

Dynamics of Consumer Demand for New Durable Goods*

Gautam Gowrisankaran
University of Arizona
HEC Montréal
NBER

Marc Rysman
Department of Economics
Boston University

February 14, 2012

Abstract

Most new consumer durable goods experience rapid declines in prices and improvements in quality, suggesting the importance of modeling dynamics. This paper estimates a dynamic model of consumer preferences for new durable goods with persistent heterogeneous consumer tastes, rational expectations and repeat purchases over time. We estimate the model on the digital camcorder industry using panel data on prices, sales and characteristics. We find that standard COLIs overstate welfare gain in later periods due to a changing composition of buyers. The one-year industry elasticity in response to a transitory industry-wide price shock is about 25% less than the one-month elasticity.

*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, Sam Kortum, John Krainer, Ariel Pakes, Minsoo Park, Rob Porter, Jeff Prince, Pasquale Schiraldi, Andy Skrzypacz, Tim Erickson, Mo Xiao and seminar participants at several institutions for helpful comments; Haizhen Lin, Ryan Murphy, Kathleen Nosal, 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 anonymous referees substantially improved the paper. We acknowledge funding from the National Science Foundation. All errors are our own.

1 Introduction

If you don't need a new set, don't rush to buy one. Prices will no doubt continue to drop over time, [and] you'll have more sets to choose from.¹

-ConsumerReports.org on 3D HDTVs

In many durable goods settings, the choice of when to buy is as important to consumers as what to buy. Particularly for consumer electronics, the quote from Consumer Reports conveys the conventional wisdom: prices will fall and new choices, often higher quality ones, will arrive. As a result, many consumers purposely delay their purchase of these goods. When they do purchase, consumers often have in mind that they will replace the product with a superior product in the foreseeable future. Thus, dynamic behavior is an important element of demand for consumer durable goods markets. This paper specifies a structural dynamic model of consumer preferences for new durable goods; estimates the model using aggregate data on digital camcorders; and uses the model to evaluate elasticities and cost-of-living indices for this market.

The digital camcorder industry is an important sector, with about 11 million units sold in the U.S. from 2000 to 2006. Between 2000 and 2006, the sector also experienced a huge evolution. Average digital camcorder prices dropped from \$930 to \$380 while average pixel counts rose from 580,000 to 1.08 million. The number of products available grew from less than 30 to almost 100. Annual sales grew by 2.6 times in 4 years. The rapidly evolving nature of the characteristics and sales – together with the fact that consumers are advised to consider the dynamic implications of their decisions – suggests that modeling dynamics is empirically very important for estimating consumer preferences in this industry. These issues have broad applicability: rapidly falling prices and improving features have been among the most visible phenomena in a large number of other new consumer durable goods markets, including computers, DVD players and HDTVs among many others.

In our model, dynamically optimizing consumers may choose among the set of available camcorders or wait. Camcorders are durable, so purchase provides flow utility into the future. The available prices, quality and variety improve over time, so waiting is valuable. In addition, while consumers can hold only one camcorder at a time, they may substitute a new camcorder for an old one, so consumers continue to evaluate the market even after purchase. Our model allows for product differentiation, endogeneity of prices, many products, changing numbers of products, persistent consumer heterogeneity, and endogenous repeat purchases over time. As our model is dynamic, we need to specify consumer perceptions over

¹<http://www.consumerreports.org/cro/electronics-computers/tvs-services/tvs/tv-buying-advice/tv-3d/tv-3d.htm>, accessed June 27, 2011.

future states of the world. We focus on a major simplifying assumption: that consumers expect 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.

The dynamics of our model build on the traditional vintage capital models (see Solow, Tobin, von Weizsacker & Yaari, 1966) in that consumers in our model hold one product at a time and endogenously reallocate to new products as the technology for capital (i.e., camcorders) improves. Our framework differs from this literature in that we model a sunk cost of acquiring new technology – namely the purchase price – which makes the consumer purchase decision dynamic. In this way, it is similar to the Rust (1987) model of bus engines, in which the agent must decide when to replace the engine.

While the Rust model concerns an industry with one homogeneous product, the camcorder industry has hundreds of different products with different prices and characteristics. To understand purchase decisions in this industry (and other consumer durable goods industries) we need to model both the “when” to buy of the dynamic literature and the “what” to buy. A different literature started by Bresnahan (1981) and Berry, Levinsohn & Pakes (1995) (henceforth BLP), has modeled static consumer decisions for differentiated products systems with many heterogeneous products. This literature has shown that incorporating consumer heterogeneity into differentiated product demand systems is important in obtaining realistic predictions. Our paper nests a BLP-style demand system within the dynamic replacement framework. By allowing for persistent consumer heterogeneity, we relax the assumption that choices are conditional independent given the observed state that is typically required for dynamic estimation, but at the cost of computational complexity. We develop a new estimation procedure that draws on the techniques of BLP for modeling consumer heterogeneity in a discrete choice model and on Rust (1987) for modeling optimal stopping decisions. Our primary methodological advance is in developing a feasible specification that allows us to combine these two separate methods.

Over the last fifteen 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. To the extent that dynamics are important for many industries, our paper may be useful in deriving better estimates for these and related questions. Indeed, recent work is using and extending our methods to examine the importance of software in the video game industry (Lee, 2010), scrapping subsidies for automobiles (such as cash-for-clunkers programs; see

Schiraldi, 2011), markups for digital cameras (Zhao, 2008), switching costs between cable and satellite television (Shcherbakov, 2009) and switching costs in consumer banking (Ho, 2011), among other research questions.

Before moving forward, we briefly consider the complex question of how estimates from dynamic and static models might differ. There are two principal causes of dynamics: durability and forward-looking behavior. The implications of durability are straightforward: the sales increase from price declines will be moderated in the durable goods model since demand endogenously falls as high-value consumers accumulate the good. Thus, the static estimation will understate the importance of price and quality, since it appears that the sales response to improvement is relatively small. The implications of accounting for forward-looking behavior are more complex. Static models predict that sales increase when prices are low, whereas dynamic models predict that sales increase when the expected change in price is low.² A price decline may not lead to a sales increase in a dynamic model if consumers expect a large price decline in the next period. It is difficult to sign the bias from estimating a static model related to this issue since it depends on the relationship of low prices, price declines and sales in the data. Finally, note that we observe both dynamic and cross-sectional variation. It is possible for the dynamic model to generate substitution patterns within time periods and across time periods that appear inconsistent from the perspective of a static model.³ Overall, it is difficult to predict what a static model will estimate when the data-generating process is dynamic, because it depends on the interaction of these many factors.

We use our results to analyze the difference between short-run and long-run price elasticities. We focus on the case in which price increases for one period and then returns to its previous level, and consumers anticipate this. When considering the industry as a whole, the long-run elasticity is substantially smaller than the short-run elasticity, since many consumers delay purchase during the price spike. However, the short- and long-run product elasticities are very close to each other. That is because we find that camcorders are relatively close substitutes, so when the price of a single product increases, consumers switch to another product rather than delaying purchase. We consider several other cases, such as

²More formally, dynamic models predict that sales increase when the user cost of capital falls, and the implicit user cost of capital is (roughly) the change in price. This point is explicit in Gandal, Kende & Rob (2000), who consider a similar but simpler dynamic model with one product in each time period. Their model generates the result that sales are a linear, decreasing function of the forward difference in prices after discounting future price, $(p_t - \beta p_{t+1})$.

³Consider the case in which prices fall and then level off. In a dynamic model, a consumer may prefer the high quality product (particularly if the consumer recognizes that purchasing the high quality good delays future purchases) but wait to purchase until the price stops dropping. Thus, the data will show that consumers bought when prices were low, but that within a period, they purchased the high price, high quality product. The static model would interpret the first sort of variation to imply that price sensitivity is high but the second form of variation to imply that price sensitivity is relatively low.

permanent price changes, and different consumer expectations.

We also 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. The Consumer Price Index and other government-computed indices are often used as COLIs, and have important implications for wage growth at many firms, government transfer programs and a variety of other government policies. The Bureau of Labor Statistics is particularly concerned about the development of accurate COLIs for consumer electronics and camcorders in particular (see Shepler, 2001). If high-value consumers purchase early and low-value consumers purchase late, standard approaches that assume demand is homogenous over time overstate the welfare gains later on (see Aizcorbe, 2005). We show that this effect, which Aizcorbe terms the “new buyer” problem, is empirically important. Standard COLI measures show welfare substantially improving in this market since prices fall and quantities rise. However, our dynamic model shows that welfare gains are attenuated, since later buyers are either low value relative to early buyers, or they already hold the good.

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 to unobserved product characteristics. 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.

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. 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 & Strebels, 2005; Prince, 2008; Keane & Wolpin, 1997). Extending these types of models to aggregate data is important for two reasons. First, in many cases, aggregate data are all that are 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.⁴ Accurate market shares are important for estimating the supply side. For instance, BLP and Goldberg (1995) use aggregate market share data to estimate pricing first-order conditions.

Finally, a number of recent papers (Gandal et al., 2000; Esteban & Shum, 2007; Melnikov, 2001; Song & Chintagunta, 2003; Gordon, 2009; Nair, 2007; Carranza, 2007; 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 aggregate data. We use a similar reduced-form approximation of the supply side as he proposed. 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. An 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 product sizes. A final comparison is to Goettler & Gordon (2012) who estimate supply and demand in the microprocessor industry. Given their focus on endogenous product quality (i.e., innovation), they solve a full model of the interaction between dynamically optimizing consumers and the two firms in the industry, whereas we use a reduced-form approximation of the supply side. Their approach is computationally costly though, and would be difficult to apply to industries with many firms or products, random coefficients, or unobserved product characteristics. In contrast, we address a context with a very large and changing number of products, random coefficients, and endogenous prices.

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 identification of the parameters.

⁴Our 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.

2.1 Model

This subsection specifies the purchase decisions of one consumer; the next discusses aggregating across heterogeneous consumers. The industry starts at time $t = 0$. The consumer has an infinite horizon and discounts the future at rate β . The consumer can benefit from at most one camcorder in a period: camcorder usage technology is Leontief. We further assume that there are no resale markets.⁵ Thus, if the consumer purchases a new camcorder to replace an old one, she costlessly discards the old camcorder and obtains flow utility from the new camcorder.

The consumer starts time 0 holding the outside good, which gives mean flow utility 0. At some point in time, she may purchase a camcorder. From this point on she will use the purchased camcorder (while keeping an eye on new products and prices). Eventually, she might upgrade her camcorder at which point she scraps the one she had. She continually repeats the process of using her latest model purchased while looking to potentially upgrade.

To formalize payoffs, at each time period t , there is a set of products $j = 1, \dots, J_t$. Each product has a net flow utility f_{jt} and a disutility from price P_{jt} (resulting from less consumption of the money good). The consumer chooses one of the available products or chooses to purchase no product. If she buys product j at time t , then at time t she receives utility of

$$u_{jt} = f_{jt} - P_{jt} + \varepsilon_{jt}, \text{ for } j = 1, \dots, J_t, \quad (1)$$

where ε_{jt} is an idiosyncratic type 1 extreme value term, distributed *i.i.d.* across products and time periods, and is meant to capture random variations in the purchase experience that do not persist across month, due to sales personnel, weather, etc.

A consumer who does not purchase any product at time t receives $u_{0t} = f_{0t} + \varepsilon_{0t}$, the mean flow utility from the good she currently owns plus another idiosyncratic type 1 extreme value term. The current flow utility f_{0t} is determined by past purchases. Thus, if time \hat{t} was her first purchase occasion and \hat{j} was her first purchase, then for $t < \hat{t}$, $f_{0t} = 0$; and for $t > \hat{t}$, $f_{0t} = f_{\hat{j}\hat{t}}$ until after she upgrades to a new product. The consumer does not need to remember the identity or characteristics of the product she purchases, only its flow utility.

Consider now the consumer dynamic optimization decision. At time t , the consumer is faced with $J_t + 1$ choices, and chooses the option that maximizes the sum of the expected discounted value of future utilities conditional on her information at time t . We assume that the consumer knows all time t information when making her time t decision but that she has no information about the future values

⁵We believe that resale markets are small for the new consumer durable goods that we examine given the speed of technological progress and price reduction.

of her ε shocks beyond their distribution. Furthermore, the set of products and their prices vary across time due to entry and exit and changes in prices for existing products. The consumer has expectations about the evolution of future products, but in most specifications lacks perfect knowledge about them.

We now define the state variables and use them to exposit the Bellman equation. Let $\vec{\varepsilon}_t \equiv (\varepsilon_{0t}, \dots, \varepsilon_{J_t t})$, where we use the right arrow to denote a vector. Then, the purchase decision for consumer i depends on $\vec{\varepsilon}_t$, endowment f_{0t} , attributes of currently available products and expectations about future product attributes. Future product attributes will depend on firm behavior which is a function of consumer endowments and supply-side factors such as technological progress. Let Ω_t denote the current industry state; Ω_t includes the number of products J_t , the price disutility and mean flow utility for each product, and any other factors that influence future product attributes. We assume that Ω_t evolves according to some Markov process $g(\Omega_{t+1}|\Omega_t)$ that accounts for firm optimizing behavior. Thus, the state vector at time t is $(\vec{\varepsilon}_t, f_{0t}, \Omega_t)$.

Let g denote densities in general; let the prime symbol (e.g., Ω') denote next period's value of a variable (which allows us to drop the t subscript); let $V(\vec{\varepsilon}, f_0, \Omega)$ denote the value function; and let $EV(f_0, \Omega) = \int_{\vec{\varepsilon}} V(\vec{\varepsilon}, f_0, \Omega) dg_{\vec{\varepsilon}}$ denote the expectation of the value function, integrated over $\vec{\varepsilon}$. Then the Bellman equation is

$$V(\vec{\varepsilon}, f_0, \Omega) = \max \left\{ f_0 + \beta E [EV(f_0, \Omega') | f_0, \Omega] + \varepsilon_0, \max_{j=1, \dots, J} \{ f_j - P_j + \beta E [EV(f_j, \Omega') | f_j, \Omega] + \varepsilon_j \} \right\} \quad (2)$$

where “ E ” denotes the expectation operator, a conditional expectation in this case, and P_j , and f_j , $j = 1, \dots, J$ and J are functions of Ω . From (2), the first element of the max operator indicates that the consumer keeps the camcorder she already has (or none), which means we pass f_0 into the next period's value function. In the second argument of the max operator, we pass f_j – the flow utility of her contemporaneous choice – into the next period's value function.

To estimate our model, we will ultimately need to solve (2) many times. But Ω has a very large dimension resulting in a curse of dimensionality. We proceed by using the aggregation properties of the extreme value distribution to express the Bellman in a relatively simple form, and then make assumptions on the perceptions of industry evolution based on this form. Specifically, we define the *logit inclusive value* to be

$$\delta(\Omega) = \ln \left(\sum_{j=1, \dots, J} \exp(f_j - P_j + \beta E [EV(f_j, \Omega') | f_j, \Omega]) \right). \quad (3)$$

The logit inclusive value is the ex-ante present discounted lifetime value of buying the preferred cam-

order, as opposed to holding the outside option.⁶ Using this definition, the expectation Bellman can be written as a simpler problem where the consumer makes a one-time purchase of a product with mean utility $\delta(\Omega)$ or continues to hold her existing product. Applying the logit aggregation once more to the binary choice between all camcorders and holding the existing product, we obtain

$$EV(f_0, \Omega) = \ln[\exp(f_0 + \beta E[EV(f_0, \Omega') | f_0, \Omega]) + \exp(\delta(\Omega))].^7 \quad (4)$$

Equation (4) shows that Ω only affects EV through its impact on the current $\delta(\Omega)$ and predictions of future values of $\delta(\Omega)$. A simplifying assumption on how consumers form these predictions then can greatly reduce the complexity of our problem. We assume that consumers predict future values of δ based only on the current δ , rather than on the full information set Ω .

Assumption 1 *Inclusive Value Sufficiency (IVS)* If $\delta(\Omega) = \delta(\tilde{\Omega})$, then $g(\delta(\Omega') | \Omega) = g(\delta(\tilde{\Omega}') | \tilde{\Omega}')$ for all $\Omega, \tilde{\Omega}$.

The IVS assumption and (4) imply that all states with the same $\delta(\Omega)$ have the same expected value; see Proposition 1 in the Appendix for details. Thus, it is sufficient for the consumer to track only two scalar variables, f_0 and δ in order to form their dynamic decisions. This assumption is restrictive: consider that δ could be high in a period 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, the consumer will assume these states share the same future. While our assumption of IVS can be interpreted as a literal assumption on how the industry evolves, it is perhaps more attractive to think of it as an assumption on how boundedly rational consumers perceive this market. The IVS assumption is valuable, since we can now replace Ω with δ in the state space, rewriting (4) as

$$EV(f_0, \delta) = \ln[\exp(f_0 + \beta E[EV(f_0, \delta') | \delta]) + \exp(\delta)], \quad (5)$$

and (3) as

$$\delta = \ln \left(\sum_{j=1, \dots, J} \exp(f_j - P_j + \beta E[EV(f_j, \delta') | f_j, \delta]) \right), \quad (6)$$

providing a tractable two-dimensional state space.

⁶Formally, the ex-ante distribution of the preferred camcorder is now distributed $\delta(\Omega)$ plus a new type 1 extreme value term. Anderson, De Palma & Thisse (1992); Rust (1987) provide proofs of this statement for static and dynamic models respectively.

⁷We omit Euler's constant from this equation as it does not affect decisions since it is constant.

Within this context, we assume rational expectations, that the consumer is on average correct about the future. We choose two functional forms. The simplest is perfect foresight, where the consumer knows all future values of f_j and P_j . This functional form is straightforward: the industry state is t . Moreover, it is a special case of IVS provided that δ is different at every time period (as would occur if quality were improving, prices were non-decreasing and the set of products were non-decreasing), as in this case, there is a one-to-one mapping from t to δ .⁸ We believe that it is more realistic to assume that the consumer has only a limited ability to predict future product attributes. Thus, for most of our specifications, we let consumer perceptions about next period's δ , $g(\delta'|\delta)$, be its actual empirical density fitted to a simple linear autoregressive specification:

$$\delta_{t+1} = \gamma_1 + \gamma_2 \delta_t + \nu_{t+1}, \quad (7)$$

where ν_{t+1} is normally distributed with mean 0 and unobserved at time t and γ_1 and γ_2 are incidental parameters. This assumption will ensure that the consumer is correct, on average, about the improvement in industry value, as embodied in δ .⁹

An implication of (7) is that, for $0 < \gamma_2 < 1$, consumers expect the δ process to converge towards an asymptote $\gamma_1/(1 - \gamma_2)$. This asymptote can be interpreted as a long-run steady state of the model. The assumption of an eventual arrival of a steady state is how we capture an evolving industry with a stationary dynamic model.¹⁰ In our results for the camcorder industry, we find that the market will eventually arrive at the steady state, but that the arrival will occur well after the end of our data.

Similar assumptions have been used in the existing dynamic literature. Most other papers (Melnikov, 2001; Hendel & Nevo, 2006) specify their analog of δ as a function of the flow utilities of available products, whereas ours also incorporates the continuation values of holding those products. Since flow utilities are taken as exogenous and continuation values incorporate endogenous decision-making, the assumption of (7) may be less palatable in our implementation. Our approach differs because we allow the characteristics of the products that δ is defined over to affect the continuation value, which means we cannot separate the continuation values from δ . In this way, our model allows a consumer that purchases a high-quality product (high f_j) to hold it for longer, which affects consumer decision-making in anticipation.

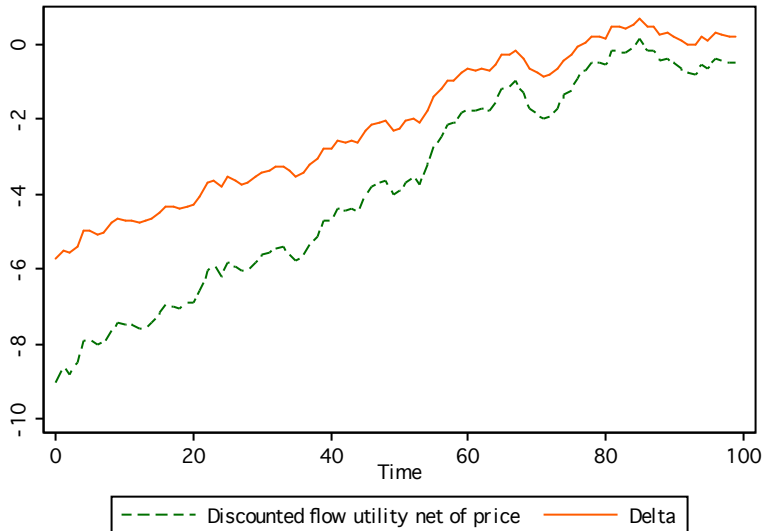
Importantly, with IVS and rational expectations, the optimal consumer decisions given an industry environment are defined by the joint solution to the expectation Bellman (5), the logit inclusive value

⁸In practice, it is more straightforward to compute this specification using t instead of δ as the industry state variable.

⁹Because of the larger dimensionality of flow utility and prices relative to δ , we can always rationalize the assumption in (7) with an assumption on fundamentals. See Proposition 2 in the Appendix for details.

¹⁰Goettler & Gordon (2012) provide an alternative approach to address a changing environment. They specify a model that is stationary relative to an endogenously advancing technology frontier.

Figure 1: Simulated evolution of flow utility and delta in a simple example

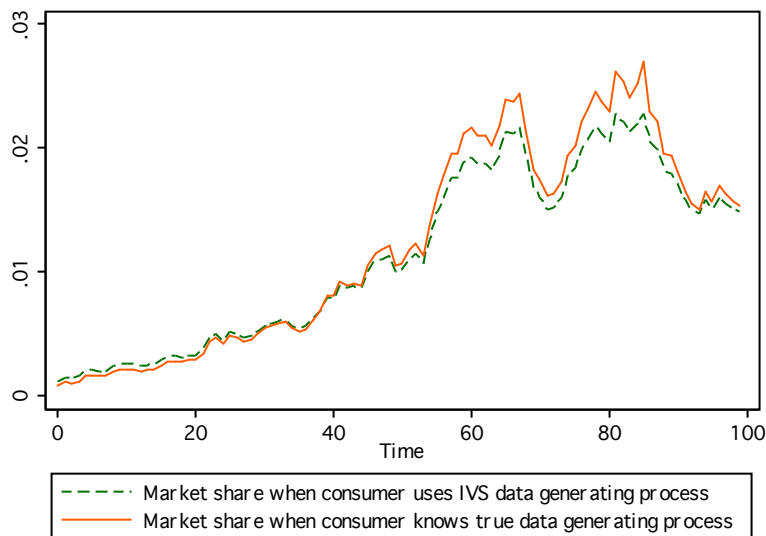


(6), and the industry evolution regression (7). While the following subsection discusses computation in detail, note for now that to compute optimal consumer decisions – for instance, to evaluate counterfactual firm policies – it is necessary to jointly solve these three equations and not just the Bellman equation. The reason for this is that a different Bellman equation (as would occur under a counterfactual policy environment) implies different values of δ which imply different γ coefficients which in turn imply a different Bellman equation.

To better understand the roles of both δ and IVS, we first evaluate δ in the context of a simple numerical example of our model, with one product every period and a constant price. We simulate flow utility to evolve according to an AR(1) process. We choose the AR(1) process and price to obtain aggregate shares that roughly match the camcorder industry.¹¹ If the consumer were forced to hold a product indefinitely once she bought it, the mean discounted flow utility net of price from purchasing the product would be $f/(1 - \beta) - P$. Figure 1 shows the evolution of this value and δ over 100 month-long periods. Note that δ is greater than the discounted flow utility net of price because δ incorporates the ability to upgrade to a new camcorder when features improve. The gap between them shrinks over time as the option value diminishes. Importantly, δ captures the path of quality increases (and, if they existed,

¹¹The AR(1) process is $f_{t+1} = .0005 + .99f_t + \nu_{t+1}$, where $\nu \sim N(0, .002)$; the price term is $P = 5$; the discount factor is $\beta = .99$; and first period flow utility is $f = -0.04$. We assume that consumers know price, the AR(1) parameters, and current f .

Figure 2: Simulated market share with approximate and true data generating process in simple example



price decreases) that occur in the data. Note also that discounted flow utility net of price will asymptote towards a long-run mean that is higher than its starting value, in this case to .05.

A further concern is how well the IVS assumption approximates δ and consumer decision-making. By assumption, consumers assume that δ evolves according to an AR(1) process. In the example, flow utility evolves according to an AR(1) process as in Melnikov (2001) and Hendel & Nevo (2006) but δ does not. To investigate how much bias might occur from this misspecification, Figure 2 shows the evolution of market share for a consumer who uses the true δ data generating process to make decisions, and for a consumer who optimizes assuming that δ follows an AR(1) process, jointly solving (5), (6), and (7). The two time paths of market shares are very close. We take from this the heuristic that if an industry is evolving rapidly (as in our example and the camcorder industry) then the exact specification of future expectations will not hugely influence purchase decisions so long as it captures the general process of evolution. Note that if we consider cases with larger numbers of products, the state space for the consumer tracking the true flow utilities grows larger whereas the state space for the consumer using our approximation stays the same, which makes our model more attractive computationally.

We believe that the assumptions that we make capture the first-order feature of the camcorder industry, that prices are dropping and quality is rising. In addition to the evidence in Figure 2, we implement the following empirical tests of the empirical validity of IVS:

1. We test the impact of adding additional state variables. This is similar to the macroeconomics general equilibrium literature on the impact of heterogeneity (see Krusell & Smith, 1998), which adds additional moments of the income distribution. Given the potential importance of the number of products in determining industry evolution, we use J as an additional predictor, so that both δ and J predict δ' and J' .¹²
2. We test the validity of the AR(1) specification in (7). We construct a moment based on the assumption of no autocorrelation $E[\nu_t \nu_{t+1}] = 0$. We do not impose this moment in estimation but rather test its validity at the estimated parameters.
3. Using the estimated parameters, we graph ν and δ over time for different simulated consumers. This allows us to examine graphically whether perceptions of their future values look systematically biased or not.

2.2 Inference

The previous subsection considered the decision of a single consumer, who is faced with a vintage capital problem in the spirit of Rust (1987). This subsection discusses aggregation across consumers, which follows BLP closely. We now assume that there is a continuum of consumers indexed by i . Consumers differ in their mean flow utility, disutility from price, idiosyncratic shocks, logit inclusive values, expectation Bellman, and expectations processes for the future. All of these terms should now be indexed by i , i.e., f_{ijt} , P_{ijt} , ε_{ijt} , $\delta_i(\Omega_t)$, EV_i and $(\gamma_{1i}, \gamma_{2i})$, respectively. In this subsection, we refer to $\delta_i(\Omega_t)$ as δ_{it} .

We assume that flow utility and price fit in the random coefficients framework developed by BLP. Specifically, we let $f_{ijt} = x_{jt} \alpha_i^x + \xi_{jt}$. Here, x_{jt} are observed characteristics of the camcorder (e.g., size, zoom, LCD screen size); ξ_{jt} is the unobserved (to the econometrician) characteristic and α_i^x are consumer i 's coefficients on observed characteristics. The ξ_{jt} plays the role of our econometric error term, and explains why market shares deviate from those predicted by observable elements of the model.¹³ We interpret it as unobserved quality. For a given model, observed characteristics are fixed over time in our data, but we do not constrain the unobserved characteristic to remain fixed over time. Note that a consumer who buys at time t obtains the same flow utility f_{ijt} every period thereafter, implying that she is receiving ξ_{jt} as the unobserved characteristic in every future period in which she uses the product.

¹²We specify two linear regressions for the state evolution that are similar to (7). The linear regressions have δ' and $\ln(J')$ as dependent variables and include both aggregate state variables as regressors.

¹³Note that ϵ is not an econometric error term since our assumption of a continuum of consumers implies that the data are assumed to reflect the integral over ϵ values.

Also, we let $P_{ijt} = \alpha_i^p \ln(p_{jt})$, where p_{jt} is price. The parameters α_i^x and α_i^p are time-invariant.

Let α_i denote consumer coefficients α_i^x and α_i^p . We assume that α_i has mean $\alpha \equiv (\alpha^x, \alpha^p)$ and variance matrix Σ . Our empirical implementation uses a diagonal Σ matrix, although correlated matrices fit within this framework. Define also the mean flow utility of product j in period t as

$$F_{jt} = x_{jt}\alpha^x + \xi_{jt}, j = 1, \dots, J_t. \quad (8)$$

Note that $f_{ijt} = F_{jt} + (\alpha_i^x - \alpha^x)x_{jt}$ where $(\alpha_i^x - \alpha^x) \sim N(0, \Sigma)$.

The structural parameters of our model are thus (α, Σ, β) . 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 $\beta = 0.99$ at the level of the month leaving only α and Σ to estimate.

Following BLP, we specify a GMM criterion function

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

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

$$(\hat{\alpha}, \hat{\Sigma}) = \arg \min_{\alpha, \Sigma} \{G(\alpha, \Sigma)' W G(\alpha, \Sigma)\}, \quad (9)$$

where W is a weighting matrix. Thus, to estimate (α, Σ) , we need to solve for $\vec{\xi}(\alpha, \Sigma)$.

In order to solve for $\vec{\xi}(\alpha, \Sigma)$, we first have to solve for market shares. As in Section 2.1, the consumer decision problem is defined by the fixed points to three equations: the expectation Bellman equation (5), the logit inclusive value (6), and the industry evolution regression (7). Based on these equations, we solve for market shares by starting at time 0 with the assumption that all consumers hold the outside good.¹⁴ Iteratively for subsequent time periods, we solve for consumer purchase probabilities given the distribution of flow utility of holdings using (10), and update consumer holdings at each period. Using (5), the probability that a consumer with holding f_{i0t} purchases good j is the aggregate probability of purchase (the first term) times the probability of purchasing product j conditional on purchase (the second term):

$$\hat{s}_j(f_{i0t}, \delta_{it}) = \frac{\exp(\delta_{it})}{\exp(EV_i(f_{i0t}, \delta_{it}))} \times \frac{\exp(f_{ijt} - P_{ijt} + \beta E[EV_i(f_{ijt}, \delta'_{it}) | f_{ijt}, \delta_{it}])}{\exp(\delta_{it})}. \quad (10)$$

¹⁴A 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 (2011).

We then integrate over consumers i to compute predicted market shares for each j and t . We perform the integration via simulation, as in BLP.

These equations define the market share conditional on the vector of mean flow utilities \vec{F} . To define the moment condition (as opposed to computing market shares), we need to recover $\vec{\xi}$. This requires solving a fourth equation, which finds \vec{F} that make shares for each product at each time period match those in the data. This equality is justified by the assumption of a continuum of consumers and is analogous to the contraction mapping defined by BLP:

$$F_{jt}^{new} = F_{jt}^{old} + \psi \cdot \left(\ln(\bar{s}_{jt}) - \ln \left(\hat{s}_{jt} \left(\vec{F}^{old}, \alpha^p, \Sigma \right) \right) \right), \forall j, t, \quad (11)$$

where $\hat{s}_{jt} \left(\vec{F}^{old}, \alpha^p, \Sigma \right)$ is the predicted market share of the model, \bar{s}_{jt} is observed market share in the data, and ψ is a tuning parameter used in the computation that we generally set to $1 - \beta$.

In practice, we solve these four equations via successive approximations, beginning with guesses of δ_{it} , $EV_i(f_{i0t}, \delta_{it})$, $(\gamma_{1i}, \gamma_{2i})$ and F_{jt} and circulating between the implicit equations until we find a fixed point in all four. We found computational benefits to taking only a small number of iterations of each equation before moving to another, although it would be possible to “nest” them, for instance, by solving (5), (6), and (7) to convergence before taking a step in (11).¹⁵

Using the \vec{F} that is the solution to these four equations, we compute $\vec{\xi}(\alpha, \Sigma)$ from (8), and then construct our objective function in (9). As we discuss in more detail in the appendix, we discretize the state space and use importance sampling to draw from the distribution of consumers. Also, as in Nevo (2000b), we can solve for the optimal α^x as a function of the other parameters using matrix algebra techniques. Thus, we perform non-linear search only over α^p and Σ . All computer code is available from the authors upon request.

An important issue is whether these implicit functions 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 always obtained convergence to the same solution. 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 with future products. As an example of a complementary between a product in the current period and two periods forward, if we exogenously lower price of a current period

¹⁵Judd & Su (2010) and Dube, Fox & Su (2011) suggest using MPEC procedures to solve this problem rather than successive approximations. This would involve numerically solving a constrained minimization problem where we minimize (9) subject to (5), (6), (7) and (11). Conlon (2010) used MPEC to successfully solved a variant of our model.

product, we increase sales at the expense of sales in the next period. However, with innovation, it may lead to higher sales in two periods as more consumers will value an upgrade. Hence, we employ:

Assumption 2 For any vector of parameters (α^p, Σ) , there is a unique vector \vec{F} such that $\ln(\vec{s}) = \ln\left(\vec{s}(\vec{F}, \alpha^p, \Sigma)\right)$.

2.3 Identification and instruments

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.

For a given model (dynamic or static), our identification strategy is similar to BLP and the literature that follows. Heuristically, the increase in market share of product j associated with a change in a characteristic of j identifies the mean of the parameter distribution α . The products that j draws market share from identifies Σ . For instance, if product j draws market share from all products, then consumer heterogeneity is captured by the *i.i.d.* term and Σ is estimated to be small. If j draws market share mostly from products with similar characteristics, then consumers differ in their value of characteristics, which leads us to estimate Σ to be large. For the dynamic model, substitution patterns across periods in addition to within periods identify parameters. For instance, a particularly low price on a low-quality product early in the sample draws in consumers interested in low quality products, and reduces their demand later in the sample. Thus, the endogenous determination of demand across periods aids in identification.

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.

Note that our model allows for consumers to purchase products repeatedly over time, but one cannot identify from aggregate sales whether purchases are by new consumers or upgraders. Formally, the model is still identified since it does not introduce any new parameters over the static model (except for the discount factor β , which we do not estimate). However, in order to identify the extent of repeat purchase from data, we incorporate household survey data on penetration in some specifications. The change in the number of households owning a camcorder over a given period relative to sales will identify the extent of repeat purchasing in a dynamic model.¹⁶

3 Data

We estimate our model principally using a panel of aggregate data for digital camcorders.¹⁷ 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 set of models, price and quantity data were collected by NPD Techworld which surveys major electronics retailers and covers 80% of the market.¹⁸ Models in our data have the same observed characteristics over time. We create market shares by dividing sales by the number of U.S. households in a year, as reported by the U.S. Census.

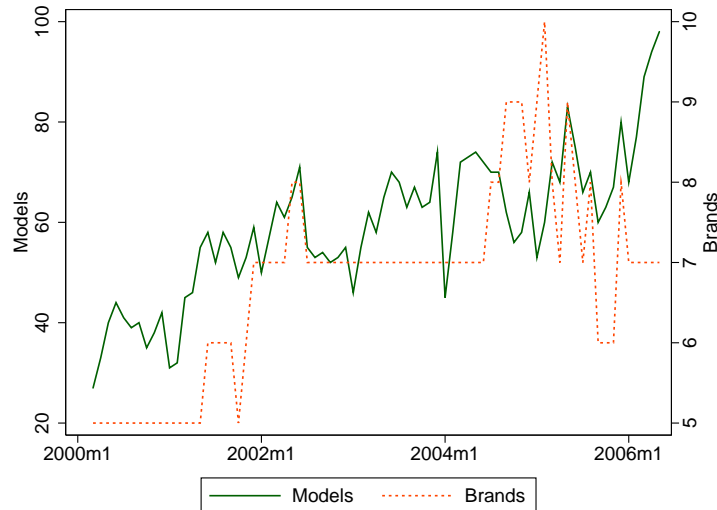
To create our final data set, we exclude from the choice set in any month all models that sold fewer

¹⁶Several recent papers relax important assumptions in our model and address the related concerns about identification. Schiraldi (2011) allows for a resale market, transactions costs, and for the depreciation of the good after purchase. Shcherbakov (2009) includes a switching cost between products. Gowrisankaran, Park & Rysman (2011) models a network effect with a complementary good, as well as households that hold multiple products.

¹⁷We 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.

¹⁸NPD 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.

Figure 3: Number of brands and models over time



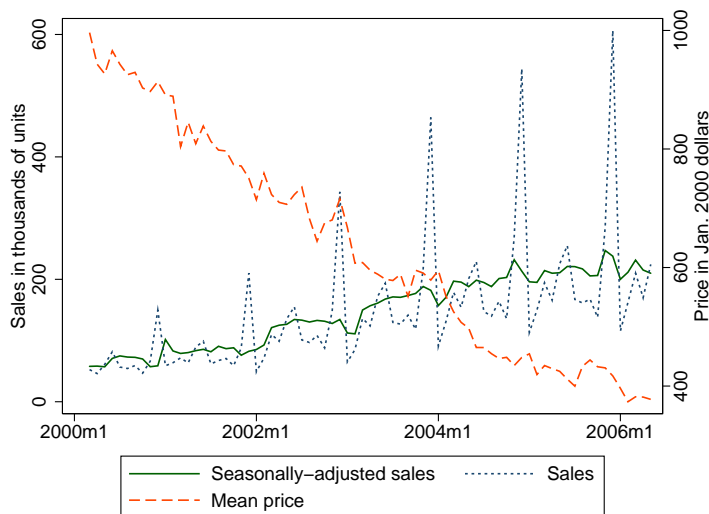
than 100 units in that month. This eliminates about 1% of sales from the sample. We also exclude from the choice set in any month all products with prices under \$100 or over \$2000 as these products likely have very different usages. This eliminates a further 1.6% of sales from the sample. Our final sample has 4,436 observations and includes 343 models and all 11 brands.

Figure 3 shows the number of models and brands over time. Both variables have a clear upward trend. The number of models varies from 29 in March 2000 to 98 in May 2006. There is substantial entry and exit – the median length of time in the data for a model is 14 months, while the mean and standard deviation are 12.9 and 7.3 months respectively. Among the brands, Canon, JVC, Panasonic and Sony are available in every month, with Hitachi, Samsung and Sharp available in most months.

In order to understand the dynamics of prices and quantities, Figure 4 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 the number of models.

Our model needs to explain the huge impact of the Christmas season on sales, which is challenging in a dynamic context. We proceed with two different methods. For most of our specifications, we address the Christmas spike issue by seasonally adjusting our data: we multiply sales by a separate constant for

Figure 4: Prices and sales for camcorders



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 4 also shows the seasonally-adjusted sales data, which are, by construction, much smoother than the unadjusted data.

For one of our specifications, we add a monthly characteristic to each product. It is unlikely that products bought over Christmas are inherently more valuable in the future. Thus, the monthly characteristic adds to utility at the time of purchase, rather than adding to f_{ijt} . This specification modifies the Bellman equation to have the month of the year as an additional state variable, and modifies the regressors in the industry evolution regression (7) to allow for month-of-year dummies instead of just a constant term. This specification adds 11 parameters, one for each month but January, all of which are estimated non-linearly. The extra state variables and parameters vastly increase the computational complexity of our estimation which is why most of our specifications seasonally adjust quantities instead of adding this characteristic.

We collected data on several important product 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.¹⁹ We observe the width and depth of each camera in inches (height was often unavailable), which we multiply together to create a “size” variable. We also record indicators for whether the camera has a lamp, whether it can take still photos and whether it has “night shot” capability, an infrared technology for shooting in low light situations. Finally, we observe the recording media the camera uses – there are four mutually exclusive media (tape, DVD, hard drive and memory card) – which we record as indicators.

Many characteristics improve on average over time. Two of the most important are the size of the camcorder and the pixel count. Figure 5 graphs simple averages of these characteristics across models by month in our final data set. Both show dramatic improvement, with pixel counts roughly doubling and size (in terms of square inches of footprint) falling by more than half. We also have several characteristics that display less dramatic improvement. Figure 6 exhibits three dummy variables over time: the presence of a lamp, the presence of night shot, and the ability to take still photographs. While all grow over time, still photo capability is widely available from the start and an included lamp never becomes widely available. Nightshot makes the largest gain. Overall, the industry appears to deliver the most significant improvement over time through decreased camcorder size and increased pixel count (notably, in our results section, we find these to be the most important drivers of consumer preferences as well).

Figure 7 presents the evolution of recording media. Early camcorders all recorded to tape, typically small tapes under the DV standard. An important innovation appearing about halfway through was the ability to record directly to DVD instead, so consumers could easily watch their recordings on their TV sets. Flash drive camcorders allow for smaller camcorders, but the memory capacities of these were very low during the time of our data set. We also observe some camcorders that use hard drives near the end of our data set. These could also be small and they had large capacity, but were expensive. By the end of the data set, less than 60% of camcorders use tape. We take from these figures that quality improvement is potentially as important as the price declines for this industry.²⁰

Finally, 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.²¹ Figure 8 shows our

¹⁹In estimation, we log all continuous variables and treat any screen of less than .1 inch as equivalent to a screen of .1 inch.

²⁰Two characteristics that we did not graph are zoom and LCD screen size. The average camera has optical zoom capability of 14X, and has an LCD screen with a 2.5 inch diameter.

²¹Data 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

Figure 5: Pixel count and camera size over time

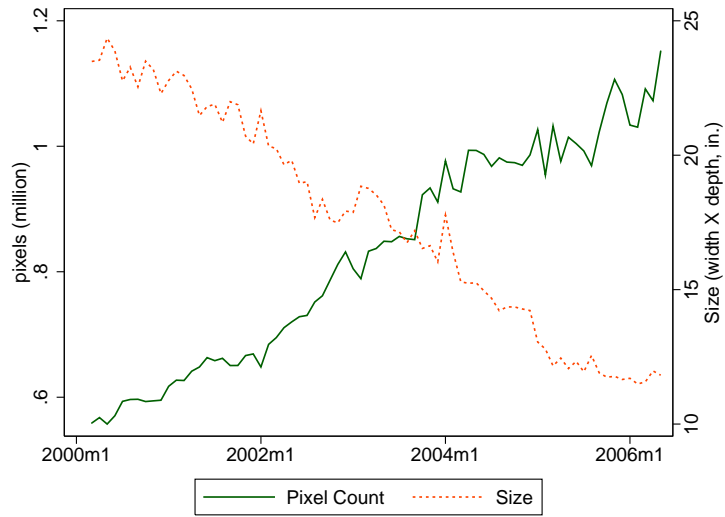


Figure 6: Indicator variables over time

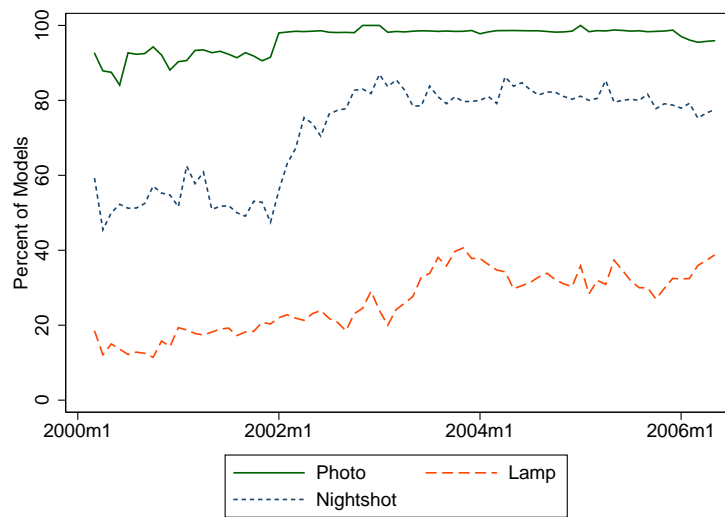
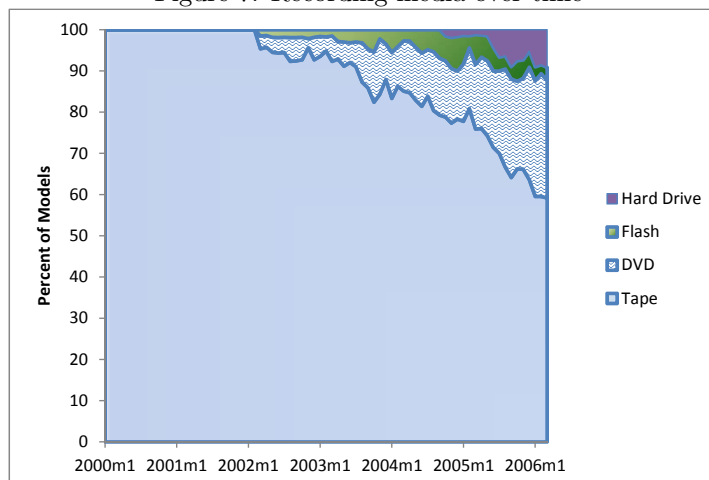


Figure 7: Recording media over time



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, as reported by NPD, the new sales of camcorders.

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

4 Results and implications

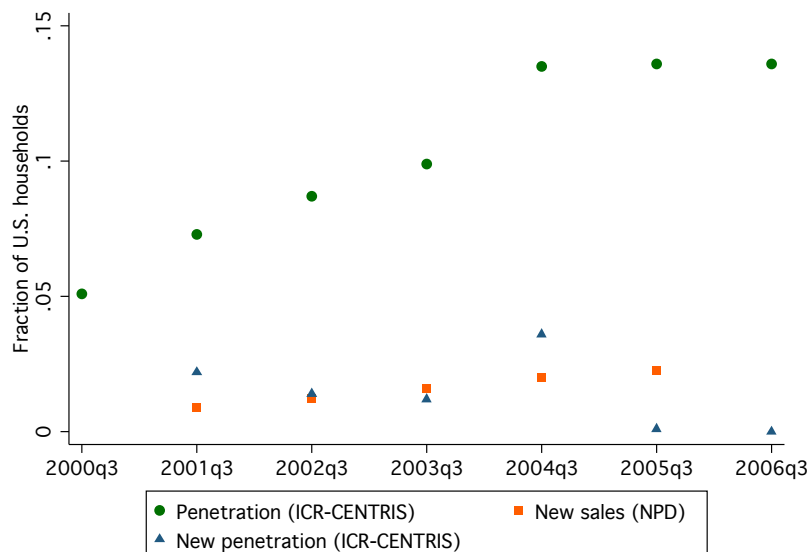
We first exposit our results, then provide evidence on the fit of the model, discuss the implications of the results and finally use our results to analyze dynamic COLIs.

Table 1: 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.

Figure 8: Penetration and sales of digital camcorders



4.1 Parameter estimates

We present our parameter estimates in Table 1. Table 1 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.

any such information.

The three included media dummies are all positive. These are relative to the flash drive technology, which is generally considered the worst during the time period of our data set. 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.

All of the estimated parameters on characteristics are smaller in absolute value than the parameter on the constant term. These characteristics either are indicators or have a standard deviation less than 1, implying 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 (and most *i.i.d* error terms) 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 F_{jt} . 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 (2007) 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).²² The value of the test statistic is 5.55, which strongly rejects the single purchase model in favor of our base model. The base model finds different results than the

²²Following Jaumandreu & Moral (2008), we base our test statistics for the non-nested test on the consistent first-stage GMM estimates.

single-purchase model, although we show below that our base model implies very few repeat purchases. This follows because in the single-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. 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 specification. To compare these coefficients with the base specification, one would have to multiply all the coefficients from this specification, except for the coefficients on price, by 1/100. The static model yields many unappealing results, including a 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. The introduction discusses possible explanations for why the static model does so poorly, such as differences between dynamic and static responses to price changes, and tension between cross-sectional and dynamic predictions in the two models.²³

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.²⁴ We chose to use only this one difference across many 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 large parameter changes on size, zoom and pixel count and small changes on the presence of a lamp or a photograph option. Hence, the model generates repeat purchase by creating a set of price insensitive consumers and increasing the importance of characteristics that improve over time.

²³An issue with comparing our model to “the static model” is that different researchers would implement the static model in different ways. Perhaps alternative specifications would perform better.

²⁴See Berry, Levinsohn & Pakes (2004) and Petrin (2002) for details on calculating weighting matrices when combining micro moments with aggregate moments.

Table 2: Robustness

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

In Table 2, 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 mean price coefficient and virtually no heterogeneity around this mean, 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.

From 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.

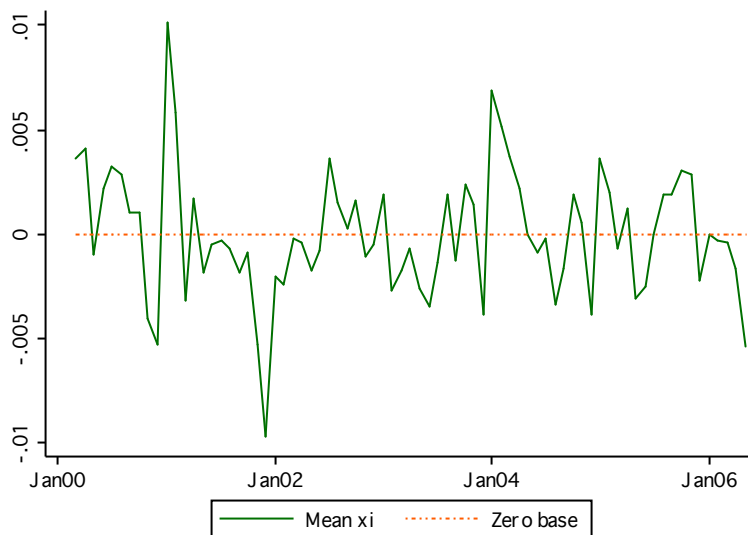
While models with the log of price tend to fit data better, it is easier to theoretically justify a model with linear price from the perspective of consumers with heterogeneous incomes. Column 4 estimates our model with a linear price. Again, the qualitative results look similar.

Column 5 estimates a 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. We find similar (unreported) results using Melnikov.

Column 6 estimates the model with monthly effects as described in Section 3. The utility function for this specification includes (unreported) 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 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.

Figure 9: Average estimation error (ξ_{jt}) by month



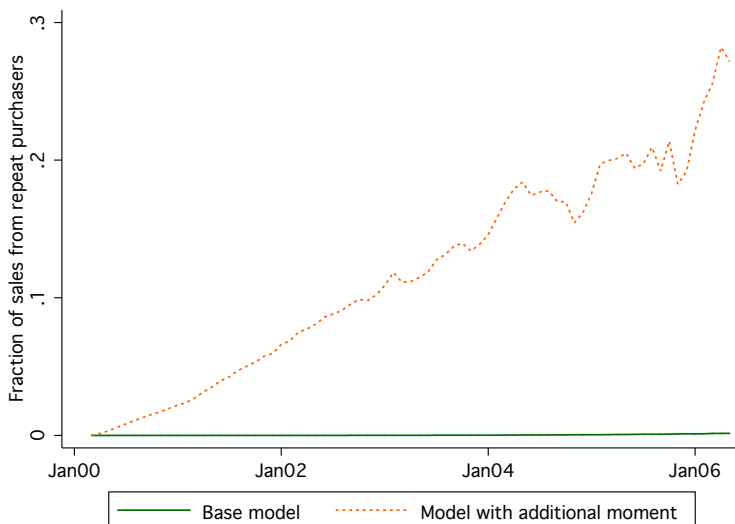
4.2 Implications of the results

We first assess the implications of the results by reporting the simple average of the unobserved quality ξ_{jt} for each month in Figure 9 using the estimated parameters from the base specification, Table 1 column 1. Note that ξ_{jt} is the econometric unobservable. 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, a product with a typical S-shaped diffusion path with a simple linear regression, we would expect to have systematic autocorrelation in ξ_{jt} . However, Figure 9 does not indicate any such pattern.

We also look at the extent to which the model generates repeat purchases. Figure 10 plots the fraction of shares due to repeat purchases for the base model as well as for the model with the micro-moment, Table 1, column 4. Under the base model, repeat purchases account for a very small fraction of total sales. Even in the final period, which has the largest fraction, repeat purchases account for only about .25% of new sales. The underlying reason why there are not more repeat purchases is that the coefficients on characteristics other than the constant term are small relative to the utility contribution from the price and the constant terms, implying that the net benefit to upgrading is low.

This finding is not consistent with the evidence, albeit imperfect, from the ICR-CENTRIS household penetration survey, that new sales are higher than new penetration. Figure 10 also plots the share of

Figure 10: Evolution of repeat purchase sales

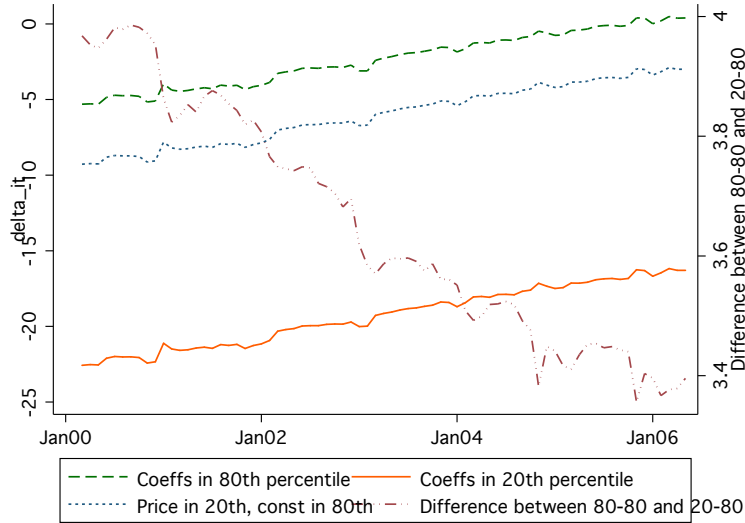


repeat purchases for the specification with the micro-moment. Since this model fits both the increase in penetration of 4.9% from Sep. 2002 to Sep. 2005 and the new sales of 5.85% over the same time period, it predicts much higher repeat purchases than the base model. In particular, it predicts that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

Finally, we plot the evolution of the logit inclusive value δ_{it} in order to compare the sources of heterogeneity in our results. Figure 11 plots δ_{it} for 3 sets of the random coefficients. We choose draws for the price and constant terms that are at the 80-80, 20-20, and 80-20 percentiles of their respective distributions. For all consumers, values are increasing close to linearly over time. As the linearity should make evident, the estimated asymptotes of the AR(1) processes are reached in the far future for the reported draws (and indeed all draws that we use). Thus, consumers expect the market to improve for the foreseeable future.

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

Figure 11: Evolution of δ_{it} over time



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.

Figure 12 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 an additional moment based on the null hypothesis of no serial correlation in ν_{it} , $E[\nu_{it}\nu_{it+1}] = 0$, 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.

Finally, Figure 13 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 current 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, both durability and forward-looking behavior, explain a very important part of the sales path. If consumers did not assume that prices and qualities changed, then sales would be

Figure 12: Difference between $\delta_{i,t+1}$ and its period t prediction

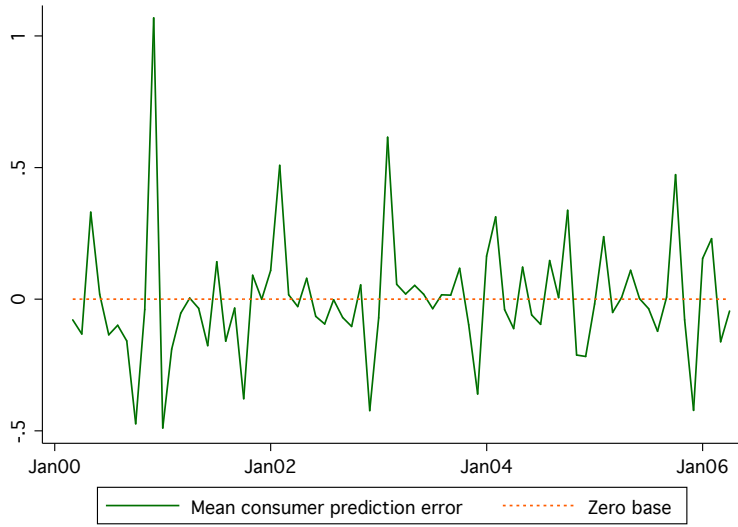


Figure 13: Evolution of digital camcorder sales under different assumptions

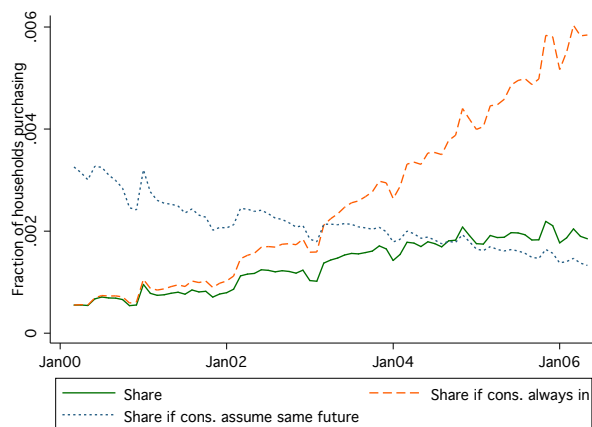
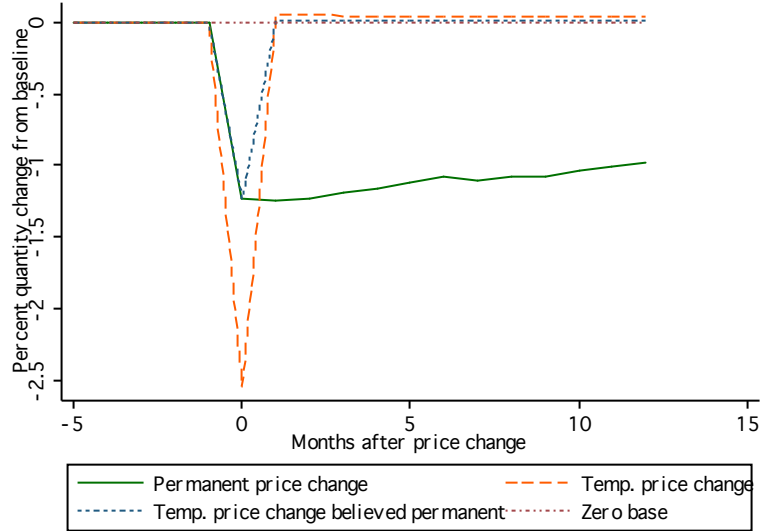


Figure 14: Industry dynamic price elasticities



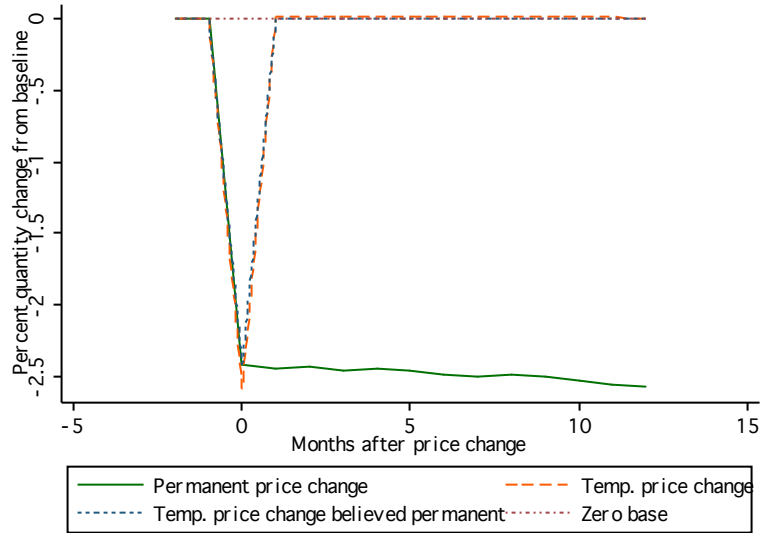
somewhat declining over time, instead of growing rapidly over the sample period, as consumers would not perceive the option value from waiting, and by the end of the sample period, many high-value consumer would already own a camcorder. If instead, we eliminate consumer holdings each period but consumers do not anticipate this taking place, the sales path would be similar to the base case until roughly two years into our sample, and then grow rapidly relative to the base case. This result is due to high valuation consumers who otherwise would be out of the market. Note that roughly 90% of the market had not purchased any digital camcorder by the end of our sample period.

4.3 Price elasticities

This subsection analyzes dynamic price elasticities, also using the coefficients in Table 1 column 1. 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 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 $\delta_{i\bar{t}}$ in Equation 6; 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.²⁵ That is, we assume that the beliefs about the future evolution of products and

²⁵The price elasticities from the static model are all virtually 0, so we do not include them on the figures.

Figure 15: Dynamic price elasticities for Sony DCRTRV250



prices, as well as the value function, conditional on the space, stays the same.

Figure 14, 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 15 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.

Strikingly, we find that the temporary price elasticities are almost the same for the industry as a whole and for the Sony DCRTRV250. However, the sources of the elasticity are different: the industry elasticity is due to consumers waiting for the price to drop whereas the product elasticity is due to consumers switching to other products. Thus, the long-run industry temporary price elasticity is much

smaller than for the product. 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 play a big role in firm strategies.

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 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 2 column 3. In performing this exercise, we hope to inform the discussion of how to improve current BLS methods.²⁶

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

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

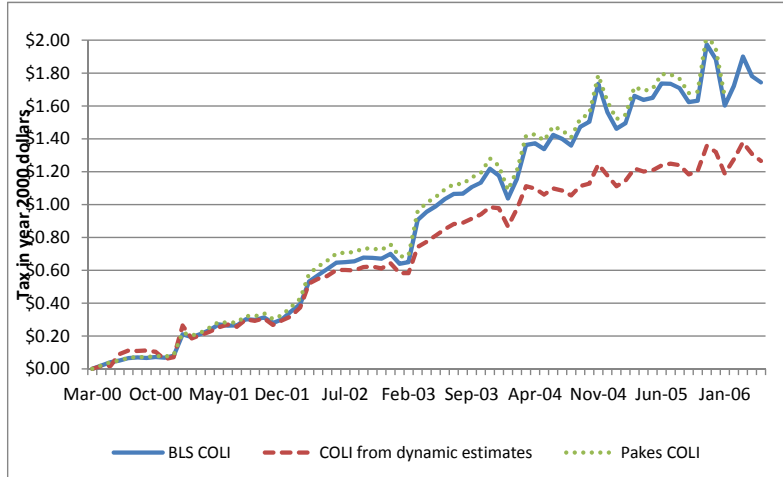
We compute a BLS-style price index from (12) using the prices and market shares (based on quantities) in our data, linked over time by model names. As is standard, we normalize the index to 100 for March 2000.²⁷ An important challenge in constructing the BLS index is determining $p_{j,t+1}$ for products that drop out of the market. One approach is to apply the average price decline for products that appear in t and $t + 1$ to products that drop out in $t + 1$. This introduces the well known “new goods” problem since the exiting product probably would have declined more quickly than 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. Interestingly, we find little evidence of a new goods problem. The indices are about the same: the BLS price index falls from 100 to 12.7 and the Pakes index falls from 100 to 14.5, slightly higher.²⁸

²⁶We 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. We focus on indices used by Pakes (2003) and the BLS but there have been other proposals for indices in dynamic settings. Reis (2009) 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).

²⁷The 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.

²⁸To implement the Pakes index, we specify a model of log price as a linear function of product characteristics (except

Figure 16: Average monthly value from camcorder market



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.²⁹ In general, one would construct COLIs by multiplying the price indices from (12) by 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). We then multiply the resulting index by the share-weighted average price in the initial period (\$961), divide by 100 and subtract the resulting term from \$961. This new index then provides the change in income relative to the first period necessary to buy the same basket of products. We plot the CPI and Pakes versions of this index in Figure 16. The indices start at \$0 by construction and end six years later at \$1.74 for the BLS COLI and \$1.72 for the Pakes COLI. That is, from the BLS COLI, a tax of \$1.74 per household in May 2006 would result in an average utility equal to the Mar. 2000 average utility, with smaller taxes necessary for earlier months. The relatively small values reflect the fact that market shares for camcorders are low.

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

for the three dummy variables on media, which show little variation within the month) and run OLS separately for each month. Using annual regressions instead of monthly regressions seems to generate more stable regressions and finds a Pakes index that falls by more (14.0 as compared to 16.6 for the BLS in December 2005), but we must drop the last four months of data because there is no following year.

²⁹Bureau of Labor Statistics (2007) states that “the concept of COLI provides the CPI [Consumer Price Index]’s measurement objective (p. 2).”

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 taxes (or subsidies) that are contingent on the aggregate industry state and 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.³⁰ 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 .³¹ 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 16.³² The dynamic index starts 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 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.

³⁰Potential 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 (2009) and Bajari, Benkard & Krainer (2005) for different approaches.

³¹If 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 unappealing.

³²We 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.

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, consumers with persistent heterogeneity over time, endogeneity of price, large and changing numbers of products, and the ability for consumers to upgrade to new durable goods as features improve. We estimate our model using a panel data set of prices, quantities and characteristics for the digital camcorder industry. We use our model to measure the welfare impact of new durable goods industries and to evaluate dynamic price elasticities for these industries.

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 short-run elasticity of demand is 2.55 for transitory price shocks, but that the long-run (12 month) elasticity is only 2.01, as many consumer delay purchase during the transitory shock. In contrast, short-run elasticity is only 1.23 for permanent price shocks. Long-run elasticities for individual products are substantially larger than market elasticities. 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.”

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. Thus, 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.

A Propositions

Proposition 1 *Assume that Assumption 1, IVS, holds. Consider states Ω and $\tilde{\Omega}$ for which $\delta(\Omega) = \delta(\tilde{\Omega})$. Then, $EV(f_0, \Omega) = EV(f_0, \tilde{\Omega})$.*

Proof We prove the proposition for the case of finite horizons and then take appropriate limits to address the case of infinite horizons. Consider first a model where the product life and market end at period T and define $EV_t^T(f_0, \Omega)$ to be the value function at time t in this case. We will prove the proposition by induction.

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

$$EV_T^T(f_0, \Omega) = \ln(\exp(\delta(\Omega)) + \exp(f_0)).$$

Since Ω_T only enters EV_T^T through δ , by the second assumption of the proposition, the valuations in period T are equal: $EV_T^T(f_0, \Omega) = EV_T^T(f_0, \delta(\Omega)) = EV_T^T(f_0, \delta(\tilde{\Omega})) = EV_T^T(f_0, \tilde{\Omega})$.

Now the inductive step. For some t such that $t \leq T$ assume that $EV_t^T(f_0, \Omega) = EV_t^T(f_0, \tilde{\Omega})$ for all f_0 and Ω for which $\delta(\Omega) = \delta(\tilde{\Omega})$. We would like to show that $EV_{t-1}^T(f_0, \Omega) = EV_{t-1}^T(f_0, \tilde{\Omega})$ when $\delta(\Omega) = \delta(\tilde{\Omega})$. We find

$$EV_{t-1}^T(f_0, \Omega) = \ln(\exp(\delta(\Omega)) + \exp(f_0 + \beta E[EV_t^T(f_0, \Omega') | f_0, \Omega])).$$

The first parts inside the ln are the same because $\delta(\Omega) = \delta(\tilde{\Omega})$ by construction. The second parts have the same conditional density because they have the same distribution of δ' by the IVS assumption and thus the same conditional density of $EV_t^T(f_0, \Omega')$ by the inductive assumption. Thus, we have proved the inductive step.

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

$$EV(f_{0t}, \Omega_t) = \lim_{T \rightarrow \infty} EV_t^T(f_{0t}, \Omega_t) = \lim_{T \rightarrow \infty} EV_t^T(f_{0t}, \tilde{\Omega}_t) = EV(f_{0t}, \tilde{\Omega}_t).$$

Because of discounting and the fact that characteristics are bounded, the limit exists and hence the equality is true. We drop the subscript t after taking the limit in T because the problem is then stationary. ■

Proposition 2 *Assume that the possible values for Ω are elements of a finite set and consider any set of contingent probabilities for the evolution of the logit inclusive value, δ . Then, there is at least one set of values $(f_j(\Omega'), P_j(\Omega') | \Omega)$, $j = 1 \dots J$ for flow utility and price respectively as a function of the state such that this set of values together with optimizing behavior imply the contingent probabilities for δ .*

Proof Let K denote the (assumed finite) number of potential values of Ω . Let k index a particular value so that $f_{j\tau k}$ denotes the realization of $f_{j\tau}$ for the k th value of Ω_τ . Let $\vec{\delta}$ denote the vector of logit inclusive values under consideration for each time period and potential state. We must define f_{jt} and P_{jt} in each of the potential states k to generate the appropriate $\delta(\Omega_t)$. There will generally be a continuum realizations of f_{jt} and P_{jt} that could generate any given $\delta(\Omega_t)$. We (arbitrarily) choose the following: let $P_{jt} = 0$ always, let product 1 at any time period τ have some contingent flow utility $f_{1\tau k}$ and let other products have a utility flow of $-\infty$.

Let \vec{f}_1 denote the vector of flow utilities for each time period and potential state. Now, define a function $m : \mathbb{R}^\infty \times \mathbb{R}^K \rightarrow \mathbb{R}^\infty \times \mathbb{R}^K$ that maps from time periods and potential states to the same space, conditional on a vector of logit inclusive values. We can write $m(\cdot|\vec{\delta})$. We define m element by element. Let the value for one particular element, $m_{\tau k}(\vec{f}_1|\vec{\delta})$, be the value of $f_{1\tau k}$ that makes the logit inclusive value for the τk th case equal to $\delta_{\tau k}$ holding constant the other flow utilities at \vec{f}_1 .

First, we show that m defines a valid function by showing that the above definition of $m_{\tau k}(\vec{f}_1|\vec{\delta})$ defines a unique value. Define the scalar-valued function $h(x|\vec{f}_1)$ as the $\delta_{\tau k}$ that would occur with flow utilities of \vec{f}_1 for every element but the τk th one and flow utility of x for the τk th one. Note that h is continuous in x : for a sufficiently low x it is unboundedly low (since the consumer will have to hold the bad product for one period); for a sufficiently high x it is unboundedly high; and it is monotonically increasing in its argument. Thus there is a unique x such that $h(x|\vec{f}_1) = \delta_{\tau k}$. This unique value, $h^{-1}(\delta_{\tau k}|\vec{f}_1)$ defines $m_{\tau k}$.

Now, we show that m has a fixed point. As m is infinite dimensional, we would like to apply Schauder's fixed point theorem. We must show that m is continuous and that it lies in a convex, compact set. The function m is continuous as h^{-1} is continuous in the argument \vec{f}_1 . To show convexity and compactness, let δ^{min} and δ^{max} denote the minimum and maximum the elements of $\vec{\delta}$ respectively. Then, no element of \vec{f}_1 will be larger than $\delta^{max}(1 - \beta)$, since purchasing a product with a flow utility of $\delta^{max}(1 - \beta)$ and never purchasing another product will already give mean expected utility δ^{max} and the actual decision allows for this option without imposing it. Thus, the elements of \vec{f}_1 are bounded above. Moreover, if the domain is bounded above by $\delta^{max}(1 - \beta)$ then the range is bounded below by $\delta^{min} - \beta(1 - \beta)\delta^{max}$, since the worst possible $f_{1\tau k}$ – which occurs if δ_τ is δ^{min} and $\delta_{\tau+1}$ is δ^{max} with certainty – yields this value. Thus, $\vec{f}_1 \in [\delta^{min} - \beta(1 - \beta)\delta^{max}, \delta^{max}(1 - \beta)]^\infty$, which is bounded and closed in \mathbb{R}^∞ and hence a compact set by Tychonov's theorem. By Schauder's fixed point theorem, g has a fixed point.

By construction, the flow utilities of a fixed point of m generate $\vec{\delta}$ as the logit inclusive values. ■

B Discretization and sampling

To perform the iterative calculation, we discretize the state space (f_{i0}, δ_i) and the transition matrix. Specifically, we compute the value function by discretizing f_{i0} into 20 evenly-spaced grid points and δ_i into 50 evenly-spaced grid points. We calculate the transition matrix by simulation as well, with 20 draws approximating the distribution of ν_t . We specify that δ_i can take on values from 20% below the minimum computed value to 20% above the maximum and assume that evolutions of δ_i that would put

it above the maximum bound simply place it at the maximum bound.³³ We have examined the impact of easing each of these restrictions and found that they have very small effects on the results.

To aggregate across draws, we need to simulate draws for α_i . A simple method is to sample standard multidimensional normal base draws $\bar{\alpha}_i \sim \phi$ and scale the base draws using $\alpha_i = \Sigma^{1/2}\bar{\alpha}_i + \alpha$. Since our estimation algorithm is very computationally intensive and computational time is roughly proportional to the number of simulation draws, we further use importance sampling to reduce sampling variance, as in BLP.

Specifically, let $\hat{s}_{sum}(\bar{\alpha}_i, \vec{F}, \alpha^p, \Sigma)$ denote the sum of predicted market shares of all camcorders at any time period for an individual with parameters (α^p, Σ) , mean flow utility \vec{F} and base draw $\bar{\alpha}_i$. Then, instead of sampling from the standard multinomial normal density we sample from the density

$$g(\check{\alpha}_i) \equiv \frac{\hat{s}_{sum}(\check{\alpha}_i, \vec{F}, \alpha^p, \Sigma) \phi(\check{\alpha}_i)}{\int \hat{s}_{sum}(\check{\alpha}, \vec{F}, \alpha^p, \Sigma) \phi(\check{\alpha}) d\check{\alpha}}, \quad (13)$$

and then reweight draws by

$$w_i \equiv \frac{\int \hat{s}_{sum}(\check{\alpha}, \vec{F}, \alpha^p, \Sigma) \phi(\check{\alpha}) d\check{\alpha}}{\hat{s}_{sum}(\check{\alpha}_i, \vec{F}, \alpha^p, \Sigma)},$$

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 g in (13) by sampling from the standard normal density ϕ and using an acceptance/rejection criterion. We compute \hat{s}_{sum} using a reasonable guess of (α^p, Σ) and the accompanying \vec{F} that solve equations 5, 6, 7 and 11 as described in Section 2.2.

In our estimation, we use 40 importance sampled draws; results for the base specification do not change substantively when we used 100 draws. Finally, instead of drawing *i.i.d.* pseudo-random normal draws for ϕ , we use Halton sequences to further reduce the sampling variance (see Gentle, 2003).

References

Ackerberg, D. A. (2003). Advertising, learning, and consumer choice in experience good markets: A structural empirical examination. *International Economic Review*, 44, 1007–1040.

³³Note that as long as the minimum of the discretization of δ_i is below the minimum observed δ_{it} and the maximum is sufficiently above the asymptotes we find from the AR1 regression from Equation 7, further expansion of the discretization does not improve our approximation.

- 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. (2007). The pure characteristics model of demand. *International Economic Review*, 48, 1193–1225.
- 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. (2007). Estimation of demand for differentiated durable goods. Unpublished manuscript, ICESI.
- Conlon, C. T. (2010). A dynamic model of costs and margins in the lcd tv industry. Unpublished Manuscript, Columbia University.

- Dube, J.-P. H., Fox, J. T., & Su, C.-L. (2011). Improving the numerical performance of BLP static and dynamic discrete choice random coefficient demand estimation. Available at SSRN: <http://ssrn.com/abstract=1338152>.
- 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*, Vol. 38, 332–354.
- 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.
- Goettler, R. & Gordon, B. (2012). Does AMD spur intel to innovate more? *Journal of Political Economy*, forthcoming.
- 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. (2009). A dynamic model of consumer replacement cycles in the PC processor industry. *Marketing Science*, 28, 846–867.
- Gowrisankaran, G., Park, M., & Rysman, M. (2011). Measuring network effects in a dynamic environment. Unpublished Manuscript, Boston University.
- Hendel, I. & Nevo, A. (2006). Measuring the implications of sales and consumer stockpiling behavior. *Econometrica*, 74, 1637–1673.
- Ho, C.-Y. (2011). Switching cost and the deposit demand in china. Unpublished Manuscript, Shanghai Jiao Tong University.
- Jaumandreu, J. & Moral, M. J. (2008). Identifying oligopoly pricing behavior: Incumbents reaction to tariffs dismantling. Unpublished Manuscript, Boston University.
- Judd, K. L. & Su, C.-L. (2010). Constrained optimization approaches to estimation of structural models. Unpublished Manuscript, Booth School of Business, University of Chicago.
- 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 Economics*, 106, 867–896.
- Lee, R. (2010). Dynamic demand estimation in platform and two-sided markets. Unpublished Manuscript, Stern School of Business.
- 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*, 5, 239–292.
- 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.
- 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. (2008). Repeat purchase amid rapid quality improvement: Structural estimation of demand for personal computers. *Journal of Economics and Management Strategy*, 17, 1–33.
- Reis, R. (2009). A dynamic measure of inflation. Unpublished manuscript, Columbia University.
- 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. (2011). Automobile replacement: A dynamic structural approach. *RA*, 42, 266–291.
- Shcherbakov, O. (2009). Measuring switching costs in the television industry. Unpublished Manuscript, University of Mannheim.

- Shepler, N. (2001). Developing a hedonic regression model for camcorders in the U.S. CPI. Bureau of Labor Statistics.
- Solow, R., Tobin, J., von Weizsacker, C., & Yaari, M. (1966). Neoclassical growth and fixed factor proportions. *Review*, 33, 79–115.
- 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.