

Dynamics of Consumer Demand for New Durable Goods*

Gautam Gowrisankaran
John M. Olin School of Business
Washington University in St. Louis
and NBER

Marc Rysman
Department of Economics
Boston University

May 11, 2006

Preliminary draft, Comments welcome

Abstract

This paper specifies and estimates a dynamic model of consumer preferences for new consumer durable goods. Most new durable consumer goods 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 specify a dynamic model of consumer preferences with persistent heterogeneous consumer tastes and estimate the model on the DVD player industry, using a panel data set on prices, sales and characteristics. Consumers in our model choose between purchasing a current product and waiting for future products, making rational forecasts about the future distribution of prices and qualities. Our model allows consumers who have already purchased a durable good to upgrade to a new model as desired. We find that dynamics are a very important determinant of consumer preferences. We use the estimates to investigate the value of new consumer goods and the determinants of price declines.

*We thank Igal Hendel, Rob Porter and seminar participants at several institutions for helpful comments, Calixte Ahokpossi for research assistance and Firat Inceoglu and Minsoo Park for assistance in providing access to these data. 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 increasing 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 cameras and DVD players. For instance, for DVD players, from their introduction in 1997 to 2003 average prices dropped from about \$730 to \$180 while many important features diffused rapidly, such as progressive scan (which enhances picture quality) which first appeared in 1999 and was in almost 50 percent of the models by 2003. The evolving nature of product attributes for new durable consumer goods suggests the importance of modeling dynamics when estimating consumer preferences. A rational consumer who contemplated purchasing a DVD player in 1997 would need to form expectations about the future path of price and quality, in order to decide whether to purchase a model then or wait for future models. Rational expectations almost certainly implied a high probability that prices would drop dramatically. A model of consumer preferences would need to incorporate dynamics to capture the fact that many people chose not to purchase a DVD player initially because they thought that prices would fall dramatically. Moreover, a dynamic model is necessary to account for the fact that the set of potential customers, and in particular the set of *high value* potential customers, is not stable over time, because an individual who has purchased a DVD player is not likely to replace the player until features improve substantially.

Measuring dynamic consumer preferences accurately for durable goods allows for the investigation of a variety of research questions that are of interest to both researchers and policymakers. One such question is the extent to which new industries have resulted in consumer welfare gains, which 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.¹ Yet, most empirical papers that examine similar industries have used static models of consumer demand. Welfare measures that are not based on demand models of demand may be biased. 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 realizing that consumers hold off from making purchases because of the expectation of lower prices and better features.

Another important research question is understanding the extent to which dynamic price discrimination explains the observed pattern of declining prices. In particular, firms may price new products high in order to extract value from

¹These include Goolsbee & Petrin (2004) for satellite cable, Park (2004) and Ohashi (in press) for VCRs, Clements & Ohashi (in press) for video games, Chintagunta, Dube & Nair (2004) for personal digital assistants and Einav (2004) for movie-going.

high value customers and gradually decrease prices in order to sell to a greater share of the market. This strategy is sometimes called “cream-skimming” and has been widely discussed in the literature. Cream-skimming may be useful if there is substantial heterogeneity in consumer preferences and that heterogeneity is persistent over time. Yet, dynamic price discrimination is certainly not the only explanation for the observed decrease in prices. Equally plausible explanations include declines in factor input costs and increases in competition through entry.² By estimating the dynamics of demand for durable goods, one can then evaluate the relative importance of these explanations in observed price decreases.

This paper specifies and estimates a structural dynamic model of consumer preferences for new consumer durable goods. We use the model to better understand the welfare gains from innovation and dynamic price discrimination for the DVD player industry. Our methods are also potentially applicable to other industries and other questions that require uncovering the dynamics of consumer preferences.

In order to capture the relevant features of new consumer durable goods industries, we require a consumer model with several important elements. First, the model must be dynamic, and allow consumers to have reasonable expectations over the attributes of future products. The model must also allow for persistent consumer heterogeneity. Berry, Levinsohn & Pakes (1995), henceforth BLP, and the literature that follows have shown the importance of incorporating consumer heterogeneity into demand systems for obtaining realistic predictions of elasticities and welfare estimates. This is even more true in the context of a dynamic model, where, for example, dynamic price discrimination strategies would never be optimal without some persistent heterogeneity. In addition, we require a model that allows for consumers to replace their durable good if features improve sufficiently to make this optimal. Finally, we seek a model that allows for the endogeneity of price, as in BLP, and that can be applied to aggregate data, which is now available for many industries.

Given these criteria, we specify a dynamic model of consumer preferences for differentiated products that allows for persistent random coefficients and repeat purchases over time and develop a method to estimate the model using aggregate data. Much of our model of consumer preferences is essentially the same as BLP: consumers in our model make a discrete choice from a set of products; they have random coefficients over observable product characteristics; the econometric unobservable is interpreted as an unobserved product characteristic; consumers receive *i.i.d.* Type 1 extreme value random components of utility at the time of purchase; and prices are endogenous. As in BLP, provided that one can recover the unobserved product characteristics, it is possible to construct a generalized method of moments (GMM) estimator using moment conditions

²These are all classic explanations for falling prices. For instance, Stokey (1979) lists these three explicitly to explain falling prices for pocket calculators. Learning-by-doing is another explanation for falling prices that is sometimes suggested. However, firms that rationally predict learning-by-doing will not typically lower prices over time. See Spence (1981) for a theoretical analysis and Benkard (2004) for a simulation analysis.

constructed from the orthogonality between unobserved product characteristics and exogenous variables.

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. The dynamics of our model are as follows. All consumers start out as potential customers at the introduction of the product. Consumers who have not purchased any product obtain some base flow utility that we normalize to zero. Once a given consumer purchases a product, she can use the product as long as she wants and in so doing would obtain the same flow utility from the purchased product in every future period. A consumer that has previously purchased a product can replace it with a different product in the future, which will be optimal if the flow utility from the product is sufficiently high to justify the current expenditure.

Consumers do not know the future set of products, but they instead perceive a distribution for these future products, and use this information to make a rational choice between purchasing now and waiting. Rather than modeling the supply side explicitly, we make a major simplifying assumption: that consumers perceive that the evolution of product characteristics will follow a simple one-dimensional Markov process, where the distribution of next period's product characteristics is a polynomial function of a simple statistic (the logit inclusive value) of the expected discounted utilities of current products.³ We assume rational expectations within the context of this simple expectations framework, in the sense that consumers' expectations will be the actual empirical distribution of quality changes.

Demand for these durable goods may be affected by characteristics of the environment in addition to physical product characteristics. Most notably, the number of DVD titles may affect the demand for DVD players. These characteristics are different than physical characteristics in that they are not constant over the life of the product. With a simple extension of our specification, we allow for utility to depend on changing environmental characteristics, assuming perfect foresight about the values of future environmental characteristics.

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. Unlike more stable industries (e.g., automobiles) 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. 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. Recent work on the second-hand automobile market by Esteban & Shum (2005) estimates a model

³Hendel & Nevo (2003) make a similar simplifying assumption to estimate the dynamics of demand for the laundry detergent market.

with forward-looking consumers and firms, using a simple vertical model where consumers must purchase a car every period. 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 can only ever purchase one product.

Several recent papers generalize the Melnikov (2001) idea. Carranza (2004) extends Melnikov's model to allow for a random coefficient on the constant term, effectively on the option to wait, for the digital camera market. 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 (2005) 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. Nair (2005) estimates the demand and supply for video games allowing for random coefficients and endogenous prices, treating each video game as a monopoly. In addition, some recent papers have estimated the dynamics of durable goods, but have focused instead on modeling the dynamics of the supply side and not of the demand side.⁴

Our paper builds on the literature on estimating dynamic demand, in that our model allows for unobserved product characteristics 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 methodology 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 computes the shares for each product conditional on a vector of mean product characteristics, and then uses these shares to define a new vector of mean product characteristics. We use a similar process to invert the share equation. However, for a set of mean product characteristics, we explicitly evaluate the

⁴See Aizcorbe & Kortum (2004) for CPUs, Copeland, Dunn & Hall (2005) for automobiles and Nair (2005) for video games.

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 nesting these two separate methods. The use of this inversion method to recover the mean product characteristics allows us to avoid the computational and modeling burden associated with estimating a full equilibrium model.

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, explain the extension of our model to allow for time-varying environmental characteristics 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 $t = 0$, at which time there is a continuum of heterogeneous potential consumers indexed by i . The unit of observation is a month. Consumers and firms 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.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 her expected discounted value of future expected utility, conditional on her information at time t . Product j at time t is characterized by observed physical product characteristics x_{jt} and price p_{jt} and an unobserved (to the econometrician) product characteristic ξ_{jt} . For instance, for DVD players, physical product characteristics include progressive scan, Dolby sound technology and built-in recorders. 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 product characteristics of a product j purchased at time t , x_{jt} and ξ_{jt} , stay constant over the infinite life of the product.⁶ We assume that consumers

⁵We 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.

⁶Subsection 2.3 adds environmental characteristics, notably the number of DVD titles, which may provide an important network benefit, and which will vary over time for a given product.

and firms know all product characteristics of current products at the time of purchase. Let

$$\delta_{ijt}^f = \alpha_i^x x_{jt} + \xi_{jt} \quad (1)$$

denote the gross flow utility at time t or greater from product j purchased at time t . At any time period, some consumers may have not previously purchased a product. We assume that these consumers are using an outside good (e.g., a VCR instead of a DVD player), and normalize the flow utility from the outside good to 0. We do not model any explicit linkage between products offered for sale at different time periods. Thus, even if two products j, t and $j, t+1$ have the same make and model and the same observable characteristics x_{jt} and $x_{j,t+1}$, we do not restrict their unobserved characteristics ξ_{jt} and $\xi_{j,t+1}$ to be the same.

We assume that a consumer purchasing product j at time t would receive a net utility at time t of

$$u_{ijt,t} = \delta_{ijt}^f - \alpha_i^p \ln(p_{jt}) + \epsilon_{ijt}. \quad (2)$$

The first component of (2) is the gross flow utility and the second component is the price term, which is only paid once, at the time of purchase. The last component, ϵ_{ijt} , is an idiosyncratic unobservable that is distributed as the difference between two Type I extreme value draws that are independent across each other, consumers, products and time. This idiosyncratic unobservable only accrues to the consumer at the time of purchase, and will capture random variations, for instance in sales personnel and store displays.⁷ When making her time t decision, the consumer knows the values of all the time t variables in (2). Because the latter two components of (2) are non-recurring, the net flow utility from the time t purchase at time $\tau > t$, expressed in period t utils is

$$u_{ijt,\tau} = \beta^{\tau-t} \delta_{ijt}^f. \quad (3)$$

We let α_i be distributed normally with mean $\alpha \equiv (\alpha^x, \alpha^p)$ and variance matrix Σ , where α and Σ are parameters to estimate. We let α_i remain constant over time for a given consumer. Note that the functional form for utility at the time of purchase, (2), fits within the random coefficients discrete choice framework of BLP.

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 rationally expect them to evolve based on the current market environment, as we explain more formally below.

⁷The logit error assumption is common but certainly not innocuous. In contexts where consumers face different numbers of products (as our consumers do over time), the logit error can imply unrealistic assumptions about product crowding. See Berry & Pakes (2005); Bajari & Benkard (2005) and Akerberg & Rysman (ress) for solutions that we could potentially implement in our setting.

We now define the state variables and use them to exposit the dynamic decision process. Let ω_t represent the set of current product attributes x_{jt} , p_{jt} and ξ_{jt} for all products available in period t . Let $\epsilon_{i,t} \equiv (\epsilon_{i1t}, \dots, \epsilon_{iJ_t t})$ denote the set of idiosyncratic unobservable draws for consumer i at period t . Then, the purchase decision for consumer i depends on the following: her coefficients α_i ; her flow utility δ_0^f for the product that she already owns; her idiosyncratic unobservable draws $\epsilon_{i,t}$; and the current and future realizations of product attributes. Future product attributes will depend on firm behavior which is a function of the current market environment, including current product attributes and consumer product holdings, and other factors such as technological progress and consumer characteristics. Group together the current product attributes and all other market characteristics (such as consumer holdings) at time t as Ω_t . We assume that Ω_{t+1} evolves according to some Markov process $P(\Omega_{t+1}|\Omega_t)$ that is consistent with all relevant industry factors including technological process, optimal pricing decisions, and the entry of new products and the exit of current products. Thus, the state vector for consumer i is $(\epsilon_{i,t}, \delta_0^f, \Omega_t)$.

We can now define the Bellman equation for consumer i as

$$V_i(\epsilon_{i,t}, \delta_0^f, \Omega_t) = \max \left\{ \delta_0^f + \beta E \left[V_i(\epsilon_{i,t+1}, \delta_0^f, \Omega_{t+1}) \middle| \Omega_t \right], \right. \\ \left. \max_{j=1, \dots, J_t} \left\{ u_{ijt,t}(\omega_t(\Omega_t)) + \beta E \left[V_i(\epsilon_{i,t+1}, \delta_{ijt}^f, \Omega_{t+1}) \middle| \Omega_t \right] \right\} \right\}, \quad (4)$$

where “ E ” denotes the expectation operator, a conditional expectation in this case. From (4), the consumer can choose to wait and keep her current product (the first option), or purchase any of the available products (the next J_t options). Note that the value of waiting is greater than the discounted stream of flow utilities $\delta_0^f/(1-\beta)$ because waiting encapsulates the option to buy a better product in the future.

We now introduce three further definitions that are necessary to explain our method of inference and later assumptions. First, define the expectation of the value function, integrated over realizations of $\epsilon_{i,t}$, as

$$EV_i(\delta_0^f, \Omega_t) = \int_{\epsilon_{i,t}} V_i(\epsilon_{i,t}, \delta_0^f, \Omega_t) dP_\epsilon. \quad (5)$$

Next, for each product $j = 1, \dots, J_t$ let

$$\delta_{ijt} = \delta_{ijt}^f - \alpha_i^p \ln(p_{jt}) + \beta E \left[EV_i(\delta_{ijt}^f, \Omega_{t+1}) \middle| \Omega_t \right]. \quad (6)$$

Note that δ_{ijt} is the part of the expected discounted utility for consumer i purchasing product j at time t that is not due to the idiosyncratic unobservable ϵ_{ijt} . Last, define the *logit inclusive value* for consumer i at time t as:

$$\delta_{it} = \ln \left(\sum_{j \in J_t} \exp(\delta_{ijt}) \right). \quad (7)$$

The inclusive value δ_{it} captures the expected value of buying in period t . Because of the logit error assumption, if consumer i knew δ_{it} for each period, that would be sufficient for the consumer to optimally choose when to purchase (see Rust, 1987; Melnikov, 2001; Akerberg, 2003). The consumer would not need further knowledge of individual products. In this sense, we can separate the consumer’s problem into two parts, a choice of when to buy based on δ_{it} and, given purchase, the choice of what to buy based on available products at time t . However, for both the choice of when to buy and the choice of what to buy, the consumer must make predictions of future product characteristics, that is of δ_{it} in future periods.

The future value of δ_{it} could depend on the entire state space Ω_t . In order to solve the consumer’s dynamic optimization problem, we need to specify industry evolution $P(\Omega_{t+1}|\Omega_t)$. The large potential dimensionality of Ω_t makes it difficult to solve (4). In the interest of tractability, we make an important simplifying assumption about the structure of industry evolution over time. We assume that the evolution of product attributes follows a one-dimensional Markov process in the logit inclusive value. Specifically, we assume that

$$P(\delta_{i,t+1}|\Omega_t) = P_i(\delta_{i,t+1}|\delta_{it}). \quad (8)$$

In other words, if two industry structures impose the same δ_{it} for any consumer i at time t , then they result in the same distribution of industry structures at time $t + 1$.

This assumption is not without loss of generality. 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. If these scenarios result in the same δ_{it} , our assumption implies they must imply the same expectation of δ_{it+1} . A similar discussion appears in Hendel & Nevo (2003).

The benefit of this simplifying assumption is that, because of the logit error assumption, the state space for the expectation of the value function is vastly reduced to two dimensions: one for the inclusive value, plus one for the flow utility δ_0^f . Thus, we can rewrite the expectation Bellman equation as

$$EV_i(\delta_0^f, \delta_{it}) = \ln \left(\exp(\delta_{it}) + \exp \left(\delta_0^f + \beta E \left[EV_i(\delta_0^f, \delta_{i,t+1}) \middle| \delta_{it} \right] \right) \right). \quad (9)$$

We now mention a couple of points about this simplifying assumption. First, from (8), we allow the conditional density $P_i(\delta_{i,t+1}|\delta_{it})$ to vary by type i , as is necessary given by the fact that $\delta_{i,t+1}$ incorporates both elements of future products and characteristics of consumers of type i . Clearly, with rational expectations, the conditional density $P(\Omega_{t+1}|\Omega_t)$ would not vary across types i . While (8) never imposes this invariance assumption, it is certainly consistent with it. Second, while it would be difficult to implement our method with transition probabilities based in arbitrary ways on Ω_t , it would not be difficult to expand our state space beyond one dimension. For instance, we could add in the number of products as another state variable, which implies that we would

then allow the density of $\delta_{i,t+1}$ and the number of products at time $t + 1$ to depend both on δ_{it} and the number of products at time t . While this would not pose any theoretical problems, it does increase the dimensionality of the state space from 2 to 3.

We assume that consumers have rational expectations about the future. Specifically, we assume that consumers perceive the empirical density of $P(\delta_{i,t+1}|\delta_{it})$ fitted to a simple functional form. We use a simple linear autoregressive specification with drift,

$$\delta_{i,t+1} = \gamma_{1i} + \gamma_{2i}\delta_{it} + u_{it}, \quad (10)$$

where u_{it} is assumed to be normally distributed with mean 0 and where γ_{1i} and γ_{2i} are parameters. Note that we can estimate this functional form with a simple 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 (10) to allow add a quadratic term of the form δ_{it}^2 . This would not substantially increase computation time.

As noted above, purchase probabilities at any time period can be expressed as a simple function of the logit inclusive values and expected future values. Specifically, the probability that a consumer of type i who owns a product with a vector of flow utilities δ_0^f purchases a new product in period t is

$$\hat{s}_{it}(\delta_0^f, \delta_{it}) = \frac{\exp(\delta_{it})}{\exp(\delta_{it}) + \exp\left(\delta_0^f + \beta E\left[EV_i(\delta_0^f, \delta_{i,t+1}) \mid \delta_{it}\right]\right)} \quad (11)$$

and the probability that this consumer purchases product j conditional on purchasing any product is

$$\hat{s}_{ij|t}(\delta_0^f, \delta_{it}, \omega_t) = \frac{\exp(\delta_{ijt})}{\exp(\delta_{it})}. \quad (12)$$

Note that (12) requires ω_t while (11) does not, because only (12) incorporates the choice of an individual product.

Using this specification, we can calculate the aggregate predicted share of each product at any time period. Doing so requires integrating over the random coefficients α and consumer holdings. Let $P(\alpha)$ denote the density of these coefficients, and let $P_{\alpha t}(\delta_0^f)$ denote the density of flow quality for consumer holdings conditional on having random coefficients α at the start of time t . Then, we can express the predicted share at any time period, as a function of the ξ vector and parameters, as

$$\hat{s}_{jt}(\xi, \Omega_t, \alpha, \Sigma, \beta) = \int_{\alpha} \int_{\delta_0^f} \hat{s}_{it}(\delta_0^f, \delta_{it}) \hat{s}_{ij|t}(\Omega_t) dP_{\alpha t}(\delta_0^f) dP(\alpha). \quad (13)$$

One can evaluate market shares in (13) by iterating over time. Specifically, for one vector of α coefficients, one can calculate (13) at time $t = 0$. This is straightforward since all consumers hold the outside good, which has flow quality

0 at time $t = 0$. One can use this calculation to calculate $P_{\alpha 1}(\delta_0^f)$, use this to simulate (13) at $t = 1$, and repeat this process until the terminal period. One can then integrate across α coefficients to calculate aggregate market shares.

We also note that our specification of price can be extended. In our utility function, we have not allowed for the distribution of willingness-to-pay to relate to income data in any way. Part of the reason for this is that the price of a DVD player or a digital camera is very small relative to average household income, so income effects are likely small. Yet, other discrete choice studies, such as Nevo (2001) allow for income effects for even smaller purchases, of breakfast cereal. Adding income data provides a natural way to have a richer specification of the price elasticity of demand.

A final issue is that we restrict consumers to value only one product at a time. We make this assumption because it would likely be difficult to identify the relative value of a second consumer electronic product from our data. However, we could potentially identify this type of parameter by including micro-moments (see Berry, Levinsohn & Pakes, 2004; Petrin, 2002) on penetration rates or individual purchasing behavior.⁸

We have thus far not discussed the supply side of the model. It is not necessary to fully specify the supply side in order to estimate demand; an assumption that product characteristics are exogenous is sufficient to estimate consumer preferences. Nonetheless, we would need to specify the firm side in order to understand the determinants of falling prices. While we do not estimate the supply side, we specify and calibrate a simple dynamic monopoly model, using the estimated demand parameters. In this model, the arrival of products is exogenous. The firm makes pricing decisions, taking into account the expected future evolution of the products, and the set of people who have already purchased and who are in the market. Importantly, the firm cannot commit to a future price path, even though it might want to commit. We feel that this assumption is reflective of the real world for new durable consumer goods, where any commitment to price exists only in the very short run. It is also possible to extend this type of specification to a dynamic oligopoly model.

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, instead of attempting to estimate β , we set $\beta = .99$, at the level of the month.

⁸Such data is available from the Consumer Expenditure Survey and superior data sets are available (see Prince, 2005; Karaca-Mandic, 2004).

We develop a method for estimating the remaining parameters that is based on Berry (1994) and Rust (1987) and the literatures that follow. Our estimation algorithm involves three levels of non-linear optimizations: on the outside is a non-linear search over the parameters; inside that is a fixed point calculation of the vector of unobserved product characteristics ξ , and inside that is the calculation of predicted market shares, which is based on consumers' dynamic optimization problems. While both the ξ fixed point calculation and the dynamic programming estimation are well-known, our innovation is in nesting the dynamic programming solution within the ξ fixed point calculation in order to develop a feasible estimator of dynamic consumer preferences. We describe each of the three levels of optimization in turn.

Starting with the outer loop, we specify a GMM criterion function

$$G(\alpha, \Sigma) = z'\xi(\alpha, \Sigma), \quad (14)$$

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 below. Our estimated parameters satisfy

$$\left(\hat{\alpha}, \hat{\Sigma}\right) = \operatorname{argmin}_{\alpha, \Sigma} \left\{ G(\alpha, \Sigma)' W G(\alpha, \Sigma) \right\}, \quad (15)$$

where W is a weighting matrix. Because $\alpha^x x_{jt}$ enters the utility function linearly and additively with ξ_{jt} , in BLP-type models, the optimal values of α^x can be written as a closed-form function of the other parameters.⁹ The dynamics of our model do not alter this simplification.¹⁰ Hence, we need only perform a non-linear search over (α^p, Σ) . We perform the 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, and then use our first stage estimates to approximate the optimal weighting matrix.¹¹

We now turn to the middle loop, the computation of ξ_{jt} , which we perform by exploiting the fixed point equation that appears in Berry (1994) and BLP. Given the closed-form solution to α^x , we define $\delta_{jt}^x \equiv \alpha^x x_{jt} + \xi_{jt}$ and perform the following fixed point on δ_{jt}^x :

$$\delta_{jt}^{x, new} = \delta_{jt}^{x, old} + \psi \cdot (\ln(s_{jt}) - \ln(\hat{s}_{jt}(\delta^{x, old}, \alpha^p, \Sigma, \beta))), \quad (16)$$

where $\hat{s}_{jt}(\delta^{x, old}, \alpha^p, \Sigma, \beta)$ is calculated from (13) and ψ is a tuning parameter that we generally set to $1/(1 - \beta)$.¹²

⁹See Nevo (2000) for a discussion.

¹⁰Unlike BLP, 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} .

¹¹See again Nevo (2000) for details.

¹²One issue relates to the properties of (16). Berry provides conditions under which this function is a contraction mapping, guaranteeing that $\hat{s}_{jt}(\delta^x, \alpha^p, \Sigma, \beta)$ is invertible in the vector of δ 's. In our case, we have found examples where this inversion is not a contraction mapping, evidently because 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.

We integrate the vector of market shares across random coefficients in (13) using simulation techniques similar to BLP. Specifically, we write

$$\alpha_i = \alpha + \Sigma^{1/2} \tilde{\alpha}_i, \quad (17)$$

where $\tilde{\alpha}_i \sim \phi_l(\tilde{\alpha}_i)$, the standard normal density with dimensionality l , the length of the vector α_i . We can sample over $\tilde{\alpha}_i$ and scale the draws based on $(\tilde{\alpha}, \tilde{\Sigma})$ to construct α_i using (17). However, since our estimation algorithm is very computationally intensive and computational time is roughly proportional to the number of simulation draws, we use importance sampling to reduce sampling variance, as in BLP. Let $\hat{s}_{sum}(\tilde{\alpha}_i, \alpha, \Sigma)$ denote the sum of predicted market shares of any durable good at any time period for an individual with parameters (α, Σ) and draw $\tilde{\alpha}_i$. Then, instead of sampling from the density $\phi_l(\tilde{\alpha}_i)$ we sample from the density

$$f(\tilde{\alpha}_i) \equiv \frac{\hat{s}_{sum}(\tilde{\alpha}_i, \alpha, \Sigma) \phi_l(\tilde{\alpha}_i)}{\int \hat{s}_{sum}(\tilde{\alpha}, \alpha, \Sigma) \phi_l(\tilde{\alpha}) d\tilde{\alpha}}, \quad (18)$$

and then reweight draws by

$$w_i \equiv \frac{\int \hat{s}_{sum}(\tilde{\alpha}, \alpha, \Sigma) \phi_l(\tilde{\alpha}) d\tilde{\alpha}}{\hat{s}_{sum}(\tilde{\alpha}_i, \alpha, \Sigma)}, \quad (19)$$

in order to obtain the correct expectation. As in BLP, we sample from the density $f(\tilde{\alpha}_i)$ by sampling from the density $\phi_l(\tilde{\alpha}_i)$ and using an acceptance/rejection criteria. We compute (18) using a reasonable guess of (α, Σ) and use a large number of simulation draws to obtain an accurate ξ vector, necessary to compute $\hat{s}_{sum}(\tilde{\alpha}_i, \alpha, \Sigma)$. Instead of drawing *i.i.d.* pseudo-random normal draws for $\phi_l(\tilde{\alpha}_i)$, we use Halton sequences based on the first l prime numbers, to further reduce the sampling variance (see Gentle (2003)).

Last, the inner loop solves the consumer dynamic programming problem conditional on a δ^x vector and nonlinear parameters (or equivalently, conditional on a ξ vector and parameters), separately for each α_i draw. For each draw, we iteratively update the product evolution Markov process from the regression in (10), the expectation Bellman from (9) and the logit inclusive values from (7), until convergence.¹³ To perform the computation, we discretize both the vectors δ_{it} and δ_0^f (which define the state space for (9)) and the transition matrix (which is based on the estimated coefficients and standard error from the linear regression (10)).¹⁴

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 simplified: conditional on the ξ vector and parameters, (7) can be solved in closed-form, because the undiscounted flow utility from the purchase of product j at time t from (1) is always just δ_{jt}^f . This then implies

¹³We found that computational time was significantly shortened by performing (9) much more frequently than either (10) or (7).

¹⁴Note that the model implicitly assumes that evolution process for δ_{it} continues out of sample. Thus, different maximum values of the grid may result in different parameter estimates, suggesting the need for robustness exercises.

that (10) can also be solved in closed-form at this step. Moreover, this model also results in a more computationally efficient outer loop, since the price coefficient α^p can also be solved in closed-form for this model, like α^x in the base model.

2.3 Time-varying environmental characteristics

This subsection extends the model to add environmental characteristics, which we did not consider earlier for ease of exposition. We allow the gross flow utility at time $\tau \geq t$ from product j purchased at time t to include a time-varying component, $\alpha_i^{x^e} x_\tau^e$, where $\alpha_i^{x^e}$ will be distributed normally, with mean α^{x^e} and a variance term that enters into Σ . This component is not specific to the product, and is meant to represent features such as the number of available DVD titles.

Consider a consumer who has never purchased before time t . Equivalent to the time-varying characteristic, we can renormalize utility by letting the consumer receive a one-time increase in expected utility from purchase at time t of

$$n_{it} \equiv \alpha_i^{x^e} E_{it} [x_t^e + \beta x_{t+1}^e + \beta^2 x_{t+2}^e + \dots]. \quad (20)$$

so that the net flow utility from purchase at time t is

$$u_{ijt,t}^n = \delta_{ijt}^f - \alpha_i^p \ln(p_{jt}) + n_{it} + \epsilon_{ijt} \quad (21)$$

instead of (2) but where the gross flow utility δ_{ijt}^f is exactly as in (1). We assume perfect foresight about future environmental characteristics so that n_{it} can be calculated using the observed path of x_t^e , conditional on $\alpha_i^{x^e}$, assuming that the number of DVD titles remains constant after the end of the sample.¹⁵

The renormalization vastly simplifies the computational burden of this model, because it implies that the environmental characteristics do not have to be treated as a separate state variable. To see this, note that with the renormalization, environmental characteristics are irrelevant for people who have already purchased a product at some time in the past, since they have already received n_{it} . Thus, for individuals who have already purchased a product, the decision problem is exactly equivalent to the model specified in Subsection 2.1.

For individuals who have never purchased a product, the problem is slightly different than earlier, because they will obtain n_{it} the first time that they purchase the product, and need to include forecasts of future values of this variable in their decision problem. We start with definitions for these individuals that are similar to those made in Subsection 2.1. Let the Bellman equation for individuals who have never purchased a product, similar to (4), be

$$V_i^n(\epsilon_{i,t}, \Omega_t) = \max \{ \beta E [V_i^n(\epsilon_{i,t+1}, \Omega_{t+1}) | \Omega_t], \quad (22) \\ \max_{j=1, \dots, J_t} \left\{ u_{ijt,t}^n(\omega_t(\Omega_t)) + \beta E \left[V_i \left(\epsilon_{i,t+1}, \delta_{ijt}^f, \Omega_{t+1} \right) \middle| \Omega_t \right] \right\} \},$$

¹⁵In principle, it is straightforward to generalize the assumption of perfect foresight by assuming idiosyncratic forecast errors, but it is hard to see how the variance of the forecast error would be identified.

where Ω_t and ω_t now include environmental characteristics. Let $EV_i^n(\Omega_t)$, be defined analogously to (5). Further, define

$$\delta_{ijt}^f(\Omega_t) = \delta_{ijt}^f + n_{it} - \alpha_i^p \ln(p_{jt}) + \beta E \left[EV_i \left(\delta_{ijt}^f, \Omega_{t+1} \right) \middle| \Omega_t \right], \quad (23)$$

and $\delta_{it}^n(\Omega_t)$ analogously to (7) as

$$\delta_{it}^n(\Omega_t) = \ln \left(\sum_{j \in J_t} \exp(\delta_{ijt}^n(\Omega_t)) \right). \quad (24)$$

Using these definitions, we make a similar assumption to (8) for individuals who have never purchased a product, that

$$P(\delta_{i,t+1}^n | \Omega_t) = P_i^n(\delta_{i,t+1}^n | \delta_{it}^n) \text{ and } P(\delta_{i,t+1} | \Omega_t) = P_i^n(\delta_{i,t+1}^n | \delta_{it}^n). \quad (25)$$

This then leads to a similar expectation Bellman as (9):

$$EV_i^n(\delta_{it}^n) = \ln \left(\exp(\delta_{it}^n) + \exp(\beta E [EV_i(\delta_{i,t+1}^n) | \delta_{it}^n]) \right). \quad (26)$$

We define expectations for future product quality for individuals who have never purchased a product with a specification that is similar to (10),

$$\begin{aligned} \delta_{i,t+1}^n &= \gamma_{3i} + \gamma_{4i} \delta_{it}^n + u_{2it}, \\ \delta_{i,t+1} &= \gamma_{5i} + \gamma_{6i} \delta_{it}^n + u_{3it}. \end{aligned} \quad (27)$$

Computation of this model requires only a slight extension to the methods outlined in Subsection 2.2. In particular, we now separately compute decision problems for individuals who have purchased a previous good and for those who have not. As in Subsection 2.2, for the inner loop we update on (10), (9) and (7) to solve for the decision problems of agents who have already purchased a product. Now we also update (27), (26) and (24) to solve for the decision problems of agents who have never purchased a product. Note that α^{xe} is a non-linear parameter, unlike α^x , because it does not enter additively with ξ . To aggregate market shares as in (13) we sum over the measure of the two types of individuals at any time period, noting that everyone starts out in the group of never having purchased a product. While the extension includes new notation, the dimensionality of the computation remains the same as the original problem, and the computational time is roughly double.

2.4 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. 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 the high valuation 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, one might wonder how it is possible to identify such a model. However, note that 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 even attempt to 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 new consumer goods.

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 to serve 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.

We believe that an instrumenting strategy in which product characteristics are assumed to be exogenous is a potential weakness of this paper and of the discrete choice literature in general. Available products may evolve in response to unobserved features of consumer demand,¹⁶ and in this case, the assump-

¹⁶For instance, de Figueiredo & Kyle (2004) find a broadening in product quality in both directions over time for laser printers.

Table 1: Characteristics of DVD players

Characteristic	Mean value
Price (Units: 1997 dollars)	\$304 (\$363)
Progressive scan	.241 (.428)
Component video	.804 (.397)
Optical audio	.841 (.366)
Coaxial audio	.837 (.369)
Dolby Digital audio	.308 (.462)
Digital theater system (DTS)	.126 (.332)
Plays CDR/RW disks	.463 (.499)
Plays MP3 files	.282 (.450)
Plays DVD-R disks	.140 (.346)
Includes VHS	.010 (.099)
Includes recorder	.019 (.418)
Multidisk player	.804 (.397)
Log DVD titles	6.69 (.795)

Number of observations: 9827; standard deviations in parentheses
Unit of observation: Model - Month

tion of exogeneity of product characteristics is not completely valid. While it is beyond the scope of the paper at present to fully endogenize product characteristics, a possible source of identification is to use the features of products only at the frontier for any given company as instruments, which is presumably more related to exogenous technological change than the set of all available products.

Another possibly endogenous variable is the number of DVD titles. Network effects between demand for DVD players and DVD titles suggests that unobserved DVD product characteristics may cause changes in the number of available DVD titles, although the number of DVD titles is more likely to respond to the stock of players, rather than the flow of players sold. It is possible to exploit the movie market to obtain instruments for the number of DVD titles. In particular, a 6 month lag of box office revenue presumably affects consumer valuation of titles but does not relate to unobserved DVD product characteristics.

3 Data

We estimate our model using a panel of aggregate data for DVD players. 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. For DVD players, we observe 522 models and 47 brands, with observations from March

1997 to August 2003. These data start from very early in the product life cycle of DVD players and include the vast majority of models. The data sets were originally collected by NPD Techworld which surveys major electronics retailers and covers 80% of the market.

Our characteristics include indicators for whether the model can play other formats besides standard DVDs (VHS tapes, MP3 files, DVD-R disks), indicators for features that improve the quality of audio and video (Dolby Digital audio, component video, coaxial audio cables, etc.) and an indicator for multisk players. We dropped portable DVD players from the data set, as we thought that these formed a different market, although they can be attached to a television and used as a substitute for a regular DVD player. The NPD data set also does not include substitutes to DVD players such as video game systems which can play DVDs. The data also reported some characteristics as missing for some observations. For about 200 model-months, the data did not report if the player could play CD-R/RW disks or MP3 files. For less than 100 model-months, the data were missing one of coaxial audio, Dolby/Digital, digital theater system (DTS) and progressive scan. We assumed that the models for which these features were missing did not have these features, since we thought that this was the most likely case. Table 1 summarizes our final set of price and characteristics data for DVD players at the level of the model-month.

Our base specification includes random coefficients for the constant term, which indicates the utility of a DVD player relative to using the outside option, and for price. Certain audio and video characteristics will add value only to the extent that consumers have complementary devices. For instance, progressive scan video is useful only with large-screen televisions. For this reason, it is also useful to include specifications with additional random coefficients.

Features of DVD players are generally improving over time, implying that characteristics are increasing over time. Figure 1 graphs the percentage of available models of DVD players that have various technological components. We graph three features: the fraction of models with progressive scan, component video connections, and Dolby Digital audio, over time. Some features, such as progressive scan, were only introduced in the middle of our sample, but caught on quickly after their introduction. Other features, such as component video, were available at the start of our sample period, and also rapidly increased in availability over time. A few features, such as Dolby Digital audio, never saw a huge increase in penetration, most likely because of the presence of other, superior substitute technologies.

Figure 2 shows total sales and average prices for DVD players over time. DVD players exhibit striking price declines over our sample period while sales increase correspondingly.

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. This would vastly complicate our model by adding state variables given that our demand system is dynamic. Given that prices do not change over Christmas, we

Figure 1: Fraction of DVD players with particular characteristics




Figure 2: Prices and sales for DVD players





Figure 3: Competition in the DVD player market



believed that this would vastly 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 2 also shows the seasonally-adjusted sales data, which are, by construction, much smoother than the unadjusted data.

A final 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 3 provides some evidence on the nature of competition in this market over time. The figure shows that the number of different DVD players is increasingly almost monotonically over time, suggesting that there will generally be closer and closer substitutes to a given player for later time periods. The Herfindahl index for DVD players is generally falling over time, although the relationship has several bumps, and hence is not close to monotonic. This pattern provides significant variation in the level of competition over time.

Table 2: Parameter estimates

Parameter	Estimates from dynamic model with repurchases allowed	Estimates from dynamic model with no repurchases	Estimates from static model
Mean coefficients (α)			
Constant	.316 (.117) *	-.280 (.279)	-102 (1,560)
Progressive scan	.060 (.003) *	.007 (.001) *	-.187 (.150)
Component video	.042 (.002) *	.009 (.002) *	.541 (.087) *
Optical audio	.014 (.002) *	-.005 (.001) *	-.794 (.081) *
Coaxial audio	.007 (.002) *	-.004 (.001) *	-.490 (.073) *
Dolby Digital audio	-.016 (.002) *	-.004 (.0007) *	-.307 (.069) *
Digital theater system (DTS)	.042 (.003) *	-.0017 (.0011)	-.953 (.140) *
Plays CDR/RW disks	.033 (.003) *	.012 (.0009) *	.933 (.073) *
Plays MP3 files	.002 (.008)	.013 (.0009) *	1.45 (.122) *
Plays DVD-R disks	.022 (.003) *	.0097 (.0010) *	.780 (.109) *
Includes VHS	.069 (.024) *	.028 (.0046) *	1.93 (.476) *
Includes recorder	.114 (.012) *	.011 (.003) *	-.760 (.246) *
Multidisk player	.037 (.002) *	.007 (.0008) *	.168 (.069)
Log price (1997 \$)	-7.48 (1.18) *	-2.50 (1.37)	.502 (5.12)
Log number of titles	.013 (.014)	.044 (.031)	-1.13 (.823)
Standard deviation coefficients ($\Sigma^{1/2}$)			
Constant	.082 (.014) *	.028 (.121)	34.6 (544)
Log price (1997 \$)	.979 (.381) *	1.24 (.946)	.442 (35.1)

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

4 Results and implications

We start by discussing our results, then provide evidence on the fit of the model and the implications of the results.

4.1 Results

We present our base parameter estimates in Table 2.¹⁷ Table 2 contains three columns of results. The first column of results provide the parameter estimates and standard errors from our base specification of the model presented in Section 2. The second column provides estimates from the dynamic model where indi-

¹⁷These results are extremely preliminary for several reasons: we have not completely verified that the code is free of bugs; we need to test alternate assumptions (such as the maximum quality level); we need to verify that the numerical approximations (i.e., the numbers of grid points and simulation draws) do not substantively affect the results; and we need to test that we are at the true minimum moment condition. Hence, our final parameter estimates are likely to be somewhat different.

viduals are restricted to purchase at most one DVD player. The third column provides estimates from a traditional static random coefficients discrete choice specification performed on the same data. The static results use the same specification for the static part of utility as in the base model, and essentially follow the Berry et al. (1995) model.

Starting with the first column of results in Table 2, the base specification reports results that are generally sensible in magnitude and sign. Price contributes negatively to utility for virtually everyone, with a base coefficient of -7.48 and a standard deviation on the random coefficient of $.979$. A person with mean tastes would obtain a positive gross flow utility from a base DVD player (relative to the outside option), with a mean constant term of $.316$. The standard deviation on this coefficient is $.082$, indicating that there is substantial variation in the gross flow utility from a DVD player but that, for most people, this term is nonetheless positive. In comparing the magnitudes of these coefficients, it is important to remember that price is paid once, while all the other coefficients relate to gross flow utility at the level of the month.

Most of the characteristics of DVD players enter utility positively. All of the characteristics except for the number of titles are indicator variables, implying that we can easily compare their magnitudes. Their magnitudes are all much smaller than the magnitude of the constant term, suggesting that these features are important, but not as important as owning a DVD player. Only one characteristic, Dolby Digital audio, enters utility negatively. This is likely due to the fact that this is one feature that did not diffuse widely over time, as shown in Figure 1. The coefficient on the number of DVD titles is positive and relatively large, when we consider that the standard deviation of this variable is $.795$. However, this coefficient is not statistically significant, unlike most of the other coefficients.

Turning now to column 2, the coefficients from the dynamic model that restricts a consumer to purchasing one time only are more often than not similar in sign, and sometimes similar in magnitude. In particular, the coefficient on price is negative, and barely not significant at the 5% level. Together with the random coefficient on price, these results indicate that the majority of people, but by no means all people, would have a negative utility from increased price. Nonetheless, these results show that there will be a substantial majority of people who like paying higher prices. If these coefficients had been applied to our base model which allows for multiple purchases, these individuals would generally buy a product most months, unless products have sufficiently negative values of ξ so that consumers obtain a negative gross utility flow from the product. In addition to being intuitively appealing, we found that this type of parameter resulted in a very bad fit for the moment criteria for the base model because of the negative ξ values. Thus, the multiple purchase feature of the base model essentially forces the price coefficient to be negative for everyone, to avoid counterintuitive implications that also do not fit the data well.

Corresponding to the less negative price coefficient, the constant term is more negative here than in the base model. The mean constant term here flips signs from positive to negative, although it loses its statistical significance. Together

with the small estimated standard deviation on the random coefficient for the constant term, this implies that most people would obtain a higher flow utility from the outside good than from the base DVD player with few DVD titles, though the mean utility from a base DVD player with the mean number of DVD titles is somewhat positive, similarly to the base specification. The majority of the other features (8 out of 12) enter utility positively, although 4 out of the 12 features, including Dolby Digital audio again, enter utility negatively.


In contrast to the two dynamic models, the static model in column 3 appears to give results that are not very realistic. In particular, price enters utility positively for the mean consumer, although the random coefficient indicates that there is a substantial set of people for whom price enters negatively. Many of the coefficients on the features for this model also have an unintuitive sign. In particular, 6 of the 12 physical characteristics enter utility negatively for this model, as does DVD titles. The constant term is far larger than any of these coefficients. Both the mean and the standard deviation coefficients on the constant term are very large relative to any other estimates from this specification, although neither is statistically significant. Thus, the data are indicating that little except for a variable constant term is very useful at explaining the results, given static preferences.

Our interpretation of the differences in these coefficients across the dynamic and static specifications is that the traditional static estimation has a hard time explaining why so few consumers purchase DVD players early on in the product life cycle. The best way to explain the lack of purchase within the static model is to assume that there is a large variance in the “taste” for DVD players. While this is not a perfect explanation for this phenomenon, it does imply that data with very different purchase probabilities would be plausible from the estimated model. Hence, data where the probability of purchase is very different early on from later would also not be completely implausible from the estimated model. In contrast, our dynamic model predicts that people did not purchase early on because they perceived that quality would rise and price would drop. We believe that this is a more appealing explanation of the data.

4.2 Fit of the model

Our estimates of the dynamic models of consumer preferences rely on the simplifying 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. In order to better understand the extent to which this assumption is valid, we examine the evolution of δ_{it} at the estimated parameter values. Figure 4 provides the mean value of δ_{it} across random coefficients as well as the mean value of δ_{it}^n . For this figure and all that follow, we use the estimated parameters reported in the first column of Table 2 and the vector of δ_{jt}^x that are consistent with these parameters and with observed shares (i.e., that solve the fixed point (16)). One can see that there is a general upward trend in both values throughout our sample period. Moreover, the trend looks roughly linear.

Figure 4: Evolution of mean δ_{it}



The mean δ_{it}^n always lies above the mean δ_{it} , implying that the value of being in the market is lower once one has already purchased a product. This is consistent with the finding that consumers value additional DVD titles positively.

We believe that Figure 4 suggests that our simple linear model of evolution for δ_{it} is reasonable. Of course, the results of this figure are based on values evaluated at the structural parameters, and so we cannot definitively ascertain whether a different industry evolution assumption would have resulted in different structural parameters that then would have generated that type of evolution assumption.

Another way of evaluating the industry evolution assumption is to examine the prediction error from the consumer decision problem. In Figure 5 we evaluate the mean value across random coefficients of the prediction error, which is $\delta_{i,t+1} - (\gamma_{1i} + \gamma_{2i}\hat{\delta}_{it} + \hat{u}_{it})$, where $\hat{\delta}_{it}$ and \hat{u}_{it} are the estimated parameters from the regression specified in (10). 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 5 does appear to show that the variance of the prediction errors increases somewhat over time, although our model imposes that the variance of the residual is constant over time. We did not display the mean prediction errors based on δ_{it}^n and either $\delta_{i,t+1}$ or $\delta_{i,t+1}^n$ because these are numerically very close to the prediction errors

Figure 6: Evolution of quantiles of δ_{it} from median period





Figure 7: Evolution of DVD player sales under different assumptions




The other two lines graph the time path of sales given alternate assumptions about consumer and firm behavior. The dashed line graphs the time path of sales that would occur if consumers assumed that their logit inclusive values for DVD players remained equal to its present value in all future periods. The dotted line plots the time path of sales that would occur if firms were faced with all consumers having no DVD player in each period, instead of high valuation consumers having purchased the product and hence generally having a higher reservation utility for buying, as occurs in our model. For both of these estimates, we use the estimated parameter vector from the base model, and the δ_{jt} generated by this estimated parameter vector. All of these lines are generated using the estimated parameter values and corresponding values of δ_{it}^x .

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 by several times 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 DVD players 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 have only the outside good in every period, then the sales path would be similar until about halfway in our sample. At this point, many of the high valuation consumers have started to purchase. By the end of our sample period, we predict that sales would be about 2 to 3 times as high as they were if the high valuation consumers were still in the market owning only the outside good. Note that this increase in sales is due to high valuation consumers not owning any DVD player, not simply a larger market, as only 27.1% of households purchased DVD players by the end of our sample, implying that 72.9% of the potential market was still there at that point.

Further information on the implied magnitudes of our parameter estimates can be provided by examining the extent to which we observed repeat purchase behavior in our sample. Figure 8 plots out the total shares as well as the shares due to repeat purchases. The figure shows that repeat purchases account for a very small fraction of total sales. Moreover, most of the repeat purchasing is happening in the last year of our data, 2003. The underlying reason why there are not more repeat purchases is that the coefficients on the physical characteristics of DVD players other than the constant term are small relative to the utility contribution from the price, the constant term and the term on DVD titles. Because these coefficients are small, there is relatively little benefit to buying a new DVD player. We believe that this result may be due in part to the fact that we do not yet allow random coefficients on any of these characteristics. Without random coefficients, our model would have to predict that *everyone* values features that are introduced later such as progressive scan, which would then imply a much higher ratio of DVD purchases to consumers than the 27.1% that we actually observe.

Figure 8: Evolution of new total and upgrade DVD player sales



Note that even though relatively few people purchase multiple DVD players during our sample period, the ability to purchase multiple DVD players results in substantially different estimates from the base model, as can be seen by comparing the results across the two specifications in Table 2. The reason for this is that the estimates from the model which does not allow for repurchasing would result in very different predictions if multiple purchases were allowed. Thus, allowing for multiple purchases serves as a substantive, and hopefully realistic, constraint on the parameter estimates.

A final implication of the results is to understand the extent to which DVD players have created consumer surplus in the economy. We evaluate the expected discounted consumer surplus by examining the first period, and integrating the value function across consumer random coefficients, evaluated at δ_{i0} . For each consumer random coefficient, we then divide by the marginal utility of a dollar, which again varies across consumers and prices. For this calculation, we use a price of \$180, which is the sales-weighted mean price of a DVD player in our sample.

Our results reveal that the DVD player market has contributed an average of \$300 in consumer surplus per U.S. household at the start of 1997. There is substantial variation in this number across consumers. Given our discount rate of $\beta = .99$, this suggests that the new flow utility from the DVD player industry averages \$3 per month. It would be of use to compare this number to the comparable figure from the static estimation of the DVD 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. More generally, this computation also shows how one can use these type of methods to provide more detail about the valuation of new goods.

5 Conclusion

This paper develops a method 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 understanding firm strategies regarding dynamic price discrimination. We estimate our model using a panel data set of prices, quantities and characteristics for the DVD player 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. Our results show that much of the reason why the initial market share for DVD players was not higher was because consumers were rationally expecting that the market would later yield cheaper and better players. We also find substantial heterogeneity in the overall utility from DVD players and also in the marginal disutility from price. We also find that the DVD player industry is worth an average of \$300 in expected value at the start of the industry.

We believe that much further work on this topic is necessary. This work will involve a variety of robustness checks of the basic model, examining the implication of firm side behavior as well as potentially examining the usefulness of adding micro-level data to the estimation procedure.

References

- Akerberg, D. A. (2003). Advertising, learning, and consumer choice in experience good markets: A structural empirical examination. *International Economic Review*, 44, 1007–1040.
- Akerberg, D. A. & Rysman, M. (in press). Unobservable product differentiation in discrete choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics*.
- Aizcorbe, A. & Kortum, S. (2004). Moore’s law and the semiconductor industry: A vintage model. Unpublished manuscript, University of Minnesota.
- 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.
- Benkard, L. (2004). A dynamic analysis of the market for wide-bodied aircraft. *Review of Economics Studies*, 71(3), 581–611.
- 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 demand model. Unpublished Manuscript, Harvard University.
- 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. (2004). Product innovation in durable goods markets. Unpublished manuscript, University of Wisconsin.
- Carranza, J. (2005). 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. (in press). Indirect network effects and the product cycle: Video games in the U.S., 1994-2002. *Journal of Industrial Economics*.
- Copeland, A., Dunn, W., & Hall, G. (2005). Prices, production, and inventories over the automotive model year. Unpublished manuscript, University of Minnesota Federal Reserve Bank.

- de Figueiredo, J. & Kyle, M. (2004). Product launch decisions by dominant and fringe firms. Unpublished manuscript, Duke University.
- Einav, L. (2004). Gross seasonality and underlying seasonality: Evidence from the U.S. motion picture industry. SIEPR Discussion Paper 02-36, Stanford University.
- Esteban, S. & Shum, M. (2005). Durable-goods oligopoly with a secondary markets. Unpublished manuscript, Johns Hopkins University.
- 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. (2003). Measuring the implications of sales and consumer stockpiling. Unpublished manuscript, Northwestern University.
- Karaca-Mandic, P. (2004). Network effects in technology adoption: The case of DVD players. Unpublished manuscript, RAND Institute.
- 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. (2005). Dynamics of pricing in durable goods markets: Applications to 32-bit console video games. Unpublished dissertation, University of Chicago Graduate School of Business.
- 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. (in press). The role of network effects in the U.S. VCR market, 1978-86. *Journal of Economics and Management Strategy*.
- 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. (2005). Measuring the digital divide: Structural estimation of the 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.
- 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.
- Spence, A. (1981). The learning curve and competition. *The Bell Journal of Economics*, 12, 49–70.
- Stokey, N. (1979). Intertemporal price discrimination. *Quarterly Journal of Economics*, 93, 355–371.