are not correlated with any products' unobserved demand factors which include unobserved characteristics and demand shocks. In general, an endogeneity problem can arise with non-price characteristics in two cases. First, firms may adjust product characteristics every period after observing demand shocks. Second, there may be important characteristics left in the unobserved term that are correlated with some of the observed characteristics. The first case is unlikely to arise in my application. I use monthly observations, so for this case to be relevant, PC manufacturers would have to adjust product characteristics monthly. However, I do not observe this behavior in the data other than in prices. It is true that product characteristics rapidly change over time in the PC industry, as will be shown in section 5, but the main driver behind this change is technological advances, which are unlikely to be correlated with contemporaneous demand shocks.

The second case is more relevant to my application. As mentioned above, firms frequently improve product characteristics and when they do so, they tend to change multiple characteristics at the same time. A larger memory capacity, for example, usually accompanies a faster CPU. Thus, it is important to control for all characteristics that change together (without causing multicollinearity) to consistently estimate PC demand. I include the ten characteristics listed in section 4.1 and also include the brand dummy and the month dummy variables. Given the range and

[13] The third type of the outside option represents consumers who actively search for personal computers but do not buy any products. This group may play an important role in determining a price trend over time. However, treating this dynamic aspect is beyond the scope of this paper.

the type of the characteristics included, the unobserved term is unlikely to contain other major characteristics.

However, the price is still an endogenous variable that is affected by unobserved demand shifters. In the IO literature (functions of) other products' observed characteristics are often used as instruments under the exogeneity assumption stated above. Usual instruments used in the literature are the sum of characteristics of the same firm's products except a given product and the sum of characteristics of other firms' products, which I call BLP instruments. The standard argument for these characteristics being correlated with prices is that in differentiated product markets close substitutes have similar characteristics and products that compete with more close substitutes tend to have low markups.

For my application I use the sum of the BLP instruments interacted with the time dummy variables as instruments. The characteristics I use are the CPU speed, the memory capacity, and the screen size. In addition to using these characteristics to capture the degree of competition products face, I also use them as proxies for production costs; in a given market (period) PCs with better characteristics are more costly to produce. I interact them with the time dummies because the cost of producing PCs goes down rapidly due to technological advances. PCs with the same characteristics are sold at substantially different prices depending on which period they were sold in, and the time-dummy interaction allows me to treat the same value of characteristics differently for each time period.[14]

I first estimate five versions of the logit model and the PCM, and report estimates in Table 2. The five logit models are the OLS logit (under *Logit*), the IV logit (*IV Logit*), the nested logit (*Nested*) and two versions of BLP (*BLP I* and *BLP II*). The month dummy variables are included in all models but are not reported. Demand estimates with other instrumental variables are reported in Table V-1 in Appendix V. The change in the price coefficient from the OLS logit model to the IV logit model shows that the price variable is positively correlated with unobserved demand components and that the instrumental variables mitigate this problem. However, only two product characteristics (CPU speed and DVD writer) have statistically significant coefficients in the

[14] I tried both the month dummies and the quarter dummies to interact with the BLP instruments and had similar results.

IV logit model. The coefficient on the smaller cache is positive and statistically significant in the OLS logit model but becomes statistically insignificant in the IV logit, although it is still positive. The positive coefficient of the smaller cache misleadingly suggests that consumers prefer products with a smaller cache size controlling for other characteristics and price. Nevertheless, notice that all estimates change in the "right" direction as the price coefficient becomes more negative. The coefficients on desirable characteristics such as memory, hard drive, screen size, etc. become more positive and those on undesirable ones such as smaller cache move towards zero.

In the nested logit model I group products based on the brands they belong to, so Sony products forms one nest, HP products forms another nest, and so on. The within-group correlation estimate (reported at the bottom of the table) is 0.84 and statistically significant, showing that the same brand PCs are much closer substitutes for one another than for other brands' PCs. However, this high correlation implies that a very high-end PC is a much closer substitute for a very low-end PC of the same brand than for a very high-end PC of another brand. Nevertheless, all estimates except three are statistically significant and have expected signs. The smaller cache variable is one of the three variables with an insignificant estimate but it has a negative sign as expected. Note that although the magnitude of the estimates is smaller than in *IV Logit*, the average consumer's willingness to pay for marginal improvement in the characteristics is not much different due to a smaller price coefficient.

The two versions of BLP only differ in terms of an assumption on the distribution of the random coefficient. In *BLP I* I assume that the price coefficient is distributed normal and in *BLP II* I assume it is distributed log normal. The table shows that *BLP I* is not statistically different from *IV Logit*. The variance parameter of the price coefficient is not statistically significant, while the mean of the price coefficient is no longer significant. All other estimates are not statistically different from those of *IV Logit*. CPU speed and DVD writer are still the only characteristics with significant coefficients. In *BLP II* the price coefficient is not statistically significant, and most of other estimates are not significant either. Tables V-2 and V-3 in Appendix V report BLP estimates in different specifications, and those estimates are not very different from those in Table 2. The insignificant estimates may not be surprising since I do not use any micro-level data for consumer heterogeneity. The literature shows that using such data helps estimating BLP-type models (BLP; Nevo, 2001; Petrin, 2002, etc.).

The last column reports demand estimates in the PCM. I use the US household desktop computer market excluding Dell products as the market for the PCM because estimates become too small as the market size approaches 2 million.[15] All coefficients on product characteristics are significant at a 5% significance level. The coefficient on the smaller cache variable has a negative sign, correctly suggesting that consumers prefer products with a larger cache size. Intel and AMD CPU have negative coefficients, suggesting that consumers prefer Apple computer CPUs.

However, consumers' willingness to pay for improved characteristics is unrealistically low in the PCM. For example, the average consumer is willing to pay only \$30 more for a PC with a 1 GHz faster CPU. The top 10% consumers are willing to pay \$320 more, implying a large degree of consumer heterogeneity. In *BLP I* the average consumer is willing to pay \$318 more, while the top 10% are only willing to pay \$466 more for the same improvement.

While none of the brand dummy variables have statistically significant coefficients in *IV logit* and *BLP I*, they are significant in the nested logit model and the PCM. Note that the brand dummies capture brand-level quality differences that are not explained by the observed characteristics. In the nested logit model HP products have the highest unobserved brand quality (among the brands that use Intel and AMD CPUs), followed by eMachines, Compaq, and Sony products, although the eMachines and the Compaq dummies are not statistically different.[16] On the other hand, in the PCM Sony products have the highest unobserved brand quality followed by HP, Compaq, and eMachines products. I discuss this difference in more detail below.

Table 3 reports demand estimates in the hybrid model. I assume that the price coefficient is distributed log normal with the log mean set to zero.[17] In the first column the market is defined as the US household desktop computer market so the third outside option is excluded. In the second to the fourth columns the market includes potential consumers, and its size is set to three

[15] In the logit model a larger market size does not change estimates substantially except for the constant term. In the PCM magnitudes of estimates decrease as the market size increases.

[16] Since I use the CPU type dummy variables and Apple used a different type from the others during the sample period, the Apple effect is not distinguishable from its CPU effect.

[17] The mean parameter is not identifiable as it changes the mean product quality of all products proportionally. The PCM requires the same normalization. See Song (2007).

million, five million and ten million respectively.

All coefficients except the coefficient of the memory variable are significant at a 5% level, and the signs of the coefficients are as expected. All desirable characteristics like CPU speed, hard drive, screen size, etc. have positive coefficients. Consumers prefer the DVD writer as much as a 1 GHz faster CPU. The small cache, the Intel CPU, and the AMD CPU variables have negative coefficients as in the PCM. The average consumer is willing to pay $130 more for a 1 GHz faster CPU. The top 10% are willing to pay $220 more.

The estimates for the brand dummies are all statistically significant and imply that HP products have the highest unobserved brand quality followed by Compaq, eMachines, and Sony, although differences among the last three are not statistically significant. Note that this ranking is different from that of the PCM but similar to that of the nested logit model. This difference stems from the way that the two models determine product quality. Recall that in the PCM prices determine the quality ranking, while both prices and market shares are important in the hybrid model. In the PCM the Sony dummy estimate takes the largest value because its products are relatively more expensive and the observed product characteristics do not fully explain price differences across the brands. On the other hand, in the hybrid model the HP dummy estimate takes the largest value because its market share is the largest and the observed characteristics do not fully explain its popularity in the market. The nested logit is similar in this sense as it relies on the group-level market share in identifying the brand dummies.

Notice that the random coefficient is statistically significant in both the hybrid model and the PCM. Their main difference from BLP is that both models link the value of the random coefficient to product choices deterministically such that consumers endowed with low values are certain to choose high price products and those with high values choose low price products. In BLP consumers endowed with low values are more likely to choose high price products but may also choose low price products at some probability. The estimation results show that this deterministic relationship between consumer types and product choices helps estimating the shape of the distribution more precisely.

The difference between the hybrid model and the PCM lies in how they rank products.

With one random coefficient the PCM ranks all products on a single quality dimension, while the hybrid model ranks products of each brand separately. The results show that the hybrid model produces economically more reasonable estimates by doing so.

The second to the fourth columns of the table show that the coefficients hardly change as the market size increases. This is because the preferences for product characteristics are identified from characteristics variations within the desktop computer market and the effect of having more potential consumers is absorbed by δ_0. This is similar to the logit model case where a larger market size makes the constant term larger without much changing the magnitude of the coefficients. However, the markup goes down as the market size increases as shown below.

4.4 Price Elasticity and Markup

The left panel of Table 4 compares the (product-level) price elasticity among the four models. Although the random coefficient is not statistically significant in *BLP I*, I use its estimate in all of the following analyses. The own-price elasticity is the largest in the PCM where a market share changes by a factor of 34 from one percent price change. It is much lower in the hybrid model where a market share changes by a factor of 11. It is still a big change compared to the nested logit model or BLP where a share changes by less than 5 percent.

Notice that the own price elasticity of the nested logit model is about 2.5 times larger than that of BLP. Although the price coefficient is much smaller in the former model, its high within-group correlation results in a much higher price elasticity. However, in the nested logit model high-end PCs are close substitutes for low-end PC as long as they belong to the same brand. For example, the demand estimates of the nested logit model imply that a 1 percent price increase of HP's high-end PC that was sold at $1,700 in September 2004 would increase the market share of its low-end PC sold at $600 by 0.08 percent while changing the market share of Sony's high-end PC (sold at $1,670) by less than 0.001 percent.

The hybrid model's high own-price elasticity is also driven by high substitutability among products of the same brand, but these within-brand substitutes are only a small subset of the same brand products and they are similar in their characteristics, usually having exactly the same characteristics except one or two. This difference leads to drastically different markup estimates between the two models. In the nested logit model firms have substantial market power over their brands and this results in unrealistically high markup estimates. In the hybrid model, on the other hand, markup estimates are much lower because only a small subset of products are substitutes within brands but they are not unrealistically low at the same time despite the high magnitude of the price elasticity because these close substitutes belong to the same brand. This is shown below in more detail.

The table also shows that only about 10% of products are substitutes across brands in the hybrid model and that the average cross-brand elasticity is 0.035, which is much smaller than the within-brand elasticity and more comparable to the cross elasticity in the logit models. To gain a better sense of the brand-level substitutability, I compare the brand-level elasticity on the right panel of Table 4 where it measures a percent change in a brand's share when the prices of all products of that brand change by 1 percent. The table shows that the own brand-level elasticity is similar between the hybrid model and BLP. In *Hybrid I*, the average own-brand price elasticity is -2.355 and the average cross-brand elasticity is 0.066, while they are -2.977 and 0.005 respectively in *BLP I*. This similarity suggests that the two models may not have drastically different markups.

Table 5 lists the average product-level percentage markup over time. Note first that in both the nested logit model and *BLP II* the average percentage markup is higher than 1 for all periods, which implies that the average marginal cost is negative. In *BLP II* this happens because the price coefficient estimate is not large enough. In the nested logit model it happens because products of the same brand belong to the same group and all products in the same group are very close substitutes. Although this feature results in more reasonable demand estimates and much higher price elasticities compared to BLP, it misleadingly implies that firms have substantial market power over their brand names and results in unrealistic markup estimates. The percentage markup estimate is more reasonable in *BLP I*, although it is over 0.5 for all periods.[18] Changing a market size does not change the magnitude. Other specifications in Table V-2 do not change the markup estimate significantly either. The average markup in the PCM is no higher than 0.005 in

[18] *IV Logit* has slightly higher but similar markup estimates as *BLP I*. One needs a larger estimate of the price coefficient to have lower markup in these models.

all periods, which is not surprising considering its high price elasticities and substitution pattern.[19]

The last three columns report the average markup for the hybrid model with different market sizes. Notice first that the markup in *Hybrid I* is generally lower than that of *BLP I* but not substantially different, consistent with what the brand-level elasticity implies and also with the results of the Monte Carlo exercise in Appendix III. Notice also that the average markup decreases as the market size increases as predicted in section 2.3. It ranges $0.31 \sim 0.70$ in *Hybrid I.* When potential consumers are included in *Hybrid II*, it goes down to $0.15 \sim 0.42$. When more people are included as potential consumers in *Hybrid III*, it goes down to $0.09 \sim 0.30$.

Table 6 lists the brand-level average markup, both in percentage and absolute terms, for *BLP I*, *PCM*, *Hybrid I* and *Hybrid II.* Brands include Apple, Compaq, eMachines, HP, and Sony.[20] In *BLP I* all brands have similar absolute markups and, thus, a brand with more (less) expensive products has a lower (higher) percentage markup. This pattern always appears in the logit model where the absolute markup is an inverse function of $\alpha\left(1-s_j\right)$. Because $1-s_j$ is usually close to 1 and not so much different among products, price dictates percentage markup such that a higher price results in a lower percentage markup. When α is the random coefficient as in *BLP I*, the absolute markup becomes higher (lower) for brands with more (less) expensive products, but the percentage markup still tends to be in the reverse order of the average price. The opposite pattern appears in the PCM where a brand with more expensive products tends to have higher percentage and absolute markups. Apple has the highest markup in both the percentage and the absolute terms, while the other brands have substantially lower markups. Nevertheless, the markup level is unrealistically low for all brands.

In the hybrid model both market shares and prices determine the brand-level absolute markup. If two brands have similar market shares, the (absolute) markup is higher for the one with more expensive products. If two brands are similar in prices, it is higher for the one with a larger share. Compare Sony and Compaq. Although Sony's average price is higher, both brands have similar markups because Compaq's market share is much higher. Consider eMachines next.

[19] Value Line reports that the before-tax margin of the Computers/Peripherals sector is 16.94% as of January 2013. See http://pages.stern.nyu.edu/~%20adamodar/New_Home_Page/datafile/margin.html

[20] The data set also includes Gateway computers and I use them in demand estimation. However, I exclude Gateway from discussion since its products only appear in March 2004.

Although its average market share is slightly higher than Compaq, its absolute markup is lower than Sony because its price is much lower. On the other hand, despite the highest average price Apple's absolute markup is the lowest because of its tiny share.[21] In the case of HP, its large market share combined with the mid-range price level results in the highest markup.

Unlike the brand-level markup, prices solely determine the ordering of the within-brand markup in the hybrid model, meaning that markup is higher for more expensive products within a brand. This never holds true in the IV logit model or in the nested logit model where products of the same brand, i.e., the same ownership, have the same absolute markup. BLP has this ordering with the random coefficient on the price variable, or on any variables that differentiate products vertically, but with the current variance estimate the markup difference is too small to be realistic. For example, in August of 2004 the absolute markup of the least expensive Compaq product is \$395 while that of the most expensive one is \$449 in *BLP I*. In the hybrid model the markup for the same products is \$208 and \$363 respectively.

The magnitude of the marginal cost implied by the demand estimates provides additional evidence that the hybrid model describes the PC market better than any other model. In *Hybrid II* the average marginal cost ranges from \$483 for eMachines to \$1,356 for Apple. In the PCM it is no different from prices because of small markups. In BLP, because markup is similar regardless of price differences, the implied marginal cost is unreasonably low for low-end brands such as eMachines (\$160) and Compaq (\$300) and also for low-end products of all brands.

In all three models HP earns the largest variable profit due to its large market share. The profit ranking among the other brands is similar between *BLP I* and *Hybrid I* and *Hybrid II*, and is mainly determined by the market share. Although Compaq and eMachines have similar markups and market shares, Compaq earns a higher profit because its price is higher. In the PCM Apple earns the second highest profit, followed by Sony, eMachines, and Compaq, but their profit levels are, again, unrealistically low.

[21] I do not model Apple's unique position in the market in this paper. Apple's operating system is not compatible with that of the other brands and its customers are considered more loyal. Treating it at the same level as the other brands may not be realistic and could be a reason for its small markup.

4.5 Consumer Welfare

Figure 3 compares the percentage change in consumer welfare. All three models are re-estimated using the same market size as in the PCM to avoid differences due to market size. With the smaller market size the constant term goes down to -1.67 in BLP. The mean of the price coefficient is -2.26 and the standard deviation is 0.21, but they are not statistically significant. CPU speed and DVD writer are still the only variables with significant coefficients and their magnitudes are almost the same as in Table 2. The estimates in the hybrid model do not change much either and are almost same as those in Table 3. All estimates are still significant at a 5% level except for the coefficient of the memory variable.

Despite all the differences in the structure of the models and demand estimates, the figure shows similar trends of welfare changes. The direction of changes is the same in 20 periods out of 35. In most periods the magnitude of changes is similar. However, there are a few interesting differences. First, consumer welfare seems to fluctuate more in BLP than in the other two models. The standard deviation is 0.12 in BLP, while it is 0.086 in PCM and 0.074 in the hybrid model. This is mainly due to the product-level taste shock in BLP that is added with a new product and goes away with a product exit. As a result, consumer welfare responds more sensitively to product entry and exit than in the other two models.

Secondly, BLP and the PCM are more similar to each other than to the hybrid model with respect to the direction of changes. The hybrid model has a different sign from the other two in 10 periods. The correlation coefficient is higher between BLP and PCM (0.88) than between the hybrid model and either of the two (0.52 with PCM and 0.51 with BLP). This may be because the hybrid model treats brand entry/exit differently from product entry and exit. The introduction of a new brand introduces a taste shock but the introduction of a new product does not. For example, in March 2004 consumer welfare decreased in both BLP and the PCM, but increased in the hybrid model. This is a period when another brand, Gateway, is added in the data but the total number of products decreases.

Lastly, consumer welfare is more sensitive to new product introduction in BLP and the PCM than in the hybrid model. This is surprising considering the results in the Monte Carlo simulations in Appendix III. The correlation coefficient between a percentage change in consumer welfare and the number of new products is 0.43 in BLP and 0.33 in the PCM, while it is 0.16 in the hybrid model. A high correlation in BLP is due to the taste shock as explained above. In the PCM it is because a value of the price coefficient (α) for a new product is very small as it is positioned at the far left end of the price coefficient distribution. Since welfare changes are multiplied by the inverse of the price coefficient, its value increases rapidly with the number of new products.

5 Product Repositioning in the Post-Merger Market

In this section I show how demand estimates can be used to analyze product repositioning in the post-merger market. The merger between HP and Compaq took place in February 2002, which is the fifth period of my 36-period sample. As explained in section 4.1, the price gap between HP and Compaq was less than \$100 right before the merger, but reached about \$300 by the end of the sample period mainly due to Compaq's price decline. As a matter of fact, Compaq is the only brand that lowered its price substantially in the post-merger period.

This price trend is not consistent with the usual prediction that the merged firm will increase prices, at least relative to the prices of non-merged firms, by exploiting its market power. According to this logic, a merger can only be justified if the supply-side efficiency gain is larger than a loss in consumer welfare. This motivates a post-merger simulation, such as Nevo (2000), that computes a consumer welfare loss from a hypothetical merger to estimate the supply-side efficiency gain that would justify the merger.

It may be the case that only Compaq enjoyed the efficiency gain from the merger. Although the trend and the level of the markup differ across the demand models, all models show that Compaq's (estimated) cost went down after the merger. This is not surprising considering the time trend of its price and market share.[22] Compaq's decreasing cost could have resulted from the efficiency gain, but it is also possible that it is associated with changes in product characteristics. That is, the merged firm may have repositioned Compaq by including disproportionately more

[22] The marginal cost is estimated by the difference between (observed) price and the inverse of (estimated) price elasticity. Thus, when price changes without much change in market share as in Compaq's case, any demand model will predict that the marginal cost moves in the same direction as price.

low-end products, and Compaq's cost went down as a result.

My data on product characteristics support this possibility. Compaq did not upgrade the memory and the hard disk capacities as fast as HP, and put the DVD writer in much fewer products (less than 40 percent) than HP (more than 60 percent). Note that the DVD writer was rarely installed in products of both brands before the merger. It also increased a portion of products with lower-end CPUs (Celeron-type CPUs) to over 60 percent, and decreased a portion of products with LCD displays from the beginning of 2004.[23] These multidimensional changes can be summarized into a single dimensional consumer utility using the demand estimates. The sign of each coefficient indicates whether consumers prefer having more of the corresponding characteristic, and its magnitude shows their willingness to pay for having more of it.

As shown in section 4.3, the hybrid model provides the most sensible estimates of the coefficients. All except one are statistically significant, and their signs and economic implications are most reasonable. If the BLP estimates were used, I would not be able to use other characteristics other than the CPU speed and the DVD writer. The estimates in the PCM are statistically significant, but their economic implications are unrealistic.

Thus, I use the demand estimates in the hybrid model to construct a brand-level index that translates changes in characteristics into the consumer utility.[24] I construct this index, which I call the quality index, in the following way. I first compute the average value of all observed characteristics for each brand. They include CPU speed, CPU type, memory capacity, hard drive capacity, the size of the level 2 cache, the screen size, whether the screen is LCD, whether the DVD reader is included, and whether the DVD writer is included. I use the simple average, instead of the sale weighted average, in order to reflect product characteristics that the brands offer, not characteristics that consumers choose. Then for brand j, I construct the index by

$$q_j = \sum_{k=1}^{L} \widehat{\beta}_k \overline{x}_{jk}, \tag{11}$$

[23] Around the time of the merger, both brands had about 30 percent of their products with lower-end CPUs.

[24] An underlying assumption is that the merger affects consumer preferences only through changes in product characteristics and prices.

where $\overline{x}_{jk}$ is brand j's average value of characteristic k and $\widehat{\beta}_k$ is its estimated marginal utility. If $\widehat{\beta}_k$ is not statistically significant, it is replaced with zero.

Figure 4 shows the quality index for five brands during the sample period. Apple's index is highest, followed by Sony and HP. The gap between Apple and Sony becomes narrower towards the end of the sample period. eMachines' index is the lowest. Most importantly, the figure shows that the quality indices of Compaq and HP diverge over time. HP's index keeps up with those of Sony and Apple, while Compaq's index becomes closer and almost identical to that of eMachines. This suggests that the merged firm gradually repositioned its two brands after the merger and that Compaq's decreased marginal cost is an outcome of this product repositioning.

I formally test if there are breaks in the trends using the Quandt Likelihood Ratio (QLR) test with 15% trimming. The regression is

$$y_{kt} = \alpha_k + \beta_k y_{kt-1} + \gamma_0 D_t\left(\tau_r\right) + \gamma_1\left[D_t\left(\tau_r\right) \times y_{t-1}\right] + \varepsilon_t,$$

where y_{kt} is the quality index of brand k at period t and $\varepsilon_t \sim N\left(0, \sigma^2\right)$. τ_r denotes the hypothesized break point and $D_t\left(\tau_r\right)$ is a dummy variable that equals zero before τ_r and one after. With 15 percent trimming I test for breaks at all months between April 2002 and April 2004. Unfortunately, the first period that can be tested for a break is two months after the merger. Thus if the divergence between HP and Compaq had started sometime before or right after the merger, this test misses the starting period.

The test results show that Compaq's largest break is in October 2002 and HP's largest one is in September 2002. However, both breaks are not statistically significant at a 10 percent significant level with 15 percent trimming. This result suggests that the divergence started around the merger so that the two brands' trends were already on different paths in the sample period. I apply the QLR test to the other brands as well. Sony's largest break is in July 2002 and Apple's largest break is in September 2003, but both breaks are not significant at a 10 percent significant level. Moreover, Apple's largest break is about a year after the other brands' breaks. However, eMachines' quality index shows a significant break in November 2002 with the QLR statistics 10.21.

This suggests that eMachines responded to the merged firm's product repositioning by repositioning its own products.

Recently, researchers have begun to pay more attention to product repositioning in post-merger markets. Berry and Waldfogel (2001) and Sweeting (2010) provide the empirical evidence of post-merger product repositioning in the radio industry. Sweeting (2011) constructs and estimates a rich dynamic model where multi-product firms make decisions on product positioning. Gandhi, Froeb, Tschanz and Werden (2008) provide a theoretical model of product repositioning. Although I do not formally model a relationship between the merger and product positioning, my finding is consistent with the theoretical prediction in Gandhi, *et.al.* (2008). In their model firms instantly change both price and location after the merger. They show that the merged firm moves its products away from each other to reduce cannibalization. The merged firm's incentive to raise prices is reduced by this repositioning, and the anti-competitive merger effects are mitigated as a result.

6 Conclusions

In this paper I consider a new discrete choice model of differentiated product demand where consumers receive an idiosyncratic taste shock only at a brand level. This model distinguishes a brand-level differentiation from a product-level differentiation. The model is a hybrid of BLP and the PCM with BLP features at a brand-level and the PCM features for products within brands.

This hybrid model describes markets where firms offer multiple products of different qualities under the same brand name. Its substitution pattern is realistic such that a given product is not necessarily a substitute for products that are either very superior or inferior in their characteristics. Also, when firms introduce new products in the market, their welfare contribution is limited to a subset of consumers.

I use the scanner-level PC data to show the advantages of using the hybrid model in a real world application. Consumer preferences including the distribution of consumer types are estimated without using individual-level data such as the current population survey or the consumer expenditure survey. This is an important advantage especially when a market is defined as one geographical market. Also, the estimates' economic implications are reasonable and the implied markup is more realistic than that of BLP or the PCM.

Because the hybrid model's demand estimates are statistically significant with reasonable magnitudes, I use them to show how firms changed their product positioning in a post-merger market. Using the quality index constructed with the demand estimates, I show that HP and Compaq became further differentiated from each other after the merger. I admit that one may need a dynamic demand model to formally analyze post-merger product repositioning. Nevertheless, under the assumption that consumer preferences are stable over time, my study provides empirical evidence on product repositioning in post-merger markets.

References

[1] Aizcorbe, A. and Y. Pho (2005), "Differences in Hedonic and Matched-Model Price Indexes: Do the Weights Matter?," BEA Working Papers 0025, Bureau of Economic Analysis.

[2] Bajari, P. and L. Benkard (2005), "Demand Estimation With Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach," *Journal of Political Economy*, 113, 1239-1276.

[3] Bayer, P., F. Ferreira, and R. McMillan (2007), "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 115, 588-638.

[4] Berry, S. (1994), "Estimating Discrete Choice Models of Product Differentiation," *RAND Journal of Economics*, 25, 242-262.

[5] Berry, S., A. Gandhi, and P. Haile (2013), "Connected Substitutes and the Invertibility of Demand," *Econometrica*, 81(5), 2087-2111.

[6] Berry, S., J. Levinsohn, and A. Pakes (1995), "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841-890.

[7] Berry, S. and J. Waldfogel (2001), "Do Mergers Increase Product Variety? Evidence from Radio Broadcasting," *Quarterly Journal of Economics*, 116, 969–1007

[8] Berry, S., O. Linton, and A. Pakes (2004), "Limit Theorems for Estimating the Parameters of Differentiated Product Demand Systems," *Review of Economic Studies*, 71, 613-654.

[9] Berry, S. and A. Pakes (2007), "The Pure Characteristics Demand Model," *International Economic Review,* 48, 1193-1225.

[10] Bresnahan, F. (1987), "Competition and Collusion in the American Auto Industry: The 1995 Price War," *Journal of Industrial Economics*, 35, 457-482.

[11] Byzalov, D. (2010), "Unbundling Cable Television: An Empirical Investigation," unpublished manuscript, Temple University.

[12] Chu, J., P. Chintagunta, and N. Vilcassim (2007), "Assessing the Economic Value of Distribution Channels: An Application to the Personal Computer Industry," *Journal of Marketing Research*, 44, 29-41.

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

[14] Gandhi, A.K., L. M. Froeb, S. T. Tschantz, and G. J. Werden (2008), "Post-Merger Product Repositioning," *Journal of Industrial Economics*, 56, 49-67.

[15] Ghose, A., P. Ipeirotis, and B. Li. (2012) "Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowd-Sourced Content," *Marketing Science,* 31, 493-520

[16] Goeree, M. (2008), "Limited Information and Advertising in the U.S. Personal Computer Industry," *Econometrica*, 76, 1017-1074.

[17] McFadden, D. (1981), "Econometric Models of Probabilistic Choice," *Structural Analysis of Discrete Data with Econometric Applications*, C. Manski and D. McFadden, eds., The MIT Press.

[18] Nevo, A. (2000), "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics*, 31, 395-421.

[19] Nevo, A. (2001), "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69, 307-322.

[20] Nevo, A. (2003), "New Products, Quality Changes and Welfare Measures Computed From Estimated Demand Systems," *The Review of Economics and Statistics,* 85, 266-275.

[21] Pakes, A., S. Berry, and J. Levinsohn (1993), "Applications and Limitations of Some Recent Advances in Empirical Industrial Organization: Price Indexes and the Analysis of Environmental Change," *American Economic Review, Papers and Proceedings*, 83, 240-246.

[22] Pakes, A. (2003), "A Reconsideration of Hedonic Price Indexes with an Application to PCs," *American Economic Review,* 93, 1578-1596.

[23] Petrin, A. (2002), "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110, 705-729.

[24] Shaked, A. and J. Sutton (1982), "Relaxing Price Competition Through Product Differentiation," *Review of Economic Studies*, 49, 3-13.

[25] Shaked, A. and J. Sutton (1987), "Product Differentiation and Industrial Structure," *The Journal of Industrial Economics*, 36, 131-146.

[26] 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.

[27] Song, M. (2010), "The Quality Adjusted Price Index in the Pure Characteristics Demand Model," *Journal of Business and Economic Statistics*, 28, 190-199.

[28] Sweeting, A. (2010), "The Effects of Mergers on Product Positioning: Evidence from the Music Radio Industry", *RAND Journal of Economics*, 41, 372–397.

[29] Sweeting, A. (2011), "Dynamic Product Positioning in Differentiated Product Industries: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry", unpublished manuscript, Duke University.

Table 1: Trends in the Desktop Personal Computer Market from 4Q2001 to 3Q2004

	Price (in dollars)				Quantity (in thousands)				CPU speed (in GHz)			
Time	mean	std	min	max	mean	std	min	max	mean	std	min	max
4Q01	825.88	255.78	445.50	1,936.92	13.55	19.08	1.01	175.31	1.29	0.26	0.50	2.00
1Q02	827.45	259.00	350.38	2,081.72	9.51	9.90	1.03	54.09	1.35	0.26	0.60	2.20
2Q02	788.99	240.12	421.21	1,878.22	6.34	5.75	1.05	32.10	1.43	0.28	0.70	2.26
3Q02	795.12	252.50	393.65	1,987.81	6.42	6.38	1.00	29.17	1.60	0.35	0.70	2.53
4Q02	739.19	253.96	388.62	1,988.93	10.03	13.87	1.03	156.90	1.91	0.37	0.70	2.66
1Q03	718.78	291.42	360.04	2,265.47	7.15	7.43	1.00	32.67	1.97	0.36	0.70	3.06
2Q03	733.54	307.05	355.55	2,074.25	6.84	5.79	1.00	37.50	2.15	0.33	0.80	3.06
3Q03	710.00	299.60	343.41	2,073.72	7.17	9.25	1.00	57.58	2.26	0.31	1.00	3.00
4Q03	724.18	295.39	373.02	2,169.56	9.89	13.79	1.04	149.49	2.46	0.31	1.00	3.20
1Q04	729.75	289.46	322.21	2,147.68	7.15	8.49	1.01	55.13	2.55	0.33	1.00	3.20
2Q04	707.41	271.81	357.77	2,132.67	6.33	6.60	1.00	30.75	2.53	0.37	1.00	3.40
3Q04	690.64	248.03	222.14	1,699.68	7.27	7.76	1.07	44.57	2.48	0.34	1.25	3.40

Table 2: Demand Estimates of Other Discrete Choice Models

Variable	Logit	IV Logit	Nested	BLP I†	BLP II‡	PCM
Constant	-5.889*	-4.488*	-5.142*	-4.216*	-5.821*	1.156*
	(0.238)	(0.951)	(0.364)	(1.831)	(0.291)	(0.059)
CPU Speed	0.611*	0.873*	0.325*	0.875*	0.634*	0.065*
	(0.115)	(0.205)	(0.090)	(0.219)	(0.139)	(0.009)
Memory	-0.014	0.003	0.000	0.006	-0.016	0.006*
	(0.021)	(0.023)	(0.008)	(0.030)	(0.020)	(0.001)
Hard Drive	0.093	0.363	0.163*	0.380	0.109	0.035*
	(0.096)	(0.195)	(0.076)	(0.235)	(0.105)	(0.009)
Intel CPU	-0.340	-1.164	-1.089*	-1.155	-0.374	-0.167*
	(0.415)	(0.602)	(0.286)	(0.653)	(0.408)	(0.022)
AMD CPU	0.112	-0.801	-1.038*	-0.815	0.097	-0.181*
	(0.404)	(0.656)	(0.301)	(0.722)	(0.395)	(0.027)
Small Cache	0.358*	0.165	-0.067	0.116	0.367*	-0.087*
	(0.076)	(0.156)	(0.057)	(0.311)	(0.071)	(0.009)
Screen Size	-0.004	0.005	0.009*	0.007	-0.003	0.003*
	(0.004)	(0.007)	(0.002)	(0.011)	(0.004)	(0.0004)
LCD Screen	0.003	0.331	0.145	0.336	0.011	0.032*
	(0.095)	(0.234)	(0.093)	(0.256)	(0.131)	(0.012)
DVD Reader	-0.030	0.021	0.067*	0.042	-0.031	0.050*
	(0.065)	(0.077)	(0.026)	(0.138)	(0.067)	(0.004)
DVD Writer	0.242*	0.485*	0.214*	0.507*	0.249*	0.064*
	(0.101)	(0.193)	(0.075)	(0.238)	(0.111)	(0.010)
Sony	-0.221	0.078	2.026*	0.105	-0.217	0.081*
	(0.370)	(0.270)	(0.248)	(0.311)	(0.217)	(0.012)
HP	-0.270	-0.111	3.166*	-0.090	-0.270	0.067*
	(0.365)	(0.203)	(0.322)	(0.230)	(0.200)	(0.009)
eMachines	-0.129	-0.148	2.736*	-0.158	-0.151	0.005
	(0.369)	(0.183)	(0.292)	(0.195)	(0.206)	(0.009)
Compaq	-0.391	-0.307	2.699*	-0.298	-0.401*	0.047*
	(0.367)	(0.186)	(0.302)	(0.202)	(0.200)	(0.008)
Price_mu	-0.944*	-2.228*	-1.047*	-2.751	0	0
	(0.126)	(0.855)	(0.334)	(2.997)		
Price_sigma	0	0	0	0.683	0.059	1.243*
				(1.912)	(5.899)	(0.191)
Corr. within group			0.841*			
			(0.069)			

The time dummy variables are included.

*significant at a 5% level.

†The price variable has a random coefficient that is distributed normal.

‡The price variable has a random coefficient that is distributed log normal.

Table 3: Demand Estimates in the Hybrid Demand Model

Variable	Hybrid I†	Hybrid II‡	Hybrid III♯	Hybrid IV⅃
Constant	0.528*	0.525*	0.521*	0.517*
	(0.076)	(0.070)	(0.073)	(0.073)
CPU Speed	0.135*	0.124*	0.123*	0.122*
	(0.059)	(0.020)	(0.020)	(0.020)
Memory	0.006	0.007	0.007*	0.007*
	(0.005)	(0.004)	(0.004)	(0.004)
Hard Drive	0.102*	0.103*	0.102*	0.101*
	(0.022)	(0.017)	(0.018)	(0.019)
Intel CPU	-0.391*	-0.378*	-0.376*	-0.374*
	(0.081)	(0.052)	(0.055)	(0.058)
AMD CPU	-0.439*	-0.432*	-0.429*	-0.427*
	(0.060)	(0.052)	(0.056)	(0.058)
Small Cache	-0.133*	-0.134*	-0.134*	-0.133*
	(0.018)	(0.012)	(0.012)	(0.012)
Screen Size	0.007*	0.007*	0.007*	0.007*
	(0.001)	(0.001)	(0.001)	(0.001)
LCD Screen	0.092*	0.092*	0.091*	0.090*
	(0.027)	(0.026)	(0.027)	(0.027)
DVD Reader	0.060*	0.060*	0.060*	0.060*
	(0.010)	(0.010)	(0.010)	(0.010)
DVD Writer	0.139*	0.139*	0.139*	0.138*
	(0.020)	(0.018)	(0.018)	(0.018)
Sony	0.426*	0.428*	0.429*	0.427*
	(0.023)	(0.025)	(0.024)	(0.023)
HP	0.748*	0.749*	0.752*	0.753*
	(0.071)	(0.067)	(0.068)	(0.068)
eMachines	0.448*	0.450*	0.454*	0.455*
	(0.048)	(0.050)	(0.050)	(0.047)
Compaq	0.497*	0.500*	0.502*	0.504*
	(0.045)	(0.047)	(0.047)	(0.045)
Price_mu	0	0	0	0
Price_sigma	0.406*	0.407*	0.413*	0.417*
	(0.109)	(0.099)	(0.104)	(0.105)

The time dummy variables are included.

*significant at a 5% level.

†The market is defined as the US household desktop computer market and its size is estimated based on the number of computers sold to household consumers.

‡The market includes potential consumers and its size is set to three millions.

♯The market is defined in the same way as in Hybrid II, but its size is set to five millions.

⅃The market is defined in the same way as in Hybrid II, but its size is set to ten millions.

Table 4: Price Elasticity in Different Demand Models

	Product Elasticity				Brand Elasticity	
	Nested	BLP I	PCM	Hybrid I	BLP I	Hybrid I
Own elasticity	-4.733	-1.882	-3,368	-1,136	-2.977	-2.355
Cross elasticity						
Within-brand	0.015	4.1×10^{-4}	1,946	489.8	n.a.	n.a.
Cross-brand	0.002	3.6×10^{-4}	1,453	0.035	0.005	0.066
% of substitutes across brands†	1	1	0.03	0.10		

See Tables 2 and 3 for estimates in different demand models.
The medians of each period are averaged over all periods.
†The mean ratio of other brand products that are substitutes to a given brand product.

Table 5: The Average Percentage Markups in Different Demand Models

	Nested	BLP I	BLP II	PCM	Hybrid I	Hybrid II	Hybrid III
4Q01	1.289	0.514	1.275	0.0025	0.404	0.232	0.152
1Q02	1.306	0.528	1.311	0.0029	0.378	0.220	0.143
2Q02	1.252	0.526	1.282	0.0037	0.351	0.145	0.092
3Q02	1.306	0.543	1.338	0.0032	0.375	0.165	0.105
4Q02	1.439	0.576	1.455	0.0036	0.696	0.417	0.300
1Q03	1.423	0.587	1.467	0.0026	0.353	0.170	0.108
2Q03	1.348	0.574	1.422	0.0031	0.392	0.149	0.094
3Q03	1.406	0.572	1.427	0.0018	0.334	0.172	0.110
4Q03	1.382	0.555	1.393	0.0033	0.316	0.221	0.143
1Q04	1.384	0.562	1.401	0.0019	0.358	0.210	0.135
2Q04	1.415	0.583	1.450	0.0019	0.399	0.187	0.119
3Q04	1.517	0.610	1.544	0.0033	0.401	0.252	0.163

See Tables 2 and 3 for estimates in different demand models.

Table 6: The Average Markup by Brand

		BLP I	PCM	Hybrid I	Hybrid II
Apple	$(p - mc)/p$	0.354	0.0215	0.133	0.070
	$p - mc$	\$484	\$43	\$194	\$101
Compaq	$(p - mc)/p$	0.609	0.0007	0.332	0.166
	$p - mc$	\$421	\$4	\$245	\$118
eMachines	$(p - mc)/p$	0.723	0.0019	0.366	0.181
	$p - mc$	\$406	\$1	\$209	\$104
HP	$(p - mc)/p$	0.527	0.0020	0.531	0.296
	$p - mc$	\$436	\$6	\$403	\$215
Sony	$(p - mc)/p$	0.398	0.0031	0.195	0.099
	$p - mc$	\$456	\$5	\$230	\$119

See Tables 2 and 3 for estimates in different demand models.

Figure 1: The Market Share by Brand from October 2001 to September 2004

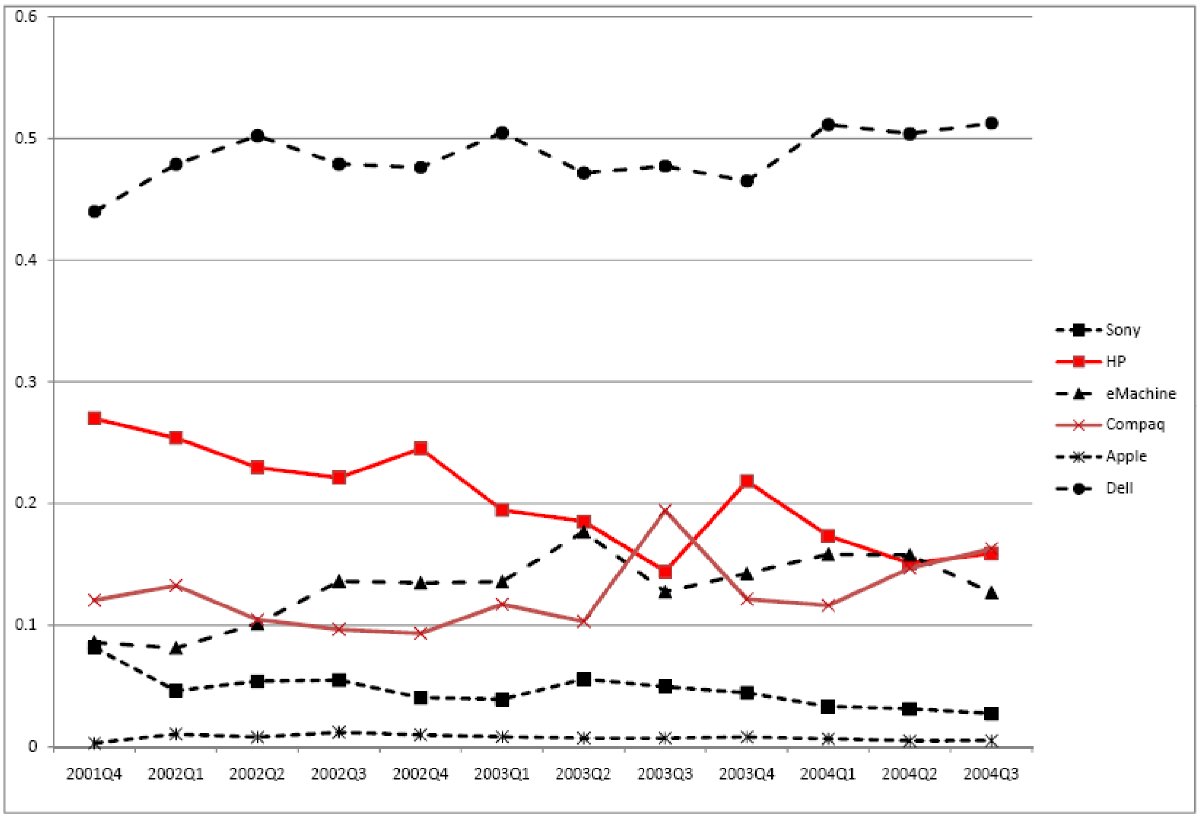

Figure 2: The Average Price by Brand from October 2001 to September 2004

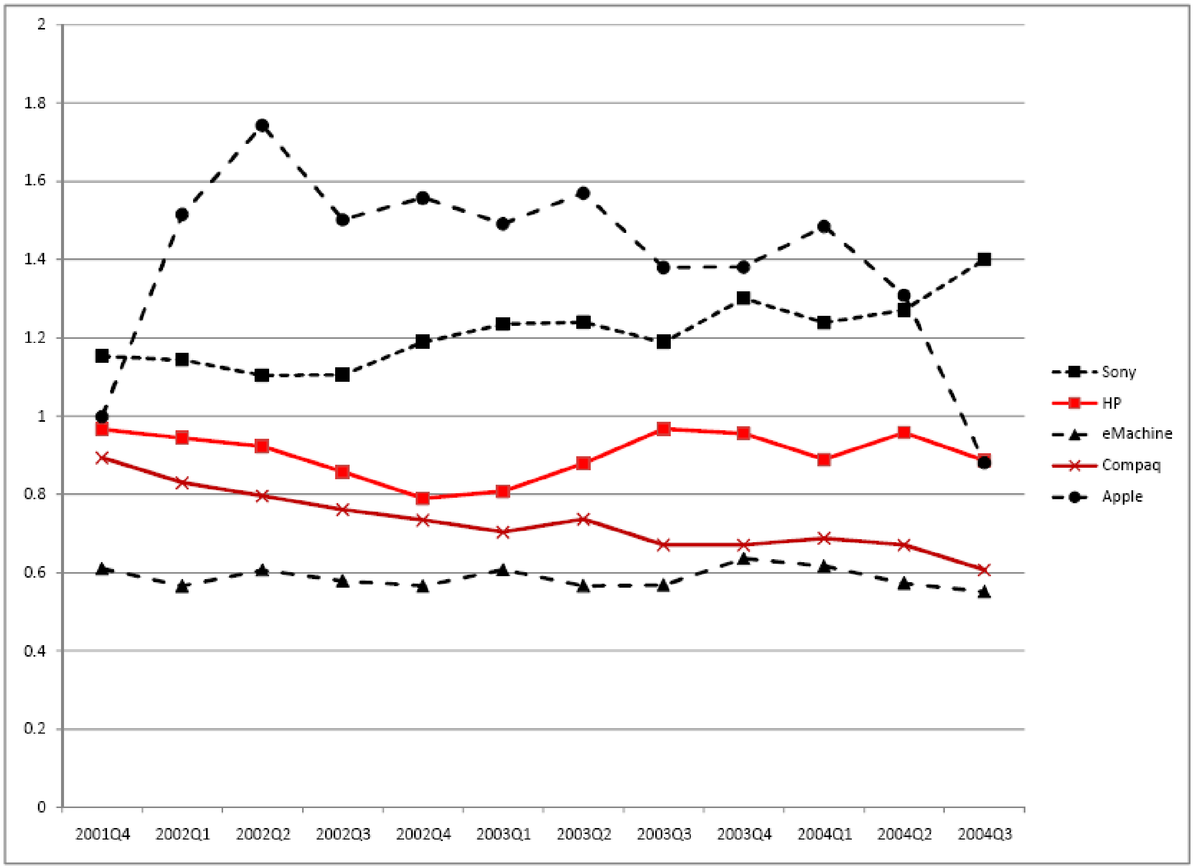

Figure 3: Consumer Welfare from October 2001 to September 2004

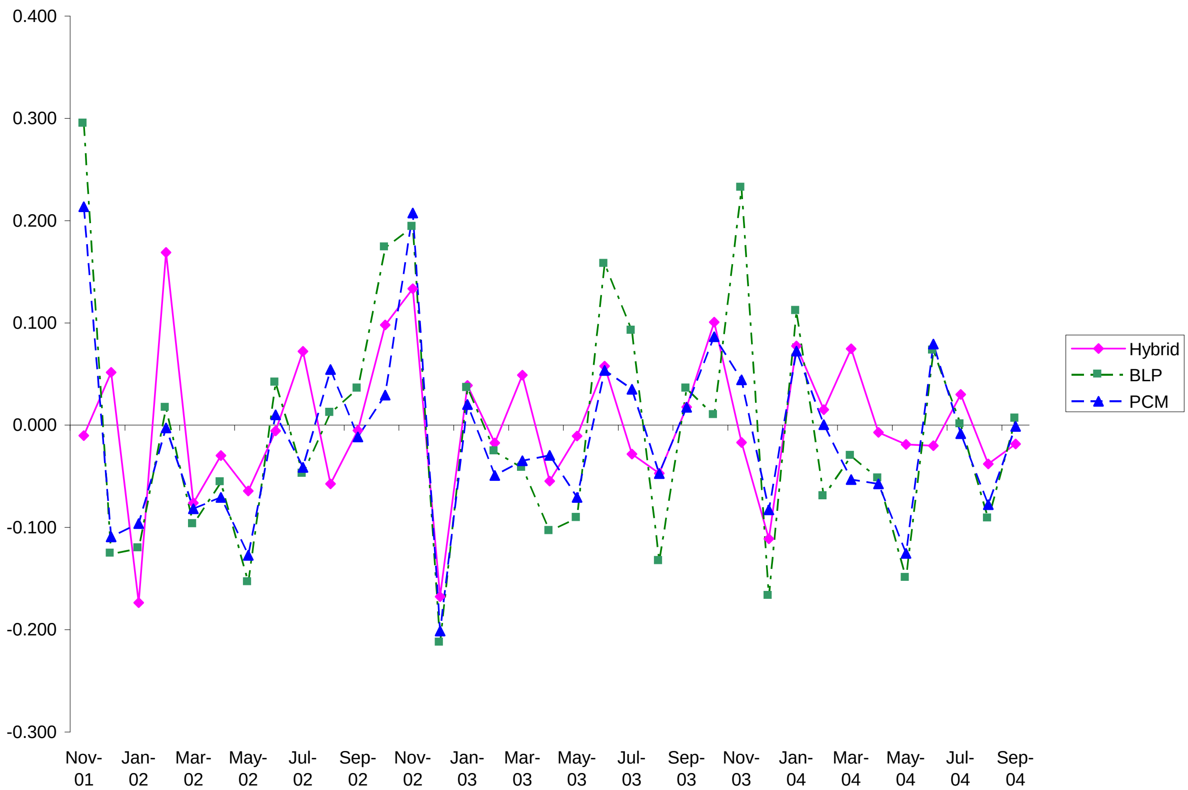

Figure 4: The Quality Index from October 2001 to September 2004

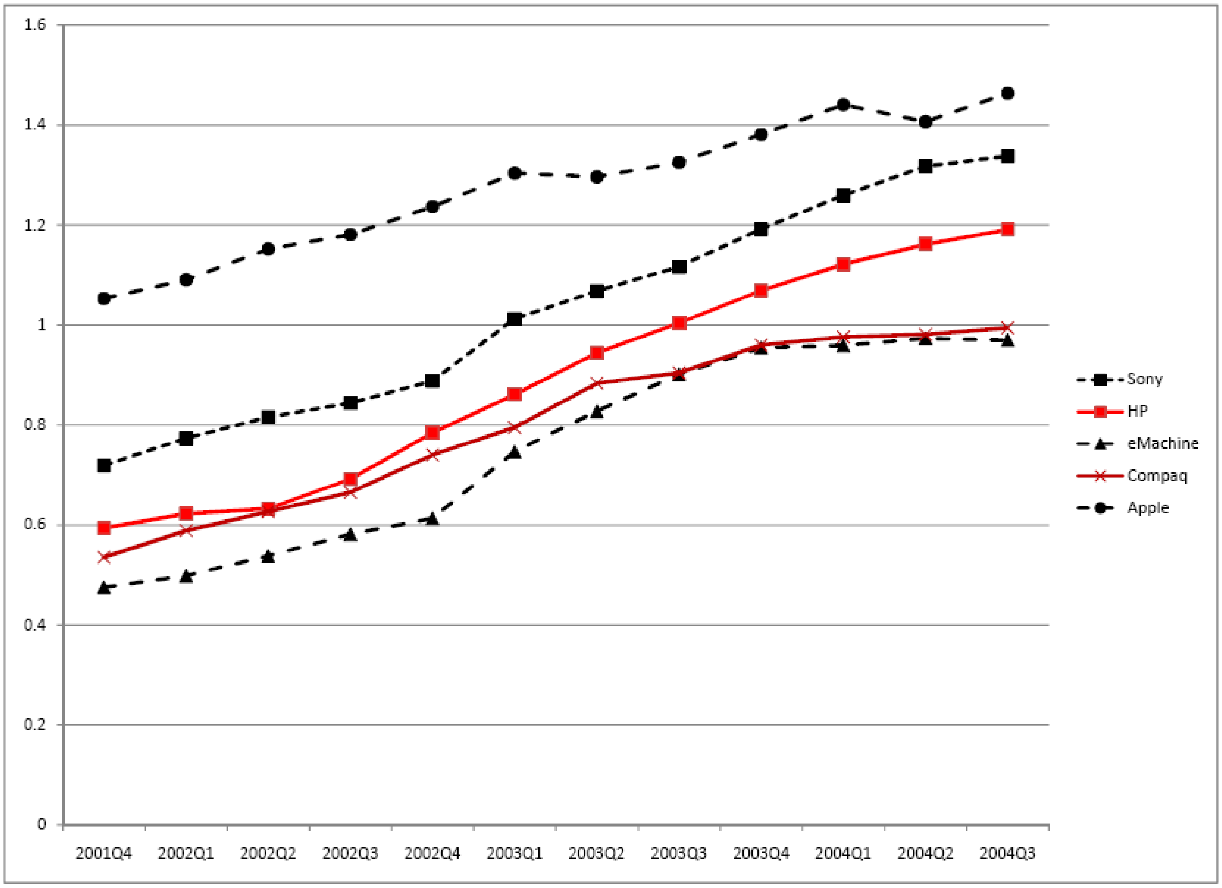

Appendix I: The Derivation of Equation (2)

Consider a consumer with $\frac{\delta_{j+1_k}-\delta_{j_k}}{p_{j+1_k}-p_{j_k}} \leq \alpha_i < \frac{\delta_{j_k}-\delta_{j-1_k}}{p_{j_k}-p_{j-1_k}}$. Her probability of choosing product j_k is the same as $\Pr\left(u_{ij_k} > u_{il_r}\right), \forall l_r \neq j_k$. Let j'_k be products of brand k other than j_k and l'_s be products of all other brands ($s \neq k$). Because $\Pr\left(u_{ij_k} > u_{ij'_k}\right) = 1$,

$$\Pr\left(u_{ij_k} > u_{ij_r}\right) = \Pr\left(u_{ij_k} > u_{ij'_k}\right)\Pr\left(u_{ij_k} > u_{il'_s}\right) = \Pr\left(u_{ij_k} > u_{il'_s}\right)$$

Now let $l^*_s = \arg\max u_{il'_s}$ for brand s, a product in brand s that maximizes this consumer's utility. Then

$$\Pr\left(u_{ij_k} > u_{il'_s}\right) = \Pr\left(u_{ij_k} > u_{il^*_s}\right)$$

where $u_{il^*_s} = \max_{l^l_s \in s} u_{il^l_s} = \max_{l^l_s \in s}\left(\delta_{l^l_s} - \alpha_i p_{il^l_s}\right) + \varepsilon_{is}$ for each brand. Because ε_i is a Type I extreme-value random variable,

$$\Pr\left(u_{ij_k} > u_{il^*_s}\right) = \frac{\exp\left(\delta_{j_k} - \alpha_i p_{j_k}\right)}{\sum_{r=1}^{K}\exp\left(\max_{l_r \in r}\left(\delta_{l_r} - \alpha_i p_{il_r}\right)\right)}$$

where the denominator includes product choices across all brands and $\max_{l_r \in k}\left(\delta_{l_r} - \alpha_i p_{il_r}\right) = \delta_{j_k} - \alpha_i p_{j_k}$.

Appendix II: Computational Details

In this appendix I explain computational details of finding product quality, $\boldsymbol{\delta}$, and show that there exists unique $\boldsymbol{\delta}$ in the one random coefficient hybrid model. In the GMM estimation procedure a computational challenge lies in finding $\boldsymbol{\delta}$ that matches the model-predicted market shares with observed market shares. The Newton-Raphson method is a natural candidate for this type of problem. For each market there are $\sum_{k=1}^{K} J_k$ number of nonlinear equations, which is equal to the number of products, and $\sum_{k=1}^{K} J_k$ variables, a vector of product quality. In my experience this method works well when the number of products per market is small, say less than twenty. When the number of products is higher, it is important to have good starting values. However, when there are more than fifty products per market, the computational time increases so much that it becomes impractical.

Thus, for markets with many products, say more than fifty, I developed a more practical yet theoretically valid method. The key idea is to use a product's observed within-brand share to approximate the conditional probability of choosing a product within a brand.[25] The within-brand share which I denote $\widetilde{s}_r$ is calculated by dividing sales units by the total unit sales of each brand, *i.e.*, $\widetilde{s}_{r_k} = q_{j_k} / \sum_{j_k=0}^{J_k} q_{j_k}$, whereas the market share, s_{j_k},is calculated by q_{j_k}/M. Note that the total unit sales of each brand should include products that were sold but are not included in data, *i.e.*, q_{0_k}. If this information is available, researchers can simply calculate within-brand market shares and recover product quality by

$$\delta_{j_k} = \delta_{j-1_k} + \left(p_{j_k} - p_{j-1_k}\right) F^{-1}\left(1 - \sum_{r=0}^{j-1} \widetilde{s}_{r_k} | \theta\right) \tag{12}$$

[25] A similar idea is used in Berry (1994) to estimate a nested logit demand model.

for product j of brand k where F is the cdf of the distribution of α.[26]

Suppose the total unit sales of each brand are not available and researchers only know the total unit sales of all brands combined. This means that no data is available for the share of the within-brand outside option, $\widetilde{s}_{0_k}$. In this case the quality of the lowest quality product in each brand, $\boldsymbol{\delta}_1$, should be found first. Given parameter values, $\boldsymbol{\delta}_1$ can be found by equating the sum of observed shares for each brand, $\sum_{j_k=1}^{J_k} s_{j_k}$, to

$$\left(1-\widetilde{s}_{0_k}\left(\boldsymbol{\delta}_1\right)\right)\int \frac{\exp\left(\max_{j_k\in k}\left(\delta_{j_k}\left(\delta_{1_k}\right)-\alpha_i p_{j_k}\right)\right)}{\sum_{r=1}^{K}\exp\left(\max_{l_r\in r}\left(\delta_{l_r}\left(\delta_{1_r}\right)-\alpha_i p_{l_r}\right)\right)} f\left(\alpha\right)d\alpha \tag{13}$$

where $\widetilde{s}_{0_k}=1-F\left(\delta_{1_k}/p_{1_k}|\boldsymbol{\theta}\right)$.[27] Because the number of elements in $\boldsymbol{\delta}_1$ is the same as the number of brands which is usually no more than twenty, the Newton-Raphson method can easily handle this computation. Once the quality of the lowest quality product is found for each brand, quality of the other products can be recovered using equation (12) where $\widetilde{s}_{r_k}$ is now a function of $\boldsymbol{\delta}_1$.

Consider a case where two brands are selling three products each (six products in total in the market) and researchers only have data on two (higher quality) products of each brand. Let p_{j_k} be the price of product j of brand k, q_{j_k} be its unit sales, and δ_{j_k} be its product quality. And let M be the market size which is the total unit sales of the two brands combined. Steps to recover product quality given prices and market shares are the following.

1. Calculate market shares by dividing q_{j_k} by the market size M. Let s_{j_k} be a market share obtained this way. Note that this share is different from the within-brand share which I denote $\widetilde{s}_{j_k}$.

2. Pick an arbitrary value for δ_{1_k} and calculate $\widetilde{s}_{0_k}=1-F\left(\frac{\delta_{1_k}}{p_{1_k}}\right)$.

3. Calculate
$$\widetilde{s}_{1_k}=\frac{s_{1_k}}{s_{0_k}+s_{1_k}+s_{2_k}}$$
where $s_{0_k}=\widetilde{s}_{0_k}\left(s_{1_k}+s_{2_k}\right)/\left(1-\widetilde{s}_{0_k}\right)$.

4. With $\widetilde{s}_{1_k}$ from step 3, calculate
$$\delta_{2_k}=\delta_{1_k}+\left(p_{2_k}-p_{1_k}\right)F^{-1}\left(1-\widetilde{s}_{0_k}-\widetilde{s}_{1_k}|\theta\right).$$

5. Evaluate whether $(s_{1k}+s_{2k})$ is sufficiently close to
$$F\left(\frac{\delta_{1_k}}{p_{1_k}}\right)\int \frac{\exp\left(\max_{j_k\in k}\left(\delta_{j_k}-\alpha_i p_{j_k}\right)\right)}{\sum_{r=1}^{2}\exp\left(\max_{l_r\in r}\left(\delta_{l_r}-\alpha_i p_{l_r}\right)\right)} f\left(\alpha|\theta\right)d\alpha \tag{14}$$

6. Repeat steps 2 to 5 until the difference in step 5 is sufficiently close to zero.

[26] Note that in the two random coefficient hybrid model this inversion is no longer possible and one should use the algorithm in Berry and Pakes (2007).

[27] $\sum_{j_k=1}^{J_k} s_{j_k}=\sum_{j_k=1}^{J_k} q_{j_k}/M=\left(1-\widetilde{s}_{0_k}\right)\sum_{j_k=0}^{J_k} q_{j_k}/M$

For a more general case where the number of products per firm is J_k, $\widetilde{s}_{j_k}$ in step 3 is calculated by

$$\widetilde{s}_{j_k} = \frac{s_{j_k}}{s_{0_k} + \sum_{1_k}^{J_k} s_{j_k}}$$

where $s_{0_k} = \widetilde{s}_{0_k} \sum_{1_k}^{J_k} s_{j_k} / (1 - \widetilde{s}_{0_k})$ and δ_{j_k} in step 4 is calculated by

$$\delta_{j_k} = \delta_{j-1_k} + (p_{j_k} - p_{j-1_k}) F^{-1} \left(1 - \widetilde{s}_{0_k} - \sum_{r_k = 1_k}^{j-1_k} \widetilde{s}_{r_k} | \theta \right).$$

In step 5 the integral in equation (14) can be approximated by drawing a large number of random variables from $F(\alpha|\theta)$ given θ and taking the sample average of $\frac{\exp\left(\max_{j_k \in k}\left(\delta_{j_k} - \alpha_i p_{j_k}\right)\right)}{\sum_{r=1}^{2} \exp(\max_{l_r \in r}(\delta_{l_r} - \alpha_i p_{l_r}))}$. This step is computationally intensive because of the max function. One needs more random draws as the number of products per brand increases, and this could increase computational burden substantially. In my application I use 35,000 draws.

This method always finds the unique $\boldsymbol{\delta}$ in the one random coefficient model**.** The existence and uniqueness of $\boldsymbol{\delta}$ is necessary for the hybrid model to be identified and can be established using the proof in Berry (1994).[28] Because the quality of the lowest quality product uniquely determines the quality of all other products, all I need to show is that the quality of the lowest quality product uniquely exists. The sufficient conditions for existence in Berry (1994) are that the market share function, $s(\boldsymbol{\delta})$, is everywhere differentiable with respect to $\boldsymbol{\delta}$, and its derivatives obey the following strict inequalities.

$$\begin{aligned} \frac{\partial s_{j_k}}{\partial \delta_{j_k}} &> 0 \\ \frac{\partial s_{j_k}}{\partial \delta_k} &< 0 \text{ for } k \neq j_k. \end{aligned}$$

Notice that the second inequality is not satisfied in the hybrid model because all products are not necessarily substitutes for one another. However, they are satisfied with the brand share function. Let s_k be the brand share function.

$$s_k(\boldsymbol{\delta}_1) = \int \frac{\exp\left(\max_{j_k \in k}\left(\delta_{j_k}(\delta_{1_k}) - \alpha_i p_{j_k}\right)\right)}{\sum_{r=1}^{K} \exp\left(\max_{l_r \in r}\left(\delta_{l_r}(\delta_{1_r}) - \alpha_i p_{l_r}\right)\right)} f(\alpha)\, d\alpha$$

where $\boldsymbol{\delta}_1$ is a vector of the lowest quality product in each brand. Notice that δ_{j_k} and δ_{j_r} are

[28] Berry, Gandhi, and Haile (2013) establish more general conditions that a demand system needs to satisfy to be identified. Their key conditions are (1) product $j's$ characteristics are excluded from the conditional indirect utilities of other products, (2) the conditional indirect utility of product j is increasing in its characteristics, and (3) there be no strict subset of products that substitute only among themselves. One can easily check that the hybrid model satisfies these conditions.

functions of δ_{1_k} and δ_{1_r} respectively. It is straightforward to show that

$$\begin{aligned} \frac{\partial s_k}{\partial \delta_{1_k}} &> 0 \\ \frac{\partial s_k}{\partial \delta_{1_r}} &< 0 \text{ for } r \neq k. \end{aligned}$$

Therefore, it follows that $\boldsymbol{\delta}_1$ exists such that $s_k(\boldsymbol{\delta}_1)$ equals brand $k's$ market share for all k. Once the existence is established, the uniqueness follows from the property that increasing all the mean product quality, including the within-brand outside option, by the same amount does not change any brand market share.

The following examples would further help readers understand how product quality is recovered from data on prices and market shares. Suppose there are two brands with each brand having two (vertically differentiated) products. Suppose further that the price of Brand 1's low-end product is the same as the price of Brand 2's high-end product, *i.e.*,

$$p_{1,high} > p_{1,low} = p_{2,high} > p_{2,low}$$

and

$$p_{1,high} - p_{1,low} = p_{2,high} - p_{2,low}$$

Note that the outside option here is to choose another product in each brand that we do not have data on and its utility is set to zero. First, assume that all four products have the same market shares. Then it should hold that

$$\frac{\delta_{1,low}}{p_{1,low}} = \frac{\delta_{2,low}}{p_{2,low}}$$

and

$$\frac{\delta_{1,high} - \delta_{1,low}}{p_{1,high} - p_{1,low}} = \frac{\delta_{2,high} - \delta_{2,low}}{p_{2,high} - p_{2,low}}$$

Because $p_{1,low} > p_{2,low}$, it follows that $\delta_{1,low} > \delta_{2,low}$, and thus $\delta_{1,high} > \delta_{2,high}$. This implies that Brand 1 can charge higher prices and still have the same market shares because the quality of their products is higher.

Now suppose the market share of Brand 1 products is larger than that of Brand 2 products. This will make the difference between $\delta_{1,low}$ and $\delta_{2,low}$ even larger, and this is consistent with the economic logic that if more consumers choose a more expensive brand, its quality must be much higher. It is also easy to see that the quality difference will be smaller as Brand 1's share becomes smaller. It is even possible that $\delta_{1,low} < \delta_{2,low}$ if the market share of Brand 1 is substantially smaller than that of Brand 2. It is not surprising that a brand with higher prices and lower quality attracts much fewer consumers. These examples also show that the zero utility normalization for the within-brand outside option does not mask the systematic quality differences across the brands.

Appendix III: Model Comparisons Using Monte Carlo Simulation

I use Monte Carlo simulations to compare the hybrid model with BLP and the PCM. I compare the models on three dimensions. First, I generate market shares in the three models and compare them. All three models have the same product quality and prices, so any differences in market shares come from the idiosyncratic error structure. Second, I take one of the three models as a "true" model at a time and use its market shares to estimate the other two models. For example, when BLP is treated as a "true" model, I use the market shares generated by BLP and estimate the PCM and the hybrid model. Third, I repeat the second exercise to compute markups and consumer welfare, using both the "true" and the estimated demand parameters. When I use BLP shares as true shares, for example, the true demand parameters are those of BLP and the estimated demand parameters are those of the other two models estimated in the second exercise. This way I show to what extent the differences in markups and consumer welfare are attributed to differences in the error structure or in the parameter values.

Throughout the simulations I consider a market where there are three brands with each brand having four products. There is one observable attribute, x, and I assume that this attribute is a desirable characteristic and distributed uniformly on $(0, 1)$. There are brand dummy variables, $\boldsymbol{\gamma} = [\gamma_1, \gamma_2, \gamma_3]$, that capture the average utility of brand-level unobserved attributes, and I set $\boldsymbol{\gamma} = [1, 0.5, 0]$. I assume that price is an exponential function of the mean product quality which is a linear combination of the observable attribute and the brand dummies. Thus, the price of product j of brand k is

$$p_{jk} = c \exp\left(\beta_0 + \beta_1 x_{jk} + \gamma_k\right)$$

where I set $c = 1/20$, $\beta_0 = 2$ and $\beta_1 = 1.5$.

I allow for an unobservable attribute, ξ_{jk} but do not allow it to be correlated with the price to make estimation simple. Thus, the price endogeneity problem is absent in this exercise. I assume that $\xi_{jk} \sim 0.001 \times N(0, 1)$. The small variance provides two benefits. First, the number of markets required for consistent estimation is small. The estimates of the "true" model are no different from true values with ten independent markets (120 products in total). Second, the smaller variance of ξ_{jk}, the easier it is to generate positive market shares in the hybrid model and especially in the PCM.

Consumers are heterogeneous on a single dimension and the heterogeneity, denoted by α_i, is captured by a random coefficient on the price variable. I assume that $\log(\alpha) \sim N(\mu, \sigma)$ where $(\mu, \sigma) = (0, 1)$. In all models I fix μ at 0 and treat σ as a parameter to estimate. Let U_{ijk} be the component of the utility function that all of the three models share,

$$U_{ijk} = \delta_{jk} - \alpha_i p_{jk}$$

where $\delta_{jk} = \beta_0 + \beta_1 x_{jk} + \gamma_k + \xi_{jk}$. Despite being simple, this utility function describes many markets including the PC market where multiple brands compete with vertically differentiated products.

Given this common component, each model generates different market shares because of the different structure of the idiosyncratic error term. Table III-1 summarizes the brand-level as well as the product-level market shares that each model generates over 100 repetitions. I compute the mean and the standard deviation across ten markets for each repetition and report their averages across 100 repetitions. For each brand, product 1 is the lowest quality and product 4 is the highest quality product.

First, note that in BLP market shares are more evenly distributed across the quality ranking and the product-level standard deviations are smaller. Because of the product-level idiosyncratic shock, the quality difference and consumer heterogeneity is less important in generating market shares. Also, note that the standard deviation is much larger in the PCM at both brand- and product-level. In the PCM all products are ranked on the single quality dimension, and the most expensive and the least expensive products have substantially larger market shares because they have one less substitute, resulting in the large product-level standard deviation. Since any brand can have these products, the brand-level standard deviation is also large. In the hybrid model, the most expensive and the least expensive products of each brand have larger market shares, resulting in smaller standard deviations.

Table III-2 shows estimation results when one of the three models is treated as a "true" model. Under the heading *Hybrid*, for example, a true model is the hybrid model so all models use the market shares generated in the hybrid model.[29] This exercise shows how close estimates are to true values when a model is mis-specified. I use the method of moments for estimation. Since μ is fixed and the price variable is not correlated with ξ, the number of moments is the same as the number of parameters. I use the grid search for the price parameter, σ, as it enters the model non-linearly.

The table shows that the hybrid model is closer to BLP than to the PCM with respect to the mean square error. Both the hybrid model and BLP do poorly when the PCM is a true model. Their estimate for σ is far above the true value and the MSEs for all parameters are the largest than in any other cases. When either of the hybrid model or BLP is a true model, the PCM substantially underestimates σ but estimates the brand parameters, $\boldsymbol{\gamma}$, closer to the true values than the other two models do.

Between the hybrid model and BLP, the hybrid estimates are more robust with respect to the model (mis-)selection. Their means are very close to the true values, and their mean square errors are one third of those of the BLP estimates for almost all parameters. The BLP estimates are close to the true values on average, but their variances are much larger, resulting in much larger MSEs.

Next I calculate the markup using the estimation results in Table III-2. In Table III-3 I report the absolute markup for Products 1, 2, 3, and 4 averaged across the brands. Both the mean and the standard deviation are computed across the markets and then averaged over 100 repetitions.

Table III-3 first shows that the markup is similar between the hybrid model and BLP, no matter which market shares (and demand estimates) are used. The latter model tends to have larger markups for Products 3 and 4, but the differences are small.[30] This is interesting, considering their differences in the substitution pattern and the magnitude of price elasticities. This similarity comes from comparable cross-brand elasticities between the two models. In the hybrid model the brand-level taste shock renders the cross-brand elasticity much lower than the within-brand elasticity and the former elasticity has a much larger impact on the magnitude of markup. The table also shows that no matter which market shares are used, the markup is substantially lower in the PCM relative to the other two models.

I also calculate consumer welfare, following the same procedure as for the markup. Rather than comparing levels, I compare how consumer welfare changes when (1) the number of brands changes and (2) the number of products changes. I compute consumer welfare without (all products of) Brand 3 for the former and without Product 4 (of each brand) for the latter, and report in Table III-4 their ratios to consumer welfare with all brands and products, which I normalize to 1. I report results with prices fixed as well as with new equilibrium prices.

[29] The estimates of the "correct" model (e.g., the hybrid model estimated with the hybrid market shares) are very close to the true values so I put the true values instead.

[30] The markup is noticeably larger when the PCM market share is used (reported in the last two columns). It is mainly because both the hybrid model and BLP predict much higher product qualities with the PCM market shares.

Consider the case with fixed prices first (the rows of *No Price Change*) where welfare changes depend solely on changes in product quality. Comparing BLP and the PCM, welfare changes are much smaller in the PCM as expected. Also, in both models welfare changes are larger when Product 4 is taken out than when Brand 3 is taken out. Since neither model distinguishes product exits from brand exits, this simply means that the welfare contribution of each brand's highest quality product is larger than that of Brand 3's four products.[31] In the hybrid model, on the other hand, consumer welfare goes down less when Product 4 is taken out because product exit does not affect the brand-level taste shock. In terms of the magnitude, the welfare change from taking out Product 4 is similar to that of the PCM, while the welfare change from taking out Brand 3 is similar to that of BLP.

Consider next the impact of price changes on welfare changes (the rows of *Price Change*). In both BLP and the hybrid model the impact of price changes on welfare changes is larger when Brand 3 is taken out, implying that prices change more substantially in response to the entry/exit of a brand than existing brands' new products. Between the two models the price impact is much more pronounced in the hybrid model. However, this is not necessarily true in the PCM where some products may not have substitutes in other brands. In the extreme case where all of Brand 3 products are cheaper than the other two brands, their absence has negligible effects on the other brands' prices.

Appendix IV: Data Description

Scanner data on PCs were collected from US retail outlets by NPD Techworld.[32] The original data set contains 6,235 computer products ranging from PDAs to desktops and has 47,923 observations for 36 months from October of 2001 to September of 2004. The main components of the data set are desktop and notebook computers. There are 2,769 desktop and 2,756 notebook computers covering 66% and 33% of the total sales respectively. The rest consists of small portable computers like tablet PCs. Dell's products are not included in the data set since they are not sold at retail outlets.

IDC estimates that about 47.6 million computers were sold in the US in 2002 and 52.7 million in 2003.[33] And about 17.1 million were sold to household consumers in 2002 and about 20 million in 2003. The total units sold in the data set is 7.7 and 8.4 million in 2002 and 2003 respectively, which means that the data set covers about 42 $\sim$ 45% of the US household market. Since products of Dell, which has about 30% market share in the US market, are not included, the data set covers 60 $\sim$ 65% of the US retail outlet market.[34]

Most observations have very small sales. 53.95% of the total observations have less than 10 units of sale, 78.64% less than 100 units, and 90% less than 1,000 units. However, the sale is concentrated on a small number of products. Observations with more than 100 units cover about 96% of the total sale and observations with more than 1,000 units cover about 90% of the sale. This means that 10% of observations cover 90% of the total sale. Moreover, when observations with less than 100 units are excluded, all of the outliers in price and product characteristics disappear.

[31] An exception is BLP with the market shares generated by BLP where the inferred quality of Brand 3 products is higher than the other cases.

[32] Aizcorbe and Pho (2005) use the same data and provide more details on the data.

[33] Gartner estimates that 43.9 million were sold in 2002 and 57.7 million were sold in 2003.

[34] Gartner estimates that the market share of Dell in the US market is 29.69% in 2002 and 28.19% in 2003.

And the summary statistics of product characteristics do not change significantly when the cutoff is raised to 1,000 units.

Appendix V: Demand Estimation in the Random Coefficient Logit Model (BLP)

In Table V-1 I report the IV logit estimation with different instrumental variables. The table shows that the price coefficient becomes positive with the BLP-type IVs (IV 1). When product characteristics are interacted with the time dummy variables, the price coefficient becomes negative. This is probably because the price does not increase over time with improving product characteristics. IV 4 is used to estimate BLP in section 4.

In Tables V-2 and V-3 I report BLP estimates with different specifications. Table V-2 reports variations of *BLP I* in Table 2. The random coefficients are added on the constant term and the speed variable in addition to the price variable, but they do not change estimation results significantly. Table V-3 reports variations of *BLP II* in Table 2. Different parameters are normalized and different instrumental variables are used for each specification.

Table III-1: Market Shares in the Monte Carlo Simulations

	Hybrid		BLP		PCM	
	mean	std	mean	std	mean	std
Brand 1	0.262	0.024	0.249	0.024	0.282	0.063
Brand 2	0.306	0.021	0.316	0.022	0.262	0.077
Brand 3	0.326	0.018	0.369	0.023	0.396	0.076
Product 1	0.122	0.027	0.092	0.014	0.113	0.079
Product 2	0.039	0.019	0.083	0.015	0.060	0.034
Product 3	0.037	0.019	0.073	0.016	0.056	0.032
Product 4	0.100	0.040	0.064	0.015	0.085	0.050

Table III-2: Estimations in the Monte Carlo Simulations

		Hybrid		BLP		PCM	
		mean	mse	mean	mse	mean	mse
Hybrid	β_1	1.5		1.517	0.010	1.407	0.205
	γ_1	1		0.872	0.018	0.978	0.030
	γ_2	0.5		0.439	0.005	0.286	0.071
	σ	1		1.074	0.013	1.687	0.519
BLP	β_1	1.587	0.046	1.5		1.554	0.161
	γ_1	1.109	0.022	1		1.161	0.055
	γ_2	0.578	0.017	0.5		0.425	0.039
	σ	0.948	0.049	1		1.856	0.757
PCM	β_1	1.600	0.011	1.632	0.021	1.5	
	γ_1	1.063	0.004	1.058	0.005	1	
	γ_2	0.529	0.001	0.525	0.001	0.5	
	σ	0.658	0.117	0.656	0.120	1	

Table III-3 : Absolute Markup in the Monte Carlo Simulations

		Hybrid		BLP		PCM	
		mean	std	mean	std	mean	std
Hybrid	Product 1	1.075	0.279	1.116	0.281	1.234	0.556
	Product 2	1.286	0.384	1.360	0.395	1.643	0.823
	Product 3	1.539	0.481	1.629	0.503	2.135	1.053
	Product 4	1.840	0.536	1.917	0.587	2.682	1.264
BLP	Product 1	1.087	0.244	1.083	0.244	1.167	0.604
	Product 2	1.301	0.364	1.337	0.406	1.658	0.942
	Product 3	1.574	0.490	1.635	0.520	2.259	1.266
	Product 4	1.887	0.582	1.934	0.617	2.923	1.566
PCM	Product 1	0.040	0.047	0.048	0.067	0.095	0.132
	Product 2	0.050	0.089	0.068	0.108	0.126	0.197
	Product 3	0.098	0.171	0.114	0.181	0.207	0.327
	Product 4	0.184	0.256	0.188	0.269	0.339	0.498

Table III-4: Consumer Welfare in the Monte Carlo Simulations

			Hybrid		BLP		PCM	
			mean	std	mean	std	mean	std
Hybrid	All Products*		1		1		1	
	No Brand 3†	No Price Change	0.909	0.007	0.926	0.005	0.970	0.010
		Price Change	0.875	0.013	0.906	0.013	0.961	0.025
	No Product 4‡	No Price Change	0.972	0.013	0.989	0.005	0.984	0.011
		Price Change	0.965	0.017	0.987	0.007	0.981	0.014
BLP	All Products*		1		1		1	
	No Brand 3†	No Price Change	0.915	0.008	0.930	0.004	0.982	0.007
		Price Change	0.863	0.013	0.901	0.014	0.949	0.068
	No Product 4‡	No Price Change	0.869	0.019	0.935	0.003	0.943	0.013
		Price Change	0.839	0.023	0.923	0.005	0.923	0.025
PCM	All Products*		1		1		1	
	No Brand 3†	No Price Change	0.996	0.003	0.994	0.005	0.994	0.005
		Price Change	0.970	0.018	0.962	0.024	0.976	0.012
	No Product 4‡	No Price Change	0.984	0.012	0.991	0.006	0.980	0.017
		Price Change	0.954	0.028	0.969	0.022	0.955	0.023

*Consumer welfare with all products is normalized to 1.
†All four products of Brand 3 are taken out.
‡Product 4 of each brand is taken out.

Table IV-1: Summary Statistics of Product Characteristics

	Memory (in Mega bytes)				Hard Drive (in Giga bytes)				Screen Size (in inches)			
	mean	std	min	max	mean	std	min	max	mean	std	min	max
4Q01	232.78	131.00	64	512	42.29	18.03	20	80	16.90	0.35	15	17
1Q02	312.66	157.13	64	512	53.30	24.78	20	120	16.36	0.83	15	17
2Q02	239.29	141.12	64	512	50.36	21.28	20	120	16.50	0.84	15	17
3Q02	291.94	174.22	64	512	53.57	19.91	20	120	16.65	0.72	15	17
4Q02	306.26	157.01	128	512	63.16	23.03	20	120	16.51	0.80	15	17
1Q03	306.21	160.91	128	1024	64.83	26.45	40	250	16.27	0.96	15	18
2Q03	344.38	152.09	128	1024	74.31	30.51	40	160	16.18	0.86	15	17
3Q03	323.26	164.24	128	1024	79.17	33.01	40	160	16.45	0.89	15	17
4Q03	366.83	139.04	128	1024	88.32	39.17	40	200	16.63	0.80	15	20
1Q04	395.06	129.20	128	1024	99.97	47.95	40	250	16.68	0.82	15	20
2Q04	460.10	116.57	256	1024	111.76	52.66	40	250	16.83	0.60	15	20
3Q04	445.36	143.47	256	1024	113.17	56.84	40	250	16.90	0.47	15	19

Table IV-2: Summary Statistics of Product Characteristics*

	Intel CPU	AMD CPU	Cache†	Screen‡	LCD screen	DVD reader	DVD writer
4Q01	0.82	0.18	0.463	0.106	0.000	0.514	0.021
1Q02	0.85	0.13	0.516	0.136	0.012	0.533	0.036
2Q02	0.78	0.20	0.452	0.073	0.014	0.526	0.041
3Q02	0.75	0.23	0.493	0.144	0.016	0.592	0.075
4Q02	0.78	0.20	0.473	0.242	0.019	0.587	0.118
1Q03	0.57	0.42	0.284	0.194	0.025	0.764	0.162
2Q03	0.58	0.41	0.277	0.217	0.051	0.745	0.216
3Q03	0.53	0.46	0.285	0.321	0.098	0.700	0.220
4Q03	0.70	0.29	0.458	0.317	0.120	0.709	0.234
1Q04	0.71	0.28	0.478	0.343	0.148	0.700	0.297
2Q04	0.58	0.41	0.408	0.233	0.046	0.821	0.422
3Q04	0.61	0.38	0.414	0.252	0.022	0.927	0.439

*Proportions of products with relevant attributes
†A proportion of products with a smaller cache size than the standard size.
‡A proportion of products with a screen included.

Table V-1: Demand Estimates in the IV Logit Model

Variable	Simple Logit	IV 1†	IV 2‡	IV 3◇	IV 4⊤
Constant	-5.889*	-9.927*	-5.498*	-4.722*	-4.488*
	(0.238)	(1.711)	(0.735)	(0.916)	(0.951)
CPU Speed	0.611*	-0.115	0.565*	0.828*	0.873*
	(0.115)	(0.342)	(0.162)	(0.198)	(0.205)
Memory	-0.014	-0.069	-0.005	-0.001	0.003
	(0.021)	(0.037)	(0.022)	(0.023)	(0.023)
Hard Drive	0.093	-0.658	0.149	0.318	0.363
	(0.096)	(0.338)	(0.157)	(0.189)	(0.195)
Intel CPU	-0.340	1.973	-0.446	-1.016	-1.164
	(0.415)	(1.049)	(0.469)	(0.582)	(0.602)
AMD CPU	0.112	2.721	-0.058	-0.637	-0.801
	(0.404)	(1.155)	(0.504)	(0.633)	(0.656)
Small Cache	0.358*	0.953*	0.310*	0.199	0.165
	(0.076)	(0.268)	(0.128)	(0.151)	(0.156)
Screen Size	-0.004	-0.030*	-0.004	0.004	0.005
	(0.004)	(0.012)	(0.006)	(0.007)	(0.007)
LCD Screen	0.003	-0.941*	0.102	0.277	0.331
	(0.095)	(0.415)	(0.187)	(0.227)	(0.234)
DVD Reader	-0.030	-0.191	-0.010	0.014	0.021
	(0.065)	(0.102)	(0.070)	(0.076)	(0.077)
DVD Writer	0.242*	-0.491	0.325*	0.446*	0.485*
	(0.101)	(0.328)	(0.158)	(0.188)	(0.193)
Price	-0.944*	2.759	-1.245	-2.012*	-2.228*
	(0.126)	(1.575)	(0.656)	(0.824)	(0.855)
Validity of IVs		Yes	Yes	Yes	Yes

† IV 1: BLP IVs. Characteristics used are CPU speed, memory, hard drive, screen size, and DVD writer
‡ IV 2: CPU speed variable interacted with time dummy variables
◇ IV 3: BLP IVs interacted with time dummy variables
⊤ IV 4: BLP IVs interacted with time dummy variables excluding hard drive and DVD writer variables.
All specifications include dummy variables for firm and time.

Table V-2: Demand Estimates in the Random Coefficient Logit Model with $\alpha_i \sim N(\mu_\alpha, \sigma_\alpha)$ and $\beta_i \sim N(\mu_\beta, \sigma_\beta)$

Variable	Spec. 1	Spec. 2	Spec. 3	Spec. 4
Means				
Constant	-4.216*	-5.660	-3.925*	-4.318*
	(1.831)	(30.606)	(1.928)	(13.357)
CPU Speed	0.875*	0.872*	0.869	0.878
	(0.219)	(0.218)	(1.102)	(0.961)
Memory	0.006	0.006	0.008	0.007
	(0.030)	(0.028)	(0.030)	(0.030)
Hard Drive	0.380	0.373	0.394	0.390
	(0.235)	(0.240)	(0.243)	(0.240)
Intel CPU	-1.155	-1.135	-1.119	-1.154
	(0.653)	(0.661)	(0.941)	(0.794)
AMD CPU	-0.815	-0.797	-0.801	-0.829
	(0.722)	(0.719)	(0.960)	(0.856)
Small Cache	0.116	0.110	0.068	0.082
	(0.311)	(0.268)	(0.297)	(0.280)
Screen Size	0.007	0.007	0.008	0.008
	(0.011)	(0.010)	(0.010)	(0.010)
LCD Screen	0.336	0.327	0.330	0.341
	(0.256)	(0.266)	(0.261)	(0.259)
DVD Reader	0.042	0.047	0.063	0.058
	(0.138)	(0.128)	(0.137)	(0.129)
DVD Writer	0.507*	0.506*	0.520*	0.523*
	(0.238)	(0.221)	(0.222)	(0.239)
Price	-2.751	-2.702	-3.322	-3.065
	(2.997)	(2.733)	(2.844)	(3.107)
Standard Deviations				
Constant	0	2.327	0	0.945
		(31.218)		(21.344)
Price	0.683	0.749	1.039	0.901
	(1.912)	(1.495)	(1.467)	(1.351)
CPU Speed	0	0	0.036	0.031
			(8.496)	(9.592)

Five thousand simulated consumers are drawn from the standard normal distribution to compute the model predicted market shares.

Table V-3: Demand Estimates in the Random Coefficient Logit Model with $\log(\alpha_i) \sim N(\mu, \sigma)$

Variable	IV 2[†]		IV 4[‡]	
	Spec. 1	Spec. 2	Spec. 1	Spec. 2
Constant	-5.772*	-6.501*	-5.821*	-2.855
	(0.226)	(1.160)	(0.291)	(1.900)
CPU Speed	0.508*	0.385*	0.634*	0.825*
	(0.150)	(0.178)	(0.139)	(0.207)
Memory	-0.009	-0.018	-0.016	0.011
	(0.020)	(0.025)	(0.020)	(0.028)
Hard Drive	0.089	-0.051	0.109	0.365
	(0.126)	(0.184)	(0.105)	(0.217)
Intel CPU	-0.262	0.130	-0.374	-1.002
	(0.437)	(0.523)	(0.408)	(0.621)
AMD CPU	0.149	0.599	0.097	-0.673
	(0.436)	(0.579)	(0.395)	(0.693)
Small Cache	0.352*	0.462*	0.367*	0.051
	(0.073)	(0.173)	(0.071)	(0.226)
Screen Size	-0.006	-0.011	-0.003	0.008
	(0.004)	(0.007)	(0.004)	(0.009)
LCD Screen	0.015	-0.154	0.011	0.275
	(0.163)	(0.212)	(0.131)	(0.234)
DVD Reader	-0.019	-0.042	-0.031	0.089
	(0.068)	(0.081)	(0.067)	(0.102)
DVD Writer	0.262*	0.136	0.249*	0.508*
	(0.121)	(0.190)	(0.111)	(0.218)
Price (μ)	0	-1.422	0	2.115*
		(3.949)		(0.744)
Price (σ)	0.381	1	0.059	1
	(1.456)		(5.899)	

Two thousand simulated consumers are drawn from the standard normal distribution to compute the model predicted market shares.

[†]CPU speed is interacted with time dummy variables.

[‡]BLP IVs are interacted with time dummy variables.