Estimation of demand for differentiated durable goods

Juan Esteban Carranza
University of Wisconsin-Madison

September 1, 2006

Abstract

This paper develops a methodology to estimate models of demand for differentiated durable goods with fully heterogeneous consumers that extends standard estimation techniques to account for the dynamic concerns of consumers. A "nested" technique is used to uncover a reduced form of the solution to the dynamic optimization problem of consumers which can be easily incorporated into a conventional multinomial discrete choice framework. The presence of such reduced form serves as the control for the dynamic optimization problem and allows a consistent estimation of the deep preference parameters. A simple application to the market for digital cameras is used to demonstrate the feasibility of the technique and its advantage over standard static techniques. The methodological framework has potential applicability in other dynamic demand problems. It is also substantially less costly than alternative estimation approaches, which in principle require the computation of the dynamic optimization problem across all consumer types and through all the steps of the estimation algorithm.

1 Introduction

The empirical literature on demand for durable differentiated goods has been dominated by static discrete choice models (Berry (1994), Berry, Levinsohn and Pakes (1995) –BLP, Goldberg (1995)) that do not account for the intertemporal incentives of market participants. This assumption seems to be particularly problematic in the case of markets for technological consumer goods, which are characterized by a very high rate of product innovation that makes the effective price of quality to decrease rapidly over time. The dynamics of quality and price and the durability of products may induce consumers to time optimally the purchase of products. Techniques based on static models implicitly assume that consumers who don't purchase any product don't value them enough and therefore lead in general to biased estimates of preference parameters. The magnitude of this effect is an empirical issue that must be addressed case by case, but is presumably significant in markets with high rates of technological change.

This paper develops a simple empirical framework for studying the demand for durable goods using product-level data. It builds partially upon an idea by Melnikov (2000) which incorporates empirical optimal stopping problems into a standard logit demand model by separating the decision on the purchase of a new product into two parts: first, consumers decide whether to buy any product solving a dynamic optimization problem which depends on the expected evolution of quality and prices, indexed by a sufficient statistic of the distribution of attainable payoffs. Then, conditional on the purchase decision, consumers choose among available products according to a static discrete choice model.

Melnikov's model inherits the limitations of the simple logit model, i.e. the unrealistic cross substitution patterns. Such feature of the model inhibits its use in complex supply models, where cross substitutions are a crucial determinant of firms' behavior. Simply extending the original technique to allow for correlation across consumers choices (e.g. via random taste parametrs) is a lot more costly, if not completely impractical, since in that case obtaining the predicted demand for any set of parameters requires the computation of the integral over the distribution of consumer characteristics of the individual demand function. This individual demand function contains the solution to a dynamic optimization problem, which itself depends on consumer characteristics [1].

The approach taken in this paper is to estimate a reduced form of the endogenous participation probability. The identification of such participation function is based on the observation over time of the total number of purchases, which is an information that is not used at all in standard static

[1] A recent working paper by Gowrisankaran and Rysman (2005) proposes such an approach.

techniques. For well defined problems, such probabilities can be conditioned on a set of known state variables and the underlying individual preference parameters. With product-level data, estimation of market shares requires the integration of the choice probabilities over the distribution of consumer tastes.

The idea that the solution to a complicated dynamic choice problem can often be approximated non-parametrically based on observed choices and states goes back at least to Hotz and Miller(1993). Some of its applications to Industrial Organization problems include the estimation of dynamic inventory problems (Aguirregabiria, 1999) and production functions (Olley and Pakes, 1996). It has become popular recently to estimate dynamic games (e.g. Aguirregabiria and Mira, forthcoming). The framework proposed in this paper further extends this basic idea by letting the choice probabilities be a random function that depends on the distribution of individual preferences.

Instead of estimating just the distribution of consumers' preferences, the goal is to estimate the joint distribution of consumers's preferences *and* the parameters of the reduced-form participation function. What identifies the distribution of the parameters of the participation function is the variation over time of the total number of purchases and its correlation with observed measures of quality. The estimated form of this function has no structural interpretation and is specific to the observed dynamic distribution of quality and prices. It therefore has predictive power, as long as the intertemporal distribution of products' characteristics doesn't change. Its incorporation into the estimation algorithm, though, allows the estimation of consumers' preferences in a manner that is consistent with an underlying dynamic optimization model. If desired, the estimated preference parameters can then be used to compute the whole dynamic problem when evaluating counterfactual equilibria, by imposing explicit structure on the dynamics of the state space.

The paper proceeds as follows: in the next section, the model of individual behavior that is assumed to be generating the data is described. Aggregation and the general estimation strategy are also discussed. The methodology is then applied to aggregate data from the digital cameras market, which is a good example of a market with rapidly changing quality and prices; details of the algorithm are discussed and salient implication of the estimates are discussed. In particular, estimates of the model are compared with estimates of a static model, which as expected yield much lower taste parameters. In the concluding section the contributions of the paper are reviewed, as well as some of its limitations.

2 A dynamic model of demand for durable differentiated products

2.1 The individual optimization problem

Given the underlying behavior of firms, observed data are assumed to be generated by a dynamic model of individual behavior, in which consumers choose optimally both the time and the quality of the products they buy. When individual i buys product j, the lifetime utility derived from this purchase is given by:

$$u_{ijt} = \xi_{ij} + \gamma_i x_j - \alpha_i p_{jt} + \epsilon_{ijt} \tag{1}$$

Following the usual representation of preferences, let ξ_{ij} be an unobserved product attribute common to all consumers who purchase product j at time t; p_{jt} is the price and x_j the vector of observed characteristics of product j. Notice that preference parameters $\{\xi_{ij}, \alpha_i, \gamma_i\}$ vary across consumers, so that resulting choices are correlated. If we let $\xi_{ij} = \xi_j + \sigma_\xi \varepsilon_{i\xi}$; $\gamma_i = \gamma + \sigma_\gamma \varepsilon_{i\gamma}$ and $\alpha_i = \alpha + \sigma_\alpha \varepsilon_{i\alpha}$, where $\varepsilon_{i.}$ is drawn from a known *iid* distribution F_ε, we can rewrite (1):

$$\begin{aligned} u_{ijt} &= (\xi_j + \gamma x_j - \alpha p_{jt}) + (\sigma_\xi \varepsilon_{i\xi} + \sigma_\beta \varepsilon_{i\gamma} x_j - \sigma_\alpha \varepsilon_{i\alpha} p_{jt}) + \epsilon_{ijt} \\ &= \delta_{jt}(x_j, p_{jt}; \theta_0, \xi_j) + \mu_{ijt}(x_j, p_{jt}; \theta_1, \varepsilon_i) + \epsilon_{ijt} \end{aligned} \tag{2}$$

where $\theta_0 = \{\gamma, \alpha\}$; $\theta_1 = \{\sigma_\xi, \sigma_\gamma, \sigma_\alpha\}$ and $\varepsilon_i = \{\varepsilon_{i\xi}, \varepsilon_{i\gamma}, \varepsilon_{i\alpha}\}$. This means that utility from purchasing any product has three components: the mean utility δ_{jt} which is common to all consumers; a term μ_{ijt} which captures the variability of tastes for quality across consumers and an idiosyncratic product-consumer shock ϵ_{ijt}.

Given the durability of products, the consumer i will buy a product at period t, if the maximum attainable utility is greater than a reservation utility R_{it} which will depend on the expected evolution of available quality and prices in the market in the near future. Let $v_{it} = \max_j\{u_{ijt}(\cdot)\}$ be the maximum lifetime utility consumer i can get from any product purchased at t; the reservation utility of the consumer is the value of not purchasing anything at time t and waiting until the next period to evaluate the problem again:

$$R_i(S_t) = 0 + \beta E\left[\max\{v_{it+1}, R_i(S_{t+1})\} \mid S_t\right] \tag{3}$$

where the instantaneous utility of the outside good has been normalized to zero and S_t is the vector of relevant states, which include all the variables that affect the consumers' purchase decision. Specifically, it should include all the variables that consumers use to predict the evolution of available products' prices and characteristics.

Potentially, the function (3) can be computed numerically, provided that the parameters in (1) and the distribution $F_{v_{it}}$ of v_{it} are known, and that we specify accordingly an intertemporal process for S_t. Not only can this computation be quite complicated given the dimensionality and non-stationarity of the state space but, more importantly, it would require knowledge of the unobserved product attributes ξ_{jt}, which are usually solved for as part of the standard estimation algorithm. As explained below, the methodology proposed in this paper circumvents this computation. For now notice that the probability that the consumer buys any product at time t, which will be denoted h_{it}, can be subsequently obtained from the known distribution of v_{it}:

$$\Pr[purchase] \equiv h_{it}(S_t) = P[v_{it} > R_i(S_t)] = 1 - F_{v_{it}}(R_i(S_t)) \tag{4}$$

where $F_{v_{it}}(R_i(S_t))$ is the pdf of v_{it} evaluated at $R_i(S_t)$; this pdf will depend on the assumed distribution of the unobservables. Notice that the solution of the dynamic optimization problem is a participation function $h_{it}(.)$ that depends on the the state vector S_t.

The probability that the consumer buys specifically product j, denoted h_{jit} is obtained by multiplying the probability of purchase by the product-specific conditional purchase probability. Let $\Im_t$ be the set of available products at time t:

$$h_{jit}(\cdot, \varepsilon_i) = h_{it}(S_t) Prob[u_{ijt} \geq u_{ikt} \forall k \in \Im_t] \tag{5}$$

2.2 The distribution of preferences and the dynamic behavior of individual consumers

Obtaining an estimable equation from (5) requires that we specify a parametric assumption for the distribution of ϵ. Throughout the remainder of the paper it will be assumed that these idiosyncratic product- and consumer-specific shocks are draws from an *iid* extreme value distribution. As a consequence the product-specific conditional choice probabilities will have a closed form solution; also, given any set of model parameters and observed demand, the unobserved product attributes are identified. These unobserved attributes are part of S_t and therefore necessary for the solution of (3).

Given this restriction, the expected value of participating in the market by purchasing *any* available product at any point in time is given by the known "inclusive value" formula:

$$r_{it}(\cdot) = \log\left[\sum_{k\in\Im_t} \exp(\delta_{kt}(\cdot) + \mu_{ikt}(\cdot))\right]$$

Therefore, r_{it} is a sufficient statistic for the distribution of v_{it} as shown by Melnikov (2000)[2]. The reservation utility (3) can then be rewritten as follows:

$$R(S_{it}) = 0 + \beta E\left[\max\{r_{it+1}, R(S_{it+1})\} \mid S_{it}\right]$$

Notice then that S_{it} comprises the set of variables used by consumer i to predict the evolution of r_{it}. In general, it should include the variables that the firms use to make their product innovation and pricing decisions. In most cases, like in the application below, the number of competing firms and products is too large and ad-hoc restrictions have to be imposed on the dimension of this set. For example, in the application below it will be assumed that r_{it} is first-order Markovian, so that its expected evolution depends only on its current value.

Given R_{it} above, the demand equations (4) and (5) can be written as a function only of r_{it} and the state variables that are used by the consumers to predict its evolution, S_{it} :

$$\Pr[purchase] \equiv h_{it}(S_{it}) = P[v_{it} > R(S_{it})] = 1 - F_{v_{it}}(R(S_{it})) \qquad (6)$$

$$h_{jit}(\cdot, \varepsilon_i) = h_{it}(S_{it})\frac{\exp(\delta_{jt}(\cdot) + \mu_{ijt}(\cdot, \varepsilon_i))}{\exp(r_{it}(\cdot, \varepsilon_i))} \qquad (7)$$

Finally and as it's usual in BLP-style models, identification of the parameters of the static utility function, which for example are crucial for the computation of elasticities, requires the specification of distributional assumptions on the joint distribution of the unobserved product characteristics ξ_j and instruments. Specifically, it will be assumed that unobserved product attributes at each point in time are uncorrelated with non-price attributes. The econometric identification of the distribution of these taste parameters, even from product-level data is made possible by the variation over time

[2]See appendix. Another distributional restriction –e.g. normality –would imply a larger state vector and, therefore, more complicated intertemporal transitions.

of purchasing behavior *across* products. The additional identification of the participation function $h(.)$ and its distribution follows from the observed variation in *aggregate* sales, which is an information that is somehow left unused in standard static models.

2.3 Aggregation and estimation

Obtaining the predicted aggregate market demand for each product, requires the integration of (7) across consumers. Besides the idiosyncratic shock ϵ_{ijt} which has already been integrated out in (7), consumers differ in their valuation of quality depending on the realization of the vector ε_i which is distributed according to F_ε. Integrating over the distribution of ε, we can obtain the predicted market share of product j, s_{jt}, given the value of θ:

$$s_{jt}(\theta_0, \theta_1) = \int \left[h_t(S_{it}(\theta_0, \theta_1, \varepsilon)) \frac{\exp(\delta_{jt}(\theta_0) + \mu_{jt}(\theta_0, \theta_1, \varepsilon))}{\exp(r_t(\theta_0, \theta_1, \varepsilon))} \right] dF_\varepsilon \tag{8}$$

Where the distribution $F(.)$ has to account for the changing distribution of preferences as consumers buy a product and leave the market (or return after a while to buy a replacement). The equation above also recognizes the fact that individual states may depend on the parameters θ and the unobserved heterogeneity.

An estimator of the parameter vector $\theta = \{\theta_0, \theta_1\}$ can be obtained by equating the predicted and observed demand:

$$M_0 s_{jt}(\theta_0, \theta_1) = Q_{jt} \tag{9}$$

where Q_{jt} is the observed demand for product j at t and M_0 the assumed initial market size[3]. Estimation requires the interaction of the unobserved product attributes ξ_j with a matrix of instruments, which depending on the case may include observed product characteristics at different points in time.

As in Berry (1994), estimation of (9) could be based on the solution for the mean utility levels $\{\delta_j\}$ in (9), so that the estimation of θ_0 can be "concentrated out" from the maximization of the criterion function. The complicated integration in (8) can be avoided by using simulation techniques. Assuming that ε is a vector of *iid* standard normal disturbances (so that σ

[3]The model can be estimated assuming that consumers who purchase a product leave the market for ever or that they return with a probability that can be estimated. In any case, the exogenous initial market size must be known.

corresponds to its standard deviation), N draws $\{\varepsilon_n\}_{n=1...N}$ of the distribution can be simulated to compute a consistent estimator of the integral:

$$s_{jt}(\theta_0, \theta_1) \approx \frac{1}{N}\sum_{n=1}^{N}\left[\psi_{nt} h_{nt}(S_{nt}(\theta_0, \theta_1, \varepsilon_n))\frac{\exp(\delta_{jt}(\theta_0) + \mu_{njt}(\theta_0, \theta_1, \varepsilon_n))}{\exp(r_{nt}(\theta_0, \theta_1, \varepsilon_n))}\right] \tag{10}$$

Where $\psi_{n,1} = 1$ and $\psi_{n,t>1} = \psi_{n,t-1}(1 - h_{n,t-1})$ is the probability that consumer n is still in the market in period t[4]. Given a value of θ and the simulated draws, the vector $\{\xi_j\}_{j\in\Im_t}$ can be solved for each t, and a numerical algorithm can the be used to find the vector θ^* that minimizes a criterion function. For each "simulated" consumer, h_{it} can be computed by solving numerically equation (3) and computing the implied probability (4).

Notice, though, that the methodology described above is computationally difficult, if not impractical. Computing the mean utility levels $\{\delta_j\}$ from (9) requires the solution of a fixed point for *each* simulated consumer. Each fixed point contains an additional fixed point algorithm to obtain the participation probability, which is the result of a dynamic problem that depends on the transition of $\{\delta_j\}$ which is what we need to compute in the first place. Moreover, even if the fixed point used to solve for the vector $\{\delta_j\}$ from (9) is computable a some points, it is not clear whether, under these conditions, it is a contraction across all the relevant parameter space.

A less demanding approach that circumvents the nested "computation" of the dynamic problem in the predicted market share can be devised by imposing a parametric structure on h_{it} and estimating its parameters as part of the whole algorithm. Given that the transition probability of r_{it} depends only on S_{it}, the participation probability, which has no closed form solution, can be approximated as a function of individual states and some parameters:

$$h_{it} = \tilde{h}_i(S_{it}(.); \tilde{\theta}_{2i}) \tag{11}$$

In (9) it is explicit that the function $\tilde{h}(.)$ varies across consumers because it is the result of the individual's underlying dynamic optimization problem. The individual participation function is indexed by the the individual parameter vector $\tilde{\theta_{2i}}$, which is to be estimated.

[4]If we allow consumers who have already bought a product to return to the market with some probability q_{it} then $\psi_{n,t>1} = \psi_{n,t-1}(1-h_{n,t-1})q_{it}$. Notice that this probability can be assumed to be an exogenous parameter to be estimated as part of the estimation algorithm or it can be further parameterized.

We can replace (11) in (10) and then solve for the mean utilities $\{\delta_{it}\}$ from (9). More specifically, for computing the predicted market shares the methodology described above requires the simulation of N draws of the vector ε_i and the computation of the simulated integral:

$$\tilde{s}_{jt}(\theta_0, \theta_1, \tilde{\theta}_2) \approx \frac{1}{N} \sum_{n=1}^{N} \left[\psi_{nt} \tilde{h}_n(S_{nt}(.); \tilde{\theta}_{2n}) \frac{\exp(\delta_{jt}(\theta_0) + \mu_{njt}(\theta_0, \theta_1, \varepsilon_n))}{\exp(r_{nt}(\theta_0, \theta_1, \varepsilon_n))} \right] \tag{12}$$

where $\tilde{h}_{nt}(\cdot)$ in (10) is given by (9). The estimation should proceed as described above with the difference that the estimation of $\tilde{\theta}_2$ will now be part of the optimization algorithm.

From (12) above it should be clear that potentially we could estimate a the consumer-specific parameters $\tilde{\theta}_{2i}$ if we had enough information. That is, we could estimate a different participation function for each simulated consumer so that we could infer the distribution of $\tilde{\theta}_2$ from the estimated $\{\hat{\tilde{\theta}}_{2i}\}$. Of course, in a realistic environment with product-level data, identifying a large number of consumer-specific participation functions is impossible. Such identification is possible if a more restrictive discrete distribution of consumers is adopted. For example, if only K consumer types are allowed, the predicted market shares would be given by:

$$\frac{q_{jt}}{M_t} = s_{jt} = \left(\sum_{k \in K} \psi_t^k \tilde{h}^k(S_t^k; \theta_2^k) \frac{\exp(\delta_{jt}^k)}{\exp(r_t^k)} \right)$$

where,

$$\psi_{t+1}^k = \psi_t^k (1 - \tilde{h}^k(r_t^k; .))$$

If K is small above, then we may be able to identify individual parameters θ_2^k for each consumer type and distribution parameters $\{\psi_0^k\}$ as in Carranza (2004)[5]. For a continuous distribution of consumer types that is approximated via a number N of simulations, what can be done is let individual participation functions be approximated non-parametrically from the underlying variation in individual taste parameters.

Specifically, notice that the shape of the participation function in (7) is determined by individual taste parameters $\{\theta_0, \theta_1\}$ that don't change across consumers and random individual taste shocks $\{\varepsilon_i\}$ that follow a known

[5]See also Berry, Carnall, Spiller (1997) for a static demand model with discrete types that can be estimated from product-level data.

parametric distribution. Therefore, we can approximate the individual participation decisions using a flexible non-parametric function of individual preference shocks, as follows:

$$h_{it} = \bar{h}(S_{it}(.); \tilde{\theta}_{2i}(\theta_0, \theta_1, \varepsilon_i)) = \bar{h}(S_{it}, \theta_0, \theta_1, \varepsilon_i; \theta_2) \tag{13}$$

Where θ_2 is a parameter vector that is common to all consumers and that captures the effects of individual taste shocks on individual participation decisions. We can, for example, approximate $h(.)$ using a logistic functional with polynomials of the variables, as in the application below.

By replacing (13) in (10), we can obtain an expression for the demand system as a function only of data and the parameters to be estimated:

$$s_{jt}(\theta_0, \theta_1, \theta_2) \approx \frac{1}{N} \sum_{n=1}^{N} \left[\psi_{nt} \bar{h}(S_{it}, \theta_0, \theta_1, \varepsilon_i; \theta_2) \frac{\exp(\delta_{jt}(\theta_0) + \mu_{njt}(\theta_0, \theta_1, \varepsilon_n))}{\exp(r_{nt}(\theta_0, \theta_1, \varepsilon_n))} \right] \tag{14}$$

$$\psi_{n,t>1} = \psi_{n,t-1}(1 - h_{n,t-1}) \tag{15}$$

$$M_t s_{jt}(\theta_0, \theta_1, \theta_2) = Q_{jt} \tag{16}$$

By setting $\psi_{n,1} = 1$, the system of equations (14)-(16) can be used to estimate parameters $\{\theta_0, \theta_1, \theta_2\}$. As usual, the equality (16) implies a vector of unobserved product attributes $\{xi_j\}$ that are assumed to be orthogonal with a matrix of instruments, so that a method of moments can be used for estimation, as will become clear in the application below.

Notice, finally, that the parameters of interest are really θ_0 and θ_1. Parameters θ_2 are secondary in the sense that after obtaining θ_0 and θ_1 –the deep preference parameters –, the "true" participation rate can be obtained for every simulated consumer by computing (3) and (4) directly for every simulated consumer[6]. Nevertheless the inclusion of θ_2 allows the estimation of θ_0 and θ_1 in a manner that is consistent with an underlying dynamic model[7]. Also, as indicated before, θ_2 are non-structural parameters that are

[6]Such computation requires specific restrictions on the transition of the states, which are not required for the estimation here. It also requires first the specification of the discount rate. On the other side, if data are rich enough, such discount rate, and any other deep parameter of the dynamic problem, can be estimated by equating predicted participation rates to observed rates.

[7]A potential iteration between the full model and its reduced form approximation is apparent. At the true value of the parameters, the predicted reduced form of the participation must be consistent with its computed version

not necessarily stable across counterfactual equilibria. Therefore, simulation of counterfactual equilibria and policy analysis require that the complete dynamic model be computed. Such computation is not trivial, given the non-stationarity of the the state variables, and is beyond the scope of this paper.

If it is assumed that consumers are homogeneous up to the extreme value preference shock, the observed participation rate, i.e. the total number of purchases over the market size, is the sample analog of the function $h(.)$. In that case, the function $\bar{h}(.)$ and the structural preference parameters can be estimated separately. In the general case, the described technique amounts to a "nested" non-parametric estimation, with a non-parametric estimation of the participation rate as part of the "wider" estimation algorithm.

The idea of the estimation is simple: if the transition of attainable payoffs can be conditioned on a known set of states, the solution to the dynamic problem can be conditioned solely on these states. Observed purchase behavior can therefore be correlated with these states to control for the dynamic concerns of consumers, within a framework that is entirely similar to the standard BLP-style techniques. Knowledge of the relevant states though, requires in general a specification of the supply model that generates the evolution of prices and quality. On the other hand, making the problem tractable requires that the number of relevant state variables be kept at the minimum. This tradeoff between tractability and formal precision is discussed in the following subsection.

2.4 Industry dynamics and consumer behavior

Estimating static models of demand usually requires no assumptions regarding the structure of supply, because consumers' actions should not be affected by the underlying interaction of firms. In contrast, estimating a dynamic model of demand requires that the expectations of consumers regarding the evolution of the industry be taken into account. In the framework discussed above, consumers should use available information to construct expectations regarding the evolution of attainable payoffs as indexed by r_{it}.

The evolution of the industry, in general, is the result of strategic decisions of firms. Rational consumers should therefore condition their decisions on whatever variables the firms are conditioning theirs. For example, in a Markov perfect equilibrium firms' strategies depend on the distribution of available products along the space of goods' prices and characteristics. Consumers should then, in principle, condition their purchase decisions on this distribution. If the number of available products and products' characteristics is large, though, this is a very complicated object and the state space becomes basically intractable.

Limiting the size of the state space requires that restrictions be imposed on the assumed model of supply, or on the way consumers construct expectations of the the evolution of attainable payoffs. In the application below, for example, it will be assumed that r_{it} is first-order Markovian, in the sense that consumers use only r_{it} to predict its expected transition. As shown afterwards, the assumption proves to be justified by the obtained results which indicate that the Markovian assumption is a very precise description of the intertemporal distribution of r_{it}.

Such simple Markovian assumption can be justified by a competitive assumption, as in Carranza (2006). Specifically, if firms actions are marginal, quality and pricing decisions depend only on the expected evolution of this "inclusive value", which is taken as given. On the other hand, it could also be justified by assuming that consumers use only a subset of the information set to construct expectations, due to the existence of costs to acquire information or some other kind of bounded rationality as discussed, for example, by Krusell and Smith (1996).

It should be noted that the need of these simplifying restrictions is not caused by the proposed estimation technique. If the complete dynamic problem of every simulated consumer was to be solved along the estimation algorithm, even stronger assumptions regarding the specifics of the state space would be required. The choice of variables and should, nevertheless, be informed by the specifics of the case and justified under the light of the obtained results.

2.5 The dynamic vs the static model

Given the extreme value assumption on ϵ, the implied individual participation rate is given by the following formula:

$$h_{it} = \frac{exp(r_{it}(.))}{exp(R(S_{it})) + exp(r_{it}(.))} = \frac{1}{1 + exp(R(S_{it}) - r_{it})}$$

so that the model above collapses to the standard BLP-style model when $R = C$, where C is a constant usually assumed to be zero. In such case, the individual participation function is given by

$$h_{it} = \frac{exp(r_{it}(.))}{1 + exp(r_{it}(.))} = \frac{1}{1 + exp(-r_{it})}$$

This static model can account for the increasing size of the market, because as r_{it} increases over time, consumers who previously gained more by staying out of the market are enticed to buy a product. Tacitly though, the model assumes that those consumers who chose not to buy don't value the characteristics of the available products much, whereas the truth may be that they value the product a lot and are waiting for better and cheaper products to be introduced later.

Given the outside static payoff normalized to zero, if consumers choose to optimally delay the purchase of a product, that means that the reservation utility is positive. As can be seen in the formulas above, given some choice probabilities, the presence of a positive reservation utility $R_{it} > 0$ implies higher values for r_{it} which would imply higher taste coefficients than the ones obtained from a model where $R_{it} = 0$[8]. How significant this bias is will depend on the specific case; presumably, if there is a significant increase in quality over time, estimates of taste parameters obtained from a static model should be biased downwards. As shown below, in the digital cameras example this bias is substantial.

3 An application: the market for digital cameras

In this section we use the described technique for estimating the demand for digital cameras using a panel of sales, prices and characteristics of digital cameras sold in the U.S. between 1998 and 2001 (for a detailed description of the data, see Carranza, 2006). Digital cameras are an almost ideal example of a durable good with rapidly improving quality and decreasing prices. In addition, quality is easily defined. During the time-span of the sample, in fact, camera resolution was by far the most important camera characteristic. Therefore, in the application below will be mostly based on a model with this one observable characteristic.

[8]This discussion presumes a given price coefficient and highlights the difficulty of identifying separately price and taste coefficients. In the application an explicit pricing equation is adopted to help with this identification.

There's a couple of salient features in the data that the model has to be able to reproduce. Specifically, there's an increasing pattern of sales over time with notorious spikes during the Christmas season. This seasonality is important because it cannot be justified by any observed variation in prices or quality. Additionally, the presence of such seasonal effect will help identifying the unobserved heterogeneity of consumers because it will end up being the result of the differential response of heterogeneous consumers to observed changes in available quality and prices.

3.1 Specification and estimation

The framework proposed in this paper is intended to be used with product-level data. Therefore, the identification of the distribution of consumers' preferences is based on a pre-specified parametric distribution. As explained below, identification of the model in this environment will require some more structure to compensate for the fact that only product-level information is available. Notice, though, that the methodology is a natural extension of BLP-style techniques, and so it accommodates extensions to environments with more detailed matching and non-matching micro data.

The lifetime utility of consumer i when purchasing model j is given by:

$$u_{ijt} = \xi_j + \gamma_i x_j - \alpha p_{jt} + \epsilon_{ijt} \tag{17}$$

Parameters have the same interpretation as before. The vector $\gamma_i \equiv \{\gamma_{0i}, \gamma_{1i}...\}$ contains a constant and taste parameters corresponding to observed characteristics. We will assume that the resolution taste coefficient γ_{1i} varies across consumers and is distributed normally, i.e. $\gamma_{1i} \sim N(\gamma_1, \sigma_\gamma^2)$; the other coefficients are assumed to be identical across consumers. We will consider alternative models with one observed attribute (resolution) and two observed attributes (resolution and optical zoom).

As explained above, under the given distributional assumptions of the unobservables, the solution to the dynamic optimization problem of consumers will depend on the individual's "inclusive value" $r_{it}(\cdot) = \log\left[\sum_{k\in\Im_t} \exp(\delta_{kt}(\cdot) + \mu_{ikt}(\cdot))\right]$, where $\Im_t$ is the set of available products at time t. Due to the heterogeneity of the consumers' taste for camera resolution, this value is different for every individual.

It is assumed that r_{it} follows a first order Markov process. The key insight is that if the expected evolution of this variable depends on its current value only, then the time-changing reservation utility will also depend only on this value. As a consequence, the participation rate will depend on this scalar-valued state variable. We will adopt a reduced form for this function; its specification should account for the correlation between the function and the structural preference parameters that generate it and that vary across consumers.

Under the stated assumptions, product j market share can be computed by simulating the errors $\varepsilon_i \sim N(0, \sigma_\gamma^2)$:

$$\tilde{s}_{jt}(\theta) = \frac{1}{N} \sum_{n=1}^{N} \left[\psi_{nt} \tilde{h}_{nt}(.) \frac{\exp(\delta_{jt}(\theta) + x_j \sigma_\gamma \varepsilon_{\gamma n}))}{\exp(r_{nt}(\theta, \varepsilon_n))} \right] \tag{18}$$

where $\delta_j = \xi_j + \gamma x_j - \alpha p_{jt}$ and $\tilde{h}_n t$ is the reduced form for the participation function $h(r_{nt})$. The parameter σ_γ is now to be estimated and $\psi_{n,t>1} = \psi_{n,t-1}(1 - h_{n,t-1})$, as defined before. As in last section problem, we can adopt a logistic approximation for h_{it}:

$$\tilde{h}_{it} = \frac{1}{1 + \exp(\pi_{0i}(\theta_0, \theta_1, \varepsilon_i) - \pi_{1i}(\theta_0, \theta_1, \varepsilon_i) r_{it})}$$

where the correlation between the reduced form and the "deep" preference parameters is accounted for. If we let the parameters of the reduced form be linear functions of the preference parameters, we get[9]:

$$\tilde{h}_{it} = \frac{1}{1 + \exp(\pi_0 + \pi_1 \varepsilon_i - \pi_2 r_{it} - \pi_3 \varepsilon_i r_{it})} \tag{19}$$

Where the participation rate ends up being just a flexible functional of individual covariates.

Parameters $\{\gamma, \alpha, \sigma_\gamma, \pi_0, \pi_1, \pi_2, \pi_3\}$ can be estimated by interacting implied mean utility levels ξ_j with instruments[10]. Following previous literature (e.g. Berry, Levinsohn and Pakes, 2004), a pricing model will be added to the demand model above to help with the identification of the price parameter α, given the lack of more detailed camera purchasing data. Specifically, prices will be assumed to be set following a static first order condition as in

[9] Notice that all we care about is the correlation of the function with the inclusive value and the error term ε_i since average preference parameters don't vary across consumers.

[10] Notice that estimation of parameters γ and α can be obtained directly from the first order conditions of the GMM optimization problem to make the algorithm faster.

Carranza (2006) and marginal cost will be assumed to be linear in product characteristics and constant in quantity:

$$p_{jt} = mc(x_{jt}, \eta_t) + \frac{1}{\alpha} = x_j \eta_t + \xi_{jt}^{mc} \tag{20}$$

where cost parameters η_t are allowed to change over time, to account for technological progress. To avoid the difficulties implied by a dynamic model of multiproduct pricing, (17) is obtained under the assumption that each product is marginal and therefore firms can disregard potential cannibalization effects (for details on this see Carranza, 2006).

3.2 Results

Estimates of the model are displayed in Table 1 with standard errors obtained using bootstrap techniques. Eight versions of the model were estimated: four specifications, each with one observed characteristic (resolution) and two characteristics (resolution and optical zoom). As a benchmark , the model with non-random coefficients ($\sigma_\gamma = 0$) was estimated (see models I). The full model corresponds to models II, while a restricted model with $\pi_3 = 0$ corresponds to models III. The later specification, in which no interaction is allowed within the participation function of average quality level and individual heterogeneity, was the preferred one, because the unrestricted model (II) implied a very unreasonable distribution of market participation decisions and had computational problems. Models IV correspond to a BLP-style static model as described in section 2.5.

Comparing the results of the model with constant taste coefficients (I) with the model with random coefficicients (II and III), it can be seen that the deep structural preference parameters γ_1 and γ_2, which are all positive and precise, are similar across specifications. The variation in parameter γ_0 can be attributed to the difference in the underlying outside option implied by each specification. Moreover, results indicate that the data are in general consistent with a very low randomness of the taste parameters γ_{1i}. The reason why a model with random coefficients is generally better to fit a panel of sales data is that it takes advantage of the correlation of market shares of similar products over time. In the digital cameras market, shares are very low across the board and therefore aren't consistent with any meaningful unobserved correlation of choice behavior.

Estimates of the parameters $\{\pi_0, \pi_1, \pi_2, \pi_3\}$ obtained from the dynamic models (I, II and III) imply that the participation function is positively correlated with average consumer's valuation of "quality", so that higher-valuation consumers adopted cameras faster. Nevertheless, when random coefficients are incorporated (II and III), estimates imply that higher than a average taste for camera resolution has a negative effect on the purchase probability. Notice that identification of the features of the participation function is based on the aggregate purchasing behavior and its covariation with individual products' purchasing behavior.

The most salient result of the estimation with respect to the existing literature, is the contrast between the estimates of the dynamic models (I, I and III) and the estimates of the static BLP-style model (IV). As indicated in the introduction, the use of a static model in markets with rapidly changing quality may lead to biased estimates of preference parameters because they implicitly assume that nonparticipating consumers don't value available products enough to buy them, whereas the truth may be that they do value them but are optimally delaying the purchase, waiting for a better deal in the future.

As illustrated in section 2.5, the widely used BLP static approach is a particular case of the general dynamic model, with a restricted participation function; in terms of the chosen specification of the model, a BLP-style model implies $\pi_0 = \pi_1 = \pi_3 = 0$ and $\pi_2 = 1$. As expected, results obtained from this static model imply a resolution parameter γ_1 that is significantly lower than the estimates obtained from *all* the dynamic specifications of the model. Notice that price parameters, which are identified from the pricing equation are the same across specifications, so that the differences in taste parameters reflect the essential difference between the static and the dynamic model. The magnitude of this bias will become clear in the following section, where a myopic version of model III will be simulated.

As discussed before, the improvement of CCD chips over time has been constant and quite dramatic, as well as the fall in prices. As a consequence of such change, consumers have strong incentives to delay the purchase of a digital camera. Its is therefore no surprise that the static model yields a lower preference for CCD resolution than the dynamic model. On the other hand, notice that the estimate of the optical zoom coefficient doesn't differ between the static and dynamic versions of the model; it is also the case that the technology of lenses hasn't changed significantly over the last decade. Results are therefore consistent with the premise, that a static specification of consumer behavior is a misleading approach in environments with rapidly changing quality that imply a nontrivial dynamic problem for consumers.

3.3 Fit of the model and additional results

The exercises presented in this subsection are mainly based on the estimates corresponding to the model with two observed characteristic and with $\omega_3 = 0$ (model III).

Notice first that in the model above the unobserved errors enter nonlinearly; also, by definition, they are correlated with prices. Therefore, evaluating the fit of the model requires the simulation of the unobserved product attributes from the empirical joint distribution of unobserved and observed attributes. As seen in Figure 1, the fit of the model is quite satisfactory. Specifically, the model can account very well for the endogenous variation in market size over time. The figures display observed and a 95% interval of the simulated aggregate market participation (i.e. number of purchases over over total number of potential costumers) obtained from the estimates described above. There is a seasonal variation in market participation in the last quarter of every year that is captured well by the model, despite the fact that the model doesn't include any specific seasonal effect.

Notice that such variation in market size as a response to available quality and prices is driven by the participation function $h(r_{it})$ which is the result of a dynamic optimization problem. Given that the quality "indices" r_{it} are stochastic, the function itself is stochastic. Figure 2 illustrates a 95% interval for the distribution of the participation function across consumer types for the model in which $\pi_3 = 0$, obtained from the estimated reduced form. Figure 3 illustrates the participation function obtained from the estimates in which $\pi_3 \neq 0$. In this case, the interaction between overall quality and the individual valuations inhibits low valuation costumers to buy at all in any period. It should be clear now that the reason why the restricted model with $\pi_3 = 0$ was estimated is that it is difficult to think how such behavior can correspond to any type of dynamic optimization.

Notice that, despite the relatively small significance of consumer heterogeneity, estimates imply that consumers with lower taste for quality increase endogenously their participation around the Christmas season. In the model –in accordance with intuitive understanding of consumers' behavior –consumers' preferences imply different intertemporal adoption patterns. As seen in Figure 4, the dispersion of costumers' valuations has a markedly seasonal pattern due to the increased participation of consumers with a lower taste for quality. As was to be expected, estimates also imply that the average valuation of quality (resolution) of buyers decrease systematically (even if very little) during the Christmas season, as illustrated in Figure 5.

The methodology has emphasized the incentives of costumers to time optimally the purchase of a new camera. It is useful to see the magnitude of these incentives on the consumers' behavior by simulating the model under the assumption that consumers buy a new camera as soon as their utility is positive[11]. As indicated above, the computation of this myopic version of the model is also useful to ascertain the magnitude of the potential bias of the static estimation of the model. Since costumers who buy any product drop off the market, the effects of such myopic behavior depend on the time at which consumers start purchasing. In Figure 6 the evolution of sales is illustrated under the assumption that costumers "start" purchasing in the first quarter of 1999. Given the high valuation of quality, most potential consumers purchase a camera right away and the market shrinks dramatically inhibiting endogenously posterior sales, which is indicative of how significantly different the estimates of the static and dynamic models are.

As indicated in the introduction, the estimation of the demand model with heterogeneous consumers yields cross substitution patterns that depend on the characteristics of the products –as opposed to models with homogeneous consumers that imply cross substitutions that end up only depending on market shares. In our present example with only one observable characteristic cross substitution elasticities depend mainly on the resolution of the camera. Moreover, changes in the price of a camera model have dynamic effects on the demand of other products as the price change has the effect of inducing intertemporal reallocation of purchase decisions. As expected cross price elasticities are positive across products and across subsequent time periods.

In this dynamic environment the way in which the elasticities are computed has to be clearly defined. For the illustrative example below, price elasticities are computed with respect to price deviations in one product in one period, keeping everything else constant, including the price of the product in subsequent periods. Take as an example a price change of the Sony DSC-S70, which is a 3.1 megapixel camera, in the second quarter of 2000; the own-price price elasticity of such price change in that period is -0.05. Cross

[11] Specifically, the simulation uses the preference parameters obtained from the dynamic model, but uses the participation function corresponding to the static BLP-style model

elasticities, which are all positive, have a very weak relationship with market shares and are closely related to the characteristics of cameras. A regression of the computed cross elasticities on the absolute values of the differences of resolutions and market shares illustrates the point: as shown in Table 2, the coefficient on such regression of the resolution difference is negative and highly significant; the coefficient of the market share difference is negative but insignificant. If we do the same exercise for the elasticities computed from the model with no consumer heterogeneity, there is no relationship between cross elasticities and product characteristics, just like in a simple logit model.

Finally, notice that the estimated preference parameters could potentially be used to compute directly the dynamic model without using the reduced form of the participation function. In fact, deep dynamic parameters, such as the discount rate, could be potentially estimated in a second stage of the estimation by matching observed data to the predictions of the model, given the taste coefficients estimated above. Nevertheless, such computation is not trivial and is beyond the scope of this paper. Specifically, it requires the adoption of a specific bounded transition for the state variables and the use of an algorithm to compute the value function of the problem, given the nonstationarity of the state space.

3.4 Summary of the estimation algorithm and further methodological remarks

Estimation of the model is based on the assumed statistical properties of the unobserved product characteristics. For a given parameter vector θ_0 this unobserved attribute ξ_{jt} for each product j at each point in time t is obtained from the solution of the set of non-linear equations implied by (10). The solution to these system can be obtained using the fixed point algorithm implied in[12]:

$$\xi = \xi + log(\bar{s}) - log(s(\theta_0, .))$$

Where $\bar{s}$ is the vector of observed market shares and $s(\theta_0)$ is the vector of predicted market shares, given θ_0.

These implied unobserved attributes can then be used to compute a set of predicted moments. Specifically, they should be interacted with a matrix

[12]See Berry, 1994. The use of logs improves the computational speed, though in some cases, depending on the form of the predicted market shares, it may fail to converge

of instruments. Usual instruments include the observed characteristics of products at each point in time which are assumed to be orthogonal to the unobserved attribute. The computed vector of predicted moments can be used to construct a metric of the mismatch of the model and the data, usually a quadratic form, as in a usual GMM, or minimum distance estimator[13]. The algorithm looks then for the vector θ^* that best matches the predictions of the model to the data, minimizing numerically the criterion function. Standard errors of the estimates are obtained numerically either using the delta method or bootstrapping the sample.

The fixed-point algorithm above can be used to solve directly for the vectors of mean utilities $\{\delta_t\}$. If the utility function is linear in the characteristics of the product, then the first order conditions of the minimization problem are linear in the mean preference parameters and therefore their estimation can be separated from the minimization algorithm.

Given the estimated parameters, the dynamic model implied by (3) can be computed and matched to the observed evolution of the market size to estimate deeper dynamic parameters, such as the discount rate. Moreover, notice that the described methodology is totally in line with demand estimation in the tradition of Berry, Levinsohn and Pakes (1995) except that a participation function is used to account for the endogenous timing of purchases. The methodology allows therefore for the utilization, when available, of non-matching information regarding the distribution of relevant variables, such as income. Specifically, when computing the simulated integral in (10), draws are taken from the empirical distribution of the non-matching observed variable. Also, aggregate information like the one in the digital cameras example can be combined with micro choice information –if available –very much in the manner described in Berry, Levinsohn and Pakes (2004).

Finally, notice that the basic estimation mechanism can also be used to estimate other type of dynamic demand models, as long as the solution to the dynamic problem has a smooth solution and the state space is identified. For example, if the demand for a product depends on time-changing aggregate adoption decisions as it's the case of network products, it seems plausible to construct a smooth dynamic adoption problem. If the solution of that problem depends on a finite set of state variables (e.g. the adoption rates), then the solution to the dynamic problem should be a function of them and an estimation algorithm like the one above can be devised.

[13]As in Nevo, 2000

4 Final remarks

The methodology discussed above allows the estimation of a dynamic model of demand for durable goods that takes into account the incentives of heterogeneous consumers to time optimally the purchase of durable goods with changing quality and/or price. In environments where products' characteristics are changing rapidly over time, accounting for these dynamic incentives is crucial to prevent the biased estimation of deep preference parameters. The main difference between the described model and the standard conventional models of demand for differentiated products is the presence of an endogenous purchase probability, which can be computed as the solution of the consumer's dynamic problem.

The difficulty of computing such problem for every consumer type along the estimation algorithm was solved by using a reduced form of the optimal participation decision, taking advantage of the structure of the problem. The technique is related to recent applications of two-step estimators *a la* Hotz and Miller (1993), where a computationally difficult problem (usually a dynamic optimization problem or a strategic choice problem) is approximated via a non-parametric form that can be estimated separately and incorporated into a second step estimation of the fully structural model. It differs with respect to this literature in that the reduced form of the dynamic problem cannot be estimated separately and is therefore "nested" in the structural estimation algorithm.

The model is estimated using a data set from the digital cameras market and yields reasonable results. Results indicate that the use of standard BLP-style models in nonstationary environments may yield estimates of preference parameters with a substantial bias. The technique is otherwise totally compatible with standard techniques and implies a similar computational burden. In addition, the same ideas discussed here can be applied to the estimation of other dynamic models of demand, such as demand for network products or products that involve learning.

References

[1] Aguirregabiria V. (1999): "The Dynamics of Markups and Inventories in Retailing Firms", *Review of Economic Studies* 66, 275-308

[2] Berry, S. (1994): "Estimating Discrete Choice Models of Product Differentiation", *RAND Journal of Economics*, 25(2), 242-262.

[3] Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium", *Econometrica*, 63(4), 841-890.

[4] Berry, S., J. Levinsohn, and A. Pakes (2004): "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market", *Econometrica*, 63(4), 841-890.

[5] Berry, S., M. Carnall and P. Spiller (1997):"Airline Hubs: Costs, Markups and the Implications of Costumer Heterogeneity", Manuscript

[6] Carranza, J.E. (2004): *Empirical Analysis of the Evolution of Quality in Durable Goods' Markets*, PhD Dissertation, Yale University, New Haven, CT.

[7] Hotz J. and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models", *Review of Economic Studies*, 60, 497-529.

[8] Krusell P. and A. Smith (1996): "Rules of Thumb in Macroeconomic Equilibrium: A Quantitative Analysis", *Journal of Economic Dynamics and Control*, 20, 527-558

[9] Melnikov, O. (2000): "Demand for Differentiated Durable Products: The Case of the U.S. Printer Market", Manuscript, Yale University.

[10] Olley S. and A. Pakes (1996): "The dynamics of productivity in the telecommunications equipment industry", *Econometrica* 64(6), 1263-1298.

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

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

A Appendix: The distribution of v_{it}

This appendix basically reproduces the proof contained in Melnikov (2000). Recall that $v_{it} = \max_j\{u_{ijt}(\cdot)\}$ is the maximum attainable utility by consumer i at time t where $u_{ijt} = \delta_{jt} + \mu_{ijt} + \epsilon_{ijt}$. Let J_t be the number of available products at time t; the distribution of v_{it} is given by:

$$\begin{aligned} F_v(z) &= P[v_{it} \leq z] = P[u_{i1t} \leq z, ..., u_{iJ_tt} \leq z] \\ &= P[\epsilon_{i1t} \leq z - \delta_{1t} - \mu_{ijt}, ..., \epsilon_{i1t} \leq z - \delta_{J_tt} - \mu_{iJ_tt}] \end{aligned}$$

Given the assumption that ϵ_{ijt} is distributed *iid* according to the Type I extreme value distribution, the probability above is:

$$\begin{aligned} F_v(z) &= \prod_{j=1}^{J_t} \exp(-\exp(-z + \delta_{jt} + \mu_{ijt})) \\ &= \exp(-\sum_{j=1}^{J_t} \exp(-z + \delta_{jt} + \mu_{ijt})) \\ &= \exp(-\exp(-z) \sum_{j=1}^{J_t} \exp(\delta_{jt} + \mu_{ijt}) \\ &= \exp(-\exp(-z) \exp(\log \sum_{j=1}^{J_t} \exp(\delta_{jt} + \mu_{ijt}))) \\ &= \exp(-\exp(-z + \log \sum_{j=1}^{J_t} \exp(\delta_{jt} + \mu_{ijt}))) \end{aligned}$$

Let $r_{it} = \log \sum_{j=1}^{J_t} \exp(\delta_{jt} + \mu_{ijt})$. Then, as seen above, v_{it} is distributed Type I extreme value with mode r_{it}:

$$F_v(z) = \exp(-\exp(-z + r_{it}))$$

Table 1: Preference estimates(standard errors in parenthesis)

	I: Constant coefficients		II: Random coefficients	
π_0	6.40 (0.92)	36.28 (8.35)	30.26	37.07
π_1			0.43	0.33
π_2	0.37 (0.04)	0.31 (0.03)	0.60	0.78
π_3			0.83	0.49
γ_1	3.04 (0.13)	3.24 (0.18)	3.13	3.03
γ_0	-3.80 (2.54)	100.92 (29.24)	7.15	14.88
γ_2	0.00	2.25 (0.17)		2.24
α	0.01 (0.00)	0.01 (0.00)	0.01	0.01
σ_x			0.07	0.22
	III: Random coefficients (restricted)		IV: Static model (BLP)	
π_0	5.79 (1.03)	5.56 (0.96)		
π_1	0.08 (0.03)	0.13 (0.05)		
π_2	0.36 (0.04)	0.30 (0.02)		
π_3				
γ_1	3.09 (0.13)	3.24 (0.15)	2.56 (0.10)	2.58 (0.15)
γ_0	-5.40 (2.93)	-5.72 (3.02)	-12.33 (0.12)	-12.18 (0.13)
γ_2	0.00	2.25 (0.18)		2.20 (0.18)
α	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)
σ_x	0.01 (0.00)	0.05 (0.01)	0.00 (0.00)	0.00 (0.00)

Table 2: Cross elasticities, product characteristics and market shares

Change in price of Sony DSC-S70 in 2:2000 (t-stat in parenthesis) Variable	Heterogeneous	Homogeneous
Constant	21.54 (2900)	21.58 (18100)
Diff. in Res.	-0.0016 (-122)	-0.00001 (-0.0005)
Market Shares	-0.0256 (-0.24)	-0.0001 (-0.0001) std

Figure 1: Observed and Predicted (95% interval) Market Participation

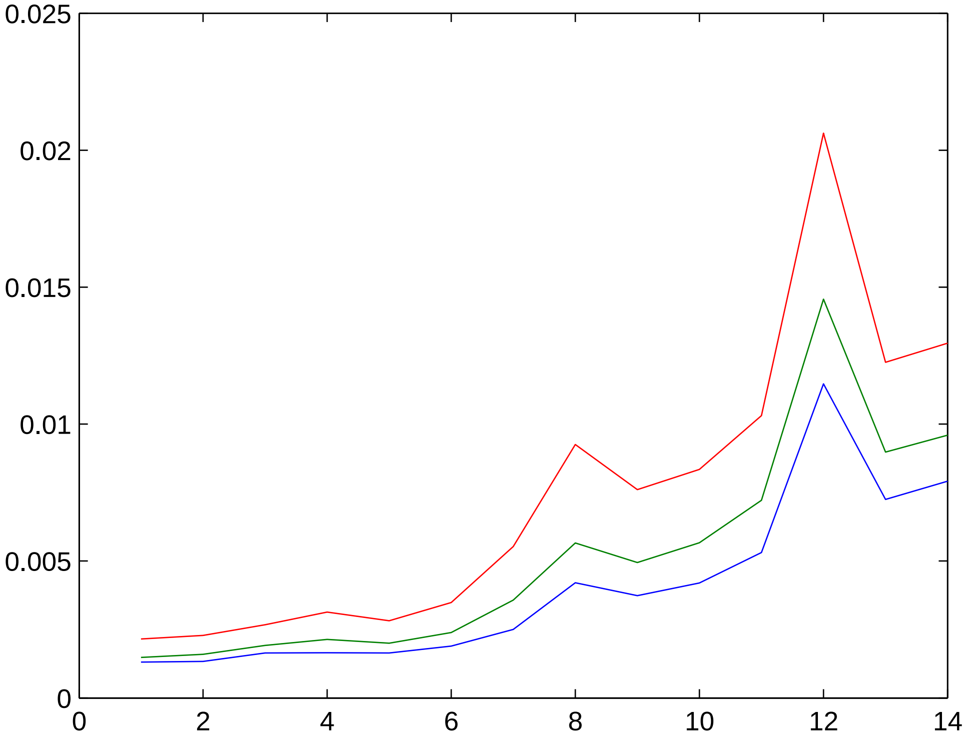

Figure 2: Observed and Predicted (95%) Participation Rate (Restricted Model)

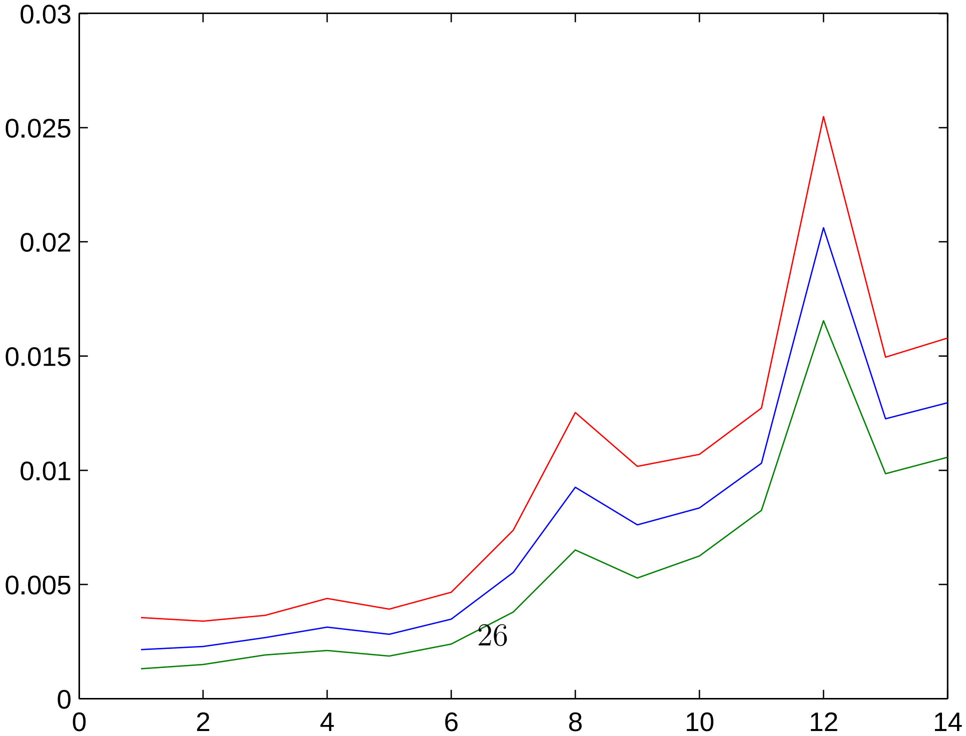

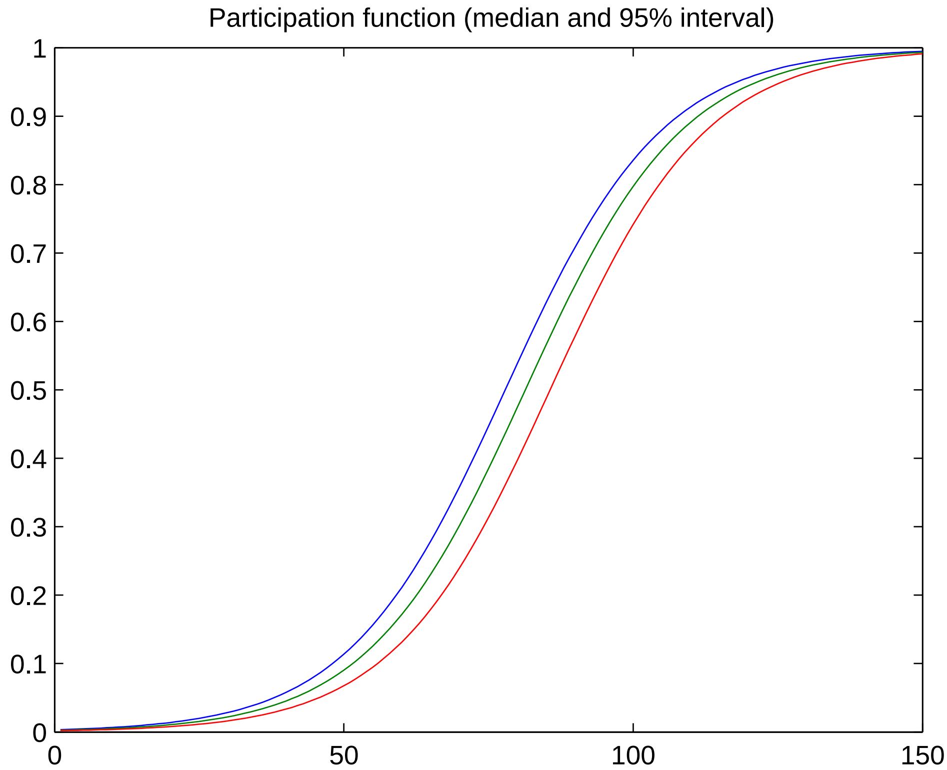

Figure 4: Participation Function (median and 95% interval, unrestricted model)

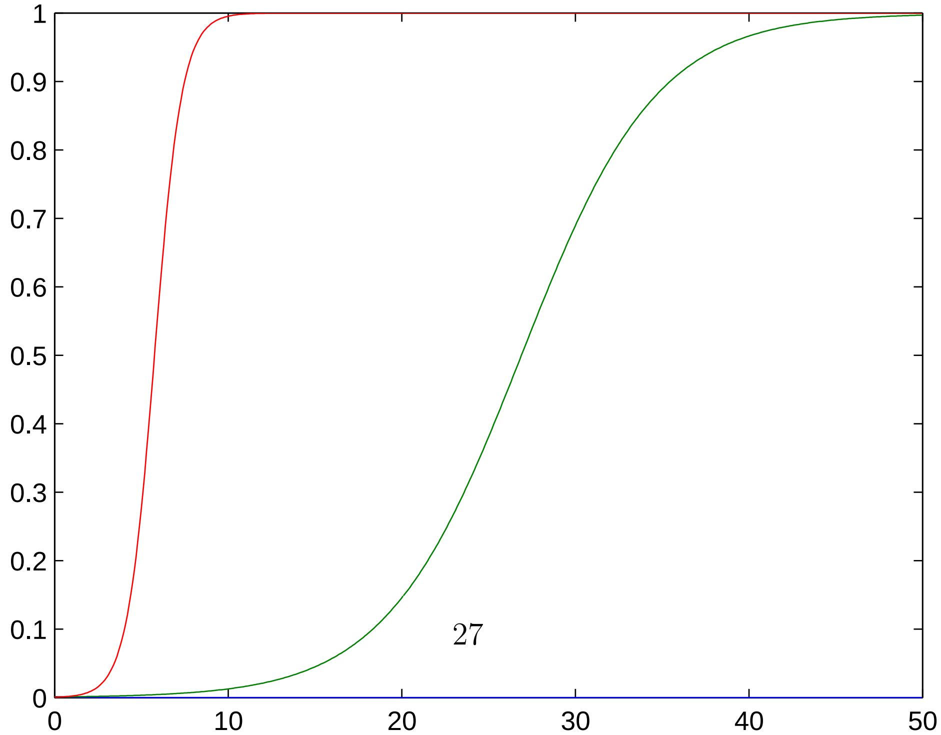

Figure 5: Dispersion of the Valuation of Quality over Time
(standard deviation of valuations)

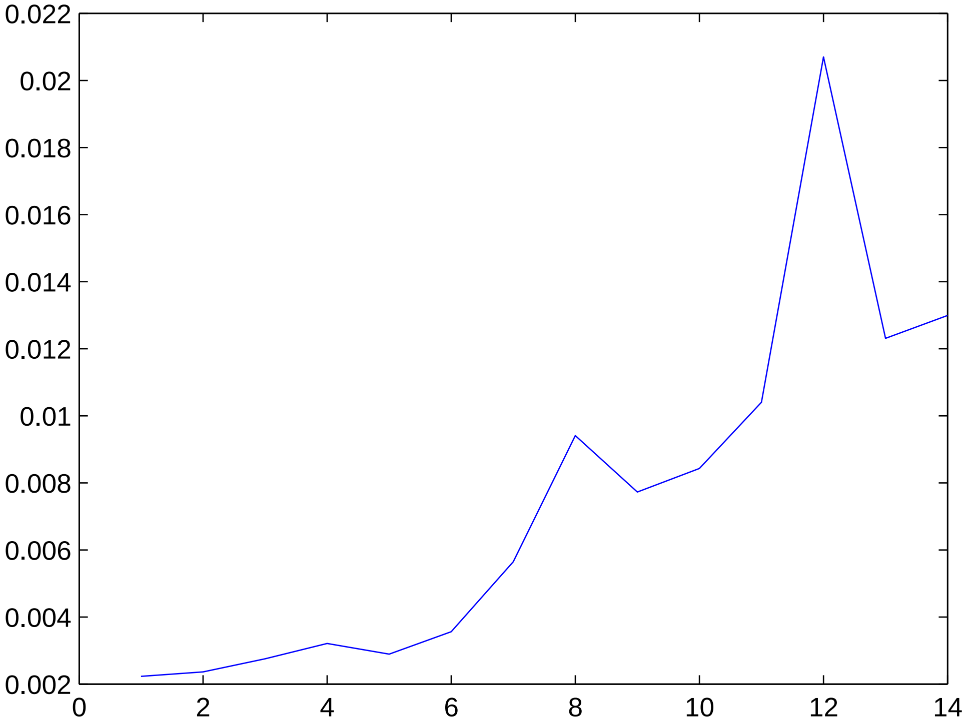

Figure 6: Valuation of Quality (resolution) of Purchasing Costumers
(average deviation from the mean)

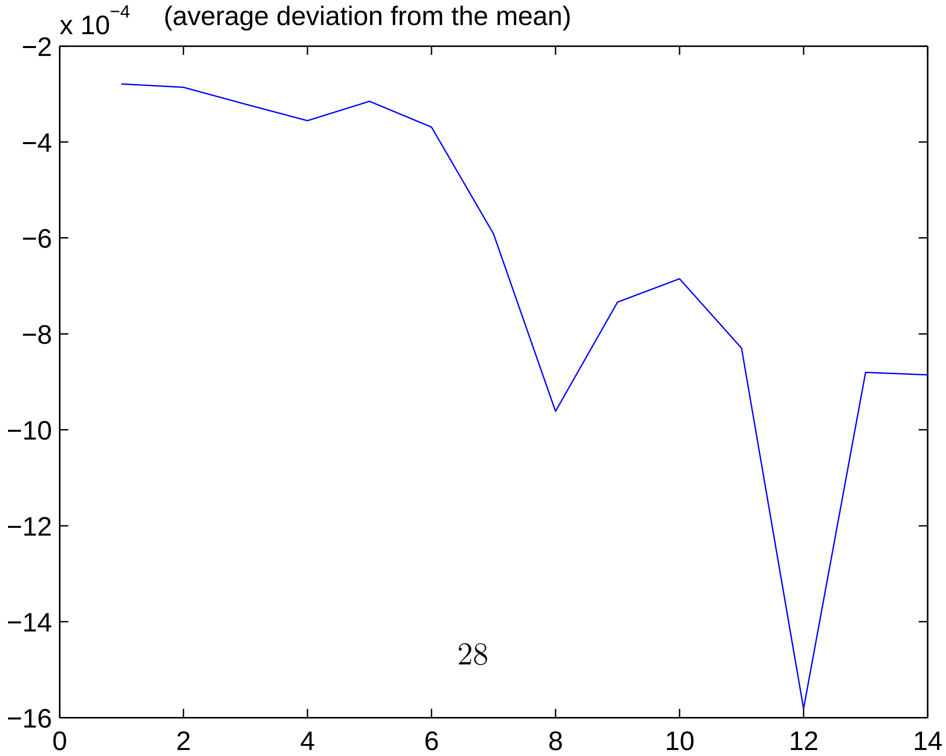

Figure 7: Simulated Participation with Myopic Costumers
(statrting in period 5) and Observed Participation