Determinants of Price Declines for New Durable Consumer Goods*

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

Marc Rysman
Department of Economics
Boston University

April 27, 2005

Abstract

Many durable consumer goods are characterized by relatively high initial prices followed by rapid declines in prices during the first few years. One possible explanation for this phenomenon is dynamic price discrimination, wherein firms price high early in order to extract surplus from high value consumers, and then gradually lower price. This paper uses data on DVD players and digital cameras to examine the extent to which dynamic price discrimination, as well as declining costs and increased competition, can explain the declining prices of consumer goods. We specify and estimate a simple dynamic model of consumer preferences with heterogeneous consumer tastes. 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. Using panels of prices, qualities and characteristics, we estimate consumer preferences separately for our two goods. We then investigate the importance of dynamic price discrimination given the estimated consumer preferences.

*We thank seminar participants at the 2005 ASSA Winter Meetings and Boston University for helpful comments, Calixte Ahokpossi for research assistance and Firat Inceoglu, Kyryl Lakishyk, Minsoo Park and Amit Pazgal for assistance in providing access to these data. Gowrisankaran acknowledges funding from the National Science Foundation (Grant SES-0318170). All errors are our own.

1 Introduction

All consumers are familiar with the dynamic problems of purchasing new consumer electronics. A purchaser today can be sure that within a short period of time, the same product will be available for less. Rapidly falling prices both within and across products 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, the prices of DVD players dropped from an average of about $700 at their introduction in 1997 to about $250 by 2003.

One possible explanation for the observed decline in prices over time is dynamic price discrimination. Firms may price high early on in order to extract surplus from high value consumers, and then gradually reduce prices towards costs in order to sell to a greater share of the market. This pattern is sometimes called "cream-skimming" and has been widely discussed in the literature. However, dynamic price discrimination is certainly not the only explanation for the observed price pattern. We typically observe large numbers of new entrants as these industries evolve; for instance, the number of brands of DVD players rose from 12 to 32 between 1997 and 2003. Entry of new firms can explain price decreases, even in a simple static oligopoly setting. Finally, prices may be falling because of decreases in marginal costs over time. Marginal costs may be decreasing because of drops in input prices (e.g. non-volatile memory for digital cameras) or because of a growing firm taking advantage of economies of scale.[1]

This paper aims to measure the determinants of rapidly falling prices for new consumer durable goods using data on the markets for DVD players and digital cameras. We develop methods that can be used to estimate dynamic demand systems for other durable goods markets. Understanding the importance of these different explanations is necessary for evaluating the impact of a variety of policies such as competition policies and subsidies for new technologies.

In order to evaluate the extent to which dynamic price discrimination and other factors explain the observed decline in prices for new consumer goods, we develop and estimate models of consumer and firm behaviors that allow for factors such as dynamic price discrimination, and use the models to understand the importance of these factors in explaining the rapid decline in prices.

To be of use in evaluating the importance of dynamic price discrimination, a model must inherently be dynamic. Myopic firms would clearly not choose to intertemporally price discriminate. Moreover, modeling dynamic *consumer* behavior is also likely important in correctly explaining purchase patterns: rational consumers in 1997 knew that they could wait and likely purchase a DVD player for much less money within a few years, and this future price path should make them less likely to purchase a DVD player initially than if prices were not expected to drop. In addition, in order for a cream-skimming strategy to be profitable, firms must expect to face different distributions of willingness to pay

[1] Another explanation for falling prices is learning-by-doing. However, firms that rationally predict that learning-by-doing will not necessarily lower prices over time. See Spence (1981) for a theoretical analysis and Benkard (2004) for a simulation analysis.

over time, as high willingness-to-pay consumers purchase the durable good early on and drop out of the market. This suggests the necessity of allowing for heterogeneous consumer preferences, with persistence in preferences over time. Last, the fact that new consumer durable goods industries include a variety of products with different features and typically only one product is purchased suggests the appropriateness of a discrete-choice differentiated products specification.

Given these criteria, we specify and estimate a model of consumer preferences with heterogeneous consumers. The model specifies consumers as rational dynamic agents, who must choose between purchasing one of the available durable goods today, or waiting until the future. Consumers are differentiated in their willingness to pay for the durable good, in their relative disutility from price and potentially on other dimensions. Consumers perceive a distribution of price and quality paths over time that allows for the observed decline in prices and improvement in quality. We also specify and calibrate a model of dynamic equilibrium pricing across firms. This model assumes an exogenous, uncertain evolution of industry quality and marginal costs, and allows for the fact that consumers and other firms act as rational dynamic agents.

The study builds on several literatures. First, a theoretical literature has discussed all of the explanations we have outlined here, which is not surprising since rapidly falling prices are one of the most striking features of some classes of new durable goods. The literature dates back at least to the Coase (1972) conjecture, which argues that a durable goods dynamic monopolist will not be able to price above cost because of competition with its future self. Stokey (1979) and Gul, Sonnenschein and Wilson (1986) formally examine the incentives of a durable goods dynamic monopolist, and show the veracity of the Coase conjecture in certain cases. Deneckere and DePalma (1988) extend this model to an oligopoly setting. Our model differs from these earlier models in several important ways, including the fact that firms are not continuously readjusting prices and that there are components to consumer utility that are idiosyncratic to one time period. These differences imply that the Coase conjecture will not hold for our model, and that firms may choose to dynamically price discriminate.

There has also been substantial empirical work on new consumer durable goods. However, most empirical papers that address situations where consumer waiting might be important have focused on other issues and have used static models of consumer demand.[2] 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.[3]

We are aware of two previous approaches to estimating aggregate demand with dynamic consumer choice. The first is Gandal, Kende and Rob (2000) which is designed for homogenous products markets. The second is Melnikov

[2]Examples are Goolsbee and Petrin (2004) for satellite cable, Park (2004) and Ohashi (2003) for VCRs, Ohashi and Clements (2004) for video games and Einav (2004) for movie-going.

[3]See Aizcorbe and Kortum (2005) for CPUs, Copeland, Dunn and Hall (2005) for automobiles and Nair (2005) for video games for examples of papers that model static demand and dynamic firm behavior.

(2001) who analyzes the dynamics of consumer choice for discrete choice differentiated products markets with durable goods.[4] Melnikov specifies a dynamic discrete choice model where the econometric unobservable is a logistic error term. The method uses the fact that logit models satisfy independence of irrelevant alternatives (IIA) to estimate most of the parameters of consumer preferences without solving for or specifying the dynamics. Unfortunately, we cannot apply Melnikov's methodology to understand dynamic price discrimination because his model makes heavy use of the IIA property from the logistic framework. IIA implies that there is no persistent heterogeneity in consumer preferences over time, which makes it unsuitable to examine our question. While we cannot use Melnikov's method of inference, our paper builds on Melnikov in that our model of consumer preferences (though not our inference given this model) is similar, though with the addition of continuous random coefficients to allow for persistent heterogeneity.

Specifically, we develop a random coefficient discrete choice model of consumer preferences with logistic (Type I extreme value) unobservables for each choice and unobserved product characteristics, following in the framework of Berry, Levinsohn and Pakes (BLP) (1995). As in Melnikov, the model is dynamic and goods are durable. Consumers have a discount factor β. All consumers start out as potential customers at the introduction of the product. Once a given consumer purchases a product, she uses that product forever, and is no longer a potential consumer for the product. Consumers obtain the same flow utility from that purchased product every period.[5] At every time period, one possible choice is the option to not purchase any product. In that case, consumers are faced with a set of new products in the subsequent period. Consumers do not know the future set of products, but they do perceive a distribution for these future products. We assume that consumers perceive the future distribution of product characteristics to be a simple polynomial function of a simple statistic (the logit inclusive value) of the current product characteristics.[6] 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.

Two assumptions of our model are relatively stark, the assumption that consumers purchase the good once in their lifetime, and the assumption that consumers' beliefs about future product characteristics depend only on simple statistics of current product characteristics. Both assumptions can be generalized to examine their robustness.

To our knowledge, this paper is the first to estimate a dynamic model of aggregate consumer demand that allows for continuous random coefficients.[7]

[4] Melnikov's methodology has been implemented to study network effects (Chintagunta, Dube and Nair, 2004 and Park, 2004) and to examine the digital camera industry by Carranza (2004).

[5] In this sense, our model is most applicable to new products in their first few years of introduction.

[6] Hendel and Nevo (2003) also make this assumption in the context of a dynamic model of stockpiling, applied to laundry detergent.

[7] A recent preliminary paper by Carranza (2005) suggests an alternative method for esti-

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 unobserved product characteristic for each product by finding the value of the unobserved product characteristic 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 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 of demand.

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 (see Petrin (2002), for a discussion) 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 implicitly creates another source of identification, which is the variation in the set of consumers who can purchase different choice sets. A central limitation of this approach is that it does not allow product characteristics to be endogenous, and hence does not allow discrimination strategies where firms optimally tailor their product mix to the set of remaining consumers.[8]

We use our estimated model of consumer preferences to understand the importance of dynamic price discrimination in explaining the observed declines in prices for DVD players and digital cameras. Specifically, we examine the impact of consumers' expectations about future price declines in explaining the share of consumer purchases. We also examine the differences between our estimates of price elasticities and random coefficients and estimates from a static model of the same industries. Last, we examine how the observed market shares would have differed had the distribution of consumers been the same across periods. By evaluating the determinants of price declines for two separate industries, we add robustness to our results, to the extent that different industries fit our model well on some dimensions but not on others.

mating this type of model.

[8] For instance, de Figueiredo and Kyle (2004) find a broadening in product quality in both directions over time for laser printers.

To directly analyze the extent of optimal dynamic price discrimination given these consumer preferences, we specify a dynamic supply side model of pricing. We build on the techniques of Ericson and Pakes (1995) and Pakes and McGuire (1994). Specifically, we specify the state as the set of consumers who have not yet purchased a product, together with the set of products that are available. We do not directly model the investment decision, because we are not directly interested in this dimension of the problem. We instead assume that the firm perceives that both the quality of its products and the aggregate industry quality evolve according to some exogenous Markov process. (This part is not yet completed. Our current dynamic pricing model restricts attention to a monopoly firm.)

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 models of consumer preferences and firm pricing and explain our method of inference. Our model is the same for DVD players and digital cameras, although we consider them separately.

2.1 Model

Our model starts with the introduction of a new consumer durable good, either DVD players or digital cameras, at time $t = 0$. At $t = 0$, there is a continuum of heterogeneous consumers indexed by i, each of whom are potential customers for a product. 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 and that each consumer purchases at most one product in her lifetime. Thus, once a consumer purchases a good, she makes no further choices.

Consider the set of consumers who have not yet purchased a product at some time t. Each of these consumers chooses one of among J_t products in period t or chooses to purchase no product in the current period in which case she is faced with a similar (though not identical) decision at time $t + 1$.

Product j at time t is characterized by observable characteristics x_{jt}, price p_{jt} and an unobservable (to the econometrician) characteristic $\hat{\xi}_{jt}$. These characteristics are known to consumers and firms.[9] Consumer preferences over x_{jt} and p_{jt} are defined by the consumer characteristics $\hat{\alpha}_i^x$ and α_i^p.

The flow utility to a consumer at time t from purchasing product j at time t is assumed to be:

$$\hat{u}_{ijt,t} = \hat{\alpha}_i^x x_{jt} - \alpha_i^p log\left(p_{jt}\right) + \hat{\xi}_{jt} + \epsilon_{ijt}. \tag{1}$$

[9]Characteristics for these goods will include physical attributes such as memory. They might also include characteristics of the environment, such as the number of available DVD titles, which might be important as network benefits in determining the utility of DVD players.

The last component of (1), ϵ_{ijt}, is an idiosyncratic unobservable that is distributed as Type I extreme value and is independent across 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 consumer tastes. We assume that the flow utility to a consumer at time t from purchasing no product at time t and instead waiting is $\hat{u}_{i0t} = \epsilon_{i0t}$, also assumed to be i.i.d. Type I extreme value.

The characteristics of a product j purchased at time t, x_{jt} and $\hat{\xi}_{jt}$, stay constant over the infinite life of the product. This implies that the flow utility to consumer i in period τ from having purchased product j in period t ($\tau > t$), expressed in period t utils is:

$$\hat{u}_{ij\tau,t} = \beta^{\tau-t} \left(\hat{\alpha}_i^x x_{jt} + \hat{\xi}_{jt} \right). \tag{2}$$

In contrast, we do not model any explicit linkage between products offered for sale at different time periods. Thus, even if two products jt and $j, t+1$ have the same model number and the same observable characteristics x_{jt} and $x_{j,t+1}$, we do not restrict their unobserved characteristics $\hat{\xi}_{jt}$ and $\hat{\xi}_{j,t+1}$ to be the same.

Let $\alpha_i^x = \hat{\alpha}_i^x/(1-\beta)$ and $\xi_j = \hat{\xi}_j/(1-\beta)$. For most of our analysis, it is sufficient to consider these "discounted value" representations of consumer tastes and product characteristics. Combining (1) and (2), the discounted net value of purchase of product j at time t expressed in period t utils becomes:

$$u_{ijt} = \alpha_i^x x_{jt} - \alpha_i^p log\left(p_{jt}\right) + \xi_{jt} + \epsilon_{ijt}. \tag{3}$$

The vector $[\alpha_i^x \quad \alpha_i^p]$ is distributed normally with mean $[\alpha^x \quad \alpha^p]$ and variance matrix Σ and remains constant over time for a given consumer. Note that our per-period model fits within the random coefficients discrete choice framework of BLP, and as such incorporates consumer heterogeneity that is persistent over time.

We now turn to the dynamic part of the consumer decision process. We refer to consumers with a given α_i^p and α_i^x as consumers of type i.[10] It will be useful to define the mean utility to consumer i who purchases product j in period t as:

$$\delta_{ijt} = \alpha_i^x x_{jt} - \alpha_i^p log\left(p_{jt}\right) + \xi_{jt}. \tag{4}$$

and also to define the *logit inclusive value* for consumer i at time t as:

$$\delta_{it} = \log\left(\sum_{j \in J_t} \exp(\delta_{ijt}) \right). \tag{5}$$

Each period t, each consumer of type i who has not yet purchased a product picks $j = 0, ..., J_t$ to maximize her expected discounted value of future expected utility, conditional on her information at time t, where $j = 0$ is taken as the

[10]This is an abuse of terminology in the sense that for consumers of type i, there are a continuum of realizations of $\epsilon_{i.t}$ also indexed by i.

decision to wait. The value of waiting is greater than the flow utility $\hat{u}_{i0t}$ because waiting encapsulates the option to remain in the market in the future.

In order to evaluate consumer i's choice at time t, we need to understand the value of waiting, which implies that we need to formalize consumer i's expectations about the future. Let the matrix Ω_t represent the values of x_j, p_{jt} and ξ_{jt} for all products available in period t. We assume that consumers know all of the currently available products but are uncertain about the product characteristics and prices that will be available to them in the future. The set and characteristics of products vary across time, due to entry and exit of products and changes in the characteristics and prices for existing products. In particular, consumers believe that Ω_t evolves according to some Markov process $P\left(\Omega_{t+1}|\Omega_t\right)$.[11] The Markov process allows for new products to enter and current products to exit. We also assume that consumers have no information about the future values of the idiosyncratic unobservables $\epsilon_{ij,t+1}$ beyond their distribution. Thus, they perceive that $P\left(\epsilon_{i.,t+1}|\epsilon_{i.t},\Omega_t\right) = P\left(\epsilon_{i.,t+1}\right)$.[12]

Given these assumptions, we can write the Bellman equation for a consumer at the beginning of period t as:

$$V_i(\Omega_t, \epsilon_{i.t}) = max\left\{max_{j=1,\ldots,J_t} u_{ijt}, \epsilon_{i0t} + \beta E\left[V_i(\Omega_{t+1}, \epsilon_{i.,t+1}|\Omega_t\right]\right\}. \quad (6)$$

From (6), a consumer that purchases a product j in period t picks j to maximize u_{ijt}. The reason for this is that this consumer will never again make a purchase during her lifetime. In contrast, by picking $j = 0$, the consumer will receive the flow utility ϵ_{i0t} and the value function in period $t+1$. Note that (6) is evaluated after the point at which consumers realize their draws of $\epsilon_{i.t}$. As in Rust (1987), it will be useful to define the expectation of the value function integrated over realizations of ϵ_{ijt},

$$EV_i(\Omega_t) = \int_{\epsilon_{i.t}} V_i(\Omega_t, \epsilon) dP_\epsilon. \quad (7)$$

Because of our assumption that the ϵ_{ijt} are i.i.d. Type I extreme value and not predictive of future states in any way, the purchase probabilities at any time period are a simple function of the logit inclusive values and expected future values (see Rust, 1987 and Melnikov, 2001). Specifically, the probability that a consumer of type i who has not purchased prior to t does so in period t is:

$$P_{it} = \frac{\exp(\delta_{it})}{\exp(\delta_{it}) + exp\left(\beta E\left[EV_i(\Omega_{it+1})|\Omega_{it}\right]\right)}. \quad (8)$$

and the probability that a consumer of type i who purchases in period t purchases product j is:

$$P_{ij|t} = \frac{\exp(\delta_{ijt})}{\exp(\delta_{it})}. \quad (9)$$

[11] Note that we assume that the distribution of Ω_{t+1} does not depend on $\epsilon_{i.t}$, since $\epsilon_{i.t}$ is idiosyncratic to a consumer.

[12] We use the notation $\epsilon_{i.t}$ to refer to $\epsilon_{i0t} \ldots \epsilon_{iJ_t t}$.

Let s^0_{it} be the share of consumers of type i who have not purchased in period t. We assume $s^0_{i0} = 1$. Then, the share of consumers of type i who purchases product j in period t is

$$s_{ijt} = P_{ij|t} P_{it} s^0_{it-1}, \tag{10}$$

and the market share for product j in period t is

$$s_{jt} = \int_{\alpha_i} s_{ijt} f(\alpha_i) d\alpha_i. \tag{11}$$

One significant remaining problem with evaluating the purchase probabilities is that the large dimensionality of Ω_t makes it difficult to solve for the consumer's dynamic optimization process. In the interest of tractability, we make an important simplifying assumptions about the structure of the industry evolution over time. We assume that the change in total product quality follows a one dimensional Markov process in the logit inclusive value. Specifically, we assume that

$$P\left(\Omega_{t+1}|\Omega_t\right) = P_i\left(\Omega_{t+1}|\delta_{it}\right). \tag{12}$$

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 has been employed in the literature.[13]

Note that 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. Even if these scenarios result in the same δ_{it}, they might imply very different expectations of δ_{it+1}.

The benefit of this simplifying assumption is that the state space is reduced to single scalar variable. The reason for this is that, as can be seen in (8), choices in any period depend only on δ_{it}. Thus, what consumers actually care about at time t is not $P\left[\Omega_{t+1}|\delta_{it}\right]$ but rather $P\left[\delta_{i,t+1}|\delta_{it}\right]$. With this simplification, we can rewrite (7), the expectation of the value function for type i at the point before the consumer realizes her $\epsilon_{i.t}$ draws as:

$$EV_i(\delta_{it}) \quad = \quad Emax\left\{\delta_{it}, \beta E\left[EV_i\left(\delta_{i,t+1}\right)|\delta_{it}\right]\right\}. \tag{13}$$

with purchase probabilities for each good as given by (9). We can then easily numerically solve the consumer decision problem by performing backward induction on (13), as there is now only one state variable.

We now mention a couple of points about this simplifying assumption. First, from (12), we allow the conditional density $P_i\left(\delta_{i,t+1}|\delta_{it}\right)$ 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 $P\left(\Omega_{t+1}|\Omega_t\right)$ would not vary across types i. While (12) never imposes this invariance assumption, it is certainly consistent

[13] For instance, Hendel and Nevo (2003) also use this assumption in the context of consumer expectations about future prices of laundry detergents.

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 both depend on δ_{it} and the number of products at time t.

Within the context of this simple functional form, we assume that consumers have rational expectations about the future. Specfically, we assume that consumers perceive the empirical density of $P(\delta_{i,t+1}|\delta_{it})$ fitted to a simple functional form. We currently choose a simple quadratic functional form (linear in some specifications),

$$\delta_{i,t+1} = \gamma_1 + \gamma_2\delta_{it} + \gamma_3\delta_{it}^2 + u_{it}, \tag{14}$$

where u_{it} is assumed to be normally distributed with mean 0. 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.

Another issue about our model is our specification of price. 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. As an alternative to using household income, we could model households as facing a budget constraint. We could posit a "budget for gadgets" b_{it} for each household i and use this as income. One could easily estimate a "budget for gadgets" from the Consumer Expenditure Survey.

We have thus far not discussed the firm side of the model. In our estimation of the demand side. It is not necessary to fully specify the firm side in order to estimate demand; an assumption that product characteristics are exogenous is sufficient to estimate consumer preferences. Nonetheless, we clearly need to specify the firm side in order to understand the optimal extent of dynamic price discrimination. We envision (but have not yet specified) a dynamic oligopoly model. In this model, firms would be faced with exogenous arrival of new products and improvements to existing products from their research labs. Entry and exit decisions would be exogenous. Firms would then simultaneously decide on pricing decisions, taking into account the expected future evolution of their products and those of competitors, and the set of people who have already purchased and who are in the market.

Currently, we are modeling a dynamic monopolist who is faced with the set of consumer preference specified above. The state variable for the monopolist is the measure of consumers of each type i who have already purchased a product and hence are out of the market. Currently, we are assuming that quality and marginal costs do not vary from period to period. The state space could easily be expanded to allow for variation in quality and marginal costs of production. Importantly, the dynamic monopolist 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.

2.2 Inference

In this subsection we discuss the estimation of the parameters of our model of consumer preferences. The parameters are (α, Σ, β), which respectively are 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 and 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.

We develop a method for estimating the remaining parameters that is based on Berry (1994), Rust (1987) and Melnikov (2001) and the literatures that follow. As in Berry, we develop a GMM estimator based on the interaction of the unobserved product characteristics ξ_{jt} with functions of the observable product characteristics x_{jt}. Also from Berry, for a given parameter vector, we recover ξ_{jt} by inverting the predicted share function to find the vector of ξ_{jt} that makes the predicted shares equal the observed shares. We use a method of repeated iteration to perform the inversion, shown by Berry to be a contraction mapping under some conditions. However, unlike in the Berry case, solving for shares in our model requires solving the dynamics of the consumer decision problem. Thus, we embed the consumer dynamic programming problem into our repeated iteration process to solve for shares. This consumer dynamic programming problem, evaluated at each iteration of the share iteration process, is similar to that of Rust. Similar to Melnikov, the structural parameters that we estimate are all essentially static preference parameters.[14]

Note that our estimation algorithm involves three levels of non-linear optimizations: on the outside is a non-linear search over the parameters; inside that is the fixed point calculation of ξ_{jt}, and inside that is the consumer dynamic optimization problem. While both the dynamic programming estimation and the inversion are well-known, our innovation is in combining them to develop a feasible estimator of dynamic consumer preferences.

While our method may seem complicated, it is actually a much easier way of

[14] In spite of this similarity and the fact that our model is very similar to Melnikov *conditional on a given consumer type*, our estimation process is quite different from Melnikov. The reason for this is that our modeling of multiple heterogeneous consumer types implies very different aggregate predictions. Thus, unlike Melnikov, sales data in any period are not sufficient to understand consumer preferences; we instead need to solve for a consumer's dynamic programming problem.

estimating a dynamic model of consumer preferences than alternative methods such as maximum likelihood. In particular, a maximum likelihood estimation that accounted for the endogeneity of price would require that we fully solve the dynamic firm and consumer problems. In contrast, we can estimate consumer preference parameters by just solving for the consumer problem. This is much less computationally and conceptually challenging than solving the whole system. The reason that we do not have to solve for the firm side is that we are using an "inversion" method where we are solving for the value of the econometric unobservables that are consistent with the underlying data, rather than actually solving for the equilibrium of the model conditional on unobservables, as would typically be done in the context of a structural maximum likelihood estimator. In addition to the computational advantage, there is another advantage of our estimation process. We can estimate the consumer preference parameters without making exact assumptions about the nature of product introductions and changes in characteristics. In contrast, a maximum likelihood method that allowed for price to be endogenous would typically need to specify the firm model.

We now explain the specifics of our method. Conditional on the solution of the vector of ξ_{jt} (which involves the dynamic consumer problem), our method of inference follows standard GMM algorithms for discrete choice models (see Nevo (2000) for instance). Specifically, we specify a GMM criterion function

$$G(\alpha, \Sigma) = z'\xi(\alpha, \Sigma), \tag{15}$$

where $\xi(\alpha, \Sigma)$ is the vector of unobserved product characteristics that sets the predicted product shares from the model equal to the observed product shares conditional on parameters, and z is a matrix of exogenous variables, described in detail below. We then choose

$$\left(\hat{\alpha}, \hat{\Sigma}\right) = argmin_{\alpha,\Sigma} \left\{ G(\alpha, \Sigma)' W G(\alpha, \Sigma) \right\}, \tag{16}$$

where W is a weighting matrix and $G(\alpha, \Sigma)$ is as defined in (15).

We employ a two-stage process 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. We first estimate our parameters using (16) with this W to obtain consistent parameter estimates $\left(\tilde{\alpha}, \hat{\Sigma}\right)$. Using these first-stage parameter estimates, we then approximate the optimal weighting matrix as $W = Var\left[G\left(\tilde{\alpha}, \tilde{\Sigma}\right)\right]^{-1}$. We then reestimate the parameters with this new weighting matrix to obtain asymptotically efficient estimates $\left(\hat{\alpha}, \hat{\Sigma}\right)$. These are the parameters that we report.

We perform our estimator with a non-linear search in Σ. However, we adopt another technique from the literature, which is to use a closed-form solution to solve for the optimal mean tastes for characteristics α for any given vector of Σ. To illustrate, we first define the mean utility for product j at time t as

$$\delta_{jt} = \alpha^x x_{jt} - \alpha^p log(p_{jt}) + \xi_{jt}. \tag{17}$$

This vector of mean utilities is a function of the two α parameters and the unobserved product characteristics. As in BLP, conditional on Σ, we recover the vector of δ_{jt} that makes predicted share match observed shares, and then solve for the moment conditions based on this. Specifically, we define $\hat{s}_{jt}(\delta, \Sigma)$ to be the share of product j at time t that would result if the vector of mean utilities was δ. We then start with an arbitrary vector δ and repeatedly evaluate

$$\delta'_{jt} = \delta_{jt} + \ln(s_{jt}) - \ln(\hat{s}_{jt}(\delta, \Sigma)) \tag{18}$$

until convergence.[15] At the fixed point of (18), $s_{jt} = \hat{s}_{jt}(\delta, \Sigma), \forall j, t$. Using the fixed point values of δ, we then solve for the α that minimizes the criterion function in (16) in closed form as

$$\hat{\alpha} = (x'ZWZ'x)^{-1}(x'ZWZ'\delta). \tag{19}$$

We then recover the vector of ξ_{jt} by applying (17) using $\hat{\alpha}$, and apply the criterion function to this vector. Thus, our non-linear search is only over Σ, which significantly reduces computational time.

As noted above, computing $\hat{s}_{jt}(\delta_{\cdot\cdot}, \Sigma)$ requires solving for the full consumer dynamic decision process. This is where we embed the techniques of Rust into the inversion routine. To calculate $\hat{s}_{jt}(\delta_{\cdot\cdot}, \Sigma)$, we solve for the dynamic programming problem given by (13) for any vector of δ_{jt} and then calculate shares by integrating across consumer types using (11).

One important detail of our estimation algorithm is that there is no closed form to (11) so we instead integrate using Monte Carlo techniques. Specifically, our estimation algorithm begins by drawing $S \cdot l$ values from a standard normal distribution, where l is the length of the vector $[\alpha_i^x \quad \alpha_i^p]$ and S is the number of simulation draws. We use these S random draws from the consumer population as our consumer types. To obtain draws with a given variance Σ, we simply transform these initial draws. Let one transformed draw of length l be denoted $d_i(\Sigma)$. Note also that the reason that we can solve for δ_{jt} instead of ξ_{jt} in (18) is because, as can be seen from (17), δ_{ijt}, the mean utility for a particular product for a consumer of type i, can be rewritten as:

$$\delta_{ijt} = \delta_{jt} + d_i(\Sigma)\left[x_{jt} \quad -log(p_{jt})\right], \tag{20}$$

an expression that does not depend on α except through δ_{jt}.

We now go through details of the estimation algorithm, starting with the innermost loop, the consumer dynamic programming problem. This is performed separately for each of the S consumer types. Solving this problem for any type i requires specifying consumer perceptions about the evolution of δ_{it}. As can be seen from its definition in (5), δ_{it} depends on δ_{jt} for all products j at time t. Hence, for each new vector of δ_{jt} from (18) we need to model a new set of consumer perceptions, which we do by performing the simple linear regression specified in (14). Using these perceptions, we can solve the Bellman equation

[15] Note that this fixed point computation is over δ_{jt} and not ξ_{jt} as we suggested earlier.

(13). Making use of our Type I extreme value variables, we follow Rust (1987) and further simplify the Bellman equation to:

$$EV_i(\delta_{it}) = log\left(exp\left(\delta_{it}\right) + exp\left(\beta E\left[EV_i\left(\delta_{i,t+1}\right)|\delta_{it}\right]\right)\right). \tag{21}$$

This is the expression that we evaluate, but in practice we discretize the vector δ_{it} and solve $EV_i(\delta_{it})$ as a step function for each "bin." For each time period t, we then plug in the aggregate purchase probability from (8), evaluating the value function at the appropriate δ_{it} as determined by the vector of δ_{ijt} for that time period. One significant issue is that our consumers must forecast distributions of quality for time periods after the end of our model, and there is no guidance from the data as to what consumers expect. This clearly suggests that we should experiment with results using different assumptions about what happens after the end of our sample. Currently, we assume that δ_{it} continues to rise at the same rate as during our sample.

Let us now turn to the middle loop, the computation of the Berry fixed point from (18). The inner loop solves the optimal decision separately for each type of consumer. This computation requires integrating across consumers to find the aggregate market share. We perform the integration by simply summing over simulation draws,

$$\hat{s}_{jt} = \sum_{i=1}^{S} s_{ijt}/S.$$

One other issue relates to the properties of (18). Berry provides conditions under which this function is a contraction mapping, guaranteeing that $\hat{s}_{jt}\left(\delta, \Sigma\right)$ is invertible in the vector of δ_{jt}. 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.

The outermost loop is the search over values of Σ. We perform the search using a simplex method. The dimensionality of Σ is currently small, which makes this process reasonably straightforward.

We now discuss our exogenous variables z and how they serve to identify the parameters. We include in z all of the product characteristics in x, which are assumed to be exogenous. We also include the mean product characteristics for a given firm at the same time period, and the mean product characteristics for all firms at the time period, for all characteristics except for the constant term. Last, we include the count of products offered by the firm and by all firms, or equivalently, the sum of the constant term across these two measures. Note that even though our model is dynamic, all the structural consumer parameters that we estimate are simple static preference parameters. The industry evolution is identified as a reduced-form by the change in the empirical distribution of the products over time.

Because the parameters that we see to estimate are the same, our identification arguments for the structural parameters are similar to those from the discrete choice literature (see Petrin, 2002, for instance). In particular, the conceptually difficult argument is in understanding how the coefficient on price and the random coefficients are identified. In our model, prices are determined by a dynamically optimizing firm, but we do not need to fully specify the firm decisions to estimate consumer preferences, and the same general properties from static models hold.In particular, consumers will be faced with different choice sets over time, and the variation in the choice sets will affect firms' pricing decisions. For instance, if two different companies offer two very similar products, then this should lower prices relative to costs, even for dynamically optimizing firms. This is the reason why we include functions of all the available products in z.

The different choice sets should also allow us to identify the random coefficients: the variation in choice sets forces substitutions across products, and different values of random coefficients will result in different substitution patterns. In addition, our model will endogenously have different distributions of consumer tastes for different time periods. For instance, consumers with high valuations for the product will likely buy early on, leaving lower valuation consumers later on. Substitution based on this aggregate variation in consumer tastes across time will further identify these parameters.

3 Data

We estimate our model using aggregate sales, characteristics and price data for digital cameras and DVD players. For both goods, we observe monthly quantity sold, average price and characteristics for a number of models. For digital cameras, our data cover 378 models and 48 brands, with observations ranging from September 1998 to September 2001. For DVD players, we observe 522 models and 47 brands, with observations from March 1997 to August 2003. For both markets, we observe sales from very early in their product life cycles and observe the vast majority of models. Both data sets were originally collected by NPD Techworld which surveys major electronics retailers and covers 80% of the market.

For DVD players, our characteristics include indicators for whether the model can play other formats besides standard DVDs (VHS tapes, MP3 files, DVDR disk), indicators for features that improve the quality of audio and video (Dolby/Digital, component video, coaxial audio cables, etc.) and an indicator for multidisk 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. We also were faced with some missing data for some of the characteristics. For about 200 model-months, we did not observe if the player could play CD-R/RW disks or MP3 files. For less than 100 model-months, our data was missing one of coaxial audio, Dolby/Digital, DTS noise reduction 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

Table 1: Characteristics of DVD Players

Characteristic	Mean value
Price (Units: 100s of 1983 $)	1.89 (2.25)
Component video	.804 (.397)
Optical audio	.841 (.366)
Coaxial audio	.837 (.369)
Dolby / Digital	.308 (.462)
DTS	.126 (.332)
Plays CDR/RW disks	.463 (.499)
Plays MP3 files	.282 (.450)
Plays DVDR disks	.140 (.346)
Progressive scan	.241 (.428)
Includes VHS	.010 (.099)
Includes recorder	.019 (.418)
Multidisk player	.804 (.397)

Number of observations: 9827
Unit of observation: Model - Month
Standard deviations of variables in parentheses

summarizes these price and characteristics data for DVD players at the level of the model-month.

For digital cameras, we observe several characteristics including picture resolution, camera weight, dimensions, storage memory, type of storage memory and lens type. We collected these data ourselves by examining old magazines and web sites. This data collection process is still in process.

Our feeling is that most of the characteristics for DVD players are relatively unimportant. In contrast, characteristics of digital cameras, such as resolution and storage memory are very important in affecting utility and their value may differ drastically across people. For this reason, we include random coefficients for DVD players only for price and the constant term (the decision to buy a DVD player versus the outside option), while for digital cameras we include more random coefficients.

Both products exhibit striking price declines both within and across models. Figure 1 shows aggregate sales and average prices for DVD players on the left and digital cameras on the right.[16] For DVD players, average nominal prices are above $700 a unit for the first few months of the data and steadily decline to less than $250 within 6 years. Sales increase in each year, with the effects

[16]For this picture, we drop digital cameras with prices over $2,500. There are several of these (some cost over $10,000) and they change the shape of this picture substantially. They likely appeal solely to a professional market.

of Christmas shopping clearly observable in the data. For digital cameras, the decline is from prices well over \$400 to prices around \$325 in 3 years. Again, sales increase with time but the Christmas season causes large spikes around this trend.

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 dynamidc. Given that prices do not change over Christmas, we were not sure what this would add to the model. Instead, we decided to seasonally adjust our data. We did this by multiplying sales by a separate constant for each month of the year. The constants were chosen so that the total sales were the same for every month of the year and so that the total sales were the same as in the raw data.

In addition to the evidence (from Figure 1) that average prices drop over time, we also find that they drop within models. To examine this feature of data, we run the following regression:

$$\ln(price_{jt}) = \alpha_j + \beta_t + \gamma(AGE_{jt}) + \epsilon_{jt}$$

In this regression, α_j and β_t represent product and month dummy variables, with the month dummy variables meant to absorb Christmas effects in particular. The function $\gamma(AGE_{jt})$ is a set of age dummy variables, where age is defined to be the number of months in t since model j first appeared the data. We run this regression with the log of sales on the left-hand side as well.

Figure 3 graphs the age dummies. The graphs for DVD players on the left and digital cameras on the right are remarkably similar. We can see that price declines for each model at a steady pace over its lifetime. Interestingly, quantity increases for a short period and then falls off. One likely reason for this is that, when initially introduced, the products are towards the high end of the price distribution and appeal to only a limited segment of the population. Another possible reason is consumer learning. Our model can easily capture the first effect, but the second one is beyond its scope.

As we suggest above, dynamic price discrimination may be an explanation for falling prices for a given model. Also, entry may be an explanation, particularly for falling prices across models. Figure 3 documents the number of models and the number of brands for DVD players and digital cameras appearing in our data by month. The increase is rapid throughout the time period. Figure 4 graphs the Herfindahl-Hirschman index of sales across brands shows that concentration has a downward trend in these industries, although the trend is very noisy.

Last, note that quality is increasing for these products. Figure 5 graphs the percentage of available models of DVD players that have various technological components. We graph three features: the presence of progressive scan (that generates a better picture), coaxial audio and optical audio (that each improve sound quality). Each of the lines steadily increase with time. Inceoglu and Park (2004) report that the number of DVD titles, which crucially affects the value of a DVD player, is also increasing rapidly with time.

Figure 1: Prices and Sales for DVD Players and Digital Cameras

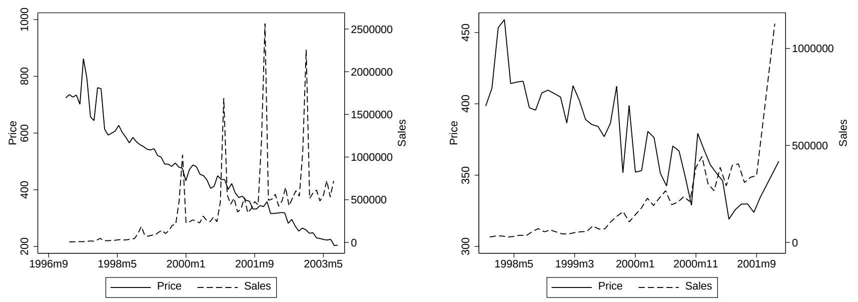

Figure 2: Coefficients on Age Dummies for DVD Players and Digital Cameras

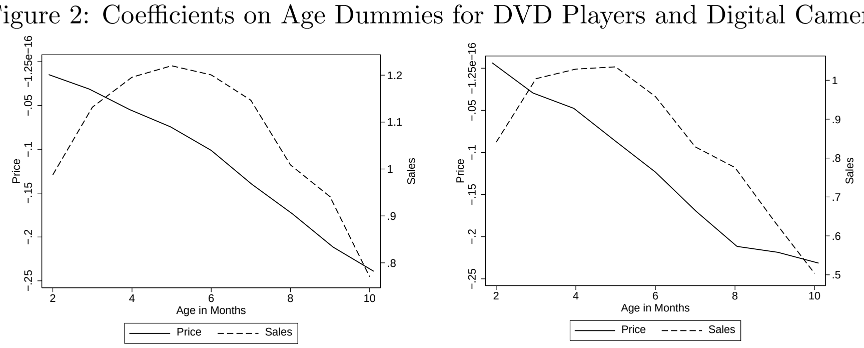

Figure 3: Counts of Models and Brands for DVD Players and Digital Cameras

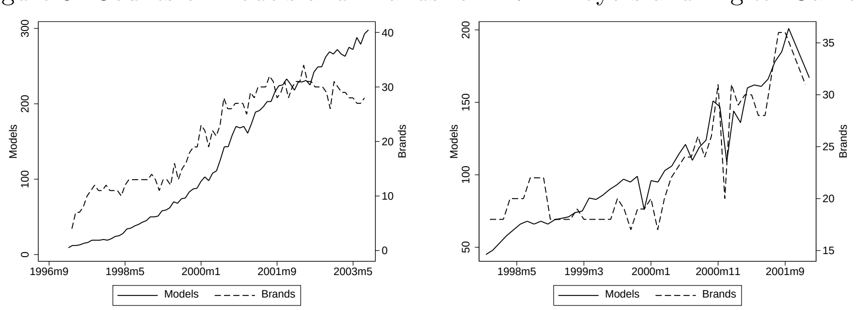

Figure 4: HHI for DVD Players and Digital Cameras

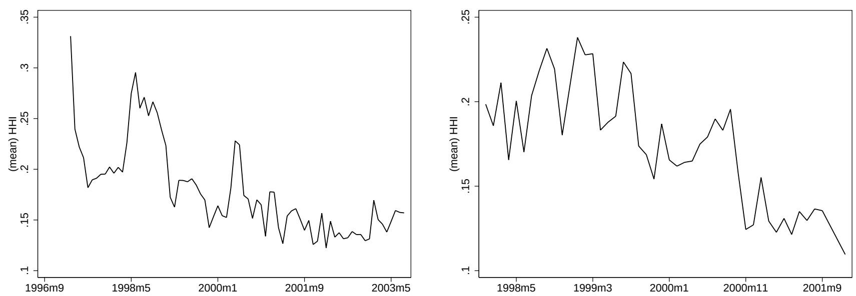

Figure 5: Percentage of Products with Features for DVD Players

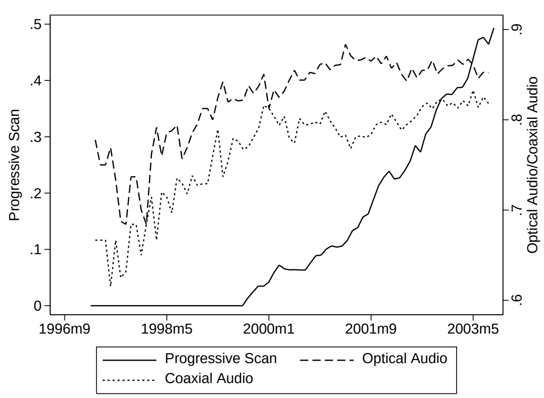

4 Results

We currently present preliminary results, and are uncertain about their robustness. As there are numerous alternate specifications we plan to try in the near future, the results that we present here only to be suggestive of what we might find, not firm conclusions about the parameters characterizing these industries. With that caveat in mind, Table 2 presents results from estimation of the DVD player industry. We specify two random coefficients, one on the constant term and the other on the price variable. Consumers are assumed to be homogenous in their preferences over the other variables, which are all indicator variables for various technologies that improve quality. We have not included the number of DVD titles in the estimation as of yet. Inceoglu and Park (2004) study this issue and the associated endogeneity problems.[17]

Table 2 contains two 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 of results provide the parameter estimates from a traditional static random coefficients discrete choice specification performed on the same data. For these results, we use the same specification for the static part of utility as in the base model.

Several of the parameter estimates are very different across the two specifications. In particular, in the base specification, the mean coefficient on $\log(p_{jt})$ is negative and statistically significant, while in the static model, this coefficient is positive and insignificant. Together, the two coefficients on price from the dynamic model imply that the vast majority of consumers will have a disutility from paying more for a DVD player, a result that makes sense. In contrast, the static model suggests that most people will have a positive utility from paying more for a DVD player, which seems unrealistic.

Another big difference is in the coefficients on the constant term, which will capture consumer preferences for owning a DVD player as opposed to using the outside option, which might involve using a VCR, going to the movies or watching TV. The static specification has a mean on the constant term that is estimated to be negative and very large in magnitude (-51.3), and a standard deviation on the random coefficient on the constant term that is correspondingly large in magnitude (-19.6). The dynamic specification has a more moderate mean coefficient on the constant term (-8.30) and on the standard deviation on the constant term (9.24). Noth of these terms are still reasonably large in magnitude and highly statistically significant, implying a market with a large heterogeneity across consumers in the willingness to pay for DVD players.

Our interpretation of the differences in these coefficients across the two specifications is that the traditional static estimation has a hard time explaining why so few consumers purchase DVD players early on. 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

[17] Some technical details are as follows. We discretize δ_{jt} into 200 bins for the purposes of the dynamics. We perform our Monte Carlo integration with 10 types of consumers.

Table 2: Results for DVD Players

Parameter	Estimate from base specification	Estimate using static model
Mean characteristics (α)		
Constant	–8.30 (.514) *	–51.3 (86.9)
Log Price (100s of 1983 $)	–1.72 (.212)	.867 (9.053)
Component video	1.44 (.101) *	.397 (.110) *
Optical audio	–.219 (.097) *	–.822 (.084) *
Coaxial audio	–.149 (.079)	–.493 (.080) *
Dolby / Digital	–.599 (.078) *	–.297 (.068) *
DTS	.413 (.128) *	–1.075 (.157) *
Plays CDR/RW disks	1.52 (.076) *	.853 (.084) *
Plays MP3 files	1.01 (.102)	1.56 (.172) *
Plays DVDR disks	1.04 (.086)	.785 (.138) *
Progressive scan	1.25 (.123)	–.319 (1.90)
Includes VHS	2.87 (.328) *	1.84 (.578) *
Includes recorder	2.18 (.277) *	–1.03 (.286) *
Multidisk player	.986 (.093) *	.081 (.085)
Std. dev. of characteristics (Σ)		
Constant	9.24 (.667) *	19.6 (39.2)
Log Price (100s of 1983 $)	.501 (.608)	.055 (7.49)

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

Figure 6: Time Path of DVD Sales Under Alternate Assumptions

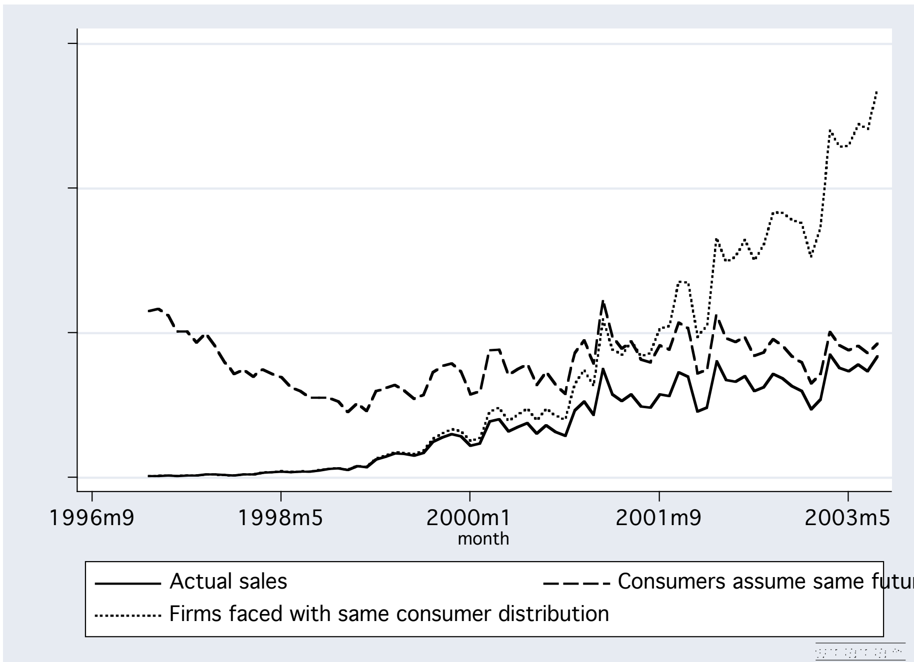

this phenomenon, it does imply that data with very different purchase probabilities would be plausible from the estimated model. Hence, data where the probability of purchase is very different early on from later would also not be completely implausible from the estimated model. In contrast, our dynamic model predicts that people did not purchase early on because they perceived that quality would rise and price would drop. We feel that this is a more appealing explanation of the data.

In contrast to the coefficients on price and the constant term, the coefficients on the other characteristics are generally similar across the two specifications. There are a couple of exceptions, for instance, the coefficients on "includes recorder" and "multidisk player" are both positive and significant in our base specification, but the first is negative and significant and the second one is insignificant in the static estimation. In the dynamic model, most of the coefficients on product characteristics are positive. Our prior is that all of these features are generally considered to be positive attributes, so this makes sense. However, one unappealing feature of our dynamic results is that three of the coefficients that measure audio quality enter negatively in the utility function, and two of these are significant. We are currently investigating the reason for this.

The magnitude of the random coefficient on the constant term from the base dynamic specification suggests substantial consumer heterogeneity and hence that dynamic price discrimination is likely to be important. We investigate this further by examining the time path of DVD player sale under a couple of different assumptions, in Figure 6.

The solid line from Figure 6 plots the actual sales of DVD players over the time period of our sample. Note that these are also the time path of sales generated by the estimated model, which matches the time path of actual sales exactly. The other two lines plot the time path of sales under alternate assumptions about consumer and firm behavior. The dashed line plots the time path of sales that would occur if consumers assumed that quality and price of DVD players (i.e.the logit inclusive value) remained equal to their present values in all future periods. The dotted line plots the time path of sales that would occur if firms were faced with same consumer distribution in the same period, instead of high valuation consumers dropping out of the market, 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.

We find that dynamics explain a huge part of the sales path. In particular, if consumers did not assume that prices and qualities changed, then sales would be roughly constant 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, but by the end of our sample period, sales would be about equal to the current sales. The implication is that consumers did not perceive huge increases in quality by the end of the sample period.

If firms were faced with the same distribution of consumers, 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. Note that this increase in sales is due to high valuation consumers being in the market, 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.

We would ideally like to further understand the importance of dynamic price discrimination by evaluating the extent of optimal price discrimination by a dynamic oligopoly industry. At this point, we perform one simple test of the implication of these results. We consider the pricing policy of a monopolist selling a single product to a population of consumers. The consumer model is as above, where consumers recognize that purchasing today means possibly forgoing lower prices in the future. For the firm side, we assume the firm has constant marginal cost and is forward looking in the sense that the firm recognizes that high prices today means consumers may delay purchase until the future. Importantly, we assume that the monopolist is not able to commit to the future path of prices.

More specifically, the firm's pricing decision is a function of the population of consumers remaining in the market. Therefore, the state space for both firms and consumers is the population of consumers remaining, i.e. the proportion of each type of consumer who has not purchased yet. We discretize the "proportion remaining" variable into equal sized bins of size Δ. Note that it is only possible for the proportion of consumers remaining to decrease, which allows us to solve the problem backwards. Given an assumption about how the firm prices when there are no consumers remaining, it is straightforward to compute how the firm prices when there are Δ consumers left of any one type. Conditional on those results, it is straightforward to compute how the firm prices when there are Δ left of any two types of consumers, or 2Δ left of any one type. We continue this way until we have solved the case when no consumers have purchased yet. In theory, the transition from one state to another conditional on the price should be deterministic but we allow some randomness. Doing so aids in computation but breaks some of the strict rules of rationality governing the process. [18] Based on these transition probabilities, we compute the probability of being in each state in period $t = 1, ..., T$. We compute the average price over time, weighted by the probability of being in each state in that period.

Figure 7 presents the price and quantity over 12 months. Marginal cost is assumed to be \$200. Price declines steadily from about \$1,600 to \$400. Unreported experiments show that as the variance in the random coefficients are increased, the difference from the highest price to the lowest increases. This decline is larger than many of the one-year declines in average price that we see in the DVD data. However, note that the price-cost mark-ups are relatively large. Competition would limit these prices and possibly limit the size of the simulated price declines. While these results are suggestive that intertemporal price discrimination may explain a large portion of the price declines in these data, much more work remains to be done before this paper can offer a firm conclusion.

5 Conclusion

We seek to understand the reasons for the large price declines that are observed for new consumer durable goods. We postulate three reasons for the declines: dynamic price discrimination, decreases in marginal costs over time, and increases in competition over time. We have developed a dynamic model of consumer preferences with heterogeneous consumers faced with an optimal purchase problem for durable goods, with the possibility changing quality and price over time. We have also started to specify a dynamic oligopoly model of firm pricing, when faced with the inability to commit to future prices and these dynamically optimizing consumers.

We estimate our model of consumer preferences for DVD players and digital cameras. Initial results suggest that dynamic price discrimination is an important part of the reason for the decreases in observed prices for DVD players

[18] Consider the case of there being Δ left of one type of consumer. The firm may price so as to sell as much as possible without switching states. Our randomness makes it more likely to switch as the firm sells more.

Figure 7: Optimal Dynamic Pricing Behavior by Monopolist

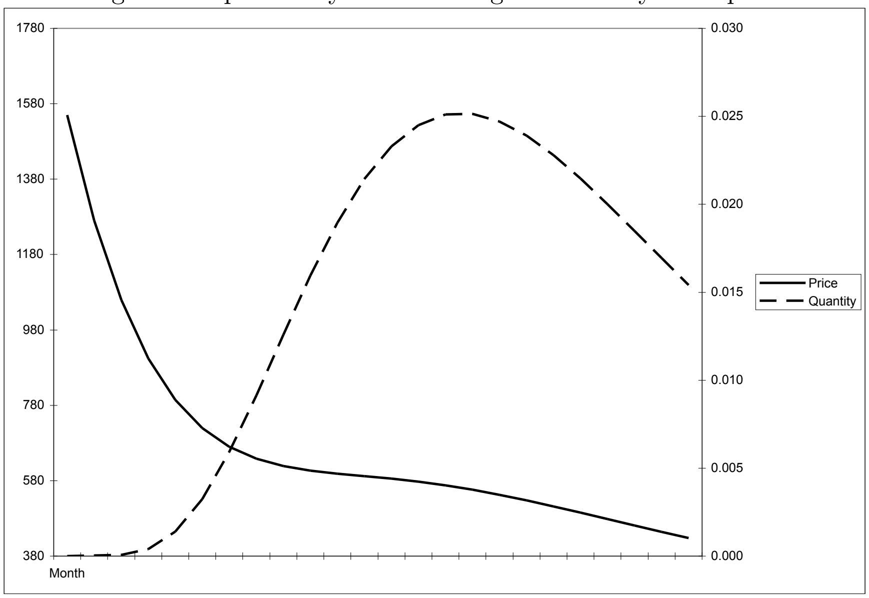

during their first few years. Much more work is needed to examine the robustness of these and other conclusions.

References

[1] Benkard, L. (2004) "A Dynamic Analysis of the Market for Wide-Bodied Aircraft," *Review of Economics Studies.*

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

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

[4] Carranza, J. (2004), "Product Innovation in Durable Goods Markets," mimeo, University of Wisconsin.

[5] Carranza, J. (2005), "Demand for Durable Goods and the Dynamics of Quality," mimeo, University of Wisconsin.

[6] Chintagunta, P., J-P Dubé and H. Nair (2004), "Empirical Analysis of Indirect Network Effects in the Market for Personal Digital Assistants, *Quantitative Marketing and Economics* 2(1), 23-58.

[7] Clements, M. and H. Ohashi (2004) "Indirect Network Effects and the Product Cycle: Video Games in the U.S., 1994-2002," *Journal of Industrial Economics*, forthcoming.

[8] Copeland, A. and W. Dunn and G. Hall (2005) "Prices, Production, and Inventories over the Automotive Model Year," mimeo, Yale University.

[9] de Figueiredo, J. and M. Kyle (2004) "Product Launch Decisions by Dominant and Fringe Firms," mimeo, Duke University.

[10] Deneckere, R. and A. de Palma (1998) "The Diffusion of Consumer Durables in a Vertically Differentiated Oligopoly," *RAND Jouranl of Economics,* 29(4), 750-771.

[11] Einav, L (2004) "Gross Seasonality and Underlying Seasonality: Evidence from the U.S. Motion Picture Industry," SIEPR Discussion Paper 02-36.

[12] Ericson, R. and A. Pakes (1995) "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies* 62(1) 53-82.

[13] Gandal, N., Kende, M. and Rob, R. (2000) "The Dynamics of Technological Adoption in Hardware/Software Systems: The Case of Compact Disk Players", *Rand Journal of Economics*, 31, 43-62.

[14] Goolsbee, A. and A. Petrin (2004), "The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable Television" *Econometrica* 72(2),351-381.

[15] Hendel, I. and A. Nevo (2003), "Measuring the Implications of Sales and Consumer Stockpiling," mimeo, Northwestern University.

[16] Magnac, T. and D. Thesmar (2002) "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70(2), 801-16.

[17] Melnikov, O. (2001) "Demand for Differentiated Products: The Case of the U.S. Computer Printer Market," mimeo, Cornell University.

[18] Nair, H. (2005), "Dynamics of Pricing in Durable Good Markets: Application to 32-bit Console Video Games," mimeo, Chicago GSB.

[19] Nevo, A (2000) "A Practitioner's Guide to Estimation of Random Coefficients Logit Models of Demand," *Journal of Economics & Management Strategy* 9(4), 513-48.

[20] Nevo, A. (2001) "Measuring Market Power in the Ready-to-Eat Breakfast Cereal Industry," *Econometrica*, 69(2), 307-42.

[21] Ohashi, H. (2003) "The Role of Network Effects in the U.S. VCR Market, 1978-86," *Journal of Economics and Management Strategy* forthcoming.

[22] Pakes, A and P. McGuire (1994) "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics,* 25(4), 555-589.

[23] Park, M. (2004), "Dynamic Demand and Network Effects in the market for DVD Players," mimeo, Boston University.

[24] Park, S. (2000), "Quantitative Analysis of Network Externalities in Competing Technologies: The VCR Case", *Review of Economics and Statistics* forthcoming.

[25] Petrin, A. (2002), "Quantifying the Benefits of New Products: The Case of the Minivan", *Journal of Political Economy*, 110, 705-29.

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

[27] Stokey, N (1979) "Intertemporal Price Discrimination," *Quarterly Journal of Economics* 93(3), 355-371.

[28] Spence, A.M. (1981) "The Learning Curve and Competition," *The Bell Journal of Economics*, 12(1), 49-70.