

Dynamics of Consumer Demand for New Durable Goods*

Gautam Gowrisankaran
University of Arizona
HEC Montréal
NBER

Marc Rysman
Department of Economics
Boston University

February 4, 2009

Abstract

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

*We thank Dan Akerberg, Victor Aguirregabiria, Ana Aizcorbe, Rabah Amir, Lanier Benkard, Steve Berry, Sofronis Clerides, Simon Gilchrist, Avi Goldfarb, Igal Hendel, Kei Hirano, Firat Inceoglu, John Krainer, Ariel Pakes, Minsoo Park, Rob Porter, Jeff Prince, Pasquale Schiraldi, Andy Skrzypacz, Mo Xiao and seminar participants at several institutions for helpful comments; Haizhen Lin, Ryan Murphy, David Rapson, Alex Shcherbakov, and Shirley Wong for research assistance; and the NPD Group, ICR-CENTRIS, Ali Hortaçsu and Jeff Prince for providing data. The comments of the editor and four anonymous referees substantially improved the paper. We acknowledge funding from the National Science Foundation. All errors are our own.

1 Introduction

This paper specifies a structural dynamic model of consumer preferences for new durable goods and estimates the model using aggregate data on the digital camcorder industry. A dynamic model is necessary to capture the fact that consumers choose not only what to buy but when to buy. Rapidly falling prices and improving features have been among the most visible phenomena in a large number of new consumer durable goods markets, such as computers, digital camcorders and DVD players. For instance, between 2000 and 2006, average digital camcorder prices dropped from \$930 to \$380 while average pixel counts rose from 580,000 to 1.08 million. The rapidly evolving nature of these industries suggests that modeling dynamics might be empirically very important in estimating consumer preferences.

Our model allows for product differentiation, persistent consumer heterogeneity, endogeneity of prices and endogenous repeat purchases over time. Berry, Levinsohn & Pakes (1995), henceforth BLP, and the literature that follows have shown that incorporating consumer heterogeneity into differentiated product demand systems is important in obtaining realistic predictions. Much of our model is essentially the same as BLP: our model is designed for aggregate data (but can incorporate consumer-level data when available); there is an unobserved product characteristic that affects equilibrium prices; consumers make a discrete choice from a set of products in a multinomial logit framework; and consumers have random coefficients over observable product characteristics. 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. As our model is dynamic, we need to specify consumer perceptions over future states of the world. We focus on a major simplifying assumption: that consumers perceive that the evolution of the value of purchase will follow a simple one-dimensional Markov process. In this sense, consumers use a reduced-form approximation of the supply side evolution to make predictions about the value of future purchases. We also examine a number of alternative specifications for perceptions, including multi-dimensional processes and perfect foresight.

Over the last 15 years, a substantial literature has used static BLP-style models to investigate questions of policy interest. This literature has analyzed questions that include (but are by no means limited to) horizontal merger policy (see Nevo, 2000a), trade policy (see Berry, Levinsohn & Pakes, 1999) and the value of new goods (see Petrin, 2002). Many of these papers investigate industries, such as automobiles, for which goods are durable. Our paper provides a framework to incorporate dynamics in BLP-style models and hence may be useful in deriving better estimates for these and related questions. Indeed, recent work is using and extending our methods to examine scrapping subsidies for automobiles (Schiraldi, 2007), markups for digital cameras (Zhao, 2008), switching costs between cable and satel-

lite television (Shcherbakov, 2008) and switching costs in consumer banking (Ho, 2008), among other research questions.

We use our results to examine the evolution of consumer value from the digital camcorder industry by calculating a cost of living index (COLI) for this sector. COLIs measure compensating variations, the dollar taxes or transfers necessary to hold welfare constant at the base level over time. Government-computed COLIs have important implications for wage growth at many firms, government transfer programs and a variety of other government policies. The BLS is particularly concerned about the development of accurate COLIs for consumer electronics and camcorders in particular (see Shepler, 2001). Systematic entry and exit of camcorders based on their characteristics may create biases, (see Pakes, 2003, and cites therein). Further biases occur when consumers act dynamically. If consumers act as rational dynamic agents and we instead assume myopic behavior, we may overstate the welfare gains later on, by assuming more high-value consumers than actually exist (see Aizcorbe, 2005).

Because we use primarily aggregate data, we develop a relatively parsimonious specification which results in the parameters that we estimate being essentially the same as in static BLP-style models: the mean and variance of consumer preferences for product characteristics. As in these 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. Our dynamic model adds to identification by making use of substitution patterns across time periods as well as within time periods and by capturing the endogenous changes in demand over time as consumer holdings evolve.

To estimate our model, we develop new methods of inference that draw on the techniques of BLP for modeling consumer heterogeneity in a discrete choice model and on Rust (1987) for modeling optimal stopping decisions. As in BLP, we solve for the vector of unobserved product characteristics by using an iterative process to find the value that makes predicted shares equal observed shares. The iterative process requires repeatedly solving for the predicted market shares. In our case, predicted market shares depend on a dynamic optimal-stopping (or purchase) problem, which we also solve with an iterative Bellman equation method. Overall, we compute a Rust-style optimal stopping problem jointly with the BLP unobserved product characteristic solution. Our methodological advance is in developing a feasible specification that allows us to combine these two separate methods.

An important feature of our model is that it is designed to be applied to aggregate data on models and market shares by month (although in some specifications, we supplement these with limited data from a survey) rather than individual household purchase data, which means we must incorporate the iterative process associated with BLP. Perhaps not surprisingly, models for household level data are substantially

more sophisticated than those for aggregate data, incorporating not only dynamics, heterogeneity and upgrading but also such features as learning, product loyalty, inventory behavior and surveys of price expectations (see Akerberg, 2003; Hendel & Nevo, 2006; Erdem & Keane, 1996; Erdem, Keane, Oncu & Strebel, 2005; Prince, 2007; Keane & Wolpin, 1997).¹ Extending these types of models to aggregate data is important for two reasons. First, in many cases, aggregate data are all that is available. Second, aggregate data are typically necessary for studying many important issues, such as oligopoly interactions. This is because household-level data sets rarely contain enough observations to measure product shares accurately.² Accurate market shares are important for estimating supply side. For instance, BLP and Goldberg (1995) use aggregate market share data to estimate pricing first-order conditions.

A number of recent papers (Gandal, Kende & Rob, 2000; Esteban & Shum, 2007; Melnikov, 2001; Song & Chintagunta, 2003; Gordon, 2006; Nair, 2007; Carranza, 2006; Park, 2008) propose dynamic consumer choice models for aggregate data. Most similar to our work is Melnikov (2001), which was the first to model dynamics in a logit-based discrete choice model with endogenous prices and a reduced-form approximation of the supply side. Our model builds on Melnikov (2001) by adding a full set of persistent random coefficients and repeat purchases over time, all modeled in an explicitly dynamic framework.

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

2 Model and Inference

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

2.1 Model

Our model starts with the introduction of a new consumer durable good at time $t = 0$. The unit of observation is a month and there is a continuum of heterogeneous potential consumers indexed by i .

¹An important comparison is to Hendel & Nevo (2006) which is also a logit-based model with endogenous repurchases and a similar approximation to the formation of expectations. By using disaggregate data, Hendel & Nevo (2006) are able to identify the parameters underlying consumer stockpiling. However, their model cannot be used with random coefficients on variables that vary within sizes.

²Our data contain 343 distinct camcorder models and 4,436 distinct model-months (a figure that is typical for new durable goods industries) implying that a survey would have to have over 100,000 purchases to measure shares accurately. The ICR-CENTRIS survey that we use for household level information interviews 4,000 individuals. By the end of our sample period, less than 15% of people had ever bought a digital camcorder, implying less than 600 total purchases.

Consumers have infinite horizons and discount the future with a common factor β . We assume that products are infinitely durable. However, if a consumer who owns one product purchases a new one, she obtains no additional utility from the old product, or equivalently, she discards the old product at no cost.³ We also do not consider resale markets because we believe that they are small for the new consumer durable goods that we examine given the speed of technological progress.

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 a similar (though not identical) decision problem at time $t + 1$. From these J_{t+1} choices, the consumer chooses the option that maximizes the sum of the expected discounted value of future expected utilities conditional on her information at time t .

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

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

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

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

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

where p_{jt} is the price of good j in period t and ε_{ijt} is an idiosyncratic unobservable meant to capture random variations in the purchase experience that do not persist across month, due to sales personnel,

³We consider the data and modeling necessary to loosen this restriction and others in Section 2.4 below.

weather, etc.⁴ We assume that ε_{ijt} is distributed type 1 extreme value, independent across consumers, products and time, and as such has mean γ , Euler's constant. We let α_i be constant over time and distributed normally with mean $\alpha \equiv (\alpha^x, \alpha^p)$ and variance matrix Σ , where α and Σ are parameters to estimate. Our empirical implementation uses a diagonal Σ matrix, although correlations can be easily added.

We also define the population mean flow utility

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

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

In our model, a consumer who does not purchase a new product at time t has net flow utility of

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

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

In order to evaluate consumer i 's choice at time t , we need to formalize consumer i 's expectations about the utility from future products. We assume that consumers have no information about the future values of the idiosyncratic unobservable shocks ε beyond their distribution. The set of products and their prices and characteristics vary across time due to entry and exit and changes in prices for existing products. Consumers are uncertain about future product attributes but have rational expectations about their evolution. We assume that each consumer is, on average over time, correct about the mean and variance of the future quality path.⁵

We now define the state variables and use them to exposit the dynamic decision process. Let $\varepsilon_{i,t} \equiv (\varepsilon_{i0t}, \dots, \varepsilon_{iJ_t t})$. Then, the purchase decision for consumer i depends on preferences α_i and $\varepsilon_{i,t}$, endowments δ_{i0t}^f , current product attributes and expectations of future product attributes. Future product

⁴Given that all camcorder prices are greater than one dollar, our specification for price is rationalized by a model where consumers have income Y_{it} ; they purchase at most two products, a money good m_{it} and possibly a camcorder; and the utility derived from the money good is $\alpha_i^p \min\{0, \ln(Y_{it} - m_{it})\}$. The specification can easily be modified to use the empirical income density, as in Nevo (2001)'s study on the breakfast cereal industry.

⁵A more general rational expectations model would allow individual consumers to have consistently biased estimates of the future logit inclusive value but let the mean expectations across consumers be accurate. While such a model would be easy to specify, it would be difficult to identify without expectations data.

attributes will depend on firm behavior which is a function of consumer endowments and supply-side factors such as technological progress. Let Ω_t denote current product attributes and any other factors that influence future product attributes. We assume that Ω_{t+1} evolves according to some Markov process $P(\Omega_{t+1}|\Omega_t)$ that accounts for firm optimizing behavior. Thus, the state vector for consumer i is $(\varepsilon_{i,t}, \delta_{i0t}^f, \Omega_t)$. Let $V(\varepsilon_{i,t}, \delta_{i0t}^f, \Omega_t)$ denote the value function, and $EV_i(\delta_{i0t}^f, \Omega_t) = \int_{\varepsilon_{i,t}} V_i(\varepsilon_{i,t}, \delta_{i0t}^f, \Omega_t) dP_\varepsilon$ denote the expectation of the value function, integrated over realizations of $\varepsilon_{i,t}$.

We can now define the Bellman equation for consumer i as

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

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

We can now use the aggregation properties of the type 1 extreme value distribution to express the expectation Bellman equation in a relatively simple form. In particular, Anderson, De Palma & Thisse (1992) and Rust (1987) show that the value of the best choice from several options in a logit model can be expressed as the logarithm of the sum of the exponents of the mean utility of each option plus a single type 1 extreme value draw.

To formally illustrate this property for our context, we require two definitions. First, for each product $j = 1, \dots, J_t$, let $\delta_{ij}(\Omega_t)$ denote the mean expected discounted utility for consumer i purchasing product j at time t ; thus

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

Second, define the *logit inclusive value* for consumer i at time t to be

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

We abbreviate $\delta_i(\Omega_t)$ as δ_{it} when it is clear to which Ω_t we are referring.

Then, Rust (1987) shows that the choice problem has the same expected value as a much simpler choice problem, where at each period t , the consumer makes a one-time purchase of a product with mean utility δ_{it} or holds the outside good with mean utility δ_{i0t}^f . Formally,

$$EV_i(\delta_{i0t}^f, \Omega_t) = \ln \left(\exp(\delta_{it}) + \exp \left(\delta_{i0t}^f + \beta E \left[EV_i(\delta_{i0t}^f, \Omega_{t+1}) \middle| \Omega_t \right] \right) \right) + \gamma. \quad (4)$$

Equation (4) shows that the state variable Ω_t only affects EV_i through its impact on the current and future values of δ_{it} . Intuitively, if two states yield the same contingent paths of δ_{it} then the expected value of either state is the same. We state this point formally:

Proposition 1 *Consider states Ω_t and Ω'_t for which the following two properties holds: (1) $\delta_i(\Omega_t) = \delta_i(\Omega'_t)$ and (2) $P(\delta_i(\Omega_{\tau+1})|\Omega_\tau) = P(\delta_i(\Omega'_{\tau+1})|\Omega'_\tau)$ if $\delta_i(\Omega_\tau) = \delta_i(\Omega'_\tau)$, for every $\tau \geq t$ and every state Ω_τ and Ω'_τ . Then, $EV_i(\delta_{i0t}^f, \Omega_t) = EV_i(\delta_{i0t}^f, \Omega'_t)$.*

The proof appears in the appendix.

Using Proposition 1, we can rewrite the state vector for EV_i so that

$$EV_i(\delta_{i0t}^f, \Omega_t) = EV_i(\delta_{i0t}^f, \delta_{it}, P[\delta_{i,\tau+1}|\Omega_\tau]). \quad (5)$$

While (5) provides a simplification, the state space is still infinite dimensional and so cannot be used for computational purposes. We wish to proceed by making simplifying assumptions on the evolution of δ_{it} . Before doing so, we first consider the relationship between δ_{it} and the underlying product characteristics. For any belief structure on the evolution of product characteristics, (2) and (3) provide an algorithm to compute current and future values of δ_{it} , showing that this mapping exists. We can also prove the converse given some regularity conditions:

Proposition 2 *Consider any set of contingent probabilities for the evolution of δ_{it} , $P_i[\delta_{i,\tau+1}|\Omega_\tau]$, $\tau > t$. Assume that the possible values for Ω_τ and $\delta_{i\tau}$ are both elements of a finite set. Then, for any vector of preferences α_i there is at least one set of conditional distributions $P_i(x_{j,\tau+1}, p_{j,\tau+1}|\Omega_\tau)$ such that this set of distributions together with optimizing behavior imply the contingent probabilities $P_i[\delta_{i,\tau+1}|\Omega_\tau]$.*

The proof appears in the appendix.

Proposition 2 shows that any set of restrictions on the evolution of δ_{it} will be consistent with some set of restrictions on the evolution of x_{jt} and p_{jt} .⁶ We proceed by making a major simplifying assumption, that the current δ_{it} is sufficient to predict future values of δ_{it} :

Assumption 1 *Inclusive Value Sufficiency (IVS)*

If $\delta_i(\Omega_t) = \delta_i(\Omega'_t)$, then $P_i(\delta_i(\Omega_{t+1})|\Omega_t) = P_i(\delta_i(\Omega'_{t+1})|\Omega'_t)$ for all t and Ω_t, Ω'_t .

The assumption of IVS and Proposition 1 together imply that all states with the same δ_{it} have the same expected value. Hence, it is sufficient to condition the value function on δ_{i0t}^f and δ_{it} rather than

⁶The proposition does not show that distributions over future states are what actually occur or that the distributions are similar across consumers with different random coefficients.

δ_{i0t}^f and Ω_t . Thus, the state space is reduced to two dimensions. We express

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

From Proposition 2, the IVS assumption together with any specification for the belief structure on the δ_{it} evolution can be generated with some belief structure on fundamentals, implying that IVS is always consistent with individual maximization.

For most of our specifications we assume that consumer i perceives $P_i(\delta_{i,t+1}|\delta_{it})$ as its actual empirical density fitted to a simple functional form and use a simple linear autoregressive specification,

$$\delta_{i,t+1} = \gamma_{1i} + \gamma_{2i}\delta_{it} + \nu_{it}, \quad (7)$$

where ν_{it} is normally distributed with mean 0 and γ_{1i} and γ_{2i} are incidental parameters specific to each consumer i . Similar functional forms have been used in the existing dynamic literature (see Melnikov, 2001; Hendel & Nevo, 2006). Our specification amounts to a bounded rationality assumption where consumers have a limited ability to predict the evolution of the future and can only predict partitions of the future state, but are correct on average about the probability of occurrence of each partition. Our functional form is of potential concern since it is not explicitly generated from supply-side dynamic optimization. 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. While dynamic profit maximization might lead these two states to have different patterns of industry evolution, consumers in our model will lump them into the same partition.⁷

In order to examine the robustness of our results, we also examine results from two other functional forms for expectations. First, we assume perfect foresight, where consumers know all future product attributes. This functional form is straightforward: the industry state is t . Moreover, it is a special case of IVS provided that the industry attributes are such that δ_{it} is never exactly the same for two time periods (as would occur if quality were improving, prices were non-decreasing and the set of products were non-decreasing): in this case, there would be a one-to-one mapping from t to δ_{it} and so δ_{it} would be sufficient to predict the future state. We do not use perfect foresight as our main specification because we believe that it is more realistic to assume that consumers have only a limited ability to predict the future.

Second, we loosen IVS by expanding the industry state space beyond δ_{it} . Given the potential importance of the number of products in determining industry evolution, we use J_t (the number of products) as an additional predictor, so that both δ_{it} and J_t predict $\delta_{i,t+1}$ and J_{t+1} . For this variant, we adjust

⁷Hendel & Nevo (2006) provide a similar discussion of the implications of (7).

IVS appropriately and we specify two linear regressions for the state evolution that are similar to (7). The linear regressions have $\delta_{i,t+1}$ and $\ln(J_{t+1})$ as dependent variables and include both aggregate state variables as regressors.⁸

Last, to examine the validity of our assumption, we also perform a test of IVS. We construct a moment based on the assumption of no autocorrelation $E[\nu_{it}\nu_{i,t+1}] = 0$ for a consumer i with mean characteristics. We do not impose this moment in estimation but rather test its validity at the estimated parameters.

An implication of (7) is that, for $0 < \gamma_{2i} < 1$,⁹ a graph of mean δ_{it} against time finds a concave line with an asymptote that is approached from below. This asymptote is important in our model since it represents a steady state in the evolution of product characteristics that the consumer expects to approach. The eventual arrival of a steady state is what allows us to treat the consumer as facing a stationary environment, even though observed choices are evolving quickly.

We now briefly discuss market shares and the supply side of the model. Using the value function resulting from IVS, we can write the probability that consumer i purchases good j as the aggregate probability of purchase times the probability of purchasing a given product conditional on purchase:

$$\begin{aligned} \hat{s}_{ij}(\delta_{i0t}^f, \delta_{ijt}, \delta_{it}) &= \frac{\exp(\delta_{it})}{\exp(EV_i(\delta_{i0t}^f, \delta_{it}) - \gamma)} \times \frac{\exp(\delta_{ijt})}{\exp(\delta_{it})} \\ &= \exp(\delta_{ijt} - EV_i(\delta_{i0t}^f, \delta_{it}) + \gamma). \end{aligned} \quad (8)$$

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

2.2 Inference

This subsection discusses the estimation of the parameters of the model, (α, Σ, β) , respectively the mean consumer tastes for product characteristics and price, the variance in consumer tastes in these variables

⁸This idea, of adding additional predictors to a limited information dynamic decision problem in order to test the impact of the limitation assumption, has been used in the macroeconomics general equilibrium literature to understand the impact of heterogeneity. See Krusell & Smith (1998).

⁹We estimate $0 < \gamma_{2i} < 1$ for all types i .

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

We develop a method for estimating the remaining parameters that is based on BLP and Rust (1987) and the literatures that follow.¹⁰ Our estimation algorithm involves three levels of non-linear optimizations: on the outside is a search over the parameters; inside that is a fixed point calculation of the vector of population mean flow utilities $\bar{\delta}_{jt}^f$; and inside that is the calculation of predicted market shares, which is based on consumers' dynamic optimization decisions.

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

For each draw, we start with initial guesses, calculate the logit inclusive values from (3), use these to calculate the coefficients of the product evolution Markov process regression in (7), and use these to calculate the expectation Bellman from (6). We repeat this process until convergence.¹¹ Using the resulting policy function $\hat{s}_{ij}(\delta_{i0t}^f, \delta_{ijt}, \delta_{it})$ and computed values of δ_{ijt} and δ_{it} , we then solve for market share for this draw by starting at time 0 with the assumption that all consumers hold the outside good.¹² Iteratively for subsequent time periods, we solve for consumer purchase decisions given the distribution of flow utility of holdings using (8) and update the distribution of flow utility of holdings based on purchases.

To perform the iterative calculation, we discretize the state space $(\delta_{i0t}^f, \delta_{it})$ and the transition matrix. Specifically, we compute the value function by discretizing δ_{i0t}^f into 20 evenly-spaced grid points and δ_{it} into 50 evenly-spaced grid points and allowing 20 points for the transition matrix. We specify that δ_{it} can take on values from 20% below the observed values to 20% above and assume that evolutions of δ_{it} that would put it above the maximum bound simply place it at the maximum bound. We have examined

¹⁰Computer code for performing the estimation is available from the authors upon request.

¹¹We also perform this three-equation fixed point to solve for new consumer decisions following counterfactual changes in product characteristics or prices. In particular, note that it is necessary to recalculate new δ_{it} , γ_{1i} and γ_{2i} values (in addition to the value function) since these are all jointly determined by optimizing behavior.

¹²A strength of our data set is that it reaches back essentially to the start of the industry, so we can assume that all consumers start with nothing. In another setting, we would have to make assumptions or estimate consumer holdings at the start of the time horizon. For an example, see Schiraldi (2007).

the impact of easing each of these restrictions and found that they have very small effects on the results.

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

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

and then reweight draws by

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

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

We now turn to the middle loop, which recovers $\bar{\delta}_{..}^f$ by performing a fixed point equation similar to that developed by BLP. We iterate until convergence

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

where $\hat{s}_{jt}(\bar{\delta}_{..}^{f,old}, \alpha^p, \Sigma)$ is the model market share (computed in the inner loop), s_{jt} is actual market share, and ψ is a tuning parameter that we generally set to $1 - \beta$. Note that it is not necessary to treat the inner loop and middle loop as separate. We have found some computational advantages to taking a step in (10) before the inner loop is entirely converged and to performing (6) much more frequently than either (7) or (3). However, we require full convergence of (3), (6), (7) and (10) before moving to the outermost loop.

An important issue is whether (3), (6), (7) and (10) have a unique fixed point, which is necessary to guarantee identification of the model. We have used a variety of different starting values and have never had a problem with not finding a solution or finding too many solutions. However, we cannot

prove uniqueness of the fixed point. Berry (1994) proves uniqueness for models where all products are substitutes. A variant of our model where consumers can only purchase once and where every current and future product attribute including the extreme value shocks are known would satisfy substitutability. In contrast, in most dynamic models, products may be complements. As one example, if we exogenously increase utility of a current period product, we increase sales at the expense of sales of other products in the current period. With innovation, if this leads to lower sales next period it may lead to higher sales in two periods as more consumers will value an upgrade. We believe that this complementarity is very unlikely to result in multiple solutions, as would occur from network effects, for instance. Nonetheless, it makes a proof of uniqueness very challenging.

The outer loop specifies a GMM criterion function

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

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

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

where W is a weighting matrix.

We minimize (11) by performing a nonlinear search over (α^p, Σ) . For each (α^p, Σ) vector, we first obtain $\bar{\delta}_{jt}^f$ from the middle loop. The fact that $\alpha^x x_{jt}$ and ξ_{jt} enter flow utility linearly (recall that $\bar{\delta}_{jt}^f = \alpha^x x_{jt} + \xi_{jt}$) then allows us to solve in closed form for the α^x that minimizes (11) given $\bar{\delta}_{jt}^f$, as in the static discrete choice literature.¹³ We perform the nonlinear search using a simplex method. We perform a two-stage search to obtain asymptotically efficient estimates. In the first stage, we let $W = (z'z)^{-1}$, which would be efficient if our model were linear instrumental variables with homoscedastic errors, and then use our first stage estimates to approximate the optimal weighting matrix.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 vastly simplified due to the fact that only consumers who have never purchased make decisions. Because of this, (2) can be simplified to $\delta_{ijt} = (\delta_{ijt}^f + \beta\gamma)/(1 - \beta) - \alpha_i^p \ln(p_{jt})$ which implies that δ_{it} in (3) does not depend on the value function. Moreover, we need only solve the expectation Bellman equation (6) for $\delta_{ijt}^f = 0$ and hence

¹³See Nevo (2000b) for details. One difference from the static model is that we cannot solve in closed form for α^p since the price term, $\alpha^p \ln(p_{jt})$, is only paid at the time of purchase, unlike ξ_{jt} .

¹⁴See again Nevo (2000b) for details.

there is effectively one state variable, δ_{it} , instead of two. The computation of the outer loop for this model is also quicker, since the price coefficient α^p can now be solved in closed-form, like α^x in the base model.

2.3 Identification and instruments

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

Note that our model allows for consumers to purchase products repeatedly over time, even though it can be estimated without any data on repeat purchase probabilities for individuals. At first glance, it might appear difficult to identify such a model. However, except for the discount factor β , which we do not estimate, this model does not introduce any new parameters over the model with one-time purchases or the static model. The reason is that we have made some relatively strong assumptions about the nature of the product which imply that the only empirically relevant reason to buy a second durable good is new features, and features are observed in the data. However, we also experiment with incorporating survey data on household purchase patterns that address this issue below.

As is standard in studies of market power since Bresnahan (1981), we allow price to be endogenous to the unobserved term (ξ_{jt}) but we assume that product characteristics are exogenous. This assumption is justified under a model in which product characteristics are determined as part of some technological progress which is exogenous to the unobserved product characteristics in any given period. As in Bresnahan and BLP, we do not use cost-shifters as instruments for price and instead exploit variables that affect the price-cost margin. Similar to BLP, we include the following variables in z : all of the product characteristics in x ; the mean product characteristics for a given firm at the same time period; the mean

product characteristics for all firms at the time period; and the count of products offered by the firm and by all firms. These variables are meant to capture how crowded a product is in characteristic space, which should affect the price-cost margin and the substitutability across products, and hence help identify the variance of the random coefficients and the price coefficient. While one may question the validity of these instruments, they are common in the literature. We consider the development of alternative instruments a good area for future research.

2.4 Extensions to the model

As noted above, we specify a base model that is relatively parsimonious but that we believe still captures the important features of the digital camcorder market. Many of the assumptions of our model could be easily relaxed provided that data to identify the resulting additional parameters exist. We now explain how to extend the model in several ways that would be of use in estimating preferences for other durable goods industries.

First, we assume that products do not depreciate. It would be easy to modify our model to allow for either stochastic or deterministic product depreciation. This would alter the Bellman equation (1) so that the future individual state would not be δ_{i0t}^f but rather a depreciated function of this value. Micro-moment data (see Berry, Levinsohn & Pakes, 2004; Petrin, 2002) on the durability of current goods or on the frequency of repeat purchases would identify the depreciation coefficients. Second, we assume that consumers can only hold one digital camcorder concurrently. We could modify the model to allow consumers to hold two products simultaneously, by allowing for two consumer state variables, one for the flow utility of each product. Each period, consumers would then choose between replacing either of their existing products or keeping them both – resulting in $2J_t + 1$ choices. Utility from the second product would be a multiple – likely less than 1 – of the utility from the first product, resulting in one extra parameter. Micro survey panel data on the stock of goods would identify this extra parameter. Third, we assume that the set of consumers in the market remains constant over time. We could model the exogenous entry of new consumers in a straightforward manner. Micro survey data on when households started considering a category for purchase would be useful to credibly identify the level of exogenous entry. In combination with micro survey data on the stock of goods, we could potentially identify differences in random coefficients by date of entry into the market.

Last, and most importantly, we assume that there are no resale markets for digital camcorders. Many durable goods industries, such as automobiles, have important resale markets. We could model an industry with resale markets using data on the quantity and price of used products. In this case,

consumers would have as choices all new and used products on the market. We would also want to allow for depreciation of products, as above. The Bellman equation would be modified to reflect the fact that a consumer who owns a product and upgrades to a new product would also receive the market price for her current product and vintage, upon trading in her current product. Thus, consumers would need to forecast the market price for their current model in addition to future values of δ_{it} . We owe this discussion to Schiraldi (2007), who extends our model to analyze the market for new and used cars in Italy in this way. He allows for transactions costs, which are necessary to explain the empirical fact that consumers do not reoptimize their choice of cars (or other products) very frequently.

3 Data

We estimate our model principally using a panel of aggregate data for digital camcorders.¹⁵ The data are at the monthly level and, for each model and month, include the number of units sold, the average price, and other observable characteristics. We observe 383 models and 11 brands, with observations from March 2000 to May 2006. These data start from very early in the product life cycle of digital camcorders and include the vast majority of models. The price and quantity data were collected by NPD Techworld which surveys major electronics retailers and covers 80% of the market.¹⁶ We create market shares by dividing sales by the number of U.S. households in a year, as reported by the U.S. Census.

We collected data on several important characteristics from on-line resources. We observe the number of pixels that the camera uses to record information, which is an important determinant of picture quality. We observe the amount of magnification in the zoom lens and the diagonal size of the LCD screen for viewing shots.¹⁷ We observe the width and depth of each camera in inches (height was often unavailable), which we multiply together to create a “size” variable. We also record indicators for whether the camera has a lamp, whether it can take still photos and whether it has “night shot” capability, an infrared technology for shooting in low light situations. Finally, we observe the recording media the camera uses – there are four mutually exclusive media (tape, DVD, hard drive and memory card) – which we record

¹⁵We have obtained similar data for digital cameras and DVD players and previous versions of this paper estimated those industries. Basic features of the results are similar across industries. We focus on camcorders because we believe this product exhibits the least amount of endogenous complementary goods or network effects (such as titles for DVD players or complementary products for producing pictures for digital cameras), which would complicate our analysis. Incorporating network effects into our framework is the subject of current research.

¹⁶NPD sales figures do not reflect on-line sellers such as Amazon and they do not cover WalMart. This could potentially bias welfare results if these vendors disproportionately sell particular types of products.

¹⁷We log all continuous variables and treat any screen of less than .1 inch as equivalent to a screen of .1 inch.

Table 1: Characteristics of digital camcorders in sample

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

as indicators.

To create our final data set, we exclude from the choice set in any month all digital camcorders that sold fewer than 100 units in that month. This eliminates about 1% of sales from the sample. We also exclude from the choice set in any month all products with prices under \$100 or over \$2000 as these products likely have very different usages. This eliminates a further 1.6% of sales from the sample. Our final sample includes 343 models and all 11 brands. The number of products varies from 29 in March 2000 to 98 in May 2006. Table 1 summarizes the sales, price and characteristics data by level of the model-month for our final sample.

Figure 1 graphs simple averages of two features over time, size and pixel count, using our final sample. Not surprisingly, cameras improve in these features over time. Weighting by sales produces similar results. Figure 2 displays a similar graph for features that are characterized by indicator variables: the presence of a lamp, the presence of night shot, the ability to take still photographs and whether the recording media is tape. The first two systematically become more popular over time. Photo ability is present in nearly every camera half-way through the sample but declines slightly in popularity by the end of our sample. Tape-based camcorders initially dominated the market but grew less popular over time relative

Figure 1: Average indicator characteristics over time



Figure 2: Average indicator characteristics over time

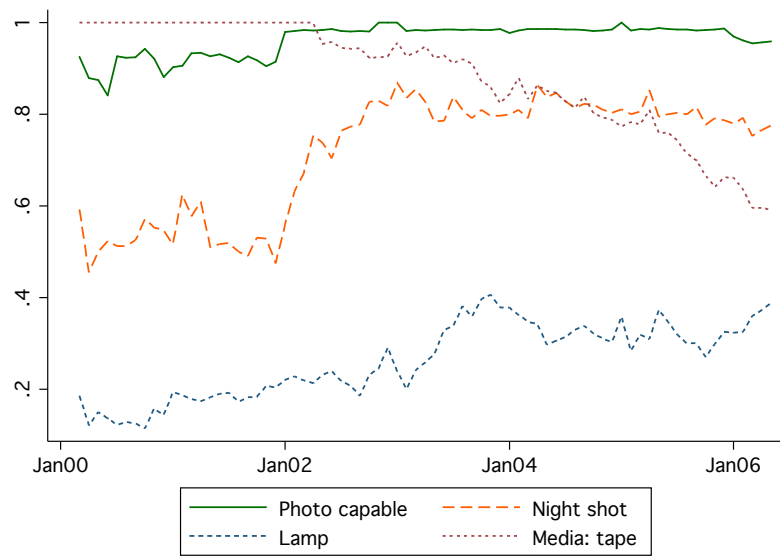
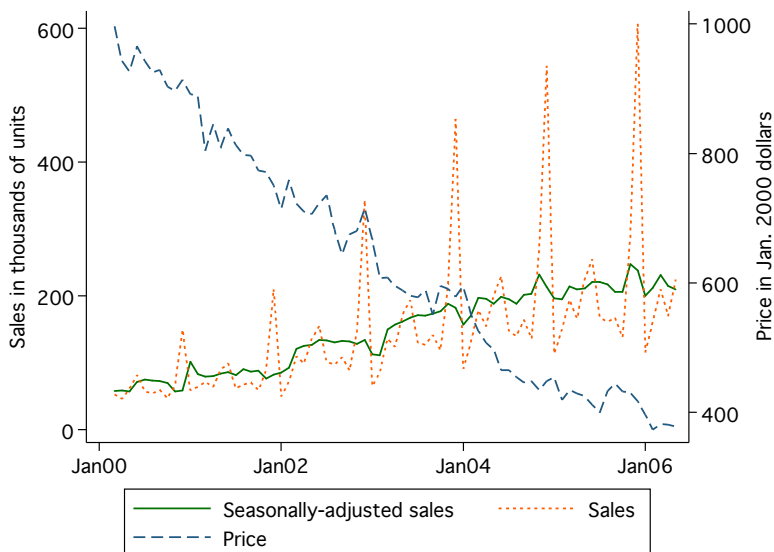


Figure 3: Prices and sales for camcorders

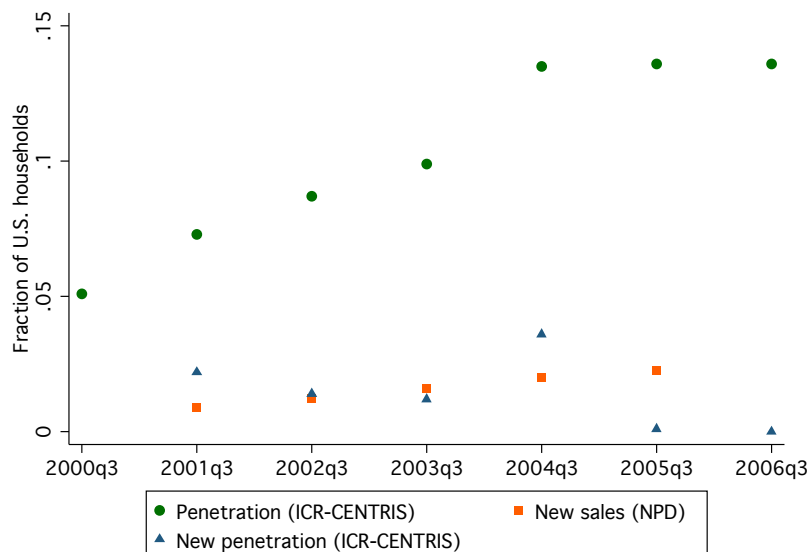


to DVD- and hard drive-based devices, representing more than 98% of devices in the first few months of the data but less than 65% in the last few.

Figure 3 shows total sales and average prices for camcorders in our final sample over time. Camcorders exhibit striking price declines over our sample period while sales increase. Even more noticeable than the overall increase in sales is the huge spike in sales at the end of each year due to Christmas shopping. Note that while quantity changes over the Christmas season, there is no visible effect on prices or characteristics.

Our model needs to have some way of explaining the huge impact of the Christmas season on sales. One way is to add a monthly characteristic to each product. Given that our demand system is dynamic, this vastly complicates our model by adding another state variable (month of the year). Moreover, it is unlikely that products bought over Christmas are inherently more valuable in the future. Thus, we believe that a reasonable model would assume that a purchase in December adds to utility at the time of purchase, rather than adding to δ_{ijt}^f , the future flow utility of the product. Thus, this results in additional parameters that are estimated non-linearly. We estimate one specification with seasonal effects. This specification adds 11 parameters, one for each month but January, that specify the additional utility at the time of purchase in that month. It also modifies the Bellman equation to have the month of the year as a state variable, and modifies the regressors in the state expectation equation (7) to allow for

Figure 4: Penetration and sales of digital camcorders



month-of-year dummies instead of just a constant term.

Given the computational complexity of this model and the fact that only sales change over Christmas, for most of our specifications, we addressed the Christmas spike issue by seasonally adjusting our data. Specifically, we multiplied sales by a separate constant for each month, constant across years. The constants were chosen so that the sales by month summed over the years in the data were the same for each month and so that total sales for each year were unchanged. Figure 3 also shows the seasonally-adjusted sales data, which are, by construction, much smoother than the unadjusted data.

In addition, in some specifications we incorporate household level data on ownership, often referred to as penetration, to better pin down repeat purchasing behavior. These data come from ICR-CENTRIS, which performs telephone interviews via random-digit dialing. ICR-CENTRIS completes about 4,000 interviews a month, asking which consumer electronics items a household owns.¹⁸ Figure 4 shows our ICR-CENTRIS data, which contain the percent of households that indicate holding a digital camcorder in the third quarter of the year for 1999 to 2006. It also shows the year-to-year change in this number and the new sales of camcorders, as reported by NPD.

The penetration data show rapid growth in penetration early on in the sample but no growth by

¹⁸Data on how many camcorders a household owns or data on the time between purchases would be even more directly useful for understanding repeat purchases. However, a lengthy search of public and private data sources did not turn up any such information.

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

4 Results and implications

We first exposit our results, then provide evidence on the fit of the model, discuss the implications of the results and finally use our results to analyze dynamic COLIs.

4.1 Results

We present our parameter estimates in Table 2. Table 2 contains four columns of results. The first column of results provides the parameter estimates and standard errors from our base specification of the model presented in Section 2 with two random coefficients, one on price and the other on the constant term. The base specification reports results that are generally sensible in magnitude and sign. As we would hope, price contributes negatively to utility for virtually everyone, with a base coefficient of -3.30 and a standard deviation on the random coefficient of $.345$. Both are precisely estimated. A person with mean tastes would obtain a negative gross flow utility from a camcorder with all characteristics zero (relative to the outside option), with a mean constant term of $-.092$. The standard deviation on the constant term in the consumer population is $.079$, indicating that there is substantial variation in the gross flow utility from a camcorder. Again, both coefficients are statistically significant. In comparing the magnitudes of these coefficients, recall that price is paid once, while all the other coefficients relate to flow utility at the level of the month, and hence the price coefficients should be roughly $1/(1 - \beta) = 100$ times the magnitude of the other coefficients as compared to a static model.

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

Table 2: Parameter estimates

Parameter	Base dynamic model	Dynamic model without repurchases	Static model	Dynamic model with micro-moment
	(1)	(2)	(3)	(4)
Mean coefficients (α)				
Constant	-.092 (.029) *	-.093 (7.24)	-6.86 (358)	-.367 (.065) *
Log price	-3.30 (1.03) *	-.543 (3.09)	-.099 (148)	-3.43 (.225) *
Log size	-.007 (.001) *	-.002 (.116)	-.159 (.051) *	-.021 (.003) *
Log pixel	.010 (.003) *	-.002 (.441)	-.329 (.053) *	.027 (.003) *
Log zoom	.005 (.002) *	.006 (.104)	.608 (.075) *	.018 (.004) *
Log LCD size	.003 (.002) *	.000 (.141)	-.073 (.093)	.004 (.005)
Media: DVD	.033 (.006) *	.004 (1.16)	.074 (.332)	.060 (.019) *
Media: tape	.012 (.005) *	-.005 (.683)	-.667 (.318) *	.015 (.018)
Media: HD	.036 (.009) *	-.002 (1.55)	-.647 (.420)	.057 (.022) *
Lamp	.005 (.002) *	-.001 (.229)	-.219 (.061) *	.002 (.003)
Night shot	.003 (.001) *	.004 (.074)	.430 (.060) *	.015 (.004) *
Photo capable	-.007 (.002) *	-.002 (.143)	-.171 (.173)	-.010 (.006)
Standard deviation coefficients ($\Sigma^{1/2}$)				
Constant	.079 (.021) *	.038 (1.06)	.001 (1147)	.087 (.038) *
Log price	.345 (.115) *	.001 (1.94)	-.001 (427)	.820 (.084) *

Standard errors in parentheses; statistical significance at 5% level indicated with *. All models include brand dummies, with Sony excluded. There are 4436 observations.

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

A potential concern in our context is the restrictiveness of the logit error assumption. Logit errors typically imply unrealistic welfare gains from new products (see Petrin, 2002). Akerberg & Rysman (2005) argue that this feature implies that logit-based models will perform poorly in contexts where consumers face different numbers of products over time. Akerberg & Rysman recommend addressing this problem by including the log of the number of products, $\ln(J_t)$, as a regressor, as if it were a linear element in $\bar{\delta}_{jt}^f$. Finding a coefficient of 0 implies the logit model is well-specified, whereas a coefficient of -1 implies “full-crowding,” so there is no demand expansion effect from variety. In unreported results, we find that other parameters change little and that the coefficient on $\ln(J_t)$ is $-.013$. Although the coefficient is statistically significant, it is very close to zero and suggests that the *i.i.d.* logit draws are a reasonable approximation. Concerns with the implications of logit draws motivate Berry & Pakes (2005) and Bajari & Benkard (2005) to propose discrete choice models that do not include logit *i.i.d.* error terms, but given this coefficient estimate, we do not further pursue this issue.

Column 2 provides estimates from the dynamic model where individuals are restricted to purchase at most one digital camcorder ever. This specification yields results that are less appealing than our base specification. In particular, the mean price coefficient drops in magnitude by a factor of 6 and loses its statistical significance. Many of the characteristics enter mean utility with an unexpected sign, including pixels, LCD screen size and lamp and many fewer mean coefficients are significant than in the base specification. The standard deviation coefficients are very small and statistically insignificant. We apply a formal test of model selection. Rivers & Vuong (2002) derive a test statistic that has a standard normal distribution under the null hypothesis that the two models fit the data equally well (in this case, in the sense of the GMM objective function).¹⁹ The value of the test statistic is 5.55, which strongly rejects the single purchase model in favor of our base model. In the one-time purchase model, the magnitude of the mean price coefficient is much smaller than the standard deviation of the extreme value distribution. Had this estimated coefficient been applied to the base model, many people would purchase a camcorder most months, which they are unable to do in the one purchase model. This sharp difference in purchase patterns between the two models explains why the coefficient estimates can be so different.

Column 3 follows BLP and estimates a traditional static random coefficients discrete choice speci-

¹⁹Following Jaumandreu & Moral (2006), we base our test statistics for the non-nested test on the consistent first-stage GMM estimates.

fication. To compare these coefficients with the base specification, one would have to multiply all the coefficients from this specification, except for the coefficients on price, by 1/100. The static model yields many unappealing results, including a barely negative price coefficient with an enormous standard error and many coefficients on characteristics that are of the opposite sign from expected. We similarly perform a non-nested test of this model against the base model and obtain a test statistic of 5.7, which strongly rejects the static model in favor our base model.

We believe that the very imprecise price coefficients in the static specifications is caused by the fact that the data cannot easily be explained by a static model. In particular, the static model cannot fit two facts that are characteristic of the data: first, many more people purchased digital camcorders once prices fell; but second, within a time period, the cheapest models were often not the most popular. Because the model cannot then estimate a significantly negative price coefficient, it also does not result in appropriate coefficients on characteristics.

The dynamic model addresses these two facts because it predicts that people wait to purchase when expected price *declines* are small, not necessarily when prices are small. Heuristically, the static price coefficient is analogous to the coefficient from a regression of market shares on prices whereas the the dynamic price coefficient is analogous to the coefficient from a regression of shares on the forward difference in price, $(p_{jt} - \beta p_{j,t+1})$.²⁰ Unlike the static explanation, the dynamic explanation for why consumers wait does not conflict with consumers buying relatively high-priced products.

As we show below, our base specification implies very little repeat purchase. Thus, we use the penetration data in the form of a micro-moment (see Petrin, 2002) as a check on our base results. Specifically, we use the penetration data to construct an additional moment that is the difference between the increase in household penetration between Sep. 2002 and Sep. 2005 predicted by the model and by the penetration data.²¹ We chose to use only these two years to mitigate the noise present in the data. Column 4 reports the result. There are two main differences between these results and the base specification. First, the standard deviation of the random coefficient on price more than doubles. That increases the set of consumers who care about price very little. Second, the coefficients on the characteristics increase, often becoming 2 or 3 times as large. The parameters that increase the most are on the characteristics that improve the most over time. For instance, there are big parameter changes on size, zoom and pixel count and small changes on the presence of a lamp or a photograph option. Hence,

²⁰Gandal et al. (2000) show that this heuristic is an exact description of the market with one product, perfect foresight, zero variance to ε_{ijt} , linearity in prices, no repeat purchase, and a concave price path.

²¹See Berry et al. (2004) and Petrin (2002) for details on calculating weighting matrices when combining micro moments with aggregate moments.

the model generates repeat purchase by creating a set of price insensitive consumers and increasing the importance of characteristics that improve over time.

In Table 3, we present a number of robustness checks. Column 1 explores the importance of the IVS assumption by including $\ln(J_t)$ as an additional state variable. The results are very similar to our base specification, lending support to the IVS assumption. The second column estimates a model with perfect foresight where the market stops evolving at the last period in our data so that the market structure available there is exactly what is available ever after. Although it leads to a smaller price coefficient and zero heterogeneity around price, the model generates mostly the same qualitative results. Hence, it does not appear that our particular specification of expectations is crucial in generating our results.

Turning now to column 3, the addition of two extra random coefficients results in parameter estimates for mean coefficients that are very similar to the base specification. In particular, the sign of the mean coefficients on price and characteristics are all the same as in the base specification, and statistical significance is similar across specifications, except that the random coefficient on price is now close to 0. Moreover, the two new random coefficients are estimated to be small and statistically insignificant.

While models with the log of price tend to fit data better, it is easier to theoretically justify a model with linear price from the perspective of consumers with heterogeneous incomes. Column 4 estimates our model with a linear price. Again, the qualitative results look similar.

Column 5 estimates a variant of the static model where we aggregate the products to the annual level.²² The results from this specification are similar to the results from the static BLP model at the monthly level. Column 6 estimates a model with single purchase and no random coefficients, which is the model considered by Melnikov (2001). We solve it based on our method rather than the multi-stage model that Melnikov proposes. The results are not particularly appealing, with an insignificant price coefficient and numerous negative coefficients on characteristics. Using Melnikov’s method finds similar results.

Column 7 estimates the model with monthly effects as described in Section 3. Not reported in the table, the utility function for this specification includes month dummies for utility at the time of purchase, which are .720 for December, .250 for November and which range from $-.146$ to $.104$ for the other months. Only the December effect is statistically significant. This model fully exploits the cross-month substitution for identification purposes since the data used in this specification does not normalize away any monthly variation. Nonetheless, the estimated coefficients on price and characteristics are remarkably similar to the base specification.

We do not address a number of issues which might be important in diffusion contexts, such as

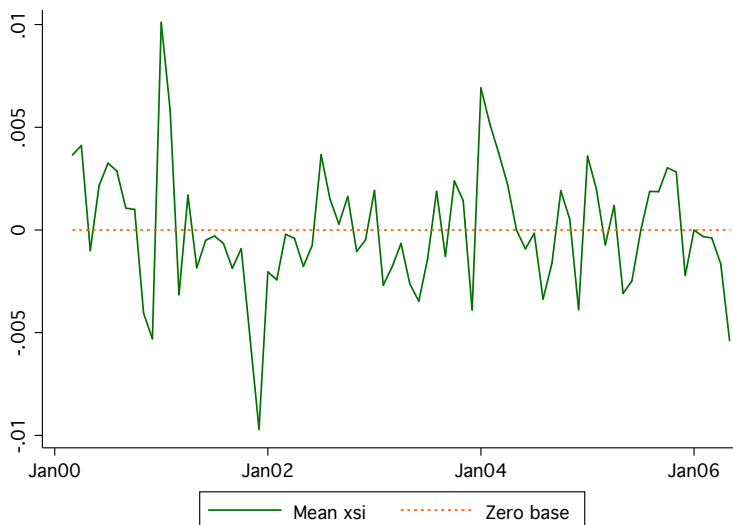
²²This specification drops the first and last year from our data, as we lack information on all months for those years.

Table 3: Robustness

Parameter	State space includes number of products		Perfect foresight	Dynamic model with extra random coefficients		Linear price	Static model aggregated to year	Meinikov's model	Month dummies
	(1)	(2)		(3)	(4)				
Mean coefficients (α)									
Constant	-.098 (.026) *	-.129 (.108)	-.103 (.037) *	-.170 (-.149)	2.26 (2325)	-6.61 (.815) *	-1.14 (.024) *		
Log price	-3.31 (1.04) *	-2.53 (.940) *	-3.01 (.717) *	-6.94 (.822) *	-304 (38.0)	-1.89 (.079) *	-3.06 (.678) *		
Log size	-.007 (.001) *	-.006 (.001) *	-.015 (.007) *	.057 (.008) *	.653 (.234) *	-.175 (.049) *	-.007 (.001) *		
Log pixel	.010 (.003) *	.008 (.001) *	.009 (.002) *	.037 (.012) *	-.863 (.822)	-.288 (.053) *	.010 (.002) *		
Log zoom	.005 (.002) *	.004 (.002) *	.004 (.002) *	-.117 (.012) *	-.779 (.455)	.609 (.074) *	.005 (.002) *		
Log LCD size	.004 (.002) *	.004 (.001) *	.004 (.002) *	.098 (.010) *	-.283 (.626)	-.064 (.088)	.003 (.001) *		
Media: DVD	.033 (.006) *	.025 (.004) *	.044 (.018) *	.211 (.053) *	.744 (2.02) *	.147 (-.332)	.031 (.005) *		
Media: tape	.013 (.005) *	.010 (.004) *	.024 (.016)	.200 (.051) *	.268 (1.43)	-.632 (-.318) *	.012 (.004) *		
Media: HD	.036 (.009) *	.026 (.005) *	.047 (.019) *	.349 (.063) *	.406 (2.63) *	-.545 (.419)	.034 (.007) *		
Lamp	.005 (.002) *	.003 (.001) *	.005 (.002) *	.077 (.011) *	-.925 (.579)	-.200 (.058) *	.004 (.001) *		
Night shot	.003 (.001) *	.004 (.001) *	.003 (.001) *	-.062 (.008) *	.855 (.275) *	.427 (.058) *	.003 (.001) *		
Photo capable	-.007 (.002) *	-.005 (.002) *	-.007 (.002) *	-.061 (.019) *	4.41 (1.71) *	-.189 (.142)	-.007 (.008)		
Standard deviation coefficients ($\Sigma^{1/2}$)									
Constant	.085 (.019) *	.130 (.098)	.081 (.025) *	.022 (.004) *	.027 (6800)		.087 (.013) *		
Log price	.349 (.108) *	2.41e-9 (.919)	1.06e-7 (.522)	1.68 (.319) *	-.006 (131)		.287 (.078) *		
Log size			-.011 (.007)						
Log pixel			1.58e-10 (.002)						

Standard errors in parentheses; statistical significance at 5% level indicated with *. All models include brand dummies, with Sony excluded. There are 4436 observations, except in the yearly model, in which there are 505.

Figure 5: Average estimation error (ξ_{jt}) by month



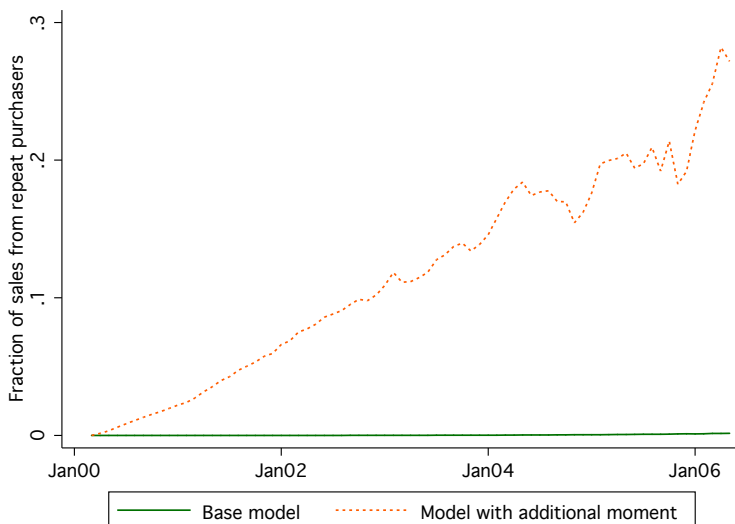
consumer learning or neighborhood effects. These would be difficult to address with aggregate data. However, a simple way to capture some of these issues is to use a time trend. In unreported results, we experimented with a quadratic time trend. The coefficients on time came out economically unimportant and the remaining parameters were very similar.

4.2 Fit of the model

We first assess the fit of the model by reporting the simple average of the unobserved quality ξ_{jt} for each month in Figure 5 using the estimated parameters from the base specification, Table 2 column 1, and the vector of $\bar{\delta}_{jt}^x$ that are consistent with these parameters. Note that ξ_{jt} is the estimation error of the model. The figure does not indicate any systematic autocorrelation or heteroscedasticity of the average error over time. This finding is important because there is no reduced-form feature such as a time trend to match the diffusion path. If one were to match, for instance, an S-shaped diffusion path with a simple linear regression, we would expect to have systematic autocorrelation in ξ_{jt} . However, Figure 5 does not indicate any such pattern.

We also look at the extent to which the model generates repeat purchases. Figure 6 plots the fraction of shares due to repeat purchases for the base model as well as for the model with the micro-moment, Table 2, column 4. Under the base model, repeat purchases account for a very small fraction of total

Figure 6: Evolution of repeat purchase sales

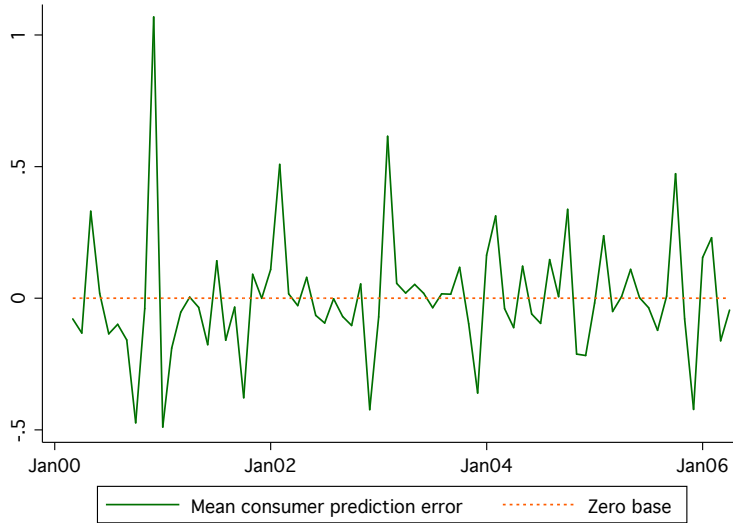


sales. Even in the final period, which has the largest fraction, repeat purchases account for only about .25% of new sales. The underlying reason why there are not more repeat purchases is that the coefficients on characteristics other than the constant term are small relative to the utility contribution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

This finding is not consistent with the evidence, albeit imperfect, from the ICR-CENTRIS household penetration survey, that new sales are higher than new penetration. Figure 6 also plots the share of repeat purchases for the specification with the micro-moment. Since this model fits both the increase in penetration of 4.9% from Sep. 2002 to Sep. 2005 and the new sales of 5.85% over the same time period, it predicts much higher repeat purchases than the base model. In particular, it predicts that over 25% of new sales are attributable to repeat purchases by the end of the sample.

Figure 7 shows the difference between $\delta_{i,t+1}$ and the period t prediction of this value, for a consumer with draws in the 50th percentile for both random coefficients. There do not appear to be any significant deviations in the AR(1) process from our assumed functional form. To verify this formally, we estimate the value of the additional moment based on serially uncorrelated values of ν_{it} using the median consumer. We find that this moment has a mean of $-.474$ with a standard deviation of 2.95 implying that we cannot reject the null hypothesis that the residuals are not serially correlated.

Figure 7: Difference between $\delta_{i,t+1}$ and its period t prediction



4.3 Implications of the results

We now analyze some implications of the results using the base specification, Table 2 column 1. We first compare the sources of heterogeneity in our results. To understand the evolution of δ_{it} in our model, Figure 8 plots δ_{it} for 3 sets of random coefficients, for which the percentiles in their draws for the price and constant terms are 80-80, 20-20, and 80-20 respectively. For all consumers, values are increasing close to linearly over time. As the linearity should make evident, the estimated asymptotes to the AR(1) processes are reached to the future of our data for the reported draws (and indeed all draws that we use).

The value that the 20-20 consumer places on the market at the end of the sample is far below the value that the 80-80 consumer places at the beginning. That is, the heterogeneity in valuation of the product swamps the changes over time. The second two lines allow us to compare consumers that differ only in their price sensitivity. Again, we see that the heterogeneity in the constant term is more important. That follows for two reasons: first, the lines are relatively close to their counterparts with different price draws and second, there is little compression over time even though prices are dropping. Because it is hard to see the level of compression, we plot the difference between the 80-80 and 80-20 lines separately; the difference decreases by 15% over the sample period.

Next, we analyze dynamic price elasticities. We compare three price changes: a temporary (one-month) 1% price increase at time \bar{t} that consumers know to be temporary; a temporary increase that

Figure 8: Evolution of δ_{it} over time

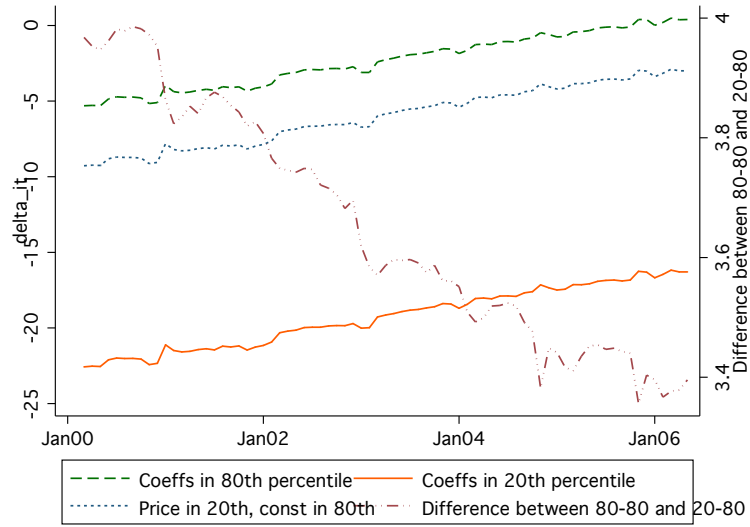


Figure 9: Industry dynamic price elasticities

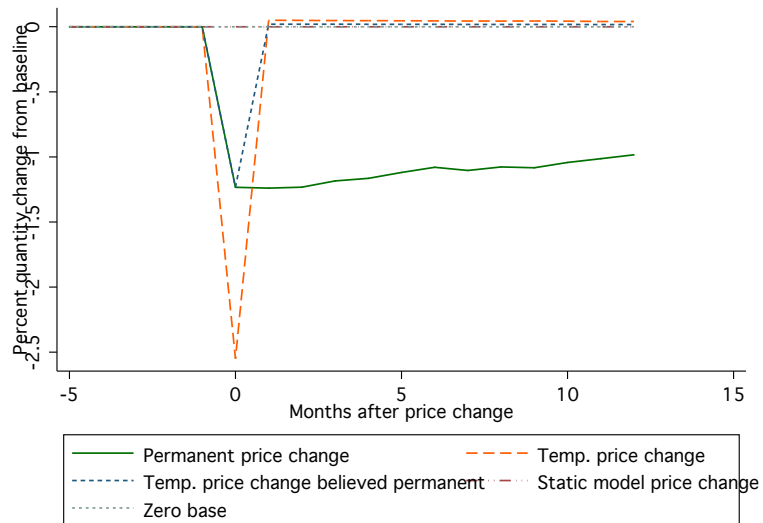
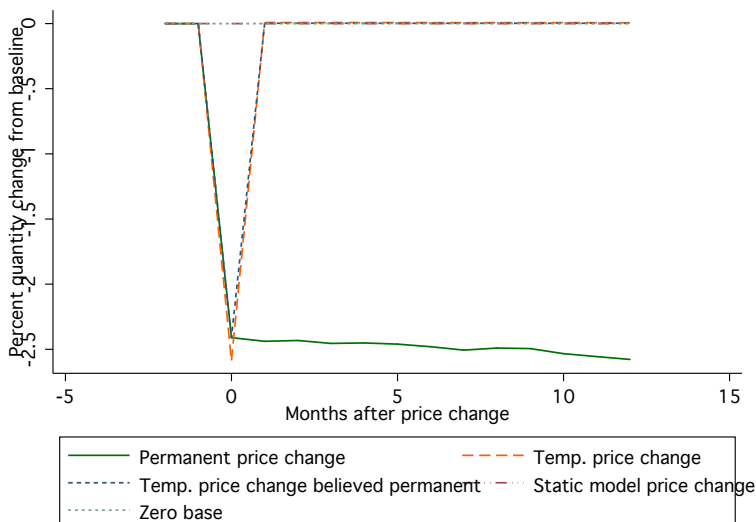


Figure 10: Dynamic price elasticities for Sony DCRTRV250



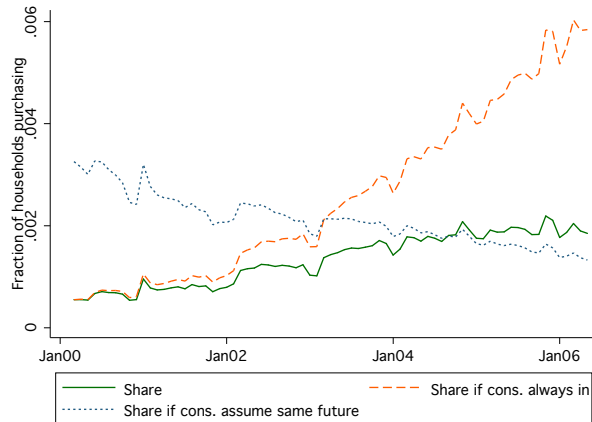
consumers believe to be permanent; and a permanent price increase. In all cases, the price increase is unexpected before time \bar{t} . When consumers believe the increase to be temporary, we compute the time \bar{t} expectations of $\delta_{i,\bar{t}+1}$ using the baseline $\delta_{i\bar{t}}$; for the perceived permanent price change, we use the realized $\delta_{i\bar{t}}$. For all specifications, we keep the estimated γ_{1i} and γ_{2i} coefficients.²³

Figure 9, which displays the industry elasticity with \bar{t} set to the median period of the sample (April 2003), shows that a temporary price change results in twice as big a response as a permanent one. Specifically, a 1% price increase leads to a contemporaneous decrease in sales of 2.55% when consumers believe it to be temporary and a decrease of only 1.23% when they believe it to be permanent. In addition, the response over the following year is also larger, but in the opposite direction: when consumers believe the temporary price change to be temporary, sales will increase by .54% of the time \bar{t} sales for the following 12 months compared to an increase of .22% under the temporary but believed-permanent price change.

Figure 10 considers the own price elasticity for the Sony DCRTRV250, which had the largest market share in the median period. Here, the difference in response between a temporary and permanent price change is small: 2.59% versus 2.41%. This result follows because consumers switch to another product rather than delay their purchase when one product changes price permanently.

²³The price elasticities from the static model are all virtually 0, so we do not include them on the figures.

Figure 11: Evolution of digital camcorder sales under different assumptions



Strikingly, we find that the temporary price elasticities are almost the same for the industry as a whole and the Sony DCRTRV250. However, the sources of the quantity change are different: a delay for the industry change but a switch to other products for the product change. Because of the difference in source, the long-run industry temporary price elasticity is much smaller than for the product, as in the industry case, consumers recover over 20% of the sales reduction in later periods, while virtually none is recovered for the product. The fact that expectations matter crucially in determining the impact of price changes suggests that expectations-setting will pay a big role in firm strategies.

Finally, Figure 11 investigates the magnitudes of the dynamic responses by examining the time path of digital camcorder sales under three different assumptions: the time path generated by the estimated model (also the actual time path of sales), the time path that would occur if consumers assumed that their logit inclusive values for digital camcorders remained equal to its present value in all future periods, and the time path that would occur if firms were faced with all consumers having no digital camcorders in each period, instead of high valuation consumers having purchased the product and hence generally having a higher reservation utility for buying, as occurs in our model.

We find that dynamics explain a very important part of the sales path. In particular, if consumers did not assume that prices and qualities changed, then sales would be somewhat declining over time, instead of growing rapidly over the sample period. At the beginning of our sample period, sales would

be huge compared to actual sales, as many consumers would have perceived only a limited option value from waiting. By the end of our sample period, sales would be significantly less than current sales, as many consumers who were likely to buy digital camcorders would have bought them early on, having assumed that quality, in the sense of the logit inclusive value, would be stable over time.

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

4.4 Cost-of-living indices

We develop a COLI for our model and compare it to widely-used COLIs. All indices are calculated with seasonally adjusted data. In performing this exercise, we hope to inform the discussion of how to improve current BLS methods. We do not mean to propose our model as a method that the BLS should consider for constructing indices as it would probably be infeasible given the time constraints under which the BLS operates.²⁴

The canonical price index I_t used by the BLS is a Laspeyres index that specifies

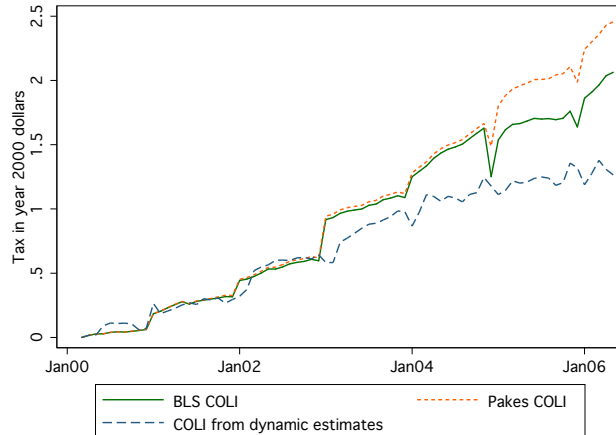
$$\frac{I_{t+1}}{I_t} = \sum_{j=1}^{J_t} s_{jt} \frac{p_{j,t+1}}{p_{jt}}. \quad (12)$$

We compute a BLS-style price index from (12) using the prices and market shares in our data, linked over time by model names. As is standard, we normalize the index to 100 for March 2000.²⁵ The main difficulty in constructing the BLS index is in determining $p_{j,t+1}$ for products that drop out of the market. A common approach used by the BLS, which we follow in our computation, is to use the average price

²⁴We focus on indices used by Pakes (2003) and the BLS but there have been other proposals for indices in dynamic settings. Reis (2005) develops a COLI from a model with durable goods. Contrary to our approach, he assumes that there are perfect resale markets, that consumers make a continuous purchase choice and implicitly, he considers established markets where diffusion is not taking place. His focus is on uncertainty in prices. He provide excellent citations on dynamics in price indices. Housing is an important area where durability has been a concern. See, for instance ,Benkard & Bajari (2005).

²⁵The BLS must deal with a number of challenging issues associated with the way enumerators collect data that we do not address here. See Pakes (2003) or more generally, Bureau of Labor Statistics (2007), Chapter 17.

Figure 12: Average monthly value from camcorder market



in the market. This introduces the well known “new goods” problem since the exiting product probably would have been priced below the market average. Pakes (2003) proposes using a prediction of $p_{j,t+1}$ from a hedonic regression, which addresses this problem. We also construct the Pakes (2003) price index, which as expected, falls further. We find that the BLS price index falls from 100 to 17.3 over our sample period and that the Pakes index falls from 100 to 13.9.

Formally, both the CPI and the Pakes index are price indices, not COLIs. However, they are both motivated by their relationship to the COLI and in practice, are used as such.²⁶ In general, one would construct COLIs by multiplying the price indices from (12) by the the expenditures in the sector. This is problematic for the camcorder sector since sales are rapidly growing and prices rapidly falling over time. Thus, we proceed by including the outside good as a product with an invariant price in (12) and by multiplying the resulting index by the share-weighted average price in the initial period (\$961) and dividing by 100. This provides the revenue savings realized in subsequent periods, which we then subtract from the March 2000 figure (of \$961) to obtain the extra value generated by the industry in any month. We plot these two indices in Figure 12. The indices starts at \$0 by construction and end six years later at \$2.06 for the BLS COLI and \$2.46 for the Pakes COLI. That is, from the BLS COLI, a tax of \$2.06 per household in May 2006 would result in an average utility equal to the Mar. 2006 average utility, with smaller taxes necessary for earlier months. The relatively small values reflect the fact that market shares for camcorders are low.

²⁶Bureau of Labor Statistics (2007) states that “the concept of COLI provides the CPI [Consumer Price Index]’s measurement objective (p. 2).”

The BLS and Pakes COLIs are designed to provide the income change necessary to buy a camcorder of equal quality in any period. However, this may deviate from the income change necessary to hold utility constant as willingness-to-pay changes due to evolving consumer holdings and expectations of the future. We use our structural model to evaluate the price changes that would hold utility constant over time.

We construct the COLI from our dynamic model as follows: we imagine a social planner who sets a sequence of aggregate-state-contingent taxes (or subsidies) that holds the average flow utility constant over time assuming that consumers follow optimizing behavior. This sequence of taxes forms a compensating variation measure because it results in the average expected value function being constant over time. This approach avoids a number of difficulties that might make a COLI for forward-looking consumers intuitively unappealing.²⁷ Note that the aggregate-state-contingent taxes do not change camcorder purchase behavior in our model. We make one final adjustment which is to assume that a consumer who buys a product that costs p_{jt} in period t pays a perpetual amortized price of $(1 - \beta)p_{jt}$ forever after, instead of paying p_{jt} at time t .²⁸ Note that a consumer pays the amortized price even after replacing the good. To eliminate this property, one could adjust β by the hazard of replacing the good.

We plot our dynamic COLI in Figure 12. In order to avoid dealing with differing marginal utilities of money based on different tax and money good quantities, our dynamic COLI is constructed from the specification with linear price, Table 3 column 3.²⁹ The dynamic index start at \$0 by construction and ends six years later at \$1.27.

The dynamic, BLS and Pakes COLI lines are very close for the first two years and then diverge substantially over the remaining four years. The dynamic COLI shows a clear concavity whereas the BLS COLI continues approximately linearly over the whole sample. Thus, we find the “new buyer problem” (Aizcorbe, 2005) to be empirically important. Sales and prices are moving linearly which causes standard COLIs to move linearly as well. However, relatively low value people are purchasing at the end of the sample and so overall, surplus is tapering off. Note that a BLS COLI that started in a later time period would have a lower slope as the average price would be lower than \$962, illustrating

²⁷Potential problems are 1) current price declines might benefit every consumer, even those who will not buy for several periods; 2) surprising price drops might affect welfare changes much more than anticipated ones; and 3) future income adjustments based on a COLI affect welfare today. See Reis (2005) and Bajari, Benkard & Krainer (2005) for different approaches.

²⁸If we measured flow utility using the entire price rather than the amortization scheme, we would find that average flow utility was less than the outside good utility throughout our sample since payments from new purchasers swamp flow from those who hold the product. Although theoretically consistent, we found this intuitively unappealing.

²⁹We also computed a COLI using the static BLP estimates. It was much larger than the other indices, and peaked at \$6.92. It did not appear reasonable.

how difficult it is to use a price index as a COLI for camcorders. Although the slope would be different, the shape would remain the same – and different from our dynamic COLI.

5 Conclusion

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

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

We find substantial heterogeneity in the overall utility from digital camcorders. Our results also show that much of the reason why the initial market share for digital camcorders was not higher was because consumers were rationally expecting that the market would later yield cheaper and better players. We find that industry elasticity of demand is 2.55 for transitory price shocks and 1.23 for permanent price shocks, with significantly larger permanent elasticities for individual products. Last, we find that the digital camcorder industry is worth an average of \$1.27 more per household per month in 2006 than in 2000 and that standard COLIs would overstate the gain in welfare due to the “new buyer problem.”

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

Appendix

Proof of Proposition 1

Proof Our approach is to prove the proposition for the case of finite horizons and then take appropriate

limits to address the case of infinite horizons. To ease notation, we omit the constant γ which enters utility every period. Consider first a model where the market ends at period $T > t$ and define $EV_i^T(\delta_{i0\tau}^f, \Omega_t)$ to be the value function when the market ends at period T . We will prove the proposition by induction.

First, the base case. In period T , we can write equation 4 as:

$$EV_i^T(\delta_{i0\tau}^f, \Omega_T) = \ln\left(\exp(\delta_i(\Omega_T)) + (1 - \beta)^{-1} \exp(\delta_{i0T}^f)\right).$$

Since Ω_T only enters EV_i^T through δ_i , by the second assumption of the proposition, we can write $P\left[EV_i^T(\delta_{i0t}^f, \Omega_T)|\Omega_{T-1}\right] = P\left[EV_i^T(\delta_{i0t}^f, \Omega_T)|\Omega'_{T-1}\right]$.

Now the inductive step. For some τ such that $t \leq \tau < T$ assume that $P\left[EV_i^T(\delta_{i0t}^f, \Omega_\tau)|\Omega_{\tau-1}\right] = P\left[EV_i^T(\delta_{i0t}^f, \Omega_\tau)|\Omega'_{\tau-1}\right]$. We would like to show that $P\left[EV_i^T(\delta_{i0t}^f, \Omega_{\tau-1})|\Omega_{\tau-2}\right] = P\left[EV_i^T(\delta_{i0t}^f, \Omega_{\tau-1})|\Omega'_{\tau-2}\right]$. We find

$$\begin{aligned} & P\left[EV_i^T(\delta_{i0t}^f, \Omega_{\tau-1})|\Omega_{\tau-2}\right] \\ &= P\left[\ln\left(\exp(\delta_i(\Omega_{\tau-1})) + \exp(E[EV_i^T(\delta_{i0t}^f, \Omega_\tau)|\Omega_{\tau-1}])\right)\middle|\Omega_{\tau-2}\right]. \end{aligned}$$

The first part has the same density under $\Omega'_{\tau-2}$ by the assumptions of the proposition while the second part has the same density by the inductive assumption. Thus, any function of them has the same density and we have proved the inductive step.

This proves the finite horizon case. The infinite horizon case holds because

$$EV_i(\delta_{i0t}^f, \Omega_t) = \lim_{T \rightarrow \infty} EV_i^T(\delta_{i0t}^f, \Omega_t) = \lim_{T \rightarrow \infty} EV_i^T(\delta_{i0t}^f, \Omega'_t) = EV_i(\delta_{i0t}^f, \Omega'_t).$$

The limit exists and hence the equality is true because of discounting and the fact that characteristics are bounded. ■

Proof of Proposition 2

Proof Let K denote the (assumed finite) number of potential values of Ω and k index a particular value so that $\delta_{ij\tau,k}^f$ denotes the realization of $\delta_{ij\tau}^f$ for the k th value of Ω_τ . We must define x_{jt} and p_{jt} in each of the potential states k to generate the appropriate δ_{it} . There will generally be many realizations of x_{jt} and p_{jt} that could generate any given δ_{it} . We (arbitrarily) choose the following: let $p_{jt} = 0$ always, let product 1 at any time period τ have some contingent flow utility $\delta_{i1\tau,k}^f$ and let other products have a utility flow of $-\infty$. For the rest of the proof, we discuss flow utility in terms of $\delta_{ij\tau k}$ rather than $x_{j\tau k}\alpha_i$ as there is a straightforward mapping between flow utilities and characteristics given preferences.

We let $\vec{\delta}_i$ and $\vec{\delta}_{i1}^f$ denote the vector of $\delta_{i\tau,k}$ logit inclusive values and $\delta_{i1\tau,k}^f$ flow utilities respectively. Now, define a function $f : \mathbb{R}^K \times \mathbb{R}^\infty \rightarrow \mathbb{R}^K \times \mathbb{R}^\infty$ that maps from potential states and time periods to the same space, conditional on the (specified) vector of logit inclusive values. We condition f on $\vec{\delta}_i$ which remains constant at the specified values, and so write $f(\cdot | \vec{\delta}_i)$. We define $f_{\tau,k}(\vec{\delta}_{i1}^f | \vec{\delta}_i)$, the value of f for a particular element, as the value of $\delta_{i1\tau,k}^f$ that makes the logit inclusive value for the τ, k th case equal to $\delta_{i\tau,k}$ holding constant the other flow utilities as $\vec{\delta}_{i1}^f$. For f to be a valid function we first show that this value is unique. We then show that f has a fixed point. By construction, the flow utilities of a fixed point of f generate $\vec{\delta}_i$ as the logit inclusive values.

To show uniqueness, for the τ, k th term consider the scalar-valued function $g(x | \vec{\delta}_{i1}^f)$, defined to be the value of $\delta_{i\tau,k}$ that results from the flow utilities of $\vec{\delta}_{i1}^f$ for every element but the τ, k th one and x for the τ, k th one. Note that g is continuous in x : for a sufficiently low $\delta_{i1\tau,k}^f$ it is unboundedly low; for a sufficiently high $\delta_{i1\tau,k}^f$ it is unboundedly high; and it is monotonically increasing in its argument. Thus there is a unique x such that $g(x | \vec{\delta}_{i1}^f) = \delta_{i\tau,k}$. This unique value, $g^{-1}(\delta_{i\tau,k} | \vec{\delta}_{i1}^f)$ defines $f_{\tau,k}$.

Now, as f is infinite dimensional, we would like to apply Schauder's fixed point theorem. We must show that f is continuous and that it lies in a convex, compact set. The function f is continuous as g^{-1} is continuous in the argument $\vec{\delta}_{i1}^f$. To show convexity and compactness, let δ_i^{min} and δ_i^{max} denote bounds for the minimum and maximum the elements of $\vec{\delta}_i$ respectively. Then, no element of $\vec{\delta}_{i1}^f$ will be larger than $\delta_i^{max}(1 - \beta)$, since purchasing a product with a flow utility of $\delta_i^{max}(1 - \beta)$ and never purchasing another product will already give mean expected utility δ_i^{max} and the actual decision allows for this option without imposing it. Thus, the elements of $\vec{\delta}_{i1}^f$ are bounded above. Moreover, if the domain is bounded above by $\delta_i^{max}(1 - \beta)$ then the range is bounded below by $\delta_i^{min} - \beta(1 - \beta)\delta_i^{max}$, since the worst possible $\delta_{i1\tau,k}^f$ occurs if the current δ_{it} is δ_i^{min} and the next period $\delta_{i,t+1}$ is δ_i^{max} with certainty and yields this value. Thus, $\vec{\delta}_{i1}^f \in [\delta_i^{min} - \beta(1 - \beta)\delta_i^{max}, \delta_i^{max}(1 - \beta)]^\infty$, which is bounded and closed in \mathbb{R}^∞ and hence a compact set by Tychonov's theorem. By Schauder's fixed point theorem, f has a fixed point. ■

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. (2005). Unobservable product differentiation in discrete choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics*, 36, 771–788.
- Aizcorbe, A. (2005). Price deflators for high technology goods and the new buyer problem. Unpublished manuscript, Board of Governors of the Federal Reserve Bank.
- Anderson, S., De Palma, A., & Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press.
- 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.
- Bajari, P., Benkard, C. L., & Krainer, J. (2005). House prices and consumer welfare. *Journal of Urban Economics*, 58, 474–487.
- Benkard, C. L. & Bajari, P. (2005). Hedonic indexes with unobserved product characteristics, and application to personal computers. *Journal of Business and Economic Statistics*, 23, 61–75.
- 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. (1999). Voluntary export restraints on automobiles: Evaluating a strategic trade policy. *American Economic Review*, 89, 400–430.
- Berry, S., Levinsohn, J., & Pakes, A. (2004). Estimating differentiated product demand systems from a combination of micro and macro data: The market for new vehicles. *Journal of Political Economy*, 112, 68–105.
- Berry, S. & Pakes, A. (2005). The pure characteristics model of demand. *International Economic Review*, in press.
- Bresnahan, T. F. (1981). Departures from marginal cost pricing in the american automobile industry: Estimates for 1997-1978. *Journal of Industrial Economics*, 17, 201–277.
- Bureau of Labor Statistics (2007). BLS handbook of methods.

- Carranza, J. (2006). Demand for durable goods and the dynamics of quality. Unpublished manuscript, University of Wisconsin.
- Erdem, T. & Keane, M. (1996). Decision making under uncertainty: Capturing dynamic brand choice in turbulent consumer goods markets. *Marketing Science*, 15, 1–20.
- Erdem, T., Keane, M., Oncu, S., & Strebel, J. (2005). Learning about computers: An analysis of information search and technology choice. *Quantitative Marketing and Economics*, 3, 207–246.
- Esteban, S. & Shum, M. (2007). Durable goods oligopoly with secondary markets: The case of automobiles. *RAND Journal of Economics*, in press.
- Gandal, N., Kende, M., & Rob, R. (2000). The dynamics of technological adoption in hardware/software systems: The case of compact disc players. *RAND Journal of Economics*, 31, 43–61.
- Gentle, J. E. (2003). *Random number generation and Monte Carlo methods* (2 ed.). New York: Springer.
- Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the U.S. automobile industry. *Econometrica*, 63, 891–951.
- Gordon, B. (2006). Estimating a dynamic model of demand for durable goods. Unpublished manuscript, Carnegie Mellon University.
- Hendel, I. & Nevo, A. (2006). Measuring the implications of sales and consumer stockpiling behavior. *Econometrica*, 74, 1637–1673.
- Ho, C.-Y. (2008). Switching cost and the deposit demand in china. Unpublished Manuscript, Georgia Institute of Technology.
- Jaumandreu, J. & Moral, M. J. (2006). Identifying behaviour in a multiproduct oligopoly: Incumbents' reaction to tariffs dismantling. Unpublished Manuscript, University of Carlos III of Madrid.
- Keane, M. & Wolpin, K. (1997). The career decisions of young men. *Journal of Political Economy*, 105, 473–522.
- Krusell, P. & Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106, 867–896.
- Magnac, T. & Thesmar, D. (2002). Identifying dynamic discrete decision processes. *Econometrica*, 70, 801–816.
- Melnikov, O. (2001). Demand for differentiated products: The case of the U.S. computer printer market. Unpublished manuscript, Cornell University.
- Nair, H. (2007). Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video games. *Quantitative Marketing and Economics*, in press.

- Nevo, A. (2000a). Mergers with differentiated products: The case of the ready-to-eat breakfast cereal industry. *RAND Journal of Economics*, 31, 395–421.
- Nevo, A. (2000b). 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.
- Pakes, A. (2003). A reconsideration of hedonic price indices with applications to PC’s. *American Economic Review*, 93, 1578–1596.
- Park, M. (2008). Estimation of dynamic demand with heterogeneous consumers under network effects. *Korean Journal of Industrial Organization*, 16, 1–38.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110, 705–729.
- Prince, J. (2007). Repeat purchase amid rapid quality improvement: Structural estimation of demand for personal computers. Unpublished manuscript, Cornell University.
- Reis, R. (2005). A cost-of-living dynamic price index, with an application to indexing retirement accounts. NBER Working Paper 11746.
- Rivers, D. & Vuong, Q. (2002). Model selection tests for nonlinear dynamic models. *Econometrics Journal*, 5, 1–39.
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, 55, 999–1033.
- Schiraldi, P. (2007). Automobile replacement: A dynamic structural approach. Unpublished Manuscript, Boston University.
- Shcherbakov, O. (2008). Measuring switching costs in the television industry. Unpublished Manuscript, Yale University.
- Shepler, N. (2001). Developing a hedonic regression model for camcorders in the U.S. CPI. Bureau of Labor Statistics.
- Song, I. & Chintagunta, P. (2003). A micromodel of new product adoption with heterogeneous and forward-looking consumers: An application to the digital camera category. *Quantitative Marketing and Economics*, 1, 371–407.

Zhao, Y. (2008). Why are prices falling fast? An empirical study of the US digital camera market. Unpublished Manuscript, Queens College, City University of New York.