

Dynamics of Consumer Demand for New Durable Goods*

Gautam Gowrisankaran
Washington University in St. Louis
and NBER

Marc Rysman
Department of Economics
Boston University

August 28, 2007

Abstract

This paper specifies and estimates a dynamic model of consumer preferences for new durable goods with persistent heterogeneous consumer tastes, rational expectations about future products and repeat purchases over time. Most new consumer durable goods, particularly consumer electronics, are characterized by relatively high initial prices followed by rapid declines in prices and improvements in quality. The evolving nature of product attributes suggests the importance of modeling dynamics in estimating consumer preferences. We estimate the model on the digital camcorder industry using a panel data set on prices, sales and characteristics. We find that dynamics are a very important determinant of consumer preferences and that estimated coefficients are more plausible than with traditional static models. We use the estimates to investigate the value of new consumer goods and intertemporal elasticities of demand.

*We thank Lanier Benkard, Igal Hendel, Kei Hirano, Firat Inceoglu, John Krainer, Minsoo Park, Rob Porter, Jeff Prince, Pasquale Schiraldi, Andy Skrzypacz and seminar participants at several institutions for helpful comments and Haizhen Lin, Ryan Murphy, David Rapson, Alex Shcherbakov, and Shirley Wong for research assistance. The NPD Group and ICR-CENTRIS graciously provided support in obtaining data, as did Jeff Prince and Ali Hortascu. We acknowledge funding from the National Science Foundation. All errors are our own.

1 Introduction

All consumers are familiar with the importance of dynamics when purchasing new consumer electronics technologies. A purchaser today can be sure that within a short period of time, a similar product will be available for less. Rapidly falling prices and improving features have been one of the most visible phenomena in a large number of interesting and important new durable goods markets, such as computers, digital camcorders and DVD players. For instance, for digital camcorders, average prices dropped from about \$930 to \$380 between 2000 to 2006, features such as night shot diffused from 53% to 77% of models, and average sizes shrank significantly. A dynamic model is necessary to capture the fact that consumers choose not only what to buy but when to buy. Moreover, the rapidly evolving nature of product attributes for new durable consumer goods suggests that modeling dynamics might be empirically very important.

This paper specifies a structural dynamic model of consumer preferences for new consumer durable goods and estimates the model using data on the digital camcorder industry. Accurately measuring dynamic consumer preferences for durable goods allows for the investigation of a variety of research questions that are of interest to both researchers and policymakers. We use the model to study the importance of dynamics in evaluating price and characteristic elasticities, intertemporal substitution and the welfare gains from innovation for this industry. Our methods are also potentially applicable to other industries and other questions that require uncovering the dynamics of consumer preferences.¹

Evaluating consumer welfare gains from new industries is necessary to develop price indices and understand how much innovation contributes to the economy. The importance of measuring these welfare gains is underscored by the substantial empirical work on the welfare of new consumer durable goods. Yet most empirical papers that examine similar industries have used static models of consumer demand.² Welfare measures that are not based on dynamic models of demand may be biased.³ Moreover, the direction of the bias is not necessarily clear. If consumers act as rational dynamic agents and we instead assume myopic behavior, we may overstate the welfare gains, by assuming more high-value consumers than actually exist, or we may understate the welfare gains, by not recognizing that consumers delay purchases because of the expectation of lower prices and better features.

Our model allows for product differentiation, for persistent consumer heterogeneity and for repeat purchases over time. Berry, Levinsohn & Pakes (1995), henceforth BLP, and the literature that follows have shown the importance of incorporating consumer heterogeneity into differentiated product demand systems for obtaining realistic predictions of elasticities and welfare estimates. Much of

¹For instance, Zhao (2007) extends our model to explain price reductions in the digital camera market and Schiraldi (2007) extends our model to study the impact of scrapping subsidies on the second-hand automobile market.

²These include Goolsbee & Petrin (2004) for satellite cable, Ohashi (2003) for VCRs, Clements & Ohashi (2005) for video games, Chintagunta, Dube & Nair (2004) for personal digital assistants and Einav (2007) for movie-going.

³See Aizcorbe (2005) for discussion on this point.

our model of consumer preferences is essentially the same as BLP: consumers in our model make a discrete choice from a set of products in a multinomial logit model and have random coefficients over observable product characteristics. Also similar is that our model is designed for aggregate data (but can incorporate consumer-level data when available) and that prices are endogenous.

Our model departs from BLP in that products are durable and consumers are rational forward-looking agents who have the option to purchase a product in the future instead of, or in addition to, purchasing one now. Consumers do not know the future set of products, but they instead perceive a distribution for the value of purchasing in future periods. Rather than modeling the supply side explicitly, we make a major simplifying assumption: that consumers perceive that the evolution of the value of purchase will follow a simple one-dimensional Markov process. In this sense, consumers use a reduced-form approximation of how the supply side evolves in order to make predictions about the value of future purchases. We assume rational expectations within the context of this simple framework in the sense that each consumer's expectations will be the actual empirical distribution of future product attributes.

As in most BLP-style models, our identification of key parameters such as price elasticities and random coefficients comes from the impact of different choice sets on purchase probabilities using the assumption that the choice sets are exogenous. We have a tremendous amount of variation in the choice sets that allow us to identify these parameters. Moreover, our dynamic model makes use of substitution patterns across time periods as well as within time periods and captures the fact that demand changes endogenously over time as consumer holdings evolve. A central limitation of this approach is that it does not allow product characteristics to be endogenous.

Related to our work, a recent empirical literature also seeks to estimate the preferences for dynamic durable goods. Gandal, Kende & Rob (2000) analyze dynamic demand for homogenous products markets. Esteban & Shum (2007) estimate a model of the second-hand automobile market with forward-looking consumers and firms using a simple vertical model where consumers must purchase a car every period. Prince (2007) estimates a demand model with upgrading using disaggregate data on purchases of personal computers. Using aggregate data on VCRs, Park (2004) interprets time dummies as capturing the value of waiting in a model without persistent consumer heterogeneity. More closely related to our paper is Melnikov (2001), who analyzes the dynamics of consumer choice for discrete choice differentiated products markets with durable goods using data on computer printers and a logit utility specification. Melnikov's framework is similar to ours but is different in that all consumer heterogeneity is captured by a term that is independently distributed across consumers, products and time, and in that consumers purchase only once in their lifetime.

Several recent papers generalize the Melnikov (2001) idea. Gordon (2006) estimates the demand for computer processors, using a logit demand specification and allowing for repeat purchases. His model does not allow for heterogeneity across consumers or for the endogeneity of price, and allows for only four products at any one time. Song & Chintagunta (2003) propose a logit utility

model of digital cameras that allows for random coefficients but does not allow for the endogeneity of price and requires the number of products to stay fixed over time. Carranza (2006) also examines digital cameras, proposing a model similar to ours although without repeat purchases, and suggests an alternative method for estimating this type of model where the dynamics are estimated through a reduced-form specification that is relatively easy to estimate. Nair (2007) estimates the demand and supply for video games allowing for random coefficients and endogenous prices treating each video game as a monopoly.

Our paper builds on the literature on estimating dynamic demand in that our model allows for unobserved product characteristics, repeat purchases, endogenous prices and multiple differentiated products and is based on an explicit dynamic model of consumer behavior. We develop new methods of inference that allow us to estimate this model. Our method draws on the techniques of Berry (1994) for modeling consumer heterogeneity in a discrete choice model and also on the Rust (1987) techniques for modeling optimal stopping decisions, where stopping corresponds to purchasing a durable good. As in Berry (1994), we solve for the vector of unobserved product characteristics for each product by finding the value of the vector that makes the predicted market share match the observed market share for each product. We then create a GMM estimator using orthogonality conditions based on the unobserved characteristics. For each parameter vector, Berry suggests finding the mean product characteristics using a contraction mapping that defines mean utilities as a function of market shares. We use a similar process to invert the share equation. However, for a set of mean product characteristics, we explicitly evaluate the dynamic demand problem in order to solve for the set of consumers that purchase the product in a given period. This Rust-style optimal stopping problem is nested within the Berry share inversion routine. Our methodological advance is in developing a specification that allows us to nest these two separate methods.

The remainder of the paper is divided as follows. Section 2 discusses the model and method of inference, Section 3 the data, Section 4 the results, and Section 5 concludes.

2 Model and Inference

In this section, we specify our dynamic model of consumer preferences, explain our method of inference and discuss the instruments and identification of the parameters.

2.1 Model

Our model starts with the introduction of a new consumer durable good at time $t = 0$. The unit of observation is a month and there is a continuum of heterogeneous potential consumers indexed by i . Consumers have infinite horizons and discount the future with a common factor β . We assume that products are infinitely durable. However, if a consumer who owns one product

purchases a new one, she obtains no additional utility from the old product, or equivalently, she discards the old product at no cost.⁴ We also do not consider resale markets because we believe that they are small for the new consumer durable goods that we examine given the speed of technological progress.⁵

Consider the decision problem for consumer i at time t . The consumer chooses one of among J_t products in period t or chooses to purchase no product in the current period. In either case, she is faced with similar (though not identical) decision problems at time $t+1$. From these J_t+1 choices, the consumer chooses the option that maximizes the sum of the expected discounted value of future expected utilities conditional on her information at time t .

Product j at time t is characterized by observed characteristics x_{jt} , price p_{jt} and an unobserved (to the econometrician) characteristic ξ_{jt} . For digital camcorders, observed characteristics include size, zoom, and the ability to take still photographs (among others), while the unobserved characteristic would encapsulate product design, ergonomics and unreported recording quality. Consumer preferences over x_{jt} and p_{jt} are defined respectively by the consumer-specific random coefficients α_i^x and α_i^p which we group together as α_i . The characteristics of a product j purchased at time t , x_{jt} and ξ_{jt} , stay constant over the infinite life of the product. We do not model any explicit linkage between products offered for sale at different time periods. We assume that consumers and firms know all time t information when making their time t decisions.

Every period, each consumer obtains a flow utility based on the product that she purchases or on the product that she already owns if she chooses not to purchase. The functional form for the flow utility fits within the random coefficients discrete choice framework of BLP. Specifically, we let

$$\delta_{ijt}^f = \alpha_i^x x_{jt} + \xi_{jt} \quad j = 1, \dots, J_t$$

denote the gross flow utility from product j purchased at time t . We assume that a consumer purchasing product j at time t would receive a net flow utility at time t of

$$u_{ijt} = \delta_{ijt}^f - \alpha_i^p \ln(p_{jt}) + \epsilon_{ijt},$$

where p_{jt} is the price of good j in period t and ϵ_{ijt} is an idiosyncratic unobservable meant to capture random variations in the purchase experience, such as the sales personnel. We assume that ϵ_{ijt} is distributed type 1 extreme value, independent across consumers, products and time, and as such has mean γ , Euler's constant. We let α_i be constant over time and distributed normally with mean $\alpha \equiv (\alpha^x, \alpha^p)$ and variance matrix Σ , where α and Σ are parameters to estimate.⁶

⁴The model could be modified to allow for a second durable good to have value. The value of the second good could be identified using micro data on penetration rates or individual purchasing behavior (see Berry, Levinsohn & Pakes, 2004; Petrin, 2002).

⁵Schiraldi (2007) extends our model to analyze the market for used cars in Italy.

⁶Our specification for price can be rationalized by a model where consumers consume two products: a money good and a camcorder or outside good, and money less than one dollar generates zero utility. (Details available upon request.) The specification can easily be modified to use the empirical income density, as in Nevo (2001)'s study on the breakfast cereal industry.

We also define the population mean flow utility

$$\bar{\delta}_{jt}^f = \alpha^x x_{jt} + \xi_{jt}, j = 1, \dots, J_t,$$

which we use to explain our method of inference in Subsection 2.2.

In our model, the value of a previously purchased product is entirely captured by its flow utility δ_{ijt}^f . We do not keep track of the identity of the product (for instance, the brand). A consumer who does not purchase a new product at time t has net flow utility of

$$u_{i0t} = \delta_{i0t}^f + \epsilon_{i0t},$$

where δ_{i0t}^f is the flow utility from the product currently owned and ϵ_{i0t} is also distributed type 1 extreme value. For an individual who has purchased a product in the past, $\delta_{i0t}^f = \delta_{ij\hat{t}}^f$, where \hat{t} is the most recent period of purchase, and \hat{j} is the product purchased at time \hat{t} . Individuals who have never purchased a product in the past use the outside good, whose mean utility we normalize to 0, so that $\delta_{i0t}^f = 0$ for those individuals.

In order to evaluate consumer i 's choice at time t , we need to formalize consumer i 's expectations about the utility from future products. We assume that consumers have no information about the future values of the idiosyncratic unobservable shocks ϵ beyond their distribution. The set of products and their prices and characteristics vary across time due to entry and exit and changes in prices for existing products. Consumers are uncertain about future product attributes but have rational expectations about their evolution. We assume that each consumer is, on average over time, correct about the mean and variance of the future quality path.⁷

We now define the state variables and use them to exposit the dynamic decision process. Let ω_t denote current product attributes x_{jt} , p_{jt} and ξ_{jt} for all products available at time t , and define $\epsilon_{i,t} \equiv (\epsilon_{i0t}, \dots, \epsilon_{iJ_t t})$. Then, the purchase decision for consumer i depends on preferences α_i and $\epsilon_{i,t}$, endowments δ_{i0t}^f , current product attributes ω_t and expectations of future product attributes. Future product attributes will depend on firm behavior which is a function of consumer endowments and supply-side factors such as technological progress. Let Ω_t denote current product attributes and any other factors that influence future product attributes. We assume that Ω_{t+1} evolves according to some Markov process $P(\Omega_{t+1}|\Omega_t)$ that accounts for firm optimizing behavior. Thus, the state vector for consumer i is $(\epsilon_{i,t}, \delta_{i0t}^f, \Omega_t)$. Let $V(\epsilon_{i,t}, \delta_{i0t}^f, \Omega_t)$ denote the value function, and $EV_i(\delta_{i0t}^f, \Omega_t) = \int_{\epsilon_{i,t}} V_i(\epsilon_{i,t}, \delta_{i0t}^f, \Omega_t) dP_\epsilon$ denote the expectation of the value function, integrated over realizations of $\epsilon_{i,t}$. Note that because $\epsilon_{i,t}$ is *i.i.d.*, it satisfies the assumption of conditional independence in Rust (1987).

⁷A more general rational expectations model would allow individual consumers to have consistently biased estimates of the future logit inclusive value but let the mean expectations across consumers be accurate. While such a model would be easy to specify, it would be difficult to identify without expectations data.

We can now define the Bellman equation for consumer i as

$$V_i(\epsilon_{i,t}, \delta_{i0t}^f, \Omega_t) = \max \left\{ u_{i0t} + \beta E \left[EV_i(\delta_{i0t}^f, \Omega_{t+1}) \middle| \Omega_t \right], \right. \\ \left. \max_{j=1, \dots, J_t} \left\{ u_{ij t} + \beta E \left[EV_i(\delta_{ij t}^f, \Omega_{t+1}) \middle| \Omega_t \right] \right\} \right\}, \quad (1)$$

where “ E ” denotes the expectation operator, a conditional expectation in this case. From (1), the consumer can choose to wait and keep her current product (option zero), or purchase any of the available products (the next J_t options). Note that the value of waiting is greater than the expected discounted stream of flow utilities $u_{i0t} + (\delta_{i0t}^f + \gamma)\beta/(1 - \beta)$ because waiting encapsulates the option to buy a better product in the future.

The large dimensionality of Ω_t makes it very difficult to compute the Bellman equation in (1). Thus, we proceed by making a major simplifying assumption that allows us to substitute a scalar variable, the logit inclusive value of purchasing in a given period, for Ω_t in the value function. We require two definitions to exposit the assumption. First, for each product $j = 1, \dots, J_t$ let

$$\delta_{ij t}(\Omega_t) = \delta_{ij t}^f - \alpha_i^p \ln(p_{j t}) + \beta E \left[EV_i(\delta_{ij t}^f, \Omega_{t+1}) \middle| \Omega_t \right] \quad (2)$$

denote the expected discounted utility for consumer i purchasing product j at time t , integrated over $\epsilon_{ij t}$. Second, define the *logit inclusive value* for consumer i at time t to be

$$\delta_{it}(\Omega_t) = \ln \left(\sum_{j=1, \dots, J_t} \exp(\delta_{ij t}(\Omega_t)) \right). \quad (3)$$

The logit assumption implies that the value of choosing from the entire set of products available in period t is the same as the value of receiving one product with mean utility δ_{it} and a single extreme value draw. In a dynamic context, a consumer can decide whether or not to purchase this period simply by comparing δ_{it} to the outside option, accounting for expectations of future values of δ_{it} . The characteristics of individual products matter only to the extent that they affect δ_{it} or expectations of its future values. Formally,

$$EV_i(\delta_{i0t}^f, \Omega_t) = EV_i(\delta_{i0t}^f, \delta_{it}, E[\delta_{it+1}, \delta_{it+2}, \dots | \Omega_t]). \quad (4)$$

Note that there is no new assumption in (4). It follows directly from the assumptions of the model, most directly the assumption of extreme value errors.

In this sense, we can separate the consumer’s decision problem into two parts, a choice of whether to buy, which is based on $\delta_{i\hat{t}}, \forall \hat{t} \geq t$, and, given purchase, the choice of what to buy, which is based on the characteristics of products available at time t . This same logic has been used to simplify the choice problem for a number of dynamic multinomial logit models (see Melnikov, 2001; Hendel & Nevo, 2007). Thus, in order to solve the consumer’s aggregate purchase problem, we need only to specify the distributions $P(\delta_{i\hat{t}} | \Omega_t), \forall i$ and $\hat{t} > t$.

Our main simplifying assumption is that the evolution of the logit inclusive value depends only on the current logit inclusive value, which we term Inclusive Value Sufficiency. Formally:

Assumption 1 *Inclusive Value Sufficiency (IVS)*

$$P(\delta_{i,t+1}|\Omega_t) = P(\delta_{i,t+1}|\Omega'_t) \quad \text{if} \quad \delta_{it}(\Omega_t) = \delta_{it}(\Omega'_t).$$

The assumption of IVS implies that if two states have the same δ_{it} for consumer i at time t , then they result in the same distribution of future logit inclusive values for that consumer. As a result of IVS, we can write (with some abuse of notation):

$$EV_i(\delta_{i0t}^f, \delta_{it}, E[\delta_{it+1}, \delta_{it+2}, \dots | \Omega_t]) = EV_i(\delta_{i0t}^f, \delta_{it}). \quad (5)$$

The IVS assumption is potentially restrictive. For example, δ_{it} could be high either because there are many products in the market all with high prices or because there is a single product in the market with a low price. We assume the same expectation of $\delta_{i,t+1}$ for these two cases even though they might lead to different outcome in reality.⁸ We could lessen the restrictiveness of Assumption 1 by expanding the state space beyond δ_{it} , for instance, by adding the number of products as another state variable. This would not pose theoretical difficulties but would substantially increase computational time. In practice, we check whether the δ_{it} evolution is well-approximated by the simple process and it appears to be so.

The benefit of the simplifying assumption is to reduce the state space for the decision problem of whether to purchase at time t from many dimensions to two, as in (5). Thus, we can write the expectation Bellman equation as

$$EV_i(\delta_{i0t}^f, \delta_{it}) = \ln\left(\exp(\delta_{it}) + \exp\left(\delta_{i0t}^f + \beta E\left[EV_i(\delta_{i0t}^f, \delta_{i,t+1}) \mid \delta_{it}\right]\right)\right) + \gamma. \quad (6)$$

We can also write the policy function, the probability that consumer i purchases good j , as the aggregate probability of purchase times the probability of purchasing a given product conditional on purchase, and then simplify:

$$\begin{aligned} \hat{s}_{ijt}(\delta_{i0t}^f, \delta_{ijt}, \delta_{it}) & \\ &= \frac{\exp(\delta_{it})}{\exp(\delta_{it}) + \exp\left(\delta_{i0t}^f + \beta E\left[EV_i(\delta_{i0t}^f, \delta_{i,t+1}) \mid \delta_{it}\right]\right)} \\ &\quad \times \frac{\exp(\delta_{ijt})}{\exp(\delta_{it})} \\ &= \exp\left(\delta_{ijt} - EV_i(\delta_{i0t}^f, \delta_{it})\right). \end{aligned} \quad (7)$$

⁸A similar assumption and discussion appears in Hendel & Nevo (2007).

To solve the consumer decision problem, we need to specify consumer expectations for $P(\delta_{i,t+1}|\delta_{it})$. Consistent with our rational expectations assumption, we assume that consumer i perceives the actual empirical density of $P(\delta_{i,t+1}|\delta_{it})$ fitted to a simple functional form. Our base specification uses a simple linear autoregressive specification with drift,

$$\delta_{i,t+1} = \gamma_{1i} + \gamma_{2i}\delta_{it} + u_{it}, \quad (8)$$

where u_{it} is normally distributed with mean 0 and where γ_{1i} and γ_{2i} are incidental parameters. We estimate (8) with an easily computable linear regression, which is useful given that this regression will be performed repeatedly in our estimation process, as noted below. It is also straightforward to extend (8) to allow a quadratic term of the form δ_{it}^2 and this would not substantially increase computation time.

We now briefly discuss the supply side of the model. We assume that products arrive according to some exogenous process and that their characteristics evolve exogenously as well. Firms have rational expectations about the future evolution of product characteristics. After observing consumer endowments and x_{jt} and ξ_{jt} for all current products, firms simultaneously make pricing decisions. Firms cannot commit to prices beyond the current period. These supply side assumptions are sufficient to estimate the demand side of the model. A fully specified dynamic oligopoly model would be necessary to understand changes in industry equilibrium given changes in exogenous variables.

2.2 Inference

This subsection discusses the estimation of the parameters of the model, (α, Σ, β) , respectively the mean consumer tastes for product characteristics and price, the variance in consumer tastes in these variables and the discount factor. We do not attempt to estimate β because it is notoriously difficult to identify the discount factor for dynamic decision models (see Magnac & Thesmar, 2002). This is particularly true for our model, where substantial consumer waiting can be explained by either little discounting of the future or moderate preferences for the product. Thus, we set $\beta = .99$ at the level of the month.

We develop a method for estimating the remaining parameters that is based on Berry (1994) and Rust (1987) and the literatures that follow.⁹ Our estimation algorithm involves three levels of non-linear optimizations: on the outside is a search over the parameters; inside that is a fixed point calculation of the vector of population mean flow utilities $\bar{\delta}_{jt}^f$; and inside that is the calculation of predicted market shares, which is based on consumers' dynamic optimization decisions. While both the $\bar{\delta}_{jt}^f$ fixed point calculation and dynamic programming estimation are well-known, our innovation is in developing a specification that allows us to nest the dynamic programming solution within the $\bar{\delta}_{jt}^f$ fixed point calculation to feasibly estimate the dynamics of consumer preferences.

⁹Computer code for performing the estimation is available from the authors upon request.

We now describe each of the three levels of optimization. The inner loop evaluates the vector of predicted market shares as a function of $\bar{\delta}_{jt}^f$ (the $\bar{\delta}_{jt}^f$ vector) and necessary parameters by solving the consumer dynamic programming problem for a number of simulated consumers and then integrating across consumer types. Let $\tilde{\alpha}_i \equiv (\tilde{\alpha}_i^x, \tilde{\alpha}_i^p) \sim \phi_l$, where l is the dimensionality of α_i and ϕ_l is the standard normal density with dimensionality l . Note that $\alpha_i = \alpha + \Sigma^{1/2} \tilde{\alpha}_i$ and $\delta_{ijt}^f = \bar{\delta}_{jt}^f + \Sigma^{1/2} \tilde{\alpha}_i x_{jt}$.

For each draw, we start with initial guesses, calculate the logit inclusive values from (3), use these to calculate the coefficients of the product evolution Markov process regression in (8), and use these to calculate the expectation Bellman from (6). We repeat this three-part process until convergence. Using the resulting policy function $\hat{s}_{ijt}(\delta_{i0t}^f, \delta_{ijt}, \delta_{it})$ and computed values of δ_{ijt} and δ_{it} , we then solve for market share for this draw by starting at time 0 with the assumption that all consumers hold the outside good. Iteratively for subsequent time periods, we solve for consumer purchase decisions given the distribution of flow utility of holdings using (7) and update the distribution of flow utility of holdings based on purchases. To perform the iterative calculation, we discretize the state space $(\delta_{i0t}^f, \delta_{it})$ and the transition matrix. We examine the impact of different numbers of grid points and different endpoints to ensure that our approximations are sufficient.

To aggregate across draws, a simple method would be to sample over $\tilde{\alpha}_i$ and scale the draws using $\Sigma^{1/2}$. Since our estimation algorithm is very computationally intensive and computational time is roughly proportional to the number of simulation draws, we instead use importance sampling to reduce sampling variance, as in BLP. Let $\hat{s}_{sum}(\tilde{\alpha}_i, \bar{\delta}_{jt}^f, \alpha^p, \Sigma)$ denote the sum of predicted market shares of any durable good at any time period for an individual with parameters (α^p, Σ) and draw $\tilde{\alpha}_i$. Then, instead of sampling from the density ϕ_l we sample from the density

$$f(\tilde{\alpha}_i) \equiv \frac{\hat{s}_{sum}(\tilde{\alpha}_i, \bar{\delta}_{jt}^f, \alpha^p, \Sigma) \phi_l(\tilde{\alpha}_i)}{\int \hat{s}_{sum}(\tilde{\alpha}, \bar{\delta}_{jt}^f, \alpha^p, \Sigma) \phi_l(\tilde{\alpha}) d\tilde{\alpha}}, \quad (9)$$

and then reweight draws by

$$w_i \equiv \frac{\int \hat{s}_{sum}(\tilde{\alpha}, \bar{\delta}_{jt}^f, \alpha^p, \Sigma) \phi_l(\tilde{\alpha}) d\tilde{\alpha}}{\hat{s}_{sum}(\tilde{\alpha}_i, \bar{\delta}_{jt}^f, \alpha^p, \Sigma)},$$

in order to obtain the correct expectation. As in BLP, we sample from the density f by sampling from the density ϕ_l and using an acceptance/rejection criterion. We compute (9) using a reasonable guess of (α^p, Σ) and computing $\bar{\delta}_{jt}^f$ from these parameters using the middle-loop procedure described below. Instead of drawing *i.i.d.* pseudo-random normal draws for ϕ_l , we use Halton sequences based on the first l prime numbers to further reduce the sampling variance (see Gentle, 2003). In practice, we use 40 draws, although results for the base specification do not change substantively when we use 100 draws.

We now turn to the middle loop, which recovers $\bar{\delta}_{jt}^f$ by performing a fixed point equation similar to that developed by Berry (1994) and BLP. We iterate

until convergence

$$\bar{\delta}_{jt}^{f,new} = \bar{\delta}_{jt}^{f,old} + \psi \cdot (\ln(s_{jt}) - \ln(\hat{s}_{jt}(\bar{\delta}_{jt}^{f,old}, \alpha^p, \Sigma))), \forall j, t, \quad (10)$$

where $\hat{s}_{jt}(\bar{\delta}_{jt}^{f,old}, \alpha^p, \Sigma)$ is the model market share (computed in the inner loop), s_{jt} is actual market share, and ψ is a tuning parameter that we generally set to $1 - \beta$.¹⁰ Note that it is not necessary to treat the inner loop and middle loop as separate. We have found some computational advantages to taking a step in (10) before the inner loop is entirely converged and to performing (6) much more frequently than either (8) or (3). However, we require full convergence of (3), (6), (8) and (10) before moving to the outermost loop.

The outer loop specifies a GMM criterion function

$$G(\alpha, \Sigma) = z' \xi(\alpha, \Sigma),$$

where $\xi(\alpha, \Sigma)$ is the vector of unobserved product characteristics for which the predicted product shares equal the observed product shares conditional on parameters, and z is a matrix of exogenous variables, described in detail in Subsection 2.3 below. We estimate parameters to satisfy

$$(\hat{\alpha}, \hat{\Sigma}) = \arg \min_{\alpha, \Sigma} \{G(\alpha, \Sigma)' W G(\alpha, \Sigma)\}, \quad (11)$$

where W is a weighting matrix.

We minimize (11) by performing a nonlinear search over (α^p, Σ) . For each (α^p, Σ) vector, we first obtain $\bar{\delta}_{jt}^f$ from the middle loop. The fact that $\alpha^x x_{jt}$ and ξ_{jt} enter flow utility linearly (recall that $\bar{\delta}_{jt}^f = \alpha^x x_{jt} + \xi_{jt}$) then allows us to solve in closed form for the α^x that minimizes (11) given $\bar{\delta}_{jt}^f$, as in the static discrete choice literature.¹¹ We perform the nonlinear search using a simplex method. We perform a two-stage search to obtain asymptotically efficient estimates. In the first stage, we let $W = (z'z)^{-1}$, which would be efficient if our model were linear instrumental variables with homoscedastic errors, and then use our first stage estimates to approximate the optimal weighting matrix.¹²

A simplified version of our model is one in which a given consumer is constrained to only ever purchase one durable good. In this case, the computation of the inner loop is vastly simplified due to the fact that only consumers who have never purchased make decisions. Because of this, (2) can be simplified to $\delta_{ijt} = (\delta_{ijt}^f + \beta\gamma)/(1 - \beta) - \alpha_i^p \ln(p_{jt})$ which implies that δ_{it} in (3) does not depend on the value function. Thus, we need only solve the expectation Bellman

¹⁰One issue relates to the properties of (10). Berry provides conditions under which this function is a contraction mapping, guaranteeing that the vector $\hat{s}_{jt}(\bar{\delta}_{jt}^f, \alpha^p, \Sigma, \beta)$ is invertible in $\bar{\delta}_{jt}^f$. In our case, we have found examples where this inversion is not a contraction mapping, implying that the dynamic demand system does not satisfy Berry's conditions. Nonetheless, we have not had any problems in ensuring convergence of this process, and have not had problems of multiple equilibria.

¹¹See Nevo (2000) for details. One difference from the static model is that we cannot solve in closed form for α^p since the price term, $\alpha^p \ln(p_{jt})$ is only paid at the time of purchase, unlike ξ_{jt} .

¹²See again Nevo (2000) for details.

equation (6) for $\delta_{ijt}^f = 0$ and hence there is effectively one state variable, δ_{it} , instead of two. The computation of the outer loop for this model is also quicker, since the price coefficient α^p can be solved in closed-form for this model, like α^x in the base model.

In practice, we compute the value function by discretizing δ_{i0t}^f into 20 evenly-spaced grid points and δ_{it} into 50 evenly-spaced grid points. We specify that δ_{it} can take on values from 20% below the observed values to 20% above and assume that evolutions of δ_{it} that would put it above the maximum bound simply place it at the maximum bound. The maximum value of δ_{it} may potentially impact our results given that we examine the dynamics of a market with an improving δ_{it} . However, we also increased the bound to 60% and found no substantive change from the base specification, suggesting that the 20% bound is sufficient. We linearly interpolate between grid points to compute the value of any given state, so arguably, we approximate the value function with a linear spline rather than use true discretization as in Rust (1987). In solving for market shares, we discretize δ_{i0t}^f into 400 evenly spaced bins, which allows for a more accurate tracking of consumer states for this important step.

2.3 Identification and instruments

Our model follows the same identification strategy as BLP and the literature that follows. Heuristically, the increase in market share at product j associated with a change in a characteristic of j identifies the mean of the parameter distribution α . The Σ parameters are identified by the set of products from which product j draws market share as j 's characteristics change. For instance, if product j draws only from products with similar characteristics, then this suggests that consumers have heterogeneous valuations of characteristics which implies that the relevant components of Σ are large. In contrast, if j draws proportionally from all products, then Σ would likely be small. Because our model is dynamic, substitution patterns across periods (in addition to within periods) identify parameters. Moreover, our model endogenously has different distributions of consumer tastes for different time periods. For instance, consumers with high valuations for the product will likely buy early on, leaving only lower valuation consumers in the market until such time as new features are introduced, which will draw back repeat consumers. Substitution based on this aggregate variation in consumer tastes across time further identifies parameters.

Note that our model allows for consumers to purchase products repeatedly over time, even though it can be estimated without any data on repeat purchase probabilities for individuals. At first glance, it might appear difficult to identify such a model. However, this model does not introduce any new parameters over the model with one-time purchases. Indeed, it does not introduce any new parameters over the static model except for the discount factor β , which we do not estimate. The reason that it does not introduce any new parameters is that we have made some relatively strong assumptions about the nature of the product: that durable goods do not wear out; that there is no resale market for them; and that there is no value to a household to holding more than one

durable good of a given type. With these assumptions, the only empirically relevant reason to buy a second durable good is new features, and features are observed in the data. While these assumptions are strong, we believe that they are reasonable for the consumer goods that we study.

As is standard in studies of market power since Bresnahan (1981), we allow price to be endogenous to the unobserved term (ξ_{jt}) but we assume that product characteristics are exogenous. This assumption is justified under a model in which product characteristics are determined as part of some technological progress which is exogenous to the unobserved product characteristics in any given period. As in Bresnahan and BLP, we do not use cost-shifters as instruments for price and instead exploit variables that affect the price-cost margin. Similar to BLP, we include the following variables in z : all of the product characteristics in x ; the mean product characteristics for a given firm at the same time period; the mean product characteristics for all firms at the time period; and the count of products offered by the firm and by all firms. These variables are meant to capture how crowded a product is in characteristic space, which should affect the price-cost margin and the substitutability across products, and hence help identify the variance of the random coefficients and the price coefficient. While one may question the validity of these instruments, they are common in the literature. We consider the development of alternative instruments a good area for future research.

3 Data

We estimate our model principally using a panel of aggregate data for digital camcorders.¹³ The data are at the monthly level and, for each model and month, include the number of units sold, the average price, and other observable characteristics. We observe 378 models and 11 brands, with observations from March 2000 to May 2006. These data start from very early in the product life cycle of digital camcorders and include the vast majority of models. The data set was collected by NPD Techworld which surveys major electronics retailers and covers 80% of the market.¹⁴ We create market shares by dividing sales by the number of U.S. households in a year, as reported by the U.S. Census.

We collected data on several important characteristics from on-line resources. We observe the number of pixels that the camera uses to record information, which is an important determinant of picture quality. We observe the amount of magnification in the zoom lens and the diagonal size of the LCD screen for

¹³We have obtained similar data for digital cameras and DVD players and previous versions of this paper estimated models for those industries. Basic features of the results are similar across industries. We focus on camcorders here because we believe this product exhibits the least amount of network effects (such as titles for DVD players or complementary products for producing pictures for digital cameras), which would complicate our analysis. Incorporating network effects into our framework is the subject of current research.

¹⁴There are major omissions from NPD's coverage. Sales figures do not reflect on-line sellers such as Amazon and they do not cover WalMart. We do not attempt to correct for these shortcomings.

Table 1: Characteristics of digital camcorders in sample

Characteristic	Mean	Std. dev.
Continuous variables		
Sales	2492	(4729)
Price (Jan. 2000 \$)	599	(339)
Size (sq. inches width \times depth, logged)	2.69	(.542)
Pixel count (logged \div 10)	1.35	(.047)
Zoom (magnification, logged)	2.54	(.518)
LCD screen size (inches, logged)	.939	(.358)
Indicator variables		
Recording media: DVD	.095	(.294)
Recording media: tape	.862	(.345)
Recording media: hard drive	.015	(.120)
Recording media: card (excluded)	.028	(.164)
Lamp	.277	(.448)
Night shot	.735	(.442)
Photo capable	.967	(.178)
Number of observations: 4436		
Unit of observation: model – month		

viewing shots.¹⁵ We observe the width and depth of each camera in inches (height was often unavailable), which we multiply together to create a “size” variable. We also record indicators for whether the camera has a lamp, whether it can take still photos and whether it has “night shot” capability, an infrared technology for shooting in low light situations. Finally, we observe the recording media the camera uses - there are four mutually exclusive media (tape, DVD, hard drive and memory card) - which we record as indicators.

To create our final data set, we exclude from the choice set in any month all digital camcorders that sold fewer than 100 units in that month. This eliminates about 1% of sales from the sample. We also exclude from the choice set in any month all products with prices under \$100 or over \$2000 as these products likely have very different usages. This eliminates a further 1.6% of sales from the sample. Table 1 summarizes the sales, price and characteristics data by level of the model-month for our final sample.

Figure 1 graphs simple averages of two features over time, size and pixel count, using our final sample. Not surprisingly, cameras improve in these features over time. Weighting by sales produces similar results. Figure 2 displays a similar graph for features that are characterized by indicator variables: the

¹⁵We log all continuous variables and treat any screen of less than 0.1 inch as equivalent to a screen of 0.1 inch.

Figure 1: Average indicator characteristics over time



Figure 2: Average indicator characteristics over time





Figure 3: Prices and sales for camcorders

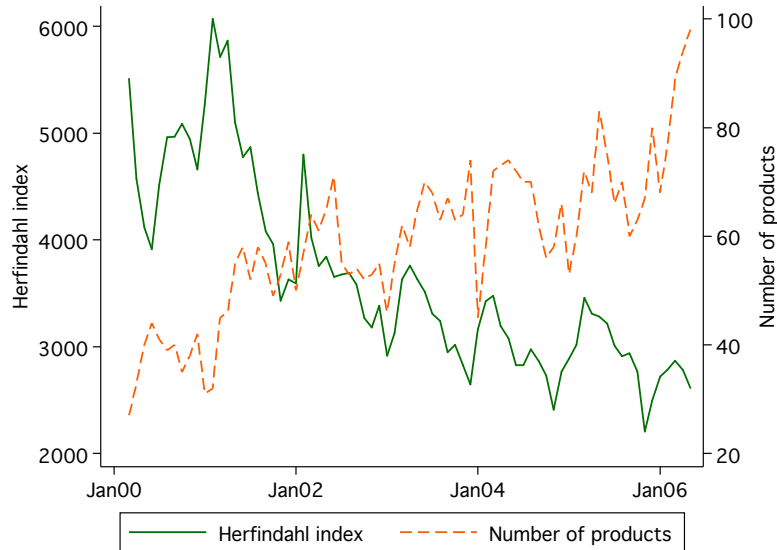


presence of a lamp, the presence of night shot, the ability to take still photographs and whether the recording media is tape. The first two systematically become more popular over time. Photo ability is present in nearly every camera half-way through the sample but declines slightly in popularity by the end of our sample. Tape-based camcorders initially dominated the market but grew less popular over time relative to DVD- and hard drive-based devices, representing more than 98% of devices in the first few months of the data but less than 65% in the last few.

Figure 3 shows total sales and average prices for camcorders in our final sample over time. Camcorders exhibit striking price declines over our sample period while sales increase. Even more noticeable than the overall increase in sales is the huge spike in sales at the end of each year due to Christmas shopping. Our model needs to have some way of explaining the huge impact of the Christmas season on sales. One way would be to add in a utility shifter for the Christmas season or general month dummies. This would vastly complicate our model by adding another state variable (months until Christmas) given that our demand system is dynamic. Because prices and features do not change over Christmas, we believe that this would increase computational time without necessarily providing any tangible benefit.

Instead, we addressed the Christmas spike issue by seasonally adjusting our data. Specifically, we multiplied sales by a separate constant for each month, constant across years. The constants were chosen so that the sales by month summed over the years in the data were the same for each month and so that total sales for each year were unchanged. Figure 3 also shows the seasonally-

Figure 4: Competition in the camcorder market



adjusted sales data, which are, by construction, much smoother than the unadjusted data.

An important note about the data concerns the identification power of our instruments. The identification from our instruments comes from the variation in the number and closeness of substitute products across locations in the product space, including time. Figure 4 provides some evidence on the nature of competition in this market over time, using our final sample. The figure shows that the number of camcorders available is steadily increasing over time roughly from 25 to 100, suggesting that there will generally be closer substitutes to a given player for later time periods. We also compute a monthly Herfindahl-Hirschman index based on brand level (not model level) market shares. This statistic declines from about 6000 near the beginning of the sample to 2500 near the end. The process is not monotonic and so provides significant variation in the level of competition over time. Interestingly, the number of brands (not shown) increases from 5 at the start to 11 in early 2005 and then declines to 7 at the end of the sample.

In addition, in some specifications we incorporate household level data on ownership, often referred to as penetration, to better pin down repeat purchasing behavior. These data come from ICR-CENTRIS, which performs telephone interviews via random-digit dialing. ICR-CENTRIS completes about 4,000 interviews a month, asking which consumer electronics items a household owns.¹⁶

¹⁶Data on how many camcorders a household owns or data on the time between purchases would be even more directly useful for understanding repeat purchases. However, a lengthy search of public and private data sources did not turn up any such information.

Figure 5: Penetration and sales of digital camcorders




Figure 5 shows our ICR-CENTRIS data, which contain the percent of households that indicate holding a digital camcorder in the third quarter of the year for 1999 to 2006. It also shows the year-to-year change in this number and the new sales of camcorders, as reported by NPD.

The penetration data show rapid growth in penetration early on in the sample but no growth by the end. The evidence from the penetration and sales data are not entirely consistent, perhaps due to differences in sampling methodology: in 3 of the 6 years, the increase in penetration is larger than the increase in new sales. We also believe the ICR-CENTRIS finding of virtually no new penetration after 2004 to be implausible. Nonetheless, the slowdown in penetration but continued growth in sales together suggest that there are substantial repeat purchases by the end of our sample. Because of the issues surrounding the penetration data, we only use it in one robustness specification, as we discuss in Section 4.2 below.

4 Results and implications

We first exposit our results, then provide evidence on the fit of the model, and finally discuss the implications of the results.

4.1 Results

We present our parameter estimates in Table 2. Table 2 contains six columns of results. The first column of results provides the parameter estimates and

Table 2: Parameter estimates

Parameter	Base dynamic model	Dynamic model with extra random coefficients	Dynamic model without repurchases	Static model	Static model aggregated to year	Dynamic model with micro moments
Mean coefficients (α)						
Constant	-.141 (.044) *	-.097 (.195) *	-.087 (1.5)	-8.90 (2e3)	-4.03 (132)	-243 (213)
Log price	-2.66 (.576) *	-2.74 (.975) *	-.056 (72.6)	.0247 (19.1)	-.089 (14.5)	-3.01 (.582) *
Log size	-.007 (.001) *	-.007 (.014)	-.002 (7e-4) *	-1.52 (.068)	-.340 (.204)	-.019 (.002) *
Log pixel	.095 (.050)	.098 (.028) *	-.002 (.027)	-2.56 (2.43)	-4.52 (5.85)	.241 (.146) *
Log zoom	.007 (.002) *	.007 (.002) *	.007 (9e-4) *	.654 (.086) *	.861 (.269)	.010 (.004) *
Log LCD size	.003 (.001) *	.003 (.001) *	-5e-4 (9e-4)	-.053 (.105)	-.361 (.325)	.011 (.004) *
Media: DVD	.025 (.005) *	.027 (.006) *	-.001 (.004)	-.177 (.344)	.229 (1.35)	.052 (.017) *
Media: tape	.007 (.005)	.007 (.005)	-.007 (.003) *	-.763 (.333) *	-.671 (1.05)	.017 (.016)
Media: HD	.020 (.006)	.023 (.008) *	-.008 (.004)	-.873 (.425) *	-1.32 (1.59)	.039 (.019) *
Lamp	.006 (.001) *	.007 (.001) *	-.002 (.001)	-.209 (.130)	-.351 (.402)	.002 (.004)
Night shot	.009 (.001) *	.009 (.001) *	.007 (6e-4) *	.646 (.073) *	1.20 (.199)	.022 (.003) *
Photo capable	-.014 (.002) *	-.015 (.003) *	-.005 (.002) *	-.431 (.205) *	-.432 (.767)	-.022 (.007) *
Standard deviation coefficients ($\Sigma^{1/2}$)						
Constant	.086 (.025) *	.058 (.130)	2e-5 (27)	.007 (4e4)	.036 (1e3)	1e-7 (.082)
Log price	7e-6 (.563)	.043 (8.06)	.0002 (817)	.001 (267)	.011 (67.7)	.651 (.233) *
Log size		5e-09 (.096)				
Log pixel		.0015 (.337)				

Standard errors in parentheses; statistical significance at 5% level indicated with *

standard errors from our base specification of the model presented in Section 2 with two random coefficients, price and the constant term. Column 2 adds random coefficients on camcorder size and pixel count. We discuss the other columns below.

Starting with column 1, the base specification reports results that are generally sensible in magnitude and sign. As we would hope, price contributes negatively to utility for virtually everyone, with a base coefficient of -2.66 that is precisely estimated and a standard deviation on the random coefficient of roughly 0, that is imprecisely estimated. A person with mean tastes would obtain a negative gross flow utility from a camcorder with all characteristics zero (relative to the outside option), with a mean constant term of $-.141$. The standard deviation on this coefficient is $.086$, indicating that there is substantial variation in the gross flow utility from a camcorder. Both the mean and standard deviation coefficients are statistically significant. In comparing the magnitudes of these coefficients, recall that price is paid once, while all the other coefficients relate to flow utility at the level of the month, and hence the price coefficients should be roughly $1/(1-\beta) = 100$ times the magnitude of the other coefficients.

Most of the characteristics of digital camcorders enter utility with the expected sign, including camcorder size, pixels, zoom, LCD screen size, night shot capability and the presence of a lamp. All of these except the pixel count are statistically significant. The three included media dummies are all positive. These are relative to the card technology, which is generally considered the worst. The one coefficient whose sign is not intuitive is photo capability, which is estimated to be negative and significant. It is hard for our utility model to generate a positive coefficient on this feature, since its diffusion reversed over time.

All of the estimated parameters on characteristics are smaller than the parameter on the constant term. In combination with the fact that these characteristics either are indicators or have a standard deviation less than 1, this implies that these features are important, but that the vertical differentiation between camcorders is small relative to the differentiation from the outside good.

Turning now to column 2, the addition of two extra random coefficients results in parameter estimates for mean coefficients that are very similar to the base specification. In particular, the sign of the mean coefficients on price and characteristics are all the same as in the base specification, and statistical significance is similar across specifications. The two random coefficients that are common across the two specifications are estimated to be similar in magnitude, although the random coefficient on the constant term loses its significance. Moreover, the two new random coefficients are estimated to be small and statistically insignificant.

A potential concern in our context is the restrictiveness of the logit error assumption. Logit errors typically imply unrealistic welfare gains from new products (see Petrin, 2002). Akerberg & Rysman (2005) argue that this feature implies that logit-based models will perform poorly in contexts where consumers face different numbers of products over time. Akerberg & Rysman recommend addressing this problem by including the log of the number of products, $\ln(J_t)$,

as a regressor, as if it were a linear element in $\bar{\delta}_{jt}^f$. Finding a coefficient of 0 implies the logit model is well-specified, whereas a coefficient of -1 implies “full-crowding,” that consumers respond to increases in the number of products as if there are no new logit draws. In unreported results, we find that other parameters change little and that the coefficient on $\ln(J_t)$ is -0.015. Although the coefficient is statistically significant, it is very close to zero and suggests that the *i.i.d.* logit draws are a reasonable approximation. Concerns with the implications of logit draws motivate Berry & Pakes (2005) and Bajari & Benkard (2005) to propose discrete choice models that do not include logit *i.i.d.* error terms, but given this coefficient estimate, we do not further pursue this issue.

Column 3 provides estimates from the dynamic model where individuals are restricted to purchase at most one digital camcorder ever. This specification yields results that are less appealing than our base specification. In particular, the mean price coefficient drops in magnitude by a factor of 5 and loses its statistical significance. Many of the characteristics enter mean utility with an unexpected sign, including pixels, LCD screen size and lamp and many fewer mean coefficients are significant than in the base specification. The standard deviation coefficients are very small and statistically insignificant.

It is useful to understand why the model with multiple purchases provides a much larger price coefficient than the dynamic model with purchases restricted to one-time only. In the one-time purchase model, the magnitude of the mean price coefficient is much smaller than the standard deviation of the extreme value distribution. Had this estimated coefficient been applied to the base model, the δ_{ijt}^f values would have to be sufficiently negative to prevent individuals from purchasing a product most months, implying that individuals dislike having their purchased camcorder. In addition to being intuitively unappealing, the negative δ_{ijt}^f values resulted in a very bad fit of the moment criteria for the base model. Thus, the multiple purchase feature of the base model essentially forces the price coefficient to be sufficiently negative to avoid implications that are counterintuitive and also do not fit the data well.

Column 4 essentially follows BLP and estimates a traditional static random coefficients discrete choice specification. To compare these coefficients with the base specification, one would have to multiply all the coefficients from this specification, except for the coefficients on price, by 1/100. The static model yields many unappealing results, including a positive mean price coefficient and many coefficients on characteristics that are of the opposite sign from expected. Column 5 estimates a variant of the static model where we aggregate the products to the annual level.¹⁷ The results from this specification are similar to the results from column 4.

We believe that the very imprecise and sometimes positive price coefficients in the static specifications is caused by the fact that the data cannot easily be explained by a static model. In particular, the static model cannot fit two facts that are characteristic of the data: first, many more people purchased digital

¹⁷This specification drops the first and last year from our data, as we lack information on all months for those years.

camcorders once prices fell; but second, within a time period, the cheapest models were often not the most popular. Because the model cannot then estimate a significantly negative price coefficient, it also does not result in appropriate coefficients on characteristics.

The dynamic model addresses these two facts because it predicts that people wait to purchase because of the expectations of price *declines* and not directly because of high prices. Heuristically, the static price coefficient is analogous to the coefficient from a regression of market shares on prices whereas the dynamic price coefficient is analogous to the coefficient from a regression of shares on the forward difference in price, $(p_{jt} - \beta p_{jt+1})$.¹⁸ Unlike the static explanation, the dynamic explanation for why consumers wait does not conflict with consumers buying relatively high-priced products.

4.2 Fit of the model

We first assess the fit of the model, by reporting the simple average of the unobserved quality ξ_{jt} for each month in Figure 6. For this figure and all that follow, we use the estimated parameters reported in the first column of Table 2 and the vector of $\bar{\delta}_{jt}^x$ that are consistent with these parameters and with observed shares. Note that ξ_{jt} is the estimation error of the model. The figure does not indicate any systematic autocorrelation or heteroscedasticity of the average error over time. This finding is important because there is no reduced-form feature such as a time trend to match the diffusion path. If one were to match, for instance, an S-shaped diffusion path with a simple linear regression, we would expect to have systematically correlation in ξ_{jt} . However, Figure 6 does not indicate any such pattern.

We next examine the reasonableness of the IVS assumption. Our estimates of the dynamic models of consumer preferences rely on the IVS assumption, that consumers perceive that next month's logit inclusive value $\delta_{i,t+1}$ depends only on the current logit inclusive value δ_{it} and only within a simple autoregressive specification with drift. Figure 7 plots δ_{it} for 3 sets of random coefficients at the estimated parameter values: individuals with random coefficients that result in them being in the 20th, 50th and 80th percentile of δ_{it} in the median month of the sample.

One can see that there is a general upward trend in these values throughout our sample period. Moreover, the trend looks roughly linear. Consistent with Table 2, Figure 7 shows that there are significant differences in valuations across coefficients. By the end of the sample period, the value of purchasing a camcorder for a 20th percentile individual is still less than the value for a median individual at the beginning of the sample period. This is consistent with the fact that total digital camcorder sales by the end of our period were only about 10% of the size of the number of U.S. households. The three paths of coefficients move in parallel, rising and dipping in the same months in response to

¹⁸Gandal et al. (2000) show that this heuristic is an exact description of the market with one product, perfect foresight, zero variance to ϵ_{ijt} , linearity in prices, no repeat purchase, and a concave price path.

Figure 6: Average estimation error (ξ_{jt}) by month




Figure 7: Evolution of δ_{it} over time





Figure 8: Difference between δ_{it+1} and its period t prediction




the introduction of new products and features and pricing changes. This is not surprising as the underlying results show that almost all of the heterogeneity is associated with the constant term rather than the valuation of prices.

Figure 7 suggests that the IVS assumption is reasonable. Of course, the results of this figure are based on values evaluated at the structural parameters, and so we cannot rule out the possibility that a more general industry evolution assumption would have resulted in different structural parameters that then would have generated a different pattern of evolution.

Another way of evaluating the industry evolution assumption is to examine the prediction error from the consumer decision problem. In Figure 8 we evaluate the mean value across random coefficients of the prediction error, which is $\delta_{i,t+1} - (\hat{\gamma}_{1i} + \hat{\gamma}_{2i}\delta_{it})$, where $\hat{\gamma}_{1i}$ and $\hat{\gamma}_{2i}$ are the estimated parameters from the regression specified in (8). The figure shows that the prediction errors fluctuate rapidly from negative to positive. There is not an overall trend where they are becoming more positive or more negative over time. In contrast, the results show that, consistent with our model, short-run changes in product attributes are the source of the difference between consumers' predictions of future values and their actual values. This provides further evidence that the evolution process that we specify is reasonable. However, Figure 8 does appear to show that the variance of the prediction errors decreases somewhat over time, although our model imposes that the variance of the residual is constant over time.

We obtain a final measure of the fit of the model by examining the extent to which we observe repeat purchasing behavior in our sample. Figure 9 plots the fraction of shares due to repeat purchases for the base model as well as

Figure 9: Evolution of repeat purchase sales




for a robustness specification that we discuss below. Under the base model, repeat purchases account for a very small fraction of total sales. Even in the final period, which has the largest fraction, repeat purchases account for only about .25% of new sales. The underlying reason why there are not more repeat purchases is that the coefficients on characteristics other than the constant term are small relative to the utility contribution from the price and the constant terms, implying that the net benefit to upgrading is low.

This finding is not consistent with the evidence, albeit imperfect, from the ICR-CENTRIS household penetration survey, that new sales are higher than new penetration. Thus, we use the penetration data in the form of a micro-moment (see Petrin, 2002) as a robustness check on our base results. Specifically, we use the penetration data to construct an additional moment that is the difference between the increase in household penetration between September 2002 and September 2005 predicted by the model and by the penetration data.¹⁹ We chose to use only these two years to mitigate the noise present in the data.

Table 2 column 6 presents results from this specification with an additional moment. The coefficient estimates on characteristics are similar to the base specification although generally somewhat larger. More importantly, the standard deviation of the random coefficient on price goes from being roughly 0 to being relatively large (more than one fifth the size of the mean price coefficient) and statistically significant, while the standard deviation of the random coefficient on the constant term changes in the opposite direction, to being roughly

¹⁹See Berry et al. (2004) and Petrin (2002) for details on calculating weighting matrices when combining micro moments with aggregate moments.

Figure 10: Evolution of digital camcorder sales under different assumptions



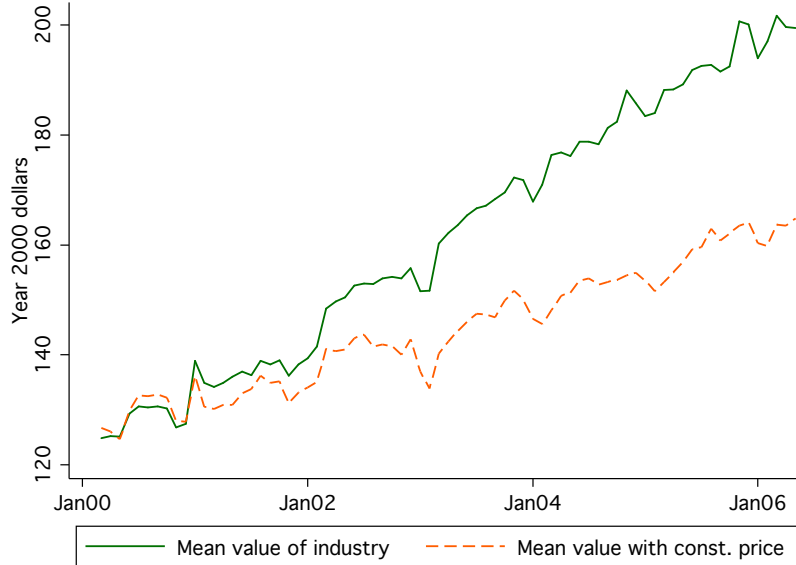
0. The larger variation in the coefficient on price allows the model to predict substantial repeat purchases, by ensuring that a substantial fraction of people care little enough about price to upgrade. The estimates fit the micro moments data almost exactly, with a difference in new household penetration of 4.9% between the two periods for both the model and data.

Figure 9 also plots the share of repeat purchases for the specification with an additional moment. Since this model fits both the increase in penetration of 4.9% from Sep. 2002 to Sep. 2005 and the new sales of 5.85% over the same time period, it predicts much higher repeat purchases than the base model. In particular, it predicts that over 25% of new sales are attributable to repeat purchases by the end of the sample.

4.3 Implications of the results

We now investigate the implications of our estimated model in terms of the importance of dynamics in consumer preferences, consumer welfare and intertemporal price elasticities. Figure 10 investigates the magnitudes of the dynamic responses by examining the time path of digital camcorder sales under three different assumptions: the time path generated by the estimated model (also the actual time path of sales), the time path that would occur if consumers assumed that their logit inclusive values for digital camcorders remained equal to its present value in all future periods, and the time path that would occur if firms were faced with all consumers having no digital camcorders in each period, instead of high valuation consumers having purchased the product and hence

Figure 11: Mean per-capita consumer surplus from digital camcorder industry




generally having a higher reservation utility for buying, as occurs in our model. All results in this subsection use the parameter estimates from our base model in Table 2 column 1.

We find that dynamics explain a very important part of the sales path. In particular, if consumers did not assume that prices and qualities changed, then sales would be somewhat declining over time, instead of growing rapidly over the sample period. At the beginning of our sample period, sales would be huge compared to actual sales, as many consumers would have perceived only a limited option value from waiting. By the end of our sample period, sales would be significantly less than current sales, as many consumers who were likely to buy digital camcorders would have bought them early on, having assumed that quality, in the sense of the logit inclusive value, would be stable over time.

If firms were faced with a situation where all consumers had only the outside good in every period, then the sales path would be similar until roughly two years into our sample. At this point, many of the high valuation consumers had started to purchase. By the end of our sample period, we find that sales in the final month would have been about 3 times as high as they actually were. Note that this increase in sales is due to high valuation consumers not owning any digital camcorders, and mostly not to having a larger market, as roughly 90% of the market had not purchased any digital camcorder by the end of our sample period.

Figure 11 examines the extent to which the digital camcorder industry has created consumer surplus. This provides evidence on the welfare gains from this new goods industry. We evaluate the expected discounted consumer surplus at

Figure 12: Industry dynamic price elasticities



each period by integrating the welfare across consumer random coefficients for the δ_{it} of that period and consumer type. We evaluate the value of the industry for each consumer at each point in time using as the basis point a world where no one owns a digital camcorder. For each consumer random coefficient, we obtain the welfare by dividing the value (measured in utility units) by the marginal utility of a dollar, which we calculate using a price of \$525, which is the sales-weighted mean price of a digital camcorder in our sample.

Our results reveal that the digital camcorder market has contributed an average of \$125 in expected discounted consumer surplus per U.S. household from the point of view of 2000, or an average of \$1.25 per month. The value of the industry rises to about \$200 by 2006. Note that the rise is less than the increase in value of the industry characteristics across these time periods, since the value in 2000 incorporates the fact that the industry will improve over time.

Figure 11 also plots the change in valuation over time with the assumption that prices for all products are equal to the weighted mean price of \$525. This plot shows that roughly half of the value gain between the start and end of our sample is due to the reduction in price. The other half is due to improvements in characteristics and increases in the number of available products.

It would be of use to compare this number to the comparable figure from the static estimation of the digital camcorder industry. However, the static estimation would provide a negative valuation number since the price coefficient is estimated to be positive. Since a negative number is clearly not plausible, we did not report the number for the static estimation.

Finally, we analyze intertemporal price elasticities. To evaluate the impor-

tance of expectations in determining elasticities, we first compare the effect of a temporary (one-month) 1% price increase at time \bar{t} when consumers believe the increase is temporary to a temporary increase when consumers believe it is permanent. In both cases, the price increase is unexpected before time \bar{t} . To compute the response when consumers think the increase is temporary, we evaluate the time \bar{t} expectations of $\delta_{i,\bar{t}+1}$ using the baseline $\delta_{i\bar{t}}$, not the one realized under the price change. To compute the response when consumers (wrongly) think the increase is permanent, we impose that consumers use the $\delta_{i\bar{t}}$ that is realized under the price increase to make predictions about $\delta_{i,\bar{t}+1}$. For all specifications, we assume that $(\gamma_{1i}, \gamma_{2i})$, the baseline transition matrix for δ_{it} , is as estimated in the baseline model.


Figure 12, which displays the industry elasticity with \bar{t} set to the median period of the sample, April 2003, shows that expectations play an important role in determining the response to price changes. Because the price change is unexpected, sales follow the baseline path prior to \bar{t} . A 1% price increase leads to a decrease in sales of 2.18% in that month when consumers believe it to be temporary and a decrease of only 1.24% when they believe it to be permanent. In addition, the response over the following year is also larger, but in the opposite direction, under the price change believed to be temporary, with an increase in sales of .57% of the time \bar{t} sales for the following 12 months compared to an increase of .27% under the believed-permanent price change. Hence, the difference between responses to temporary price changes that consumers do and do not perceive as temporary is much smaller in the long-run than the short-run.

Next we consider a permanent 1% price change from time \bar{t} onwards. Consumers know the price change is permanent in the sense that they use $\delta_{i\bar{t}}$ as realized under the price change to make predictions about future values of δ_{it} . By construction, the quantity change in the first month of 1.24% is the same as when consumers thought the temporary price change was permanent. There is little difference between the long-run and short-run elasticity, which we expect given that expectations are similar immediately after the price change and in the long-run.

We also study the effects of a price change for a single product at the median period of our sample. We consider the Sony DCRTRV250, as it has the largest market share in this period. We depict our results in Figure 13. Here again, consumers respond more strongly in the month of the price change when they believe the price change is temporary than when they believe it is permanent. However, the difference in responses is small: 2.21% versus 2.08%. This result follows because consumers switch to another product rather than delay their purchase when only one product changes price, so expectations matter less. We also consider a permanent price change that consumers recognize as permanent. As in the industry case, the long-run and short-run response to the price change appear similar in this case. Note that elasticities would be larger for most other products as they have smaller market shares.

Strikingly, we find that the short-run elasticities are almost the same for the industry as a whole and the product for the case of a temporary price change that is perceived to be temporary (2.18 and 2.21). However, the sources of

Figure 13: Dynamic price elasticities for Sony DCRTRV250



the quantity change are different: a delay for the industry experiment but a switch to other products for the product experiment. Because of the difference in source, the long-run impacts do show the expected pattern of the market elasticity being significantly smaller than the product elasticity. Note that the result that consumer primarily switch rather delay in the product experiment is in part due to the high variance of the coefficient on the constant term, which implies that many consumers view all camcorders as close substitutes relative to the outside option.

5 Conclusion

This paper develops new methods to estimate the dynamics of consumer preferences for new durable goods. Our model allows for rational expectations about future product attributes, heterogeneous consumers with persistent heterogeneity over time, endogeneity of price, and the ability for consumers to upgrade to new durable goods as features improve. Our model is of use in measuring the welfare impact of new durable goods industries and in evaluating the intertemporal price elasticities for these industries, among other economic questions. We estimate our model using a panel data set of prices, quantities and characteristics for the digital camcorder industry.

Our estimates of consumer preferences that account for dynamics are generally sensible. A variety of robustness measures show that the major simplifying assumptions about the dynamics in the model are broadly consistent with the data. In contrast, a static analysis performed with the same data yields less

realistic results.

We find substantial heterogeneity in the overall utility from digital camcorders. Our results also show that much of the reason why the initial market share for digital camcorders was not higher was because consumers were rationally expecting that the market would later yield cheaper and better players. We find that industry elasticity of demand is 2.18 for transitory price shocks and 1.24 for permanent price shocks, with significantly larger elasticities for individual products. Last, we find that the digital camcorder industry is worth an average of \$125 in expected value at the start of the industry.

We believe that our results show that dynamic estimation of consumer preferences is both feasible and important for analyzing industries with new goods. We see several avenues of future research, including evaluating firm decision problems in the presence of consumer and firm dynamics.

References

- Akerberg, D. A. & Rysman, M. (2005). Unobservable product differentiation in discrete choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics*, 36, 771–788.
- Aizcorbe, A. (2005). Price deflators for high technology goods and the new buyer problem. Unpublished manuscript, Board of Governors of the Federal Reserve Bank.
- Bajari, P. & Benkard, C. L. (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of Political Economy*, 113, 1239–1276.
- Berry, S. (1994). Estimating discrete choice models of product differentiation. *RAND Journal of Economics*, 25, 242–262.
- Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63, 841–890.
- Berry, S., Levinsohn, J., & Pakes, A. (2004). Estimating differentiated product demand systems from a combination of micro and macro data: The market for new vehicles. *Journal of Political Economy*, 112, 68–105.
- Berry, S. & Pakes, A. (2005). The pure characteristics model of demand. *International Economic Review*, in press.
- Bresnahan, T. F. (1981). Departures from marginal cost pricing in the american automobile industry: Estimates for 1997-1978. *Journal of Industrial Economics*, 17, 201–277.
- Carranza, J. (2006). Demand for durable goods and the dynamics of quality. Unpublished manuscript, University of Wisconsin.
- Chintagunta, P., Dube, J., & Nair, H. (2004). Empirical analysis of indirect network effects in the market for personal digital assistants. *Quantitative Marketing and Economics*, 2(1), 23–58.
- Clements, M. & Ohashi, H. (2005). Indirect network effects and the product cycle: Video games in the U.S., 1994-2002. *Journal of Industrial Economics*, 53, 515–542.
- Einav, L. (2007). Seasonality in the U.S. motion picture industry. *RAND Journal of Economics*, 38, in press.
- Esteban, S. & Shum, M. (2007). Durable goods oligopoly with secondary markets: The case of automobiles. *RAND Journal of Economics*, in press.
- Gandal, N., Kende, M., & Rob, R. (2000). The dynamics of technological adoption in hardware/software systems: The case of compact disc players. *RAND Journal of Economics*, 31, 43–61.
- Gentle, J. E. (2003). *Random number generation and Monte Carlo methods* (2 ed.). New York: Springer.
- Goolsbee, A. & Petrin, A. (2004). The consumer gains from direct broadcast satellites and the competition with cable television. *Econometrica*, 72, 351–381.

- Gordon, B. (2006). Estimating a dynamic model of demand for durable goods. Unpublished manuscript, Carnegie Mellon University.
- Hendel, I. & Nevo, A. (2007). Measuring the implications of sales and consumer stockpiling behavior. *Econometrica*, in press.
- Magnac, T. & Thesmar, D. (2002). Identifying dynamic discrete decision processes. *Econometrica*, 70, 801–816.
- Melnikov, O. (2001). Demand for differentiated products: The case of the U.S. computer printer market. Unpublished manuscript, Cornell University.
- Nair, H. (2007). Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video games. *Quantitative Marketing and Economics*, in press.
- Nevo, A. (2000). A practitioner's guide to estimation of random coefficients logit models of demand. *Journal of Economics & Management Strategy*, 9, 513–548.
- Nevo, A. (2001). Measuring market power in the ready-to-eat breakfast cereal industry. *Econometrica*, 69, 307–342.
- Ohashi, H. (2003). The role of network effects in the U.S. VCR market, 1978–86. *Journal of Economics and Management Strategy*, 12, 447–494.
- Park, S. (2004). Quantitative analysis of network externalities in competing technologies: The VCR case. *Review of Economics and Statistics*, 86, 937–945.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110, 705–729.
- Prince, J. (2007). Repeat purchase amid rapid quality improvement: Structural estimation of demand for personal computers. Unpublished manuscript, Cornell University.
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, 55, 999–1033.
- Schiraldi, P. (2007). Automobile replacement: A dynamic structural approach. Unpublished Manuscript, Boston University.
- Song, I. & Chintagunta, P. (2003). A micromodel of new product adoption with heterogeneous and forward-looking consumers: An application to the digital camera category. *Quantitative Marketing and Economics*, 1, 371–407.
- Zhao, Y. (2007). Why are prices falling fast? An empirical study of the US digital camera market. Unpublished Manuscript, Yale University.