DYNAMICS OF CONSUMER DEMAND FOR NEW DURABLE GOODS

Gautam Gowrisankaran
Marc Rysman

Dynamics of Consumer Demand for New Durable Goods
Gautam Gowrisankaran and Marc Rysman
NBER Working Paper No. 14737
February 2009, Revised December 2012
JEL No. C23,E31,L1,L13,L68

ABSTRACT

Most new consumer durable goods experience rapid prices declines and quality improvements, suggesting
the importance of modeling dynamics. This paper specifies a dynamic model of consumer preferences
for new durable goods with persistently 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 the one-year elasticity in response to a transitory industry-wide
price shock is about 25% less than the one-month elasticity. Standard cost-of-living indices overstate
welfare gain in later periods due to a changing composition of buyers.

Gautam Gowrisankaran
Department of Economics
University of Arizona
P.O. Box 210108
Tucson, AZ 85721-0108
and HEC Montreal
and also NBER
gowrisankaran@eller.arizona.edu

Marc Rysman
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
mrysman@bu.edu

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 model with a superior model in the foreseeable future. Thus, dynamic behavior is an important element of demand for consumer durable goods. This paper specifies a structural dynamic model of consumer preferences for new durable goods; estimates the model using data on digital camcorders that are primarily at the level of the camcorder model; 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. During this time, 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 models 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.

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 may 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, a changing number of models, 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. To make the problem tractable, 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 models 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 almost 100 models at some time periods, each 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 emphasizes 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 conditionally independent given the observed state, an assumption that is typically required for dynamic estimation. The cost is 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, 2012), scrapping subsidies for automobiles (such as cash-for-clunkers programs; see Schiraldi, 2011),

markups for digital cameras (Zhao, 2008), and switching costs in consumer banking, Medicare managed care, and subscription television (see Shcherbakov, 2009; Ho, 2011; Nosal, 2012, respectively), among other research questions.

We estimate very different price and characteristic elasticities between our dynamic model and similar static models. Moreover, the dynamic results are more intuitive, e.g., with negative price elasticities and positive elasticities on important characteristics such as LCD screen size. Thus, it is useful to consider how static estimates of durable goods models might be biased. By modeling dynamics, we have essentially transformed a consumption problem into a capital investment problem. Firms will invest in a piece of capital if the flow of services from the capital is greater than the rental cost of the capital. The rental cost of capital is essentially the *difference* in the present-value price of capital between this period and next period.[2] Thus, for the camcorder industry, static models would predict a steady increase in sales as prices drop, while dynamic models would predict that sales would increase the most only when prices stop dropping. By incorrectly using price instead of the difference in price, a static estimation applied to a durable good purchase decision with falling prices will then result in mismeasurement that may tend to bias the price coefficient towards zero. Our model adds complexity in that we have heterogeneous agents. In particular, the sales increase from price declines will be moderated over time since demand endogenously falls as high-value consumers accumulate the good. Not accounting for heterogeneity will further cause the static model to understate the importance of price and quality, since the population response to product improvement over time is smaller than the average individual response, due to the changing population of available consumers. 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, which may also lead to implausible estimates from a static estimation.[3]

We use our results to analyze the difference between a short-run elasticity, computed from the change in sales for one month, to a long-run elasticity, which measures the change in sales over 12 months. We compute these elasticities for a temporary price increase, in which prices increase for one month and

---

[2]Gandal, Kende & Rob (2000) show formally that in a simple dynamic model with one model, sales are a linear, decreasing function of the forward difference in price net of discounted future price, $(p_t - \beta p_{t+1})$.

[3]Consider the case in which prices fall and then level off. In a dynamic model, a consumer may prefer the high quality model but wait to purchase until the price stops dropping. Thus, the data will show first that consumers bought when prices were low; and second that within a period, they purchased the high price, high quality model. 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 low, leading to an imprecise and inconsistent price coefficient. The dynamic model can rationalize these data appropriately.

then return to their previous level. When considering a temporary price change to the industry as a whole, the long-run elasticity is only 56% the size of the short-run elasticity, since many consumer delay purchase during the price spike, but still purchase a camcorder within one year. However, the short- and long-run elasticities for a price change to a single camcorder model are very close to each other. That is because we find that camcorders are relatively close substitutes, so when the price of a single model increases, consumers switch to another model rather than delaying purchase. We also compute elasticities from a permanent price increase, in which prices remain at the new level. The short-run elasticity of a permanent industry-wide price change is about 48% of the elasticity of a temporary industry-wide price change.

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). Standard COLI measures show welfare substantially improving in this market since prices fall and quantities rise. However, 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. Specifically, our dynamic model finds that welfare gains are attenuated in later periods, since later buyers are either low value relative to early buyers, or they already hold the good.

Because we use primarily model-level 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 model 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 model 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 data on prices and market shares by camcorder model and 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

5

household-level data are substantially more sophisticated than those for model-level data, incorporating not only dynamics, heterogeneity and upgrading but also such features as learning, product loyalty, inventory behavior and surveys of price expectations (see Ackerberg, 2003; Hendel & Nevo, 2006; Erdem, Imai & Keane, 2003; Erdem & Keane, 1996; Prince, 2008; Keane & Wolpin, 1997). We extend dynamic estimation with model-level data. We believe that these models are important for two reasons. First, in many cases, model-level data are all that are available. Second, model-level data are typically necessary for studying many important issues, such as oligopoly interactions. This is because household-level datasets rarely contain enough observations to measure individual model shares accurately.[4] Accurate market shares are important for estimating the supply side. For instance, BLP and Goldberg (1995) use model-level data to estimate pricing first-order conditions.

A number of recent papers (Gandal et al., 2000; Esteban & Shum, 2007; Melnikov, 2012; Song & Chintagunta, 2003; Gordon, 2009; Nair, 2007; Carranza, 2007; Park, 2008) estimate dynamic consumer choice problems with model-level data. Most similar to our work is Melnikov (2012), which was the first to model dynamics in a logit-based discrete choice model with endogenous prices and model-level data. We use a similar reduced-form approximation of the supply side as Melnikov. Our model builds on Melnikov by adding a full set of persistent random coefficients and repeat purchases over time, all modeled in an explicitly dynamic framework. Another 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 household-level data, Hendel & Nevo (2006) are able to identify the parameters underlying consumer stockpiling. However, their model restricts future purchases to not depend on the current purchase except through its size, and their model does not allow for persistent heterogeneity in preferences for characteristics that vary within models of the same size, such as price. A final comparison is to Goettler & Gordon (2011), who estimate supply and demand in the microprocessor industry. Given their focus on endogenous model 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, which makes it difficult to incorporate persistent consumer heterogeneity or unobserved model characteristics, or to apply it to industries with many firms or models. In contrast, we address a context with a very large and changing number of models, persistent random coefficients, and endogenous prices.

---

[4]Our 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.

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.

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 models 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 models $j = 1, \ldots, J_t$. Each model 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 models or chooses to purchase no product. If she buys model j at time t, then at time t she receives utility of

$$u_{jt} = f_{jt} - P_{jt} + \varepsilon_{jt}, \text{ for } j = 1, \cdots J_t, \tag{1}$$

where ε_{jt} is an idiosyncratic type 1 extreme value term, distributed *i.i.d.* across models 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 model 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 ε_{0t}, which is also

---

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

i.i.d. across time periods. 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 $t \le \hat{t}$, $f_{0t} = 0$. For $t > \hat{t}$, f_{0t} will be the flow utility for the model purchased at $\hat{t}$. This level of flow utility continues until the next purchase, after which point it will be the level from this subsequent purchase, etc. Thus, the consumer does not need to remember the identity or characteristics of the model she purchases, only the flow utility of her most recent purchase.

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

We now define the state variables and use them to exposit the Bellman equation. Let $\vec{\varepsilon}_t \equiv (\varepsilon_{0t}, \ldots, \varepsilon_{J_t t})$ (where we use the right arrow to denote a vector), and let $g_{\vec{\varepsilon}_t}$ denote the joint density of $\vec{\varepsilon}_t$. Then, the purchase decision for the consumer depends on $\vec{\varepsilon}_t$, endowment f_{0t}, attributes of currently available models, and expectations about future model attributes. Future model 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 models J_t, the price disutility and mean flow utility for each model, and any other factors that influence future model attributes. We assume that Ω_t evolves according to some Markov process $g_\Omega(\Omega_{t+1}|\Omega_t)$ that accounts for firm optimizing behavior.[6] Thus, the state vector at time t is $(\vec{\varepsilon}_t, f_{0t}, \Omega_t)$.

Let the prime symbol (e.g., Ω') denote next period's value of a variable (which allows us to drop the t subscript) and let $V(f_0, \Omega)$ denote the value function prior to the realization of $\vec{\varepsilon}$. Then,

$$V(f_0, \Omega) = \tag{2}$$

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

where "E" denotes the expectation operator, a conditional expectation in this case, and P_j, and f_j, $j = 1, \ldots, J$ 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

---

[6]We assume in Section 2.2 below that there is a continuum of consumers and hence that an individual consumer's decision does not affect the evolution of Ω.

contemporaneous choice – into the next period's value function, where it becomes the new f_0 moving forward.

To estimate our model, we will ultimately need to solve the Bellman equation 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 (2) in a relatively simple form, and then make assumptions on the perceptions of industry evolution based on this form. Specifically, we can rewrite

$$V\left(f_0, \Omega\right) = \ln\left[\exp\left(f_0 + \beta E\left[V\left(f_0, \Omega'\right)|\,\Omega\right]\right) + \exp\left(\delta(\Omega)\right)\right],^{7} \tag{3}$$

where $\delta(\Omega)$, the *logit inclusive value* is defined as

$$\delta(\Omega) = \ln\left(\sum_{j=1,\dots,J} \exp\left(f_j - P_j + \beta E\left[V\left(f_j, \Omega'\right)|\,\Omega\right]\right)\right). \tag{4}$$

The logit inclusive value is the ex-ante present discounted lifetime value of buying the preferred camcorder, as opposed to holding the outside option.[8] Thus, (3) shows that, for any state (f_0, Ω), the value function is equivalent to the value function from a simpler problem where the consumer makes a one-time purchase of a model with mean utility $\delta(\Omega)$ or continues to hold her existing model. Moreover, (3) shows that Ω affects V only through its impact on the current $\delta(\Omega)$ and predictions of future values of $\delta(\Omega)$. Thus, the current δ and the predictions of future values of δ (and holdings) are sufficient statistics for the consumer's dynamic optimization problem. The consumer uses Ω only to make predictions of the future distribution of δ.[9]

A simplifying assumption on how consumers form predictions of $\delta(\Omega)$ can then greatly reduce the complexity of our problem. Let g_δ denote the density of δ. 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 $g_\delta(\delta(\Omega')|\Omega) = g_\delta(\delta(\tilde{\Omega}')|\tilde{\Omega})$ for all Ω, $\tilde{\Omega}$.*

The IVS assumption and (3) imply that all states with the same $\delta(\Omega)$ have the same expected value; see Proposition 1 in Appendix A for details. Thus, it is sufficient for the consumer to track only two

---

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

[8]Formally, 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) and Rust (1987) provide proofs of this statement for static and dynamic models respectively.

[9]Gowrisankaran & Rysman (2012) provides a formal proposition and proof for this model; Melnikov (2012) provides a similar proof for his model.

scalar variables, f_0 and δ in order to make their dynamic decisions.[10] The IVS assumption is valuable, since we can now replace Ω with δ in the state space. We can then rewrite (3) as

$$\mathcal{V}(f_0, \delta) = \ln\left[\exp\left(f_0 + \beta E\left[\mathcal{V}(f_0, \delta')\,|\delta\right]\right) + \exp(\delta)\right], \tag{5}$$

and (4) as

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

where $\mathcal{V}$ is the value function defined over (f_0, δ). Note that this provides a two-dimensional, and hence tractable, state space.

The IVS assumption is also restrictive: consider that δ could be high in a period either because there are many models in the market all with high prices or because there is a single model 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 subject to IVS 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 potentially boundedly rational consumers perceive this market. For the camcorder industry, we believe that the most important dynamics stem from consumers' expectations in early periods of the rapidly declining prices and increasing quality that later occurred. Thus, we assume rational expectations, that the consumer is on average correct about the future.

We employ one of two different specifications of the rational expectations. 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 models were non-decreasing), as in this case, there is a one-to-one mapping from t to δ.[11] We believe that it is more realistic to assume that the consumer has only a limited ability to predict future model attributes. Thus, for most of our specifications, we let consumer perceptions about next period's δ, $g_\delta(\delta'|\delta)$, be its

---

[10]Similar assumptions have been used in the existing dynamic literature. Most other papers (Melnikov, 2012; Hendel & Nevo, 2006) specify their analog of δ as a function of the flow utilities of available models, whereas ours also incorporates the continuation values of holding those models. 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. The advantage of our approach is that it captures the endogeneity of purchase decision to the quality of the model purchased, e.g. that a consumer who purchases a high quality model will be less likely to upgrade. See Gowrisankaran & Rysman (2012) for a further discussion of different assumptions that can simplify the state spaces for durable goods problems.

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

actual empirical density fitted to a simple linear autoregressive specification:

$$\delta_{t+1} = \gamma_1 + \gamma_2 \delta_t + \nu_{t+1}, \tag{7}$$

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.[12]

The assumption of IVS and (7) is crucial to our approach, since this is what allows us to proceed with dynamic estimation without jointly solving the supply side. Moreover, we believe that it captures the first-order feature of the camcorder industry, that prices are dropping and quality is rising.[13] Given its centrality, it is important to explore the plausibility of the assumption. Appendix B provides a Monte Carlo analysis of how our approach performs on simulated data. Appendix B also reviews the empirical evidence from throughout the paper on the efficacy of IVS and (7). Each of our exercises supports IVS and (7).

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.[14] 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.

Importantly, with IVS and rational expectations, the optimal consumer decisions given an industry environment are defined by the joint solution to the Bellman (5), the logit inclusive value (6), and the industry evolution regression (7). 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.

---

[12]We can always rationalize the assumption in (7) with an assumption on the evolution of the underlying exogenous model characteristics. See Proposition 2 in Appendix A for details. Intuitively, this proof makes use of the larger dimensionality of flow utility and prices relative to δ.

[13]For similar tractability reasons, the macroeconomics general equilibrium literature on the impact of income and wealth heterogeneity (see Krusell & Smith, 1998) includes moments that capture the important components of heterogeneity instead of modeling the complete distribution.

[14]Goettler & Gordon (2011) provide an alternative approach to address a changing environment. They specify a model that is stationary relative to an endogenously advancing technology frontier.

2.2 Aggregation and equilibrium

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 and the concept of equilibrium. The aggregation follows BLP closely. We 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, Bellman equation, and expectations processes for the future. All of these terms should now be indexed by i, i.e., f_{ijt}, P_{ijt}, ε_{ijt}, δ_{it}, $\mathcal{V}_i$, and $(\gamma_{1i}, \gamma_{2i}, \nu_{it})$ respectively.

We assume that flow utility and price fit in the random coefficients framework developed by BLP. Specifically, in (1), 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.[15] 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 model.[16] 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.

For the supply side, we assume that models arrive according to some exogenous process and that their characteristics evolve exogenously as well. Firms have rational expectations about the future evolution of model characteristics. After observing consumer endowments and x_{jt} and ξ_{jt} for all current models, firms simultaneously make pricing decisions in a Markov Perfect Equilibrium. 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 outcomes given changes in exogenous variables.

---

[15]Note 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.

[16]One could relax this assumption by allowing f_0 to vary over time in the absence of any purchase. This would be useful to model sectors with substantial depreciation such as automobiles (Schiraldi, 2011).

2.3   Inference

We now discuss inference of the structural parameters, (α, Σ, β), which requires bringing together the dynamic and static parts of the model. To do so, it will be helpful to define the *mean flow utility* of model j in period t as

$$F_{jt} = x_{jt}\alpha^x + \xi_{jt}, j = 1, \ldots, J_t. \tag{8}$$

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

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 camcorders. 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 characteristics (ξ_{jt}) for which the predicted shares equal the observed shares conditional on parameters, and z is a matrix of exogenous variables, described in detail in Section 2.4 below. We estimate parameters to satisfy

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

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 for each model. As in Section 2.1, the consumer decision problem is defined by the fixed points to three equations: the Bellman equation (5), the logit inclusive value (6), and the industry evolution regression (7). Based on these equations, we compute the probability of purchase of each model, for each consumer and each possible set of state variables. This *conditional probability of purchase* is defined as:

$$\frac{\exp(\delta_{it})}{\exp\left(\mathcal{V}_i(f_{i0t}, \delta_{it})\right)} \times \frac{\exp\left(f_{ijt} - P_{ijt} + \beta E\left[\mathcal{V}_i\left(f_{ijt}, \delta_{i,t+1}\right)| f_{ijt}, \delta_{it}\right]\right)}{\exp(\delta_{it})}. \tag{10}$$

In (10), the probability that a consumer who holds a model with value f_{i0t} and faces a market with value δ_{it} purchases good j is the probability of purchasing any camcorder in that time period (the first term) times the probability of purchasing model j conditional on purchase (the second term).

With the conditional probability of purchase in hand, we solve for model-level market shares by starting at time 0 with the assumption that all consumers hold the outside good.[17] Iteratively for subsequent time periods, we use purchase probabilities computed for the realized states to update consumer holdings at each period. We then integrate over consumers i to compute predicted market shares by model and time period. We perform the integration via simulation, as in BLP. This integration is further described in Appendix C.

The above process finds market shares conditional on the vector of mean flow utilities $\vec{F}$. We wish to find the values of $\vec{F}$ that make predicted shares for each model at each time period match those in the data. This requires solving a fourth equation,

$$s_{jt} = \hat{s}_{jt}\left(\vec{F}, \alpha^p, \Sigma\right), \forall j, t, \tag{11}$$

where s_{jt} is the observed market share in the data and $\hat{s}_{jt}\left(\vec{F}, \alpha^p, \Sigma\right)$ is the predicted market share of our model. We solve (11) by iterating on $\vec{F}$ analogously to the contraction mapping used by BLP:

$$F_{jt}^{new} = F_{jt}^{old} + \psi \cdot \left(\ln(s_{jt}) - \ln\left(\hat{s}_{jt}\left(\vec{F}^{old}, \alpha^p, \Sigma\right)\right)\right), \forall j, t, \tag{12}$$

where ψ is a tuning parameter used in the computation that we generally set to $1 - \beta$.

In practice, we solve these equations via successive approximations, beginning with guesses of δ_{it}, $\mathcal{V}_i(f_{i0t}, \delta_{it})$, and F_{jt} and circulating between them, (7), and (10) until we find a fixed point. 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 computing (10) and taking a step in (12).[18]

Using the $\vec{F}$ that arises from this fixed point algorithm, we compute $\vec{\xi}(\alpha, \Sigma)$ from (8), and then we construct our objective function in (9). As we detail in Appendix C, 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 a non-linear search only over α^p and Σ. All computer code is available from the authors upon request.

---

[17]A strength of our dataset 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 on holdings by consumer type at the start of the market (see, for example Schiraldi, 2011; Shcherbakov, 2009).

[18]Judd & Su (2012) and Dube, Fox & Su (2012) 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), (10), and (11). Conlon (2010) used MPEC to successfully solve a variant of our model.

An important issue is whether these equations 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 models are substitutes. A variant of our model where consumers can only purchase once and where every current and future model attribute including the extreme value shocks are known would satisfy substitutability. In contrast, in most dynamic models, models may be complements with future models. As an example of a complementary between a model in the current period and two periods forward, if we exogenously lower price of a current period model, 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 $\vec{s} = \vec{s}\left(\vec{F}, \alpha^p, \Sigma\right)$.*

2.4 Identification and instruments

For a given economic model (dynamic or static), our identification strategy is similar to BLP and the literature that follows. Heuristically, the increase in market share of model j associated with a change in a characteristic of j identifies the mean of the parameter distribution α. The models from which j draws market shares identify Σ. For instance, if model j draws market share equally from all models in proportion to their shares, 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 models 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 model early in the sample draws in consumers interested in low quality models, 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 model characteristics are exogenous. This assumption is justified under a model in which camcorder characteristics are determined as part of some technological progress that is exogenous to the unobserved model 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 model characteristics in x; the mean model characteristics for a given firm at the same time period; the mean model characteristics for all firms at the time period; and the count of models offered by the firm and by all firms.

These variables are meant to capture how crowded a model is in characteristic space, which should affect the price-cost margin and the substitutability across models, 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 sales data 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.[19]

3   Data

We estimate our model principally using a panel of model-level data for digital camcorders.[20] The data are at the monthly level and, for each model and month, include the number of units sold, the average price, and other 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.[21] 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 dataset, we exclude from the choice set in any month all models that sold fewer

---

[19]Several 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.

[20]We 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. See Gowrisankaran et al. (2011).

[21]NPD 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 models.

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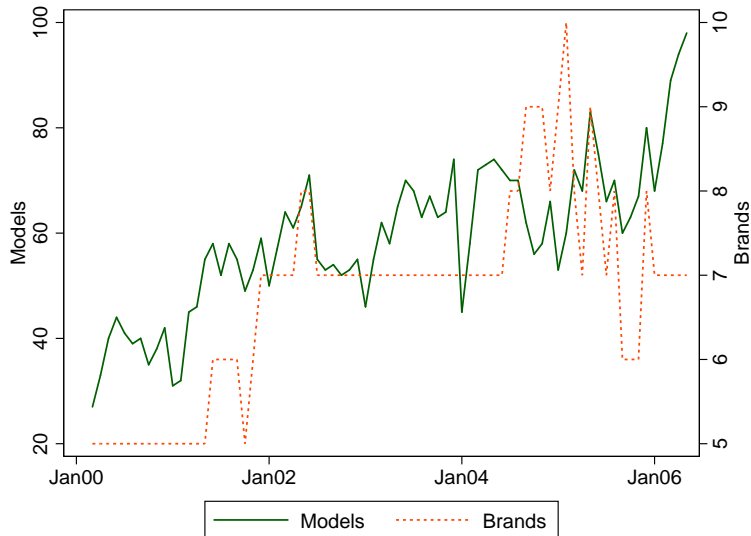
than 100 units in that month. This eliminates about 1% of sales from the sample. We also exclude from the choice set in any month all models with prices under \$100 or over \$2000 as these models likely have very different usages. This eliminates a further 1.6% of sales from the sample. Our final sample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

Figure 1 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 2 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

17

Figure 2: Prices and sales for camcorders
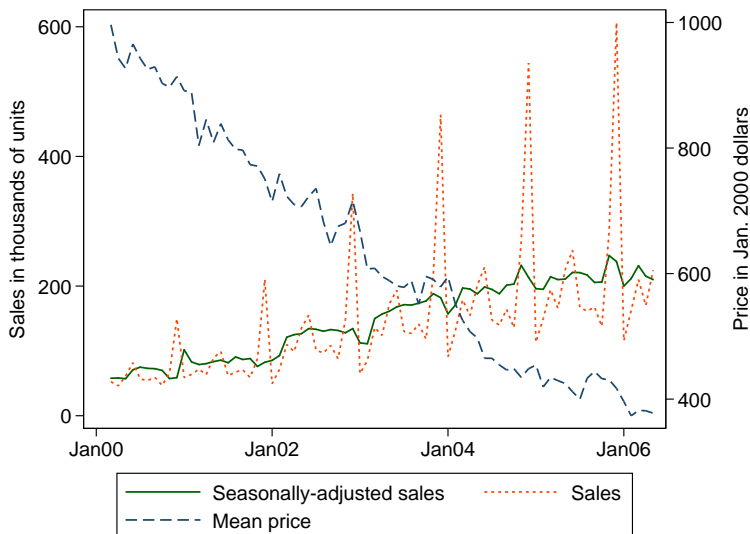
for each month of the year, chosen so that (a) the geometric mean of adjusted sales is the same across months of the year and (b) adjusted total sales over the sample period are equal to total sales in the raw data.[22] Figure 2 also shows the seasonally-adjusted sales data, which are, by construction, much smoother than the unadjusted data.

For one of our specifications, we use the raw sales data add a month-of-year indicator as a characteristic to each model. It is unlikely that products bought over Christmas are inherently more valuable in the future. Thus, the month-of-year characteristic adds to utility at the time of purchase, rather than adding to f_{ijt}. Thus, this specification adds 11 parameters, one for each month but January, all of which are estimated non-linearly. Also, this specification modifies the regressors in the industry evolution regression (7) to allow for month-of-year dummies instead of just the constant and autoregressive terms. In order to put month-of-year dummies into (7), this specification requires the Bellman equation to have the month of the year as an additional state variable. The extra parameters and state variables 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 model characteristics from on-line resources. We observe the

---

[22]We perform the seasonal adjustment by first regressing log sales on month-of-year dummies. We then construct adjusted sales by dividing raw sales for each model in each month by the exponentiated dummy for its month of the year, and then multiplying by a common factor that makes total adjusted sales match total sales in the raw data.

number of pixels that the camcorder 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.[23] We observe the width and depth of each camcorder in inches (height was often unavailable), which we multiply together to create a "size" variable. We also record indicators for whether the camcorder 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 camcorder 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 3 graphs simple averages of these characteristics across models by month in our final dataset. 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 4 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 5 presents the evolution of the 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 dataset. We also observe some camcorders that use hard drives near the end of our dataset. These camcorders could also be small and they had large capacity, but were expensive. By the end of the dataset, 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.[24]

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

---

[23]In estimation, we log all continuous variables and treat any screen of less than 0.1 inch as equivalent to a screen of 0.1 inch.

[24]Two characteristics that we did not graph are zoom and LCD screen size. The mean optical zoom capability is 14X, and the mean LCD screen size diameter is 2.5 inches.

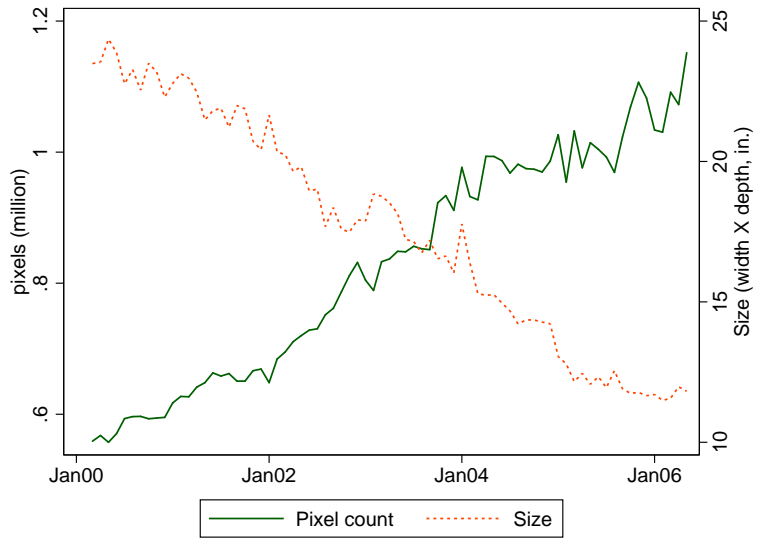Figure 3: Pixel count and camcorder size over time

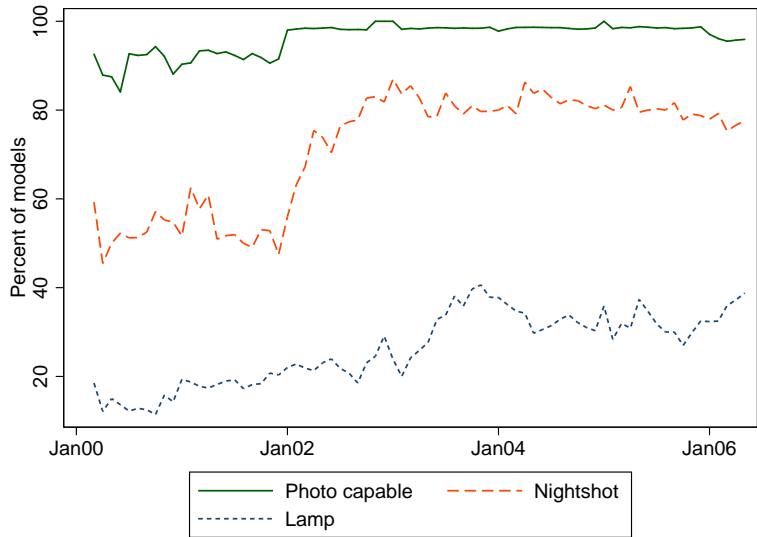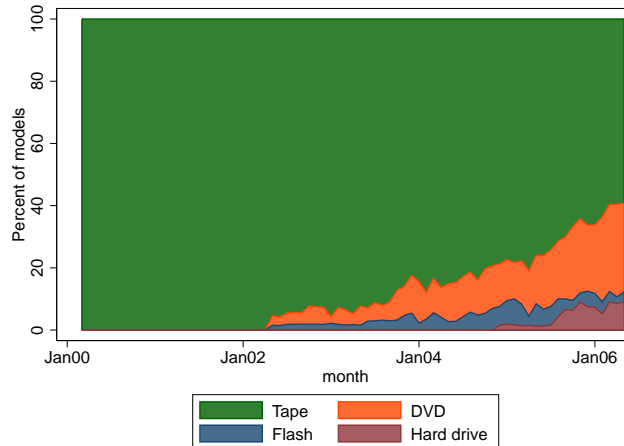Figure 4: Indicator variables over time

Figure 5: Recording media over time

interviews a month, asking which consumer electronics items a household owns.[25] Figure 6 shows our ICR-CENTRIS data, which contain the percent of households that indicate holding a digital camcorder in the third quarter of the year for 1999 to 2006. It also shows the year-to-year change in this number and, as reported by NPD, the new sales of camcorders.

These 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.

---

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

Figure 6: Penetration and sales of digital camcorders
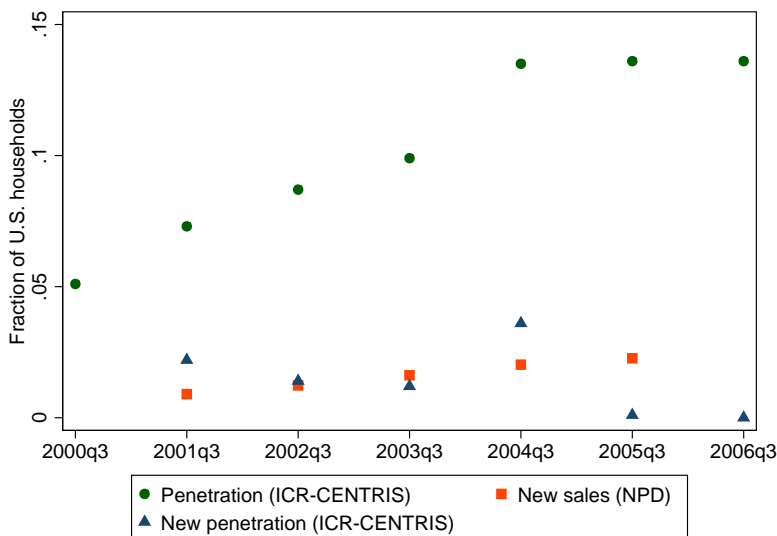
4.1 Parameter estimates

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. This specification has two random coefficients, one on price and the other on the constant term. It obtains 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 0.345. 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 -0.092. The standard deviation on the constant term in the consumer population is 0.079, 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 coefficients other than on price should be roughly $(1 - \beta) = 0.01$ times the magnitude of the comparable coefficient in a static model, while the coefficients on price should have similar magnitudes to those from 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.

22

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)/10	.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.

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 dataset. 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). Ackerberg & Rysman (2005) argue that this feature implies that logit-based models will perform poorly in contexts where consumers face different numbers of models over time. Ackerberg & Rysman recommend addressing this problem by including the log of the number of models, $\ln(J_t)$, 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 -0.013. 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).[26] The value of the test statistic is 5.55, which strongly rejects the single purchase model in favor of our base model.

---

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

Interestingly, we show below that the base model implies very few repeat purchases. Still, the base model and the single-purchase model lead to different parameter results. In particular, the magnitude of the mean price coefficient in the single-purchase model is much smaller. Had this estimated coefficient been applied to the model that allowed repurchases, the amount of repurchases would be very high, which would lead to unrealistic sales patterns. This sharp difference in predicted purchase patterns explains why the coefficient estimates can be so different.

Column 3 follows BLP and estimates a traditional static random coefficients discrete choice specification. These coefficients can be interpreted as the lifetime value of the associated characteristics, whereas characteristic coefficients in the base specification represent the flow value, so to compare the characteristic coefficients (but not the price coefficients) with the base specification, we have to multiply them by $1 - \beta = 0.01$. 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.[27]

We are concerned that repeat purchase behavior is not well-identified from data on market shares. Thus, Column 4 uses the penetration data in the form of a micro-moment (see Petrin, 2002). 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 in the penetration data and as predicted by the model.[28] We chose to use only this one difference across many years to mitigate the noise present in these data.

Our model can use two features to generate repeat purchases. First, there must be some consumers with a low disutility of price, as only these consumers will repurchase. Second, consumers must care about model features that improve over time. We find both of these elements are important. In Column 4, the standard deviation of the random coefficient on price more than doubles. That is, the model matches the repeat purchases by generating more consumer heterogeneity, which increases the set of consumers who care about price very little. Second, the coefficients on some of the characteristics

---

[27]An 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.

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

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, all characteristics that show substantial improvement over time in our data. In contrast, we find small changes for coefficients on the presence of a lamp or a photograph option, which are both characteristics that change little over the dataset. Hence, the model generates repeat purchase by creating a set of price insensitive consumers and increasing the importance in the utility function of characteristics that improve over time.

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 that consumers use to predict future realizations of the inclusive value.[29] 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 as our base specification. 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 (2012). We solve it based on our method rather than the two-step estimation algorithm 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

---

[29]In this exercise, we specify two linear regressions, analogous to (7), for the state evolution process. The dependent variables are δ' and $\ln(J')$ and both regressions include both state variables as regressors. Note that this robustness check is different from the Ackerberg & Rysman (2005) test for crowding in unobserved characteristic space. For the test based on Ackerberg & Rysman (2005), we include $\ln(J_t)$ as if it were a characteristic in the utility function. In Column 1, Table 2, we include $\ln(J_t)$ as an explanatory variable in the AR(1) regression describing the evolution of δ_{it} (7). The former does not require including $\ln(J_t)$ as a state variable, whereas the latter does.

Table 2: Robustness

Parameter	State space includes number of products (1)	Perfect foresight (2)	Dynamic model with extra random coefficients (3)	Linear price (4)	Melnikov's model (5)	Month dummies (6)
Mean coefficients (α)						
Constant	-.098 (.026) *	-.129 (.108)	-.103 (.037) *	-.170 (.149)	-6.61 (.815) *	-.114 (.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)/10	.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.

using Melnikov's two-step estimation algorithm.

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 0.720 for December, 0.250 for November and which range from -0.146 to 0.104 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 that 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.

4.2   Implications of the results

Figure 1 assesses the implications of the results by reporting the simple average of the unobserved quality ξ_{jt} for each month, 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 7 does not indicate any such pattern.

We also look at the extent to which the model generates repeat purchases. Figure 8 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 0.25% of new sales. The underlying reason why there are not more repeat purchases is that that the most important source of heterogeneity is the random coefficient on the constant term. Thus, consumers who buy camcorders are disproportionately those with high draws on the constant term. That is, they value a camcorder highly but do not distinguish very much between different models. Since consumers place low value on the improvements in characteristics, particularly relative to paying for another camcorder, the value to upgrading is low.
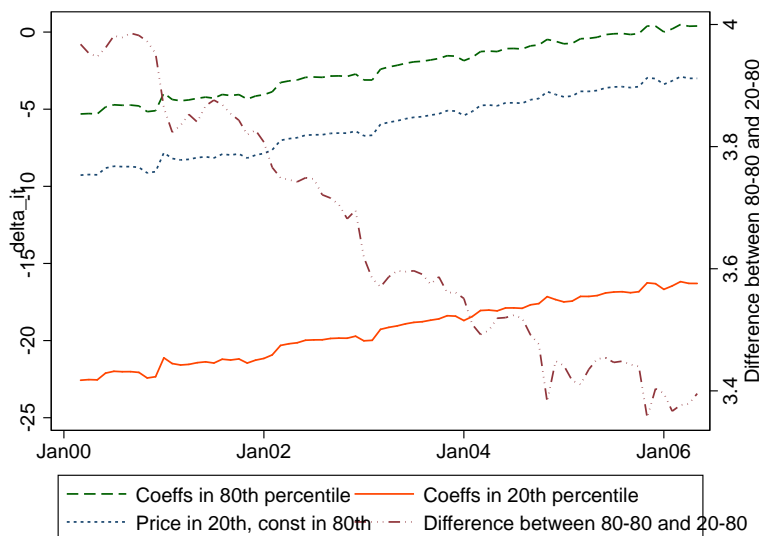
Figure 7: Average estimation error (ξ_{jt}) by month

Figure 8: Evolution of repeat purchase sales


29

Figure 9: Evolution of δ_{it} over time

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

Finally, we plot the evolution of the logit inclusive value δ_{it} in order to compare the sources of heterogeneity in our results. Figure 9 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.

Figure 10: Difference between $\delta_{i,t+1}$ and its period t prediction
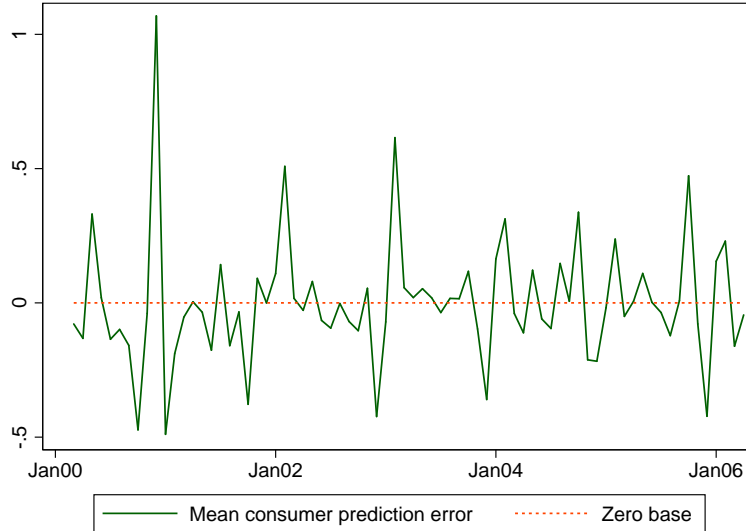
That follows for two reasons: first, the lines are relatively close to their counterparts with different price draws and second, there is little compression over time even though prices are dropping. Because it is hard to see the level of compression, we plot the difference between the 80-80 and 80-20 lines separately; the difference decreases by 15% over the sample period.

Figure 10 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 -0.474 with a standard deviation of 2.95 implying that we cannot reject the null hypothesis that the residuals are not serially correlated.

Finally, Figure 11 investigates the magnitudes of the dynamic responses by examining the time path of digital camcorder sales under three different assumptions: the time path generated by the estimated model (also the actual time path of sales), the time path that would occur if consumers in each period assumed that the market did not improve from its current state (that is, 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 all consumers started each period with no digital camcorder, but still believed at the time of purchase that the good is durable.

31

Figure 11: Evolution of digital camcorder sales under different assumptions
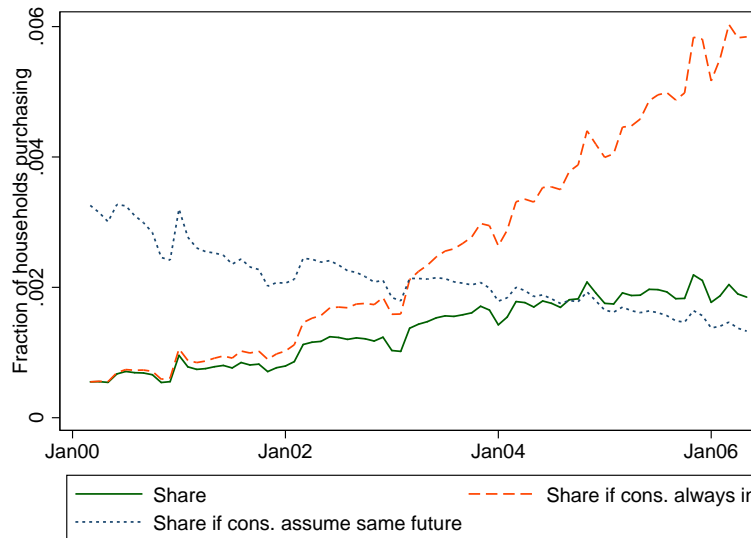
We find that dynamics, both durability and forward-looking behavior, explain a very important part of the sales path. If consumers assume that prices and qualities did not change, then sales would be somewhat declining over time, instead of growing rapidly over the sample period, as consumers would not 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 still believe at the time of purchase that they will be able to keep their camcorders, 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 specifically due to high valuation consumers remaining in the market and not just to a bigger market, as 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.[30] We compare a temporary (one-month) 1% price increase to a permanent 1% price increase. The price increase takes place at $\bar{t}$, and is unexpected before then. We assume that consumers at time $\bar{t}$ know that the temporary price increase is temporary and that the permanent price increase is permanent. For the

---

[30]As the price elasticities from the static model are all virtually 0, we focus here only on the base dynamic specification.

Figure 12: Industry dynamic price elasticities
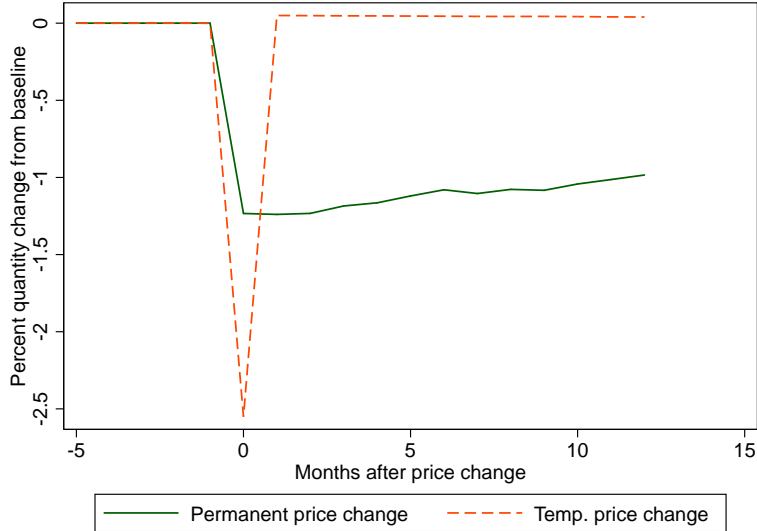
temporary increase, we compute the time $\bar{t}$ expectations of $\delta_{i,\bar{t}+1}$ using the baseline $\delta_{i\bar{t}}$ in the AR(1) equation (7); for the permanent price change, we use the realized $\delta_{i\bar{t}}$. For all specifications, we keep the estimated γ_{1i} and γ_{2i} coefficients. That is, we assume that the beliefs about the future evolution of characteristics and prices and the value function stay the same, conditional on $\delta_{i\bar{t}}$.

We start with the industry-wide elasticity, when all models change price simultaneously. Figure 12, 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 over twice as big an immediate response as a permanent one. Specifically, a temporary 1% price increase leads to a contemporaneous decrease in sales of 2.55%, while a permanent increase leads to only a 1.23% decrease. Unlike the permanent price increase, the temporary price increase causes a substantial fraction of consumers to delay purchases instead of forgoing them. Specifically, 44% of the decrease in sales at time $\bar{t}$ from the temporary price increase are recaptured over the following 12 months.

We compare these industry-wide results to the effects of a price change of a single model. Figure 13 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 modest: 2.59% versus 2.41%. The temporary and permanent single-model elasticities are closer to each other than are the industry elasticities because when only one model changes price, consumers

Figure 13: Dynamic price elasticities for Sony DCRTRV250
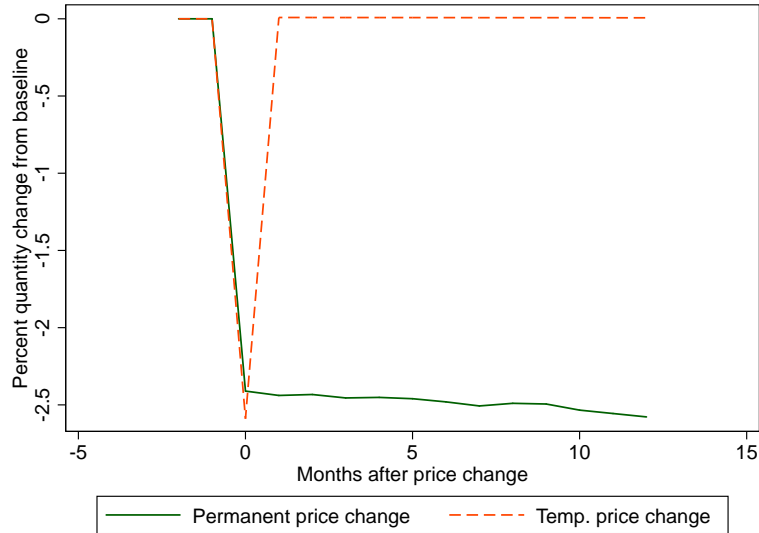
switch to another model rather than delay their purchase. Recall that in our base specification, purchasing consumers primarily value having a camcorder, as opposed to preferring particular camcorder characteristics.

Conventional wisdom is that price elasticities for individual models are higher than for the market as a whole. But strikingly, we find that the temporary short-run price elasticities are almost the same for the market and for the Sony DCRTRV250. However, the sources of the elasticity are different: when all prices increase, consumers reduce purchases by waiting for prices to drop, whereas when price rises for a single model, consumers reduce purchases by switching to other models. Thus, the long-run temporary price elasticity is much smaller for the market than for the model. 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

A classic application of demand estimation is a calculation of the value of new goods. We focus on the use of cost-of-living indices to evaluate how the value of a new good changes over time, using digital camcorders as the example. We develop a COLI from 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.[31]

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

$$\frac{I_{t+1}}{I_t} = \frac{\sum_{j=1}^{J_t} s_{jt} p_{j,t+1}}{\sum_{j=1}^{J_t} s_{jt} p_{jt}}. \tag{13}$$

We compute a BLS-style price index from (13) using the prices and market shares (based on quantities) in our data, linked over time by model names.[32] As is standard, we normalize the index to 100 for March 2000. An important challenge in constructing the BLS index is determining $p_{j,t+1}$ for models that drop out of the market. One approach is to apply the average price decline for models that appear in t and $t+1$ to models that drop out in $t+1$. This introduces the well known "new goods" problem since the price of the exiting model 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.[33]

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.[34] In general, one would construct COLIs by multiplying the price indices from (13) by the the expenditures in the sector. This is problematic for the camcorder sector since sales are rapidly growing and prices rapidly falling over time.

---

[31] We do not mean to propose our model as a method that the BLS should consider for constructing indices as it would probably be infeasible given the time constraints under which the BLS operates. We focus on indices used by Pakes (2003) and the BLS but there have been other proposals for indices in dynamic settings. For example, 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).

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

[33] To implement the Pakes index, we specify a model of log price as a linear function of model characteristics (except 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.

[34] Bureau of Labor Statistics (2007) states that "the concept of COLI provides the CPI [Consumer Price Index]'s measurement objective (p. 2)."

Figure 14: Average monthly value from camcorder market
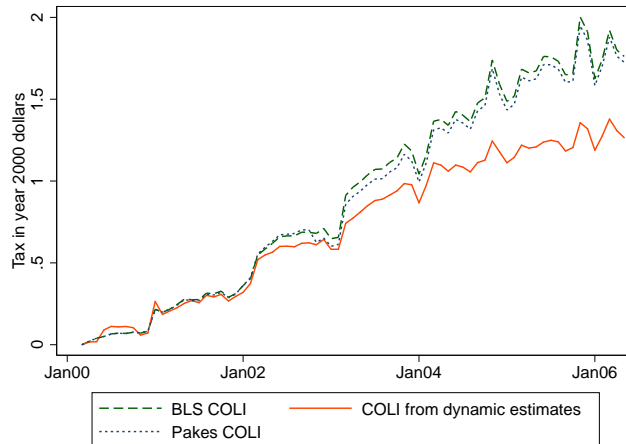
Thus, we proceed by including the outside good as a model with an invariant price in (13). We then multiply the resulting index by the share-weighted average price in the initial period ($961), divide by 100 and subtract the resulting term from $961. This new index then provides the amount that income could be reduced relative to the first period that would still allow a household to buy the same basket of models. We plot the CPI and Pakes versions of this index in Figure 14. 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 March 2006 average utility, with smaller taxes necessary for most 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 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.[35] Note that the aggregate-state-contingent taxes do not change camcorder purchase behavior in our (linear price) model. We make one final adjustment which is to assume that a consumer who buys a model 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.[36] 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 14.[37] 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. Although the slope would be different, the shape would remain the same – and different from our dynamic COLI.

5   Conclusion

This paper develops new methods to estimate the dynamics of consumer preferences for new durable goods. Our model allows for rational expectations about future model attributes, consumers with persistent heterogeneity over time, endogeneity of price, large and changing numbers of models, and the ability for consumers to upgrade to new durable goods as features improve. We estimate our model using a panel dataset 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

---

[35]Potential 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.

[36]If 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.

[37]We 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.

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 the short-run industry elasticity of demand is 2.55 for temporary price shocks but only 1.23 for permanent price shocks. Short-run elasticities for individual models for permanent price shocks are substantially larger than short-run industry-wide 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 this gain in welfare due to the "new buyer problem," which is that later adopters tend to value the product less than earlier adopters.

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, $V(f_0, \Omega) = V\left(f_0, \tilde{\Omega}\right)$.*

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 $V_t^T(f_0, \Omega)$ 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 3 as:

$$V_T^T(f_0, \Omega) = \ln\left(\exp(\delta(\Omega)) + \exp(f_0)\right).$$

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

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

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

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 $V_t^T(f_0, \Omega')$ by the inductive assumption. Thus, we have proved the inductive step.

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

$$V(f_{0t}, \Omega_t) = \lim_{T \to \infty} V_t^T(f_{0t}, \Omega_t) = \lim_{T \to \infty} V_t^T(f_{0t}, \tilde{\Omega}_t) = V(f_{0t}, \tilde{\Omega}_t).$$

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 $(f_j(\Omega'), P_j(\Omega')|\Omega), j = 1 \ldots J$ 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 kth 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 model 1 at any time period τ have some contingent flow utility $f_{1\tau k}$ and let other models 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 $m : \Re^\infty \times \Re^K \to \Re^\infty \times \Re^K$ 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 τkth 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 τkth one and flow utility of x for the τkth 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 model 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 model with a flow utility of $\delta^{max}(1 - \beta)$ and never purchasing another model 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 $\mathbb{R}^\infty$ 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.$\blacksquare$

B   Evidence on the inclusive value sufficiency (IVS) assumption

To better understand the roles of both δ and IVS, we evaluate δ in the context of a simple numerical example of our model, with one model 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.[38] If the consumer were forced to hold a model indefinitely once she bought it, the mean discounted flow utility net of price from purchasing the model would be $f/(1-\beta) - P$. Figure 15 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, 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 0.05.

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 (2012) and Hendel & Nevo (2006) but δ does not. To investigate how much bias might occur from this misspecification, Figure 16 shows the

---

[38]The AR(1) process is $f_{t+1} = 0.0005 + 0.99f_t + \nu_{t+1}$, where $\nu \sim N(0, 0.002)$; the price term is $P = 5$; the discount factor is $\beta = 0.99$; and first period flow utility is $f = -0.04$. We assume that consumers know price, the AR(1) parameters, and current f.

Figure 15: Simulated evolution of flow utility and delta in a simple example
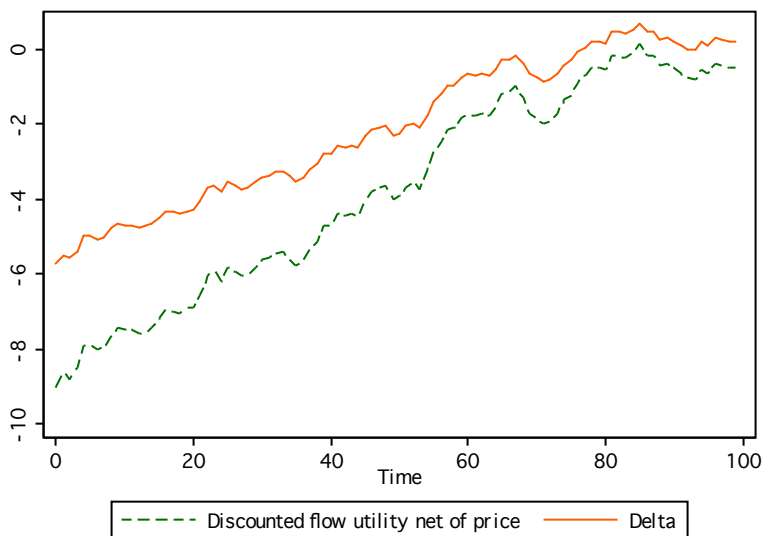
evolution of market share for a consumer who uses the true δ data generating process to make decisions, and for a consumer who optimizes assuming that δ follows an AR(1) process, jointly solving (5), (6), and (7). The two time paths of market shares are very close. We take from this the heuristic that 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 models, 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.

In addition to this Monte Carlo analysis, the paper provides several pieces of evidence in support of IVS and (7). We briefly review them here:

1. Proposition 2 in Appendix A shows that we can always rationalize the assumption in (7) with an assumption on the evolution of the underlying exogenous model characteristics.

2. We estimate a specification of our model that evaluates the impact of adding additional state variables that the consumer can use to predict future inclusive values. This exercise is similar to one performed in (see Krusell & Smith, 1998), who analyze the impact of additional moments of the income distribution as a robustness check. Given the potential importance of the number of models in determining industry evolution, we use J as an additional predictor, so that both δ and

41

Figure 16: Simulated market share with approximate and true data generating process in simple example
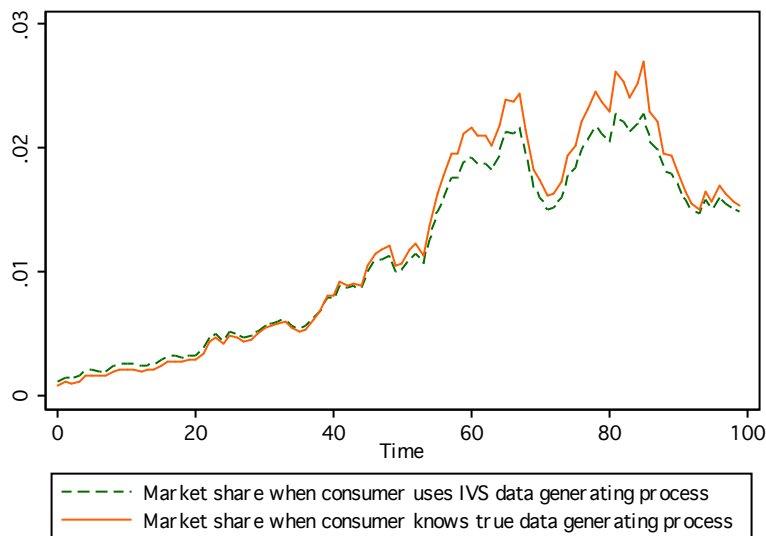
J predict δ' and J'. See Table 2, Column 1.

3. We test the validity of the AR(1) specification in (7). We construct a moment based on the assumption of no autocorrelation $E\left[\nu_t\nu_{t+1}\right] = 0$. We do not impose this moment in estimation but rather test its validity at the estimated parameters.

4. Using the estimated parameters, we graph ν_i and δ_i over time for different simulated consumers. This allows us to examine graphically whether perceptions of their future values look systematically biased or not.

Each of these exercises supports the use of IVS and (7).

C   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

42

it above the maximum bound simply place it at the maximum bound.[39] 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 $\hat{s}_{sum}\left(\bar{\alpha}_i, \vec{F}, \alpha^p, \Sigma\right)$ 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

$$g(\check{\alpha}_i) \equiv \frac{\hat{s}_{sum}\left(\check{\alpha}_i, \vec{F}, \alpha^p, \Sigma\right)\phi(\check{\alpha}_i)}{\int \hat{s}_{sum}\left(\check{\alpha}, \vec{F}, \alpha^p, \Sigma\right)\phi(\check{\alpha})d\check{\alpha}}, \tag{14}$$

and then reweight draws by

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

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 (14) 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, 10, and 12 as described in Section 2.3.

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.

---

[39]Note that so long as the minimum bound of δ_i in the computation for i is sufficiently below the minimum observed δ_i across time periods in the data and the maximum is sufficiently above the asymptotes we find from the AR(1) regression from (7), further expansion of the bounds does not improve our approximation.

Ackerberg, 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. (2012). Improving the numerical performance of BLP static and dynamic discrete choice random coefficient demand estimation. *Econometrica*, *80*, 2231–2267.

Erdem, T., Imai, S., & Keane, M. (2003). Consumer prices and promotion expectations: Capturing consumer brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and Economics*, *1*, 5–64.

Erdem, T. & Keane, M. (1996). Decision making under uncertainty: Capturing dynamic brand choice in turbulent consumer goods markets. *Marketing Science*, *15*, 1–20.

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. (2011). Does AMD spur intel to innovate more? *Journal of Political Economy*, *119*, 1141–1200.

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.

Gowrisankaran, G. & Rysman, M. (2012). A framework for modeling industry evolution in dynamic demand models. Unpublished manuscript, University of Arizona.

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. (2012). Constrained optimization approaches to estimation of structural models. *Econometrica*, *80*, 2213–2230.

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. (2012). Vertical integration and exclusivity 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. (2012). Demand for differentiated products: The case of the U.S. computer printer market. Forthcoming in Economic Inquiry.

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.

Nosal, K. (2012). Estimating switching costs for medicare advantage plans. Unpublished manuscript, University of Mannheim.

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). Autobmobile replacement: A dynamic structural approach. *RAND Journal of Economics*, *42*, 266–291.

Shcherbakov, O. (2009). Measuring switching costs in the television industry. Unpublished manuscript, Mannheim University.

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 of Economic Studies*, *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.