

Demand for Differentiated Durable Products: The Case of the U.S. Computer Printer Market*

Oleg Melnikov[†]

Department of Economics,
Cornell University

October 14, 2001

Abstract

This paper develops an estimation technique for analyzing the impact of technological change on the dynamics of consumer demand in a differentiated durable products industry. The paper presents a dynamic model of consumer demand for differentiated durable products that explicitly accounts for consumers' expectations of future product quality and consumers' outflow from the market, arising endogenously from their purchase decisions. The timing of consumers' purchases is formalized as an optimal stopping problem. A solution to that problem defines the hazard rate of product adoptions, while the nested discrete choice model determines the alternative-specific purchase probabilities. Integrating individual decisions over the population distribution generates rich dynamics of aggregate and product level sales. The empirical part of the paper takes the model to data on the U.S. computer printer market. The estimates support the hypothesis of consumers' forward-looking behavior, allowing for better demand forecasts and improved measures of welfare gains from introducing new products.

*I would like to thank Patrick Bayer, John Rust, Christopher Timmins, and especially my advisors, Steven Berry, Ariel Pakes and Martin Pesendorfer for valuable advice and general encouragement. I have greatly benefitted from discussions with Eugene Choo, Philip Haile, Jerry Hausman, Günter Hitsch, Nikolay Moshkin, Katja Seim and Nadia Soboleva. Seminar participants at Chapell Hill, Chicago, Cornell, Duke, Harvard, Toronto, and Yale provided many helpful suggestions. I am indebted to Mark Bates of PC Data, Inc. for providing me with the data without which this research would not be possible. I am grateful to Susan Olmsted for her help with administrative issues. All errors are my own.

[†]Contact information: e-mail om32@cornell.edu, phone (607) 255-6314.

1 Introduction

This paper aims to provide a theoretical and empirical description of consumers' behavior in a dynamic market for differentiated durable goods. The computer hardware industry serves as an example of the type of market I consider. Markets for computer systems and components, as well as for many other high-technology products, exhibit many striking similarities. The following features are of particular concern for an empirical economist:

- Products are differentiated in multiple dimensions in addition to price.
- Products are durables yielding consumption services over multiple periods.
- Product quality rapidly improves over time.

Empirical work on estimating differentiated product demand systems is dominated by static discrete choice models with random utility (Berry, Levinsohn, and Pakes (1995), Goldberg (1995), Nevo (200?), among others). However, markets for durable goods exhibit several features at odds with their assumption of myopic consumer behavior.

First, traditional static discrete choice models implicitly assume that consumers participate in the market every period, choosing one of the available products or an outside option. Thus, the fraction of consumers who do not buy any product in a given period could only be explained by the high value they assign to the outside alternative. For durable goods, however, purchases are made infrequently and result in a consumer's exiting the market for a significant period of time. Hence, the model needs to account for consumers' outflow from the market arising from their purchase decisions.

Second, a static framework may give a misleading picture of the aggregate sales time pattern in a dynamic industry. For many markets for durable products it is not uncommon to see rapidly improving product quality and falling prices. Given *any* reasonable set of consumer preferences over product characteristics space, static discrete choice models for such an industry will tend to predict an upward trend in aggregate sales. This fails to account for a variety of patterns commonly seen in the data.

Third, the evolution of product quality leads to the possibility of intertemporal demand substitution. By foregoing a current purchase, the consumer retains the option to buy a potentially better product in the future. Hence, the decision of when to buy may be just as important for consumers as the decision of what to buy. This trade-off cannot be accounted for in the static framework.

The purpose of this paper is to remedy these shortcomings of static discrete choice models

with respect to durable goods markets, with particular attention paid to technological change and the dynamics of adoption of new products.

The paper presents a dynamic model of consumer demand for differentiated durable products that explicitly accounts for consumers' expectations of future product quality and consumers' outflow from the market, arising endogenously from their purchase decisions. In my model, the consumer faces a sequence of static discrete choice problems over a non-stationary choice set. I show that for a subclass of random utility models (which includes the multinomial logit and the McFadden (1978) generalized extreme value model) there exists a scalar-valued sufficient statistic that determines the value of the option to postpone the purchase. This allows me to model industry evolution over a vastly reduced state space and to formalize the consumer's decision of when to buy as an optimal stopping problem. A solution to that problem defines the hazard rate of product adoptions, while the nested discrete choice model determines the alternative-specific purchase probabilities. Integrating individual decisions over the population distribution generates rich dynamics of aggregate and product level sales.

The paper proceeds to present Monte-Carlo evidence that ignoring the time dimension of consumer choice will induce a downward bias in the estimates of consumer preferences over quality, sometimes reversing the sign of the coefficients. I propose a nested three-stage estimator that allows for sequential identification of the model parameters with aggregate data from the relatively short time series.

The empirical part of the paper takes the model to data on the U.S. computer printer market. During the 90s, the industry expanded dramatically and was characterized by high rates of product entry and exit and remarkable technological innovation. These trends are indicative of the performance of the computer industry as a whole. Given the annual rate of decline in the quality-adjusted prices of more than 30%, it appears that intertemporal consumer optimization is an important feature of the market. The data set used in this study is a detailed monthly sampled scan-code level panel on more than 400 printer models. The estimates support the hypothesis of consumer's forward-looking behavior, allowing for better out-of-sample demand forecasts and improved measures of welfare gains from introducing new products.

The rest of the paper is organized as follows. Section 2 contains a review of the existing literature on the topic. Section 3 presents the model of consumer demand for durables, considers its implications, and discusses the estimation strategy. Section 4 performs a set of Monte Carlo experiments to check the viability of the estimation approach. Section 5 describes the data and reports estimation results. Section 6 concludes and suggests some directions for future research.

2 Review of the Related Literature

The model presented below builds upon three sources: discrete choice models, the durable goods literature, and stochastic growth theory.

Static discrete choice models with random utility, pioneered by McFadden's (1973) seminal paper and developed in McFadden (1978), Berry (1994), Berry, Levinsohn, and Pakes (1995), McFadden and Train (200?), among others, have been a staple of empirical research on estimating differentiated product demand systems. In these models, products are treated as portfolios of characteristics. Preferences are defined over the product characteristic space and consumers select the product that yields the highest utility value. The presence of the random utility component, which accounts for idiosyncratic differences in perceived assessment of quality, ensures that the aggregate demand is stochastic. The postulated functional form for the utility function and the error term correlation structure jointly determine product-level purchase probabilities. Thus, the random discrete choice model accounts for consumers' heterogeneity in an elegant way consistent with the utility-maximizing behavior. Recent empirical applications of the discrete choice framework include Goldberg (1995), Nevo (200?), Davis (2000).

The fundamental assumption of discrete choice models is that consumers choose a single unit from the set of available alternatives. This assumption seems to be particularly justified for purchases of durable products. Hence, most empirical IO applications of this framework deal with durable goods markets. There is also a growing literature that uses this approach to analyze various aspects of industry conduct in highly dynamic markets (Bresnahan, Stern, and Trajtenberg (1997) for office computers, Himmelberg and Olley (1998) for hard disc drives, etc.). But the same "single unit only" assumption seems to imply that consumers' behavior should be forward looking, which is at odds with the static nature of the discrete choice setting.

In the micro level consumer data, postulating a reduced form that includes current ownership proxies may help to alleviate the problem. I.e., Goldberg (1995) in her study of the U.S. auto industry uses the age of the current vehicle as an individual level demand shifter. However, for many markets, detailed consumer level data are not available and the researcher must (at best) rely on the aggregate product level data. In the latter case, ownership variables are inherently unobservable.

For stable markets where the set of products does not change much over time, the applicability of static models may be justified under assumption of a stationary distribution of consumer holdings. This effectively explains all industry sales as replacement demand. Again, for many practical applications the object of interest is the entry of new products and/or firms and other structural

industry changes, which may lead to shake-ups of consumer holdings.

The myopic assumption seems to be particularly ill-suited for studying markets for new products and industries with a rapid rate of technological change, such as the computer industry. For such diverse applications as marketing, welfare measures, and policy analysis it is essential to separate new demand from the replacement demand. The main contribution of this paper is an attempt to address these issues in a structural way consistent with a discrete choice framework.

The impact of product durability on the evolution of prices and sales has received much attention in the durable goods literature. The bulk of this research is devoted to pricing new products, a topic that attracted a considerable theoretical interest since the celebrated Coasian conjecture (Coase (1972)). Briefly, this result states that in a market for durable goods, current sales cannibalize future sales prospects and induced competition with the future will force an efficient (marginal cost) pricing even in a monopoly case. The Coasian conjecture turned out to be extremely sensitive to assumptions on the distribution of unobserved consumer preferences, as has been shown by Sobel and Takahashi (1983) and Bagnoli, Salant, and Swierzbinski (1989) in various settings.

Unfortunately, most of this literature focuses on the homogenous products markets. This can probably be attributed to the complexity of dynamic equilibrium models. Effectively, one has to solve for the entire equilibrium price path to determine prices at any given time. Moreover, consumer expectations must be self-fulfilling along this price path in any viable equilibrium. Not surprisingly, few results have been obtained for differentiated durable goods markets. To the best of the author's knowledge, the only paper that deals with pricing in the differentiated durable products oligopoly is Deneckere and de Palma (1998), but their set of assumptions seems too rigid to resemble real world markets.

There are also some theoretical and empirical results on the equilibrium prices and holdings in secondary markets for used durables (Stokey (1981), Rust (1985), Stolyarov (2000)). These papers also deal mainly with homogenous products. While secondary markets are undoubtedly an important feature of many durable goods industries, for our empirical application there is very little trade in the used computer equipment. Thus, the model presented below assumes away the market for used products.

3 The Model

This section provides an overview of the consumer choice model used in this paper. I start by considering the decision problem faced by a representative consumer in the evolving market for

durable goods. Integrating individual decisions over the population distribution generates a system of equations for aggregate sales which can be taken to the data. The section completes by discussing econometric identification of model parameters from the aggregate data.

3.1 Consumer Decision Problem

Environment. Time is discrete and denoted by $t = 0, 1, \dots$. Consumers, indexed by $i = 1, \dots, M$, buy at most one unit of a durable product in a lifetime¹. Consumer i ' state at time t is given by $S_{it} \in \{0, 1\}$, where $S_{it} = 0$ if i does not own any product at time t and $S_{it} = 1$ otherwise. The effective market size at time t , M_t , is determined by the cardinality of the set $\mathfrak{M}_t = \{i \mid S_{it} = 0\}$. Consumers with $S_{it} = 1$ are considered to be out of the market, so the discussion below only concerns first-time buyers. The set of products available to consumers at time t is denoted by \mathfrak{J}_t and may change over time.

Payoffs. In each period t , consumer $i : S_{it} = 0$ has two options: (a) to buy one of the products $j \in \mathfrak{J}_t$ available on the market this period or (b) to postpone the purchase by choosing an outside alternative. If option (b) is chosen, the consumer obtains a one period utility payoff of c . If consumer i buys a product $j \in \mathfrak{J}_t$, he obtains a terminal payoff

$$u_{ijt} = f(\mathbf{x}_j, \mathbf{y}_{jt}; \theta_i) + \varepsilon_{ijt}, \quad (3.1)$$

where:

\mathbf{x}_j is a $(1 \times K)$ vector of static product attributes;

\mathbf{y}_{jt} is a $(1 \times L)$ vector of product characteristics that may change over time (which includes price, p_{jt} , and potentially other dynamic parameters like the product age);

θ_i is a $((K + L) \times 1)$ vector of (stationary) consumer preferences over \mathbf{x} and \mathbf{y} ;

ε_{ijt} is an individual-specific random utility component drawn from some underlying distribution F_ε .

Variable u_{ijt} is interpreted as a lifetime utility value of owning a durable. The exact functional form of $f(\cdot)$ and the distinction between observed and unobserved product characteristics are not important at this point and will be discussed in detail in Section 3.3. In addition, we assume that:

[A1] $\theta_i = \theta \forall i = 1, \dots, M$ (homogenous consumer preferences over the product characteristics space);

¹This assumption will be somewhat relaxed later on to allow an exogenous inflow of consumers into the market.

[A2] ε_{ijt} terms are independent across individuals and time periods and are drawn from the generalized extreme value (GEV) distribution with the cumulative joint distribution function given by

$$F_\varepsilon(\varepsilon, \dots, \varepsilon_{Jt}) = \exp(-G(e^{-\varepsilon_1}, \dots, e^{-\varepsilon_{Jt}})), \quad (3.2)$$

where the function $G(x_1, \dots, x_J) = G(\mathbf{x})$ satisfies the following properties:

- (i) $G(\mathbf{x}) \geq 0 \forall \mathbf{x} \in \mathfrak{X}^J$;
- (ii) $G(\mathbf{x})$ is homogenous of degree one in \mathbf{x} ;
- (iii) $\lim_{x_j \rightarrow \infty} G(x_1, \dots, x_J) = \infty \forall j = 1, \dots, J$;
- (iv) for any distinct sequence (j_1, \dots, j_k) , $\partial^k G / \partial x_{j_1} \dots \partial x_{j_k} \geq 0$ if k is odd and ≤ 0 if k is even².

If $G(x_1, \dots, x_J) = \sum_j x_j$, then in a static framework assumptions A1 and A2 generate a well-known multinomial logit demand system (see McFadden (1973), Berry (1994), and Berry, Levinsohn, and Pakes (1995)). The GEV assumption also accommodates the nested logit model developed in McFadden (1981). Given assumption A1, it is convenient to decompose (3.1) as $u_{jt} = \delta_{jt} + \varepsilon_{jt}$, where δ_{jt} is the mean utility level associated with choosing product j ³.

Consumer problem. The consumer learns realizations of his idiosyncratic preferences at the beginning of each time period. Consumer decision problem is a two stage process. First, based on the realizations of ε_{ijt} 's, the consumer chooses the product j_t^* that maximizes (3.1) from among \mathfrak{J}_t . Second, the consumer decides whether to buy this product or to postpone his purchase until the next period. The latter decision generates an optimal stopping problem on the consumer's side. Denoting the time period in which consumer decides to buy a product by τ , the consumer's optimization problem can be written as

$$J(I_t) = \max_{\tau} \left[\sum_{k=t}^{\tau-1} \beta^{k-t} c + \beta^{\tau-t} E_t \max_{j \in \mathfrak{J}_\tau} u_{j\tau} \right] \quad (3.3)$$

where $\beta \in [0, 1)$ is a common discount factor and $E_t[\cdot] = E[\cdot | I_t]$ denotes a conditional expectation given the information set I_t available to the consumer at time t .

To cast (3.3) in a standard dynamic programming framework, we need to define a distribution of $\max_{j \in \mathfrak{J}_t} u_{jt}$ and to specify an information set available to consumer at time t . To this end, let

$$v_t = \max_{j \in \mathfrak{J}_t} u_{jt}. \quad (3.4)$$

²See McFadden (1978) for an extensive discussion.

³For the logit model it is easy to verify that u_{jt} is distributed type I extreme value with the mode of δ_{jt} . The mean utility from choosing product j across the population is given by $\delta_{jt} + \gamma$, where γ is the Euler's constant.

The distribution of v_t is characterized by the following proposition.

Proposition 1. Under assumption A2, v_t is distributed type I extreme value with the cumulative distribution and density functions given by

$$\begin{aligned} F_v(u; r_t) &= \exp(-e^{-(u-r_t)}); \\ f_v(u; r_t) &= e^{r_t} \exp(-e^{-(u-r_t)} - u) = e^{-(u-r_t)} F_v(u; r_t), \end{aligned} \tag{3.5}$$

where r_t is the mode of the distribution given by

$$r_t = \ln G(\exp(\delta_{1t}, \dots, \delta_{J_t,t})) = \ln R_t. \tag{3.6}$$

Proof. See Appendix A.

In other words, r_t is a scalar-valued sufficient statistic for the distribution of future payoffs.

Given equations (3.4)-(3.6), the recursive formulation for the consumer decision problem (3.3) is

$$J(v_t, I_t) = \max \{v_t, c + \beta E_t J(v_{t+1}, I_{t+1})\}, \tag{3.7}$$

where v_t is drawn from (3.5).

Industry Evolution. From Proposition 1, although the expectation in (3.7) should be taken over the changing set of products, attributes, and prices, the distribution of the maximum utility that can be obtained from participating in the market is completely characterized by a scalar parameter r_t , the mode of the extreme value distribution. r_t effectively serves as an index of the consumption value of the market as a whole. Given that I_t is a common knowledge, assumption A1 ensures that all consumers have identical expectations on the future distribution of r_t .

Given the dimensionality of the product characteristic space and the number of competing products, for a typical market it is computationally infeasible to model the process generating r_t 's directly. We assume, therefore, that this process can be adequately described by a reduced set of state variables describing the state of the market at time t . Using an approach taken in most stochastic growth models, we assume that the evolution of the mean utility level can be characterized by a homogenous Markov process with the transition density $\Phi(r_{t+1} | r_t, \theta_r)$. Correspondingly, E_t is the conditional expectation operator generated by Φ .

Postulating a functional form for the transition density $\Phi(r_{t+1} | r_t, \theta_r)$ completes specification of the problem. Below we assume that r_t follows a diffusion process

$$r_{t+1} = \mu(r_t) + \sigma(r_t)\nu_{t+1}, \tag{3.8}$$

where ν_t are i.i.d. standard normals and functions $\mu(r)$ and $\sigma(r)$ satisfy the following:

Model	Diffusion Type	Functional Form
stochastic linear growth	random walk with drift	$r_{t+1} = r_t + \gamma + \sigma\nu$, $\gamma \geq 0$
mean-reverting process	stable AR(1) process	$r_{t+1} = f + \alpha r_t + \sigma\nu$, $ \alpha < 1$
bounded returns to technological progress	generalized Wright-Fisher	$r_{t+1} = r_t^\gamma a^{1-\gamma} + \sigma(r_t)\nu$, $0 < \gamma < 1$, $\lim_{r \rightarrow a} \sigma(r) = 0$

Table 3.1: Admissible Functional Forms for the Quality Generating Process

[F1] $\mu(r)$ and $\sigma(r)$ are continuous and almost everywhere differentiable;

[F2] $0 < \sigma(r) < \infty \forall r \in \mathfrak{R}$;

[F3] r_t is a weak submartingale: $\mu(r_t) \geq r_t$;

[F4] $\lim_{n \rightarrow \infty} \beta^n \mu^n(r) < \infty$ where $0 \leq \beta < 1$, $\mu^0(r) = \mu(r)$ and $\mu^n(r) = \mu(\mu^{n-1}(r))$.

Alternative specifications for the quality-generating process r_t are discussed in the extensions section. Notice, however, that the functional form (3.8) is quite flexible and encompasses many specifications used in the literature on economic growth and technological change. Table 3.1 lists several growth models compatible with the formulation (3.8).

Bellman Equation and Policy Function. Under the above assumptions, we can rewrite (3.7) as

$$J(v_t, r_t) = \max \{v_t, c + \beta E [J(\tilde{v}_{t+1}, \tilde{r}_{t+1}) | r_t]\}, \quad (3.9)$$

where v_t is distributed extreme value with the mode of r_t .

(3.9) is a standard optimal stopping problem, similar to the classic job search model of Stigler (1961) and McCall (1970). From the optimal stopping theory (Stokey and Lucas (1989), Bertsekas (1995)), a general solution to this class of problems is given by a subset of state space where it is optimal to stop. In our case, the stopping set is given by $\mathcal{S} = \{(v, r) | v \geq c + \beta E [J(\cdot) | r]\}$. Clearly, $J(\cdot)$ is increasing in its first argument. From (3.5), the distribution of v is stochastically increasing in r . By assumptions F1 and F3, r_{t+1} is monotone and stochastically non-decreasing in r_t . Assumption F4 ensures that the second term inside the brackets in (3.9) is bounded away from infinity for any r_t ⁴. It follows that:

(i) $J(v, r)$ is monotone and (weakly) increasing in both arguments;

(ii) $E [J(\cdot) | r]$ is monotone and non-decreasing in r ;

(iii) the stopping set $\mathcal{S} = \{v | v \geq W(r)\}$, where $W(r)$ is the reservation utility level defined by

⁴Notice that it does not imply that $J(\cdot)$ is bounded; it suffices to require that expected future returns do not grow faster than geometric series with modulus β^{-1} .

$$W(r_t) = c + \beta E [J(\tilde{v}_{t+1}, \tilde{r}_{t+1}) | r_t]. \quad (3.10)$$

Due to the structure of the optimal policy,

$$J(v, r) = \begin{cases} v & \text{if } v \geq W(r) \\ W(r) & \text{otherwise} \end{cases}.$$

Thus, the reservation utility level must satisfy the functional equation

$$W(r) = c + \beta \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \max(v, W(s)) dF(v | s) d\Phi(s | r), \quad (3.11)$$

where

$$\begin{aligned} F(v | s) &= \exp(-\exp(-(v - s))), \\ d\Phi(s | r) &= \phi\left(\frac{s - \mu(r)}{\sigma(r)}\right) dr, \end{aligned}$$

and $\phi(\cdot)$ is the standard normal density.

3.2 Sales Dynamics and Aggregation

We now consider the dynamics of sales generated by the model.

Demand Structure. The consumer will buy some product at time t if and only if $v_t > W(r_t)$, hence, the probability of postponing the purchase until the next period π_{0t} implied by the model is simply $F_v(W(r_t), r_t)$:

$$\pi_{0t}(r_t) = \mathbb{P}\{S_{i,t+1} = 0 | S_{it} = 0, r_t\} = \exp(-\exp(-(W(r_t) - r_t))). \quad (3.12)$$

The complement to the search probability π_{0t} , $h(r_t) = 1 - \pi_{0t}(r_t)$, defines the individual hazard rate of the product adoption. Product-specific purchase probabilities are determined by

$$\begin{aligned} \pi_{jt}(r_t, \cdot) &= \mathbb{P}\{u_{jt} \geq u_{kt} \forall k \neq j; u_{jt} \geq W(r_t)\} \\ &= \mathbb{P}\{u_{jt} \geq W(r_t) | u_{jt} \geq u_{kt} \forall k \neq j\} \mathbb{P}\{u_{jt} \geq u_{kt} \forall k \neq j\} \\ &= \mathbb{P}\{v_t \geq W(r_t)\} \mathbb{P}\{u_{jt} \geq u_{kt} \forall k \neq j\} \\ &= h(r_t) \frac{\exp(\delta_{jt}) G_j(e^{\delta_{j1}}, \dots, e^{\delta_{j_t, t}})}{G(e^{\delta_{j1}}, \dots, e^{\delta_{j_t, t}})} = h(r_t) \frac{\exp(\delta_{jt}) G_j(\cdot)}{R_t}, \end{aligned} \quad (3.13)$$

where the last line follows from the functional form for unconditional choice probabilities in the GEV model (see McFadden (1981)), and $G_j(\cdot)$ denotes the partial derivative of $G(\cdot)$ with respect to j -th argument (for instance, in the multinomial logit model $G_j(\cdot) = 1 \forall j$).

Hazard rate. From (3.13), the consumer's reservation utility level only affects his individual choice probabilities through the hazard rate $h(r_t)$. From (3.12) the hazard rate is a monotone transformation of $Y(r_t) = W(r_t) - r_t$. Thus, it is convenient to derive an explicit functional equation for $Y(r_t)$. From Proposition I, $v_t \stackrel{d}{=} r_t + \varepsilon$, where ε is a zero mode extreme value deviate. Using this property, one can rewrite (3.10) as

$$\begin{aligned} W(r_t) &= c + \beta E[\max(r_{t+1} + \varepsilon, W(r_{t+1})) | r_t] \\ &= c + \beta E[r_{t+1} + \max(\varepsilon, W(r_{t+1}) - r_{t+1}) | r_t], \end{aligned} \quad (3.14)$$

Subtracting r_t from both sides of (3.14) yields

$$Y(r_t) = c + \beta E[r_{t+1} | r_t] - r_t + \beta E[\max(\varepsilon, Y(r_{t+1})) | r_t]. \quad (3.15)$$

Letting $M(z) = E \max(z, \varepsilon)$ and utilizing the postulated functional form (3.8) for the law of motion of r_t produces

$$Y(r) = c + \beta \mu(r) - r + \beta \int_{-\infty}^{+\infty} M(Y(x)) \phi\left(\frac{x - \mu(r)}{\sigma(r)}\right) dx. \quad (3.16)$$

It is easy to show that under assumptions made above $Y(r)$ is a monotone, decreasing function of r . Numerically, $Y(r)$ can be computed as a fixed point of the mapping (3.16). Notice that the function $M(z)$ does not depend on the state variable r and its computation can be factored out from the value function iterations. Hence, each stage of the iterative procedure only requires a computation of the one-dimensional integral.

With $Y(r)$ at hand, the hazard rate $h(r; c, \beta, \theta_r)$ can be computed from (3.12). The properties of the hazard rate function are illustrated on Figure 1. For this example, r_t is specified as a homoscedastic random walk with drift, i.e. $\mu(r) = r + \gamma$, $\sigma(r) = \sigma$. The probability of adoption follows an S-shaped pattern familiar from the empirical marketing literature on the diffusion of innovations (see Mahajan, Muller, and Bass (1990) for a review). At low levels of r , the value of participating in the market is low relative to the outside alternative and the probability of adoption is close to zero. As time goes by, the quality of the mean product on the market improves, thus increasing the probability of the purchase. Since in this example the quality generating process is unbounded, eventually the average quality rises to the level where the gain from waiting does not justify its cost (as the quality growth is linear, but discounting is geometric). Therefore, once the market quality exceeds a certain level, adoption is nearly immediate. Products' diffusion in the marketplace occurs while r_t traces the interval in between these two extremes.

Model parameters affect the speed of the diffusion process. The discount factor, β , has the most significant impact on the adoption rate. Patient consumers are more likely to postpone their




Figure 3.1: Hazard Rate of Product Adoptions $h(r; c, \beta, \gamma, \sigma)$.

purchase in anticipation of a higher value in the future. Higher variance, σ^2 , lowers the probability of adoption since the option value of waiting increases with the market volatility. Higher drift has a similar effect, but its impact on the time pattern of sales is ambiguous: the probability of adoption at every given level of r_t is lower, but r_t itself grows faster. Finally, c essentially serves as a scaling parameter, shifting the whole hazard rate function.

Aggregation. Integrating individual purchase probabilities (3.13) over the population distribution of (c, β) yields the aggregate dynamic demand system. For the sake of simplicity, I assume that $c_i = c$, $\beta_i = \beta \forall i$ ⁵.

Recall that consumer i 's state at time t is $S_{it} \in \{0, 1\}$, where $S_{it} = 0$ if i does not own any product at time t and $S_{it} = 1$ otherwise. From (3.12), the transition of consumer states is governed by a Markov matrix $H : \{0, 1\} \rightarrow \{0, 1\}$ specified as

$$H_1(r_t) = \begin{bmatrix} \pi_{0t}(r_t) & h(r_t) \\ 0 & 1 \end{bmatrix}. \quad (3.17)$$

Notice that the model can be readily extended to incorporate exogenous replacement decisions, that is, the product "breaks down" with an exogenous probability q in which case consumer $i : S_{it} = 1$

⁵In principle, it is straightforward to generalize the model to allow for heterogeneity in consumers' dynamic parameters. Aggregate market shares can be obtained by numerical integration of individual consumer's decisions. The difficulty arises due to the fact that with heterogeneity, the distribution of consumers' holdings is a function of their types. It appears that the data set that I use in the empirical part of this paper is not rich enough to reliably identify this function.

returns to the market. For this case, the state transition matrix is

$$H_2(r_t) = \begin{bmatrix} \pi_{0t}(r_t) & h(r_t) \\ q & 1 - q \end{bmatrix}. \quad (3.18)$$

Let $\mu_t = \mathbb{P}\{S_{it} = 0\}$, which will be referred below as the participation rate. Given μ_t , the set of search probabilities $\pi_{0t}(r_t)$, and the transition matrix H_2 , the participation rate evolves over time according to the Kolmogorov-Chapman equation

$$\mu_{t+1} = \mu_t \pi_{0t}(r_t) + q(1 - \mu_t). \quad (3.19)$$

Adjusting conditional choice probabilities (3.13) for the participation rate produces the following equations for the market shares

$$s_t = \mu_t h(r_t); \quad (3.20)$$

$$s_{jt} = \mu_t h(r_t) \frac{\exp(\delta_{jt}) G_j(\cdot)}{R_t}. \quad (3.21)$$

Combining (3.19)-(3.21) and the definition of the hazard rate

$$h(r_t) = 1 - \exp(-\exp(-Y(r_t))) \quad (3.22)$$

generates a system of equations that describes the evolution of aggregate sales over time.

The above system of equations implies a hierarchical two-level aggregate demand system. Borrowing the terminology from the almost ideal demand system literature (Deaton and Muellbauer (1980), Hausman (1997)), the lower-level demand specification determines relative market shares of individual products based on the differences in their quality. From (3.21),

$$\ln s_{jt} - \ln s_t = \delta_{jt} + \ln G_j(\cdot) - r_t.$$

which is similar to the traditional discrete choice demand equation with the fixed time effects. The upper-level demand equations (3.19), (3.20), and (3.22) are generated by the solution to the consumers' intertemporal optimization problem and describe the dynamics of aggregate industry-wide sales.

3.3 Econometric Specification

In this section we outline the estimation of the model from the aggregate data.

We assume that econometrician observes the following:

Q_{jt} , aggregate sales of product j in period t ;

M_t , total market size in period t (the parsimonious assumption in the empirical literature is to identify M_t with the total number of households in a given period);

x_j , vector of static product attributes;

y_{jt} , vector of time varying choice-specific factors (which includes the average prices, p_{jt} , and potentially some other observable demand shifters).

Also, let $Q_t = \sum_j Q_{jt}$ and define the absolute market shares s_{jt}^a as Q_{jt}/M_t , relative market shares as $s_{jt}^r = Q_{jt}/Q_t$, and the total "industry share" s_t as Q_t/M_t . Our task is to identify consumer preferences, parameters of the transition density θ_r , value function parameters $\theta_v = (c, \beta)$, initial participation rate μ_0 , and replacement probability q from a panel of observations on s_{jt} , x_j , and y_{jt} .

For simplicity, I will assume a fixed effects specification for the consumer utility function and a logit error structure, i.e.

$$u_{jt} = \xi_j + y_{jt}b + \varepsilon_{jt} = \delta_{jt} + \varepsilon_{jt},$$

where ε_{jt} are i.i.d. extreme value deviates.

The model can be potentially estimated using Rust (1987) nested fixed point maximum likelihood algorithm. However, given the hierarchical nature of the demand system, it seems easier to employ the following three-stage procedure.

First, we identify static parameters. Combining (3.21) with the logit assumption, the probability of buying a product j at time t is

$$\pi_{jt} = \mu_t h_t(r_t; \theta_r, \theta_v) \frac{e^{\delta_{jt}}}{R_t}, \quad (3.23)$$

where $R_t = \sum_k e^{\delta_{kt}}$, while the aggregate probability of making a purchase at time t is

$$\pi_t = \mu_t h_t(r_t; \theta_r, \theta_v). \quad (3.24)$$

Taking log of the ratio π_{jt}/π_t and π_t yields⁶

$$\ln(\pi_{jt}/\pi_t) = \delta_{jt} - r_t = \xi_j + y_{jt}b - r_t, \quad (3.25)$$

Replacing the odds ratio π_{jt}/π_t by its sample analog s_{jt}^r , b , ξ_j , and r_t can be estimated from (3.25) by OLS up to a single fixed effect. I take OLS estimates of the time effects, \hat{r}_t , as the estimates of the actual realizations of the quality-generating process.

⁶Alternatively, one can use

$$\ln \pi_{jt} - \ln \pi_{1t} = \delta_{jt} - \delta_{1t} = (d_j - d_1) + (y_{jt} - y_{1t})b,$$

which identifies consumer preferences up to a single fixed effect d_1 and construct r_t time series using its definition.

Second, given the \hat{r}_t 's, $\hat{\theta}_r$ can be obtained from the postulated functional form of $\Phi(\hat{r}_{t+1}|\hat{r}_t; \theta_r)$ using maximum likelihood.

Finally, for identification of dynamic parameters, we turn to the aggregate data. Given θ_v and $\hat{\theta}_r$ from the second stage, the value function (3.16) and implied search probabilities $\pi_{0t}(\hat{r}_t; \hat{\theta}_r, \theta_v)$ can be computed from (3.12). Denote the number of consumers who are shopping for a new product by $N_t = \sum_i^{M_t} (1 - S_{it}) = M_t \mu_t$. Given θ_v , q , and N_t , the expected level of aggregate sales \hat{Q}_t is given by

$$\hat{Q}_t(\theta_v, q, N_t; \hat{r}_t, \hat{\theta}_r) = N_t h(\hat{r}_t; \hat{\theta}_r, \theta_v).$$

From the Kolmogorov-Chapman equation (3.19)

$$N_{t+1} = N_t \pi_{0t}(\hat{r}_t; \hat{\theta}_r, \theta_v) + q(M_t - N_t) + (M_{t+1} - M_t),$$

where the first term gives the number of consumers who have chosen to continue searching in the previous period, the second term is the replacement demand inflow, and the last term accounts for exogenous change in the market size. Combining two equations above, remaining parameters $\theta_d = \{\theta_v, q, \mu_0\}$ can be found by fitting predicted sales to the data with the moment condition

$$\hat{Q}_t(\hat{r}_t; \theta_d, \hat{\theta}_r) = Q_t. \tag{3.26}$$

4 Monte Carlo Experiments

In this section we perform a set of Monte-Carlo experiments to illustrate the applicability of the method.

As was shown above, demand in a given period nontrivially depends on consumer expectations of the future mean product quality. Correspondingly, correct specification of the quality generating process is important. Unfortunately, for a realistic market it is unlikely that any specification of the industry dynamics will lead to an analytical form for the density $\Phi(r_{t+1}|r_t)$. Hence, we would like to check the robustness of estimates to specification errors.

The simplest (and perhaps most intuitive) specification for the quality generating process r_t is the stochastic linear growth model which assumes that r_t follows a random walk with drift:

$$r_{t+1} = r_t + \gamma + \sigma_r \nu_{t+1}, \tag{4.1}$$

where $\nu_{t+1} \sim N(0, 1)$. A prime example of such a growth is "Moore's law" for microprocessor industry, an empirically stable relationship showing that the capacity (and, consequently, speed)

of computer microchips doubles roughly every 18 months⁷.

As a benchmark, consider the following highly stylized model of the Moore's Law industry where the quality generating process abides by the linear stochastic growth model (4.1) almost exactly. Products have two characteristics, the fixed effect ξ_j and the composite price-adjusted quality index X_{jt} . The number of products does not change over time ($J_t = J$). X_{jt} (exogenously) obeys Moore's law and evolves as

$$X_{j,t+1} = (1 + l)X_{jt}e^{\nu_{jt}},$$

where l is the growth rate and $\nu_{jt} \sim N(0, \sigma_x^2)$. Consumers have a diminishing marginal utility of quality specified as

$$u_{jt} = \xi_j + \alpha x_{jt} + \varepsilon_{jt}, \quad (4.2)$$

where $x_{jt} = \ln X_{jt}$ and ε_{jt} is the i.i.d. logit error term. Correspondingly, x_{jt} follows a random walk:

$$x_{j,t+1} = x_{j,t} + l + \nu_{j,t+1}. \quad (4.3)$$

Market size is equal to $M_t = M$ and is observed by the econometrician, so the expected sales level of product j is $Q_{jt} = M\pi_{jt}$. Our goal is to recover ξ_j , α , and dynamic parameters (l, σ, c, β, q) from the panel of observations on Q_{jt} and x_{jt} .

First, we need to translate (4.2) and (4.3) into model primitives. The state of the industry at time t is given by $r_t = \ln R_t$, where $R_t = \sum_j \exp(\xi_j + \alpha x_{jt})$. If random shocks to product quality levels were identical across products (i.e. $\nu_{j,t} = \nu_t$), the model would abide by (4.1) exactly as

$$\begin{aligned} R_{t+1} &= \sum_j \exp(\xi_j + \alpha x_{j,t+1}) = \sum_j \exp(\xi_j + \alpha(x_{jt} + l + \nu_{t+1})) \\ &= \exp(\alpha l + \alpha \nu_{t+1}) R_t. \end{aligned} \quad (4.4)$$

Unfortunately, in this case the model would not be identified since quality differentials between products (and, accordingly, relative market shares) would not change over time. If shocks to the product quality are independent across the products, observed relative market shares would change over time, but the quality generating process would not have an exact analytical representation,

⁷In 1965, future Intel's co-founder Gordon Moore has noticed that the the number of transistors per square inch on integrated circuits had doubled every year since they were invented and predicted that this trend would continue for the foreseeable future. This rate of increase, which later became known as Moore's Law, is possibly the best known stylized fact about the computer industry. Industry experts expect that Moore's law will continue to hold for at least another decade. Similar relationships were documented in manufacturing of computer memory and hard disk drives, telecommunications, consumer electronics, and other high-tech industries.

as the sum of lognormals does not have an explicit functional form. The best that can be said (as easy to verify) is that

$$\begin{aligned} E[R_{t+1} | R_t] &= \exp(\alpha l + \alpha^2 \sigma_x^2 / 2) R_t; \\ \text{Var}[R_{t+1} | R_t] &\leq \exp(2\alpha l + \alpha^2 \sigma_x^2) (\exp(\alpha^2 \sigma_x^2) - 1) R_t^2. \end{aligned} \tag{4.5}$$

To justify using representation in (4.1) as an approximation, one could appeal either to the large sample distribution theory or to numerical arguments. Appendix B shows that in a large sample (J and T jointly converge to infinity) $\alpha l + \alpha^2 \sigma_x^2 / 2$ is a consistent estimator for $E[\ln R_{t+1} - \ln R_t]$. In a finite sample, if l and σ_x^2 are small (close to zero, which is to be expected in the high frequency data), (4.1) is reasonably accurate due to the approximate linearity of the exponential function near zero.

The simulated data set for Monte Carlo experiments below was generated as follows. The number of products was set at $J = 3$. Fixed effects were generated as independent draws from $N(0, \sigma_\xi^2)$. Consumers' preference parameter α and dynamic parameters $(l, \sigma_x, c, \beta, q)$ were calibrated and an implied hazard rate function $h(r_t; \cdot)$ was computed. Initial values for x_{jt} were obtained as independent draws from $N(m_x, \sigma_x^2)$. In principle, the choice of m_x is arbitrary. However, an "ideal" data set would trace the evolution of the industry from its inception until maturity. From Section 3.2, the linear growth rate model (4.1) has somewhat extreme predictions concerning the relationship between the level of r_t and the hazard rate of product adoptions. Specifically, within given tolerance levels (ε, δ) the industry exhibits nontrivial dynamics on the compact subset of the state space $[r_{\min}, r_{\max}]$ defined by $h(r_{\min}) = \varepsilon$, $h(r_{\max}) = 1 - \delta$. Thus, we want the industry to initiate at r_0 near r_{\min} . Given the postulated distribution for ξ and \mathbf{x}_0 , $E r_0 = \sum_j \exp(\xi_j + \alpha x_{j0}) = J \exp(\alpha m_x + (\sigma_\xi^2 + \alpha^2 \sigma_x^2))$. I choose m_x that solves $E r_0 = r_{\min}$.

Given the set of initial values and calibrated model parameters, the evolution of individual product characteristics x_{jt} and consumers' demand Q_{jt} was simulated for $T = 200$ periods in accordance with (4.3) and (3.19)-(3.22), correspondingly. Since precise aggregate data will rarely be available to the econometrician, as they are commonly derived by extrapolating measurement results for a subsample of the market, I also subject the simulated data to the heteroscedastic measurement error specified as $Q_{jt} = Q_{jt}^* \exp(\omega_{jt})$, where $\omega_{jt} \sim N(0, \sigma_m^2)$.

After obtaining the simulated panel of observations on (Q_{jt}, x_{jt}) , model parameters can be recovered by applying the algorithm described in the previous section. Estimation results are reported in Table 4.1. Preference parameter estimates are contrasted with the static logit model

$$\ln Q_{jt} - \ln Q_{0t} = \xi_j + \alpha x_{jt} + \varepsilon_{jt}, \tag{4.6}$$

Parameter	True Value	Logit		Full Model		
		estimate	std.err.	estimate	std.err.	
<i>Static Parameters</i>						
ξ_1	(fixed effects)	1.3802	-14.3611	0.0874	.	.
ξ_2		1.6651	-13.5349	0.0871	.	.
ξ_3		4.1457	-2.5378	0.1039	.	.
$\xi_1 - \xi_3$		-2.7654	-11.8233	0.1358	-2.8072	0.0387
$\xi_2 - \xi_3$		-2.4806	-10.9971	0.1356	-2.5087	0.0365
α	(preference on quality)	1.0000	-0.0590	0.0077	0.9963	0.0044
R^2				0.947		0.998
<i>Dynamic Parameters (full model)</i>						
γ	(r_t drift, $\alpha l + \alpha^2 \sigma_x^2 / 2$)	0.1312	.	.	0.1258	0.0191
σ_r^2	(r_t variance)	0.0625	.	.	0.0727	0.0073
Starting Value:			Near θ_0		Random	
β	(discount factor)	0.9000	0.8901	0.0352	0.8699	0.0645
q	(replacement probability)	0.0500	0.0507	0.0012	0.0507	0.0015
c	(outside good payoff)	0.0000	-0.0086	0.0530	-0.0085	0.0322
μ_0	(participation rate at $t = 0$)	0.9500	0.8836	.	0.8934	.

Table 4.1: Monte Carlo Parameter estimates

which obtains after normalizing utility of the outside good to zero (see Berry (1994) for an excellent discussion).

The first part of Table 4.1 reports estimates of the static preference parameters. Notice that the static logit model potentially identifies all fixed effects (since it normalizes utility of the unobserved outside good consumption), while (3.19)-(3.22) requires normalizing one of the fixed effects to zero. The static logit correctly estimates that $\xi_1 < \xi_2 < \xi_3$, but apart from ranking alternatives it performs poorly. The preference parameter α is badly underestimated and its sign is reversed. Essentially, the logit model attributes all the variance in observed market shares to differences in fixed effects, missing the market dynamics. In contrast, first stage estimates from (3.25) are sharp and close to true parameter values. This is hardly surprising: in the dynamic model, relative market shares and total industry sales could be stable while the quality of all products steadily improves over time. But given any positive value of α , the static logit must predict that eventually the share of the outside good should converge to zero. Since this does not happen in the panel, the logit explains the data by driving the consumer taste for quality towards zero.

After obtaining static parameters, estimates of r_t realizations can be computed up to a constant. Figure 4.1 compares true realizations of r_t in the simulated data and their estimates \hat{r}_t . Since the dynamic model correctly identified preference parameters, it is not surprising that \hat{r}_t tracks realizations of the quality generating process rather well. The difference in levels of r_t and \hat{r}_t is due to the omitted fixed effect. For the above data, $T^{-1} \sum (\hat{r}_t - r_t) = 4.1727$ (standard error




Figure 4.1: Simulated quality generating process r_t : realizations vs. estimates

0.0842), while the true value of ξ_3 is 4.1457. This shift does not affect the estimates of transition density parameters, since they are defined over differences in r_{t+1} and r_t . The estimates of r_t drift γ and variance σ_r^2 are reported in 4.1. Projection of the model primitives α , l and σ_x^2 onto γ and σ_r^2 poses some difficulties as was described above. From Appendix B, as $J \rightarrow \infty$, $\text{plim } T^{-1} \sum (r_{t+1} - r_t) = \gamma = \alpha l + \alpha^2 \sigma_x^2 / 2$ but in a finite sample this is a biased estimate of the process drift. For variance, I can only derive an upper bound for σ_r^2 . Correspondingly, "true values" for γ and σ_r^2 were obtained by bootstrapping r_t process with fixed values for x_0 , ξ , and (l, σ_x^2) .

The bottom line of Table 4.1 reports third stage estimates of remaining dynamic parameters obtained by applying a Nelder-Mead optimization routine for minimizing the norm of (3.26). To check whether the GMM objective function is numerically well-conditioned, two sets of estimates are reported. The first set of estimates is obtained by initiating the minimization procedure in the vicinity of true values. For the second set, starting values were generated at random. Both minimizations were performed using identical convergence tolerance levels. Both sets of estimates are close to each other and to true parameter values, which suggests that the estimator is numerically well behaved. The only troublesome parameter is the initial participation rate μ_0 . Recall that $1 - \mu_t$ denotes the fraction of consumers who bought the product in the past. As such, μ_t is an integrated process whose dependence on μ_0 diminishes with time, until we arrive to the stationary distribution (towards the tail of the simulated data), which, by definition, does not depend on μ_0 at all. Thus, regardless of the length of time series, very few observations carry the information that helps to




Figure 4.2: Simulated industry sales Q_t : realizations vs. fitted values

identify μ_0 . By the same token, however, the estimates for remaining parameters should not be strongly affected by the error in μ_0 , which is indeed what we find. As an alternative to optimizing over μ_0 , the estimates reported in Table 4.1 were obtained by choosing the level of μ_0 which fits the first observation exactly given the vector of other parameters.

As an overall measure of the model fit, figure 4.2 compares simulated total industry sales Q_t with their fitted values.

5 An Empirical Example: The U.S. Computer Printers Market

In this section we estimate the model described above using the data on the U.S. computer printer market. Section 5.1 provides a brief description of the market and analyzes relevant industry trends. Section 5.2 describes the data and gives descriptive statistics. Section 5.3 reports the estimates and contrasts their implications with these of the static model.

5.1 The Industry

Description. The U.S. computer printer industry comprises a large segment of the overall computer market and may serve as a good example of the fast technical change typical for computer industry as a whole. In 1998, more than 14 millions general purpose personal computer (PC) printers were

sold. The industry generated \$4.5 billion in sales of new PC printers alone⁸ and accounted for 21% of total retail sales of computer products and accessories, trailing only desktop systems sales. With the advent of the inexpensive ink jet technology and following the general growth of PC sales, the industry expanded dramatically. During the last five years, the estimated annual unit sales growth rate was at 12%⁹.

Product Differentiation. PC printers are differentiated in multiple dimensions. Qualitatively, the most visible source of printer differentiation is the technology they are based upon. Two main types of PC printers are inkjets and lasers. Inkjet printers, the dominant type, generate images by spraying an ultrafine pattern of ink dots on the paper. In recent years, virtually all inkjets available on the marketplace support color printing. Different hues are obtained by combining cyan, magenta, and yellow inks from what is often termed a CMY cartridge, for the three basic colors. All inkjets can use a separate black cartridge for black-and-white printing, but only so-called CMYK models (K for black) hold both color and black cartridges at the same time. Some inkjets include even more inks for handling complex graphic images and color photographs. Laser printers form an image by transferring toner (powdered ink) to paper as it passes over a photoconductor drum¹⁰. Although the average price of color lasers has fallen tremendously over the past few years, they remain too expensive for the mass market, so laser printers aimed at home use are strictly black-and-white.

For monochrome images, as a rule, lasers outperform inkjets in terms of both printing speed and quality. However, the performance gap between two competing technologies has narrowed. This is due in part to advances in the computer technology. Since the laser technology requires forming an image of the whole page before it can be transferred to paper, most laser printers supply their own on-board microprocessor and memory to process the data. Inkjets, on the other hand, print text line-by-line and usually utilize their host computer resources to handle the data stream. Hence, inkjet performance improves as computers themselves become more powerful. Other factors contributing to the inkjet speed are chemical advances in fast-drying inks and better interfaces for moving data rapidly from PC to the printer.

There is also a small but stable market for older impact technology. Impact printers produce

⁸Sales of replacement components, such as ink cartridges, laser toners, refill kits, trays and feeders, account for another major source of revenue for printer manufacturers. It also should be noted that a significant amount of printer products are not targeted towards the retail sector. In particular, sizeable revenues are generated by sales of special-purpose products, such as point-of-sale and label printers.

⁹Sources: PC Data *Hardware Monthly Report* described below, *Windows* (June 1999), *Consumer Reports* (March 1998, March 1999).

¹⁰A variation of this method is Okidata's proprietary LED technology, which uses an array of light-emitting diodes (LED) instead of the laser beam to form an image on the drum.

characters and graphics by striking an appropriately shaped combination of pins (dot matrix) or hammers (line printers) against an ink ribbon. Impact printers are not well suited for printing graphics and their resolution is low, but they can be faster than lasers and inkjets for tasks that do not require high quality output¹¹. Their other unique feature is the possibility of producing several copies of the document simultaneously. Accordingly, retail models of impact printers are positioned to attract home business users for tasks such as printing forms, labels, and invoices¹².

Another source of printer differentiation is their functionality. From a purely engineering perspective, printers have many common components with other office equipment. Starting in mid-90s, several manufacturers have realized this fact and started to offer integrated multifunction devices (MFD). A representative MFD combines a printer, scanner, copier, and possibly fax machine under the same cover. MFDs were initially heralded as "the consumer device for the nineties" but their acceptance has proven to be lukewarm. Partially this may be attributed to the widespread opinion that MFDs achieve their functionality by sacrificing performance of individual components¹³. But perhaps the main factor preventing MFDs from gaining the market share was the increasing reliance on computers for electronic data transfer, reducing demand for more traditional means of information exchange.

Finally, there is clear separation between personal and workgroup printers. The latter are targeted towards the small office/home office (SOHO) users and typically include the network interface card (NIC) for sharing the printer between several linked computers. Workgroup models are designed for heavy printing volume and are much faster than personal printers, but bear a significantly higher price tag. It appears that as inexpensive inkjet printers improve in speed and quality, laser printer manufacturers attempt to differentiate their products by including more workgroup printers and MFDs in their product lines.

Stylized Industry Facts. The industry is highly concentrated. Table 5.1 reports retail market shares of leading manufacturers. Hewlett-Packard (HP) is an undisputable industry leader, controlling nearly half of the market. 4-firms concentration ratio is 91%. HP is a dominant firm across the board, offering a complete product line from portable inkjets to powerful color lasers. Other manufacturers tend to specialize. Canon and Epson¹⁴ aggressively compete against HP in the inkjet

¹¹Some "industrial-strength" line printers offer print speed up to 100 pages per minute.

¹²Few models use the thermal transfer technology, similar to the one used in fax machines. Thermal transfer models are comparable to impact printers in terms of the output quality, but they have none of their advantages. This technology seems to disappear from the market.

¹³One reviewer has sarcastically suggested to integrate MFDs with a coffeemaker so that the waiting for a printed page would be less painful.

¹⁴Epson historically is a leading supplier of impact printers, but, as was mentioned earlier, significance of this market has declined in last few years. Epson also has a dominant position on the market for point-of-sale printers chiefly

Manufacturer	Unit Share	Dollar Share
Hewlett-Packard	49.94%	62.07%
Epson	18.19%	12.90%
Canon	15.54%	11.45%
Lexmark	7.81%	4.63%
Total	91.48%	91.05%

Table 5.1: Market Shares of Leading Manufacturers

sector. On a market for personal laser printers, HP is close to a monopoly status, although even in this segment it has some competition from Lexmark and Okidata. Brother is particularly active on the market for MFDs. Nevertheless, all leading manufacturers have extensive product lines which include dozens of models with comprehensive option packages, which suggests that product differentiation plays an important role in the printer market.

5.2 The Data

Sources and Description. The data set used in this paper comes primarily from *Hardware Monthly Report* electronically distributed by PC Data, Inc., a leading marketing research agency in the field of computer software and hardware¹⁵. PC Data collects point-of-sale, scan-code level data on (essentially) all computer products and accessories sold in the U.S. through major retailers, mail order firms, and e-commerce sites (CompUSA, Best Buy, Circuit City, OfficeMax, CDW, eBay, etc.). My subset of this database contains monthly data on sales and average prices of computer printers and multifunction devices. The sample covers 14 months from January 1998 until February 1999. PC Data estimates their market coverage at 55%.

This panel is supplemented by the data on the static product attributes compiled by the author from several sources, primarily from CNET (<http://computers.cnet.com>), CompUSA (<http://compusanet.com>), and manufacturer's on-line documentation. Product characteristics for which we have data include technology type and various measures of product performance, such as speed, printing quality, paper handling capabilities, and size. Exact definitions of printer attributes used in our analysis are given in Table 5.2.

The original PC Data sample included sales data on more than 800 printer models. Products with extremely low sale volumes that were found to be discontinued before January 1998, models not generally available through retail channels, printers incompatible with PC or Apple, demo and

based on the impact technology.

¹⁵This data set was graciously made available to us by Mark Bates of PC Data.

Variable ^a	Definition	Avail.
<i>Printer Class Dummies: equal to one if</i>		
bLaser	Laser/LED (light emitting diodes) printer (L)	all
bInkJet	Inkjet/micro-dry printer (J)	all
bImpact	Impact (dot matrix) printer (I)	all
bTT	Thermal transfer (only used in multifunction devices) (T)	all
bColor	Color printer (C)	all
bMF	Multifunction device (MF)	all
<i>Performance Measures</i>		
nSpeedBW	Maximum speed in black and white mode, pages per minute (ppm) ^b	all
nRAM	Memory buffer size, Kb	all-12
nCPU	On-board CPU speed, mHz	L
nFonts	Number of resident fonts	L, I
bPS	Supports Adobe PostScript language	all
bUSB	Universal Serial Bus (USB) interface	all
bSCSI	Small Computer System Interface (SCSI)	all
<i>Printing Quality Measures</i>		
nResBW	Maximum resolution in black and white mode, in dots per inch (dpi); $\sqrt{DpiX * DpiY}$	all
nColors	Number of colors used to generate an image	all
<i>Paper Handling Measures</i>		
nPHC	Paper handling capacity, in pages	L, I, T
bA3	Supports printing on A3 paper size	all
bA2plus	Supports printing on A2 paper size and above	all
bDuplex	Supports duplex (double-sided) printing	all
<i>Miscellaneous</i>		
nFootPrt	Footprint, width×depth, sq.in.	all
bPort	Portable printer	all
bNIC	Network Interface Card (usually Ethernet) included	all
nExtras	1+sum of the following dummies: scanner, color scanner, copier, color copier, fax, color fax, PC fax, message center	all
bRefurb	Refurbished and factory-serviced models	all
Notes:		
^a General convention for naming variables used throughout this section is that dummy (boolean) variables start with the prefix <i>b</i> , continuous variables - with the prefix <i>n</i> , variables transformed to logs - with the prefix <i>l</i> .		
^b For impact printers, this variable was constructed from the characters per second statistic assuming 4000 characters per page.		

Table 5.2: Definitions of Printer Characteristics

pilot models were eliminated from the sample. Models differing only in minor characteristics (body color, bundled software available for free elsewhere, etc.) were aggregated. As a result, our final sample contains data on 462 models from 27 manufacturers. High rates of technological innovation and product obsolescence in the industry translate into high rates of product entry and exit. Thus, the panel is unbalanced; treating each month/model pair as a single observation, the total sample size is 4344.

The sample identifies the product by its stock-keeping unit (SKU) code. In general, products are assigned distinct SKU codes if they differ in at least one (potentially) observable characteristic. Thus, the products in our sample are "transaction" models rather than "base" models, i.e. the printers with different option packages were considered to be distinct products. Using SKU-level data is clearly advantageous to the more common practice of identifying the product with the base model since, particularly for upscale lasers, the cost of included options may well exceed the price of the base model.

Extensive as it is, the PC Data set does have certain limitations. First, it covers only retail sales. Institutional purchases and direct merchandising are not included in the sample. Since I only model consumer-level demand, this should not create a problem for demand side estimates, but it might lead to an incomplete picture of the overall industry dynamics. Second, on the retail level printers are sometimes bundled with new computer system packages. There is some limited evidence which seems to indicate that this practice is not too common. Original PC Data sample also had the information on retail sales of computer systems which included the list of bundled components. According to these data, only about 8% of new computer systems come bundled with a printer. In any case, it seems reasonable to think that the purchase of a bundled system is chiefly determined by consumers' preferences for a personal computer, and in this sense bundled printer sales must constitute a separate market.

To check the accuracy of the data, I have matched total 1998 sales with the aggregate annual statistics from independent sources. For inkjets, the PC Data figures agreed with the aggregate data within 5-7%, which suggests that most of inkjet sales in fact occur on the retail level. On the other side, lasers appear to be undersampled in the data. This is to be expected, since many high performance lasers are targeted towards business and institutional users. Modeling institutional demand is beyond the scope of this paper, but I should stress that any references to "market shares" below should be interpreted as retail market shares, especially for lasers. I have also checked selected price series against the data from PriceSCAN.com, which agreed well.

Descriptive Statistics. Tables 5.3 and 5.4 provide some descriptive statistics on the industry

Time	No. of models			Q_t , mil.	Market Shares			Average Prices, \$			
	all	entry ^a	exit		Laser	Jet	MF	All	Laser	Jet	MF
98.1	244	.	.	1.127	0.127	0.836	0.107	335	747	272	436
98.2	248	11	0	1.158	0.115	0.853	0.104	326	752	268	434
98.3	265	21	2	1.287	0.116	0.851	0.105	326	750	267	429
98.4	279	22	5	0.965	0.134	0.828	0.114	337	769	267	439
98.5	292	19	3	0.809	0.145	0.816	0.124	349	800	269	438
98.6	304	13	5	0.984	0.151	0.813	0.108	339	794	255	463
98.7	310	20	10	0.877	0.142	0.815	0.110	331	805	248	432
98.8	326	25	5	1.155	0.115	0.852	0.103	303	793	235	428
98.9	340	12	6	1.453	0.116	0.854	0.096	294	786	226	414
98.10	342	22	14	1.057	0.120	0.847	0.120	304	788	234	393
98.11	344	13	13	1.355	0.092	0.884	0.087	265	767	212	408
98.12	349	19	13	2.092	0.077	0.906	0.076	266	834	216	404
99.1	353	12	14	1.371	0.109	0.865	0.091	283	788	218	402
99.2	348	9	24	1.372	0.111	0.863	0.091	282	799	215	404
All	462	218	114	17.062	0.119	0.849	0.103	310	784	243	423

Notes:

^aI define entry in period t as $\sum_{k=1}^{t-1} Q_{jt} = 0$ and $Q_{jt} > 0$; similarly exit in period t is defined by $Q_{j,t-1} > 0$ and $\sum_{k=t}^T Q_{jt} = 0$. Therefore, J_t need not equal $J_{t-1} + n_t^{entry} - n_t^{exit}$ since some products may have gaps in the sales record. Mostly such gaps are caused by extremely low market shares for some "exotic" products, but this is not necessarily the case. Indeed, during the sample period one of the most popular Hewlett-Packard ink jets, HP720c was discontinued and then reintroduced several months later "back by the popular demand".

Table 5.3: Dynamics of Prices and Market Shares

dynamics. Table 5.3 traces the motion of aggregate prices and market shares over time. During the sample period, printer sales were growing at the average monthly rate of 1.5%, although there is a high degree of variation around this trend, perhaps due to the influence of seasonal factors. It appears that ink jets are slowly gaining ground against the other technology types, but this trend is statistically very weak. Multifunction devices occupy a relatively stable niche of the overall market.

From January 1998 to January 1999 the average sales-weighted price of a printer has fallen by nearly 16%, although this downward price trend is far from homogenous for different printer types. In fact, lasers became more expensive, by roughly 5%. The latter tendency may be partially explained by the higher rates of quality improvement in the laser technology. However, the product-level data seem to indicate that as cheap ink jets approach near-laser printing quality, laser printer manufacturers try to differentiate their products by including more upscale, small business-targeted models in their marketing mix. Ink jet prices have declined by more than 20%. Notably, the prices for multifunction devices have also fallen despite the massive transition from thermal transfer printing to laser and color ink jet technologies witnessed in 1998 by this sector.

Time	nSpeedBW			nResBW			nRAM			nCPU	nColors
	All	Laser	Jet	All	Laser	Jet	All	Laser	Jet	Laser	Jet
98.1	5.84	9.14	5.43	638	720	641	606	2634	317	37.73	3.78
98.2	6.00	9.07	5.66	638	715	640	611	2720	344	37.05	3.81
98.3	6.01	9.04	5.68	632	713	635	615	2756	340	36.51	3.79
98.4	6.11	9.21	5.70	641	721	644	697	2906	362	38.04	3.75
98.5	6.13	9.41	5.65	648	731	651	749	3144	354	40.07	3.78
98.6	6.10	9.34	5.58	658	739	660	760	3213	333	40.35	3.78
98.7	6.09	9.57	5.59	660	750	665	770	3393	345	41.77	3.85
98.8	5.93	9.53	5.51	657	747	660	680	3440	325	41.40	3.83
98.9	5.86	9.70	5.41	667	754	669	694	3496	331	41.39	3.83
98.10	5.98	9.83	5.51	674	754	679	759	3725	360	42.70	3.82
98.11	5.76	10.02	5.36	675	744	678	692	4186	342	44.26	3.79
98.12	5.93	10.77	5.55	702	768	705	730	5262	356	50.19	3.85
99.1	6.22	10.54	5.73	691	753	696	943	5226	427	48.01	3.92
99.2	6.26	10.60	5.77	687	796	686	1059	5680	493	50.69	3.92
All	6.02	9.70	5.58	662	743	665	740	3699	359	42.15	3.82

Table 5.4: Technological Trends

Table 5.4 reports sales-weighted means for selected technical parameters reflecting printers performance. All of them show clear upward trends. This trend is particularly pronounced for the size of installed RAM, which have grown by 55% for ink jets and more than doubled for lasers. Printers have noticeably improved in terms of speed and resolution and have become more "colorful". Lasers still offer the best printing speed and quality, but the performance gap between two competing technologies has narrowed. This can be seen better on Figure 5.1. The left two graphs on Figure 5.1 project unit market shares on the subset of the product characteristics space (speed and resolution) in January 1998 and January 1999. Two graphs on the right similarly compare distributions of manufacturers' revenues. In 1998, the consumer demand is concentrated around two clusters roughly associated with lasers and inkjets. In 1999, these "twin peaks" nearly merged and the density drifted towards the northeast corner. It is also clear that in 1999, laser printer manufacturers positioned more of their products in the northeast corner of the distribution, perhaps in response to the inkjets' assault.

The downward trend in prices becomes even more pronounced when we control for quality changes. Table 5.5 reports results from the hedonic pricing regression for the whole sample (columns 2-3), black and white lasers (columns 4-5), and color ink jets (columns 6-7). For all three regressions, the dependent variable is $\ln p_{jt}$. The empirical specification also includes time dummies, whose values together with the confidence intervals are shown on the graph. With a single exception of CPU speed, for which I expected a significant positive effect, all variables have the expected sign

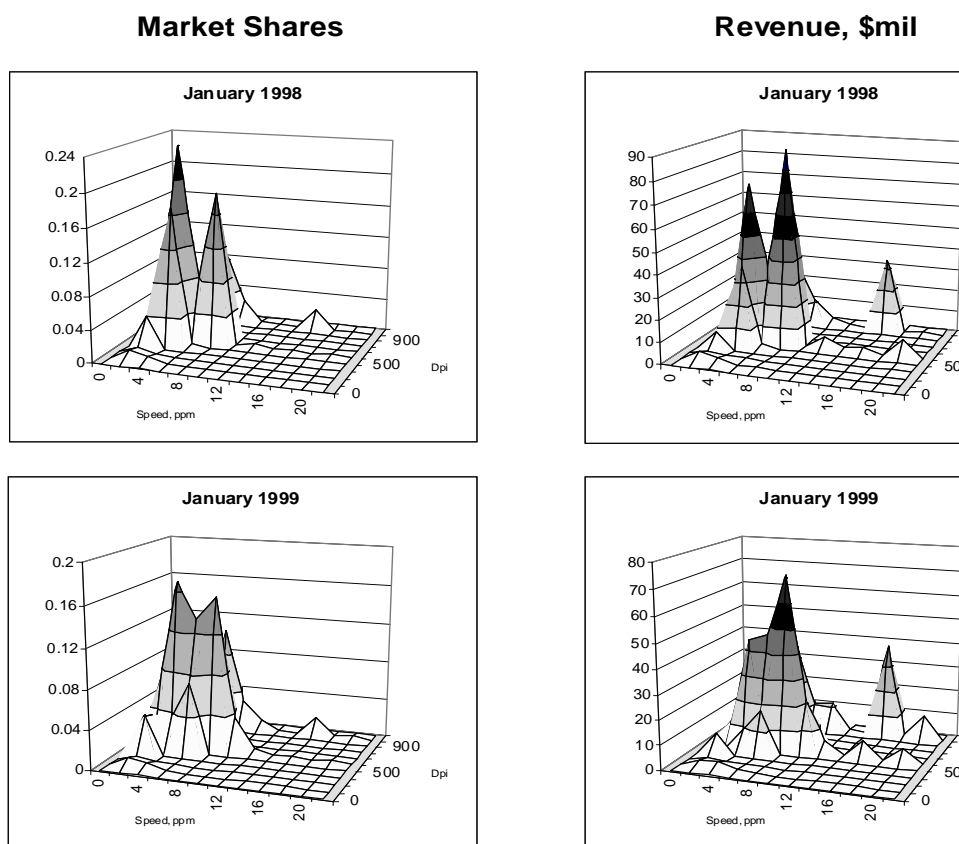


Figure 5.1: Market Shares and Revenues as Functions of Continuous Demand Characteristics

and are sharply estimated. Overall, differences in speed seem to be the main source of dispersion in printers' prices. Laser prices are also significantly influenced by the printer's resolution, amount of RAM, and paper handling capacity. For inkjets, resolution and the number of colors both appear to be more important than speed, while RAM essentially has no effect on the price. Fast interfaces (SCSI and USB), Postscript-compatibility, and network support features all have high and significant impact on the printers' price. Factory refurbished models sell at heavy discounts, especially inkjets. Such a steep depreciation is typical for many computer products and probably explains why there is little trade in secondary markets for SOHO computer equipment. It is also interesting to note that Hewlett-Packard printers sell at a considerable premium with respect to models from other manufacturers.

The bottom part of Table 5.5 plots fixed time effects for hedonic regressions. Remarkably, within a year quality-adjusted prices have fallen by more than 30%¹⁶. For SOHO printer types,

¹⁶For the pooled regression, $30.5\% \pm 9.2\%$ within 95% confidence interval assuming normally distributed disturbances

this effect is even higher. For inkjets, prices have dropped by 32%, and for lasers - by nearly 40%. Comparing these results with the descriptive price statistics, we see that for lasers this decline is mainly due to an increase in quality attributes of comparably priced printers; for inkjets - due to decline in nominal prices.

Considering the industry trends reported above, it seems natural to think that intertemporal aspects of a consumer's choice should play an important role in the demand evolution. We now report estimates of demand parameters for the model in Section 3.

5.3 Empirical Results

The demand system for the full model is derived from the consumers' optimization problem described in Section 3. We make the following parametric specifications:

(i) for the static demand side, we adopt the logit model with unobservable random effects described in detail in Berry (1994):

$$u_{jt} = \xi_j + x_j\beta - \alpha p_{jt} + \varepsilon_{jt},$$

where ξ_j is an unobserved random quality component of the product quality and ε_{jt} is the logit error term;

(ii) the quality generating process follows a random walk with drift

$$r_{t+1} = r_t + \gamma + \sigma\nu_{t+1}.$$

These functional forms are admittedly restrictive; more flexible specifications were discussed in Section 3. The choice of functional forms was motivated mainly by data constraints and, to lesser extent, by computational considerations. The stochastic linear growth specification is particularly troublesome since, with positive drift, it implies that participating in the market will eventually become infinitely more attractive to the consumer than the outside alternative. It seems natural that the quality growth must eventually slow down and the difference between the value of the market and the outside alternative must be bounded. Specifying r_t process as bounded would require a nonlinear functional form for $\mu(r_t)$. However, in short time series we observe only a limited subset of the state space over which this process appears to be approximately linear. Correspondingly, it may be difficult to identify nonlinearities in the data generating process.

The above model requires the market size to be observed by the econometrician. We identify the total market size with the number of U.S. households who own at least one computer. The data

and using an exact sample distribution formula.

Variable	Pooled		Lasers, B&W		Inkjets, Color	
	Est.	Std.Err.	Est.	Std.Err.	Est.	Std.Err
lSpeedBW	0.3409	0.0179	0.2514	0.0293	*0.0356	0.0273
lRAM	*-0.0083	0.0038	0.1503	0.0085	*-0.0092	0.0049
lCPU	*-0.0289	0.0109	*0.0009	0.0086	.	.
lFonts	0.0512	0.0105	-0.0415	0.0112	.	.
bPS	0.5265	0.0355	0.2042	0.0278	1.6767	0.0810
bUSB	0.1933	0.0242	0.2380	0.0645	0.0640	0.0280
bSCSI	0.7559	0.0572	0.3071	0.0689	1.1777	0.1174
lResBW	0.0565	0.0235	0.1896	0.0174	0.3833	0.0266
lColors	0.3667	0.0328	.	.	0.2408	0.0344
lPHC	0.0815	0.0139	0.2164	0.0194	0.0827	0.0209
bA3	0.4249	0.0302	0.2767	0.0324	0.4139	0.0655
bA2+	0.5450	0.0595	.	.	-0.4494	0.0957
bDuplex	0.2898	0.0619	0.2503	0.0460	.	.
lFootPrt	0.6032	0.0262	0.3919	0.0271	0.3567	0.0326
bPort	1.4815	0.0730	.	.	1.0409	0.0767
bNIC	0.3734	0.0297	0.2529	0.0262	1.0092	0.0646
lExtras	0.2749	0.0169	0.2445	0.0305	0.3536	0.0186
bRefurb	-0.6383	0.0350	-0.3062	0.0399	-0.8580	0.0418
bHP	0.3009	0.0287	0.1899	0.0279	0.5986	0.0438
bEpson	-0.2882	0.0383	.	.	0.2069	0.0587
bCanon	-0.2885	0.0401	.	.	*-0.0171	0.0442
bLexmark	0.1300	0.0320	0.1392	0.0332	*-0.0945	0.0583
bBrother	-0.2113	0.0475	-0.5881	0.0408	.	.
bOkidata	*0.0152	0.0383	*-0.0616	0.0369	.	.
# Obs.		4344		1598		1582
\bar{R}^2		0.868		0.906		0.781

*Not significant at the 95% level.

Fixed Time Effects:

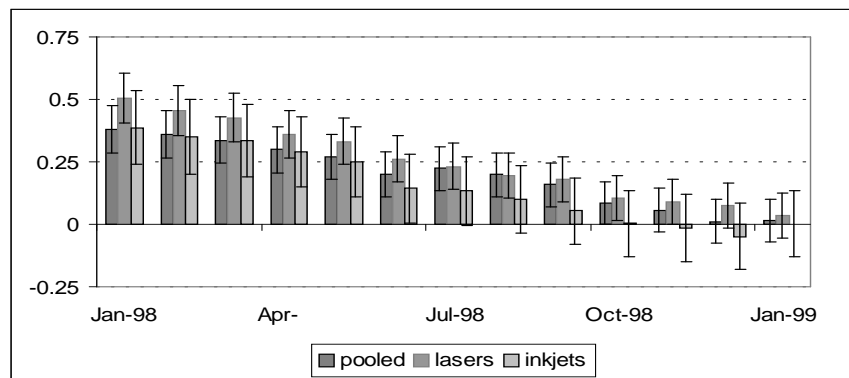


Table 5.5: Hedonic Pricing Regression

on computer ownership were obtained from 1998 Internet and Computer Use Supplement to the Current Population Survey available from the Bureau of Labor Statistics. According to the survey, in 1998 42.1% of U.S. households owned a computer, which translates into the average market size of 42 million. Similar surveys were conducted in 1997 and 2000 (August). Unfortunately, these data are only available on the annual basis. Since purchases of new computers may be an important factor beyond the growth in printer sales, I convert annual data to monthly frequency assuming the uniform growth rate in computer ownership. This translates into the monthly market expansion rate of approximately 0.5mln. households.

Table 5.6 reports static demand parameters. Most coefficients are precise and of the expected sign. Speed, resolution, added functionality (IExtras), and CPU speed all have a strong and significant effect on relative market shares. The strongly negative coefficient on the "refurbished" dummy shows that there is little demand for used equipment. There are also a few anomalies. The coefficient on memory is negative albeit close to zero. I have argued above that an increase in the host computer memory reduces the need for printers' on-board RAM, so the estimate may reflect this substitutability. I also expected to find a positive preference for portable printers, but perhaps there is little demand for mobile printing.

Estimated manufacturers' dummies show that brand names play an important role in structuring demand. The estimates suggest that after controlling for observable quality differentials, a representative consumer is more likely to purchase a model manufactured by one of the three leading vendors (HP, Canon, and Epson). Not surprisingly, Hewlett-Packard is viewed by consumers as a quality leader, while Canon and Epson models are essentially close substitutes. In the pooled regression, Okidata occupies the fourth place despite its small market share, but this effect goes away if we exclude impact printers from the sample. This is perhaps due to the fact that Okidata (together with Epson) has a strong presence on the impact printers market.

The price dispersion in the sample is quite high (the least expensive model in the sample, refurbished black and white Canon-30, had the lowest price of \$40, while several color lasers have the average price in the vicinity of \$10,000). To allow for different price elasticities for different price groups, a spline in price is used instead of nominal price. Spline cutoff points were initially chosen at deciles of the price distribution. I then merged categories that resulted in nearly identical elasticities. Using the price spline, as well as manufacturers' fixed effects, also should alleviate the potential endogeneity between the price and the unobservable component of product quality (see Berry (1994) and Goldberg (1995) for different opinions on this subject). Berry (1994) argues that price may be correlated with the latent product-specific quality and suggests instrumenting for price to correct for resulting endogeneity. Since I do not have data on the product-level supply shifters,

Variable	Pooled		Excluding Impacts	
	Est.	Std.Err.	Est.	Std.Err.
ISpeedBW	0.1525	0.0599	0.1705	0.0680
IRAM	-0.0474	0.0124	-0.0458	0.0143
ICPU	0.1199	0.0293	0.1305	0.0314
bPS	*0.2244	0.1198	*0.2337	0.1321
bUSB	0.3065	0.0787	0.2200	0.0853
bSCSI	*-0.2541	0.1873	*-0.2829	0.1988
lResBW	0.2814	0.0764	0.3233	0.0926
lColors	0.2439	0.1080	0.4069	0.1247
bA3	*-0.1870	0.1012	*0.0893	0.1274
bA2+	0.6600	0.1949	0.6003	0.3025
bDuplex	-0.4096	0.1988	*-0.2999	0.2112
lFootPrt	*-0.0353	0.0869	*-0.1065	0.0976
bPort	-0.7841	0.2389	-0.8212	0.2591
bNIC	*-0.1093	0.0996	*-0.1182	0.1070
lExtras	0.3276	0.0555	0.3906	0.0593
bRefurb	-2.5912	0.1180	-2.7121	0.1268
lPrice Spline:				
$p_{jt} < \$125$	-0.7359	0.1686	-0.9042	0.2053
$\$125 \dots \199	-0.5156	0.1442	-0.6745	0.1773
$\$200 \dots \299	-0.4664	0.1326	-0.6295	0.1633
$\$300 \dots \399	-0.5254	0.1266	-0.6817	0.1563
$\$400 \dots \999	-0.6234	0.1181	-0.7709	0.1453
$p_{jt} \geq \$1000$	-0.7296	0.1068	-0.8890	0.1348
Manufacturers Dummies:				
bHP	3.0924	0.0932	3.0972	0.1013
bEpson	2.0110	0.1265	2.3700	0.1890
bCanon	2.3970	0.1304	2.3681	0.1395
bLexmark	0.8812	0.1038	0.8228	0.1192
bBrother	*0.2462	0.1518	*0.1044	0.1715
bOkidata	1.2187	0.1224	*0.3159	0.1713
# Obs.	4344		3577	
\bar{R}^2	0.474		0.471	
Notes:				
*Not significant at the 95%level.				

Table 5.6: Static Structural Parameters Estimates

finding appropriate instruments is not straightforward. An alternative to the instrumental variable technique, which is pursued here, is including fixed effects. One can interpret the intercepts and spline coefficients in my specification as the estimates for the unobservable quality level which is localized to the technology type, manufacturer, and price group.

The estimates suggest that the demand for printers is relatively elastic. Spline coefficients for different price groups exhibit a U-shaped pattern. Not surprisingly, we find that demand in the lowest price category is the most elastic. High price elasticity for expensive printers is also to be expected: there are many models available in this category which compete for a relatively low demand, and price may be a decisive demand factor in this segment. In contrast, for medium-priced models I would expect that the vendor's reputation and brand image are ultimately more important than price, which is consistent with our findings.

Figure 5.2 shows first-stage estimates of r_t process realizations in the data together with the fitted random walk and conditional confidence intervals \hat{s}_t defined as $\mathbb{P}\{r_{t+1} \in [\mu(r_t) - \hat{s}_t, \mu(r_t) + \hat{s}_t] \mid r_t\} = 0.95$. Notice that the approximation quality deteriorates around the Christmas shopping season. Since PC Data sample covers a relatively short time interval, I cannot separately identify seasonal effects from the trend component, which is an obvious disadvantage of having a short time series. However apart from Christmas, the stochastic linear growth model fits the data reasonably well. It is clear from the graph that the mean quality level in January 1999 was substantially higher than in January 1998.

Table 5.7 reports dynamic demand parameters for the full model and several constrained estimators. Column 2 shows estimates for the full model. Estimates imply that consumers' forward-looking behavior plays an important role in the market. The estimated discount factor has a reasonable magnitude of 0.9366. The replacement probability from the full model is 0.0205, which implies that the average lifetime of a consumer printer is slightly above four years. The estimate for the initial participation probability suggest that at the beginning of the sample, 59% of households with computers had a printer. By the end of the sample period, this figure has risen to 65%.

To show the effect of dynamic parameters on the demand, columns 3-5 report estimates for several constrained estimators. Column 4 shows demand coefficients for myopic consumers. According to these estimates, 98% of U.S. households owned a printer in the beginning of 1998. Correspondingly, for myopic consumers the model attributes all printer purchases to the replacement demand. To explain the magnitude of sales, the model is forced to adjust the replacement probability, which is estimated close to the ratio of average monthly sales to the total market size. Column 5 reports estimates for very "patient" consumers. Here, the magnitude of the new demand resulting from




Figure 5.2: Estimated r_t Realizations in the Data.

the incomplete penetration in the beginning of the sample is insufficient to explain sales levels and so the constrained estimator drives the replacement probability up to 0.0759, suggesting the annual replacement frequency.

Figure 5.3 shows the fit of different estimators to the data. All three estimators are unbiased, but significantly differ in the decomposition of variance. The "myopic" estimator fits the aggregate data to the average monthly sales levels. Patient consumers are more likely to absorb stochastic quality shocks by postponing the purchase and consequently smooth out their consumption more. Accordingly, the "patient" estimator oversmooths the data, although it picks up an upward trend in

<i>Second Stage</i>					
		Estimate	Std. Err.		
γ	(r_t drift)	0.0834	(0.0622)		
σ_r^2	(r_t variance)	0.0503	(0.0197)		
<i>Third Stage</i>					
Estimator:		Unconstrained	Myopic	$\beta = 0.99$	
β	(discount factor)	0.9366	(0.0499)	0.0000	0.9900
q	(replacement probability)	0.0205	(0.0182)	0.0251	0.0759
c	(outside good payoff)	0.7726	(0.2878)	4.3326	0.0187
μ_0	(participation rate at $t = 0$)		0.4091	0.0224	0.5112
μ_T	(participation rate at $t = 14$)		0.3489	0.0245	0.4887
Replacement frequency, months			49	40	13

Table 5.7: Dynamic Structural Parameters Estimates

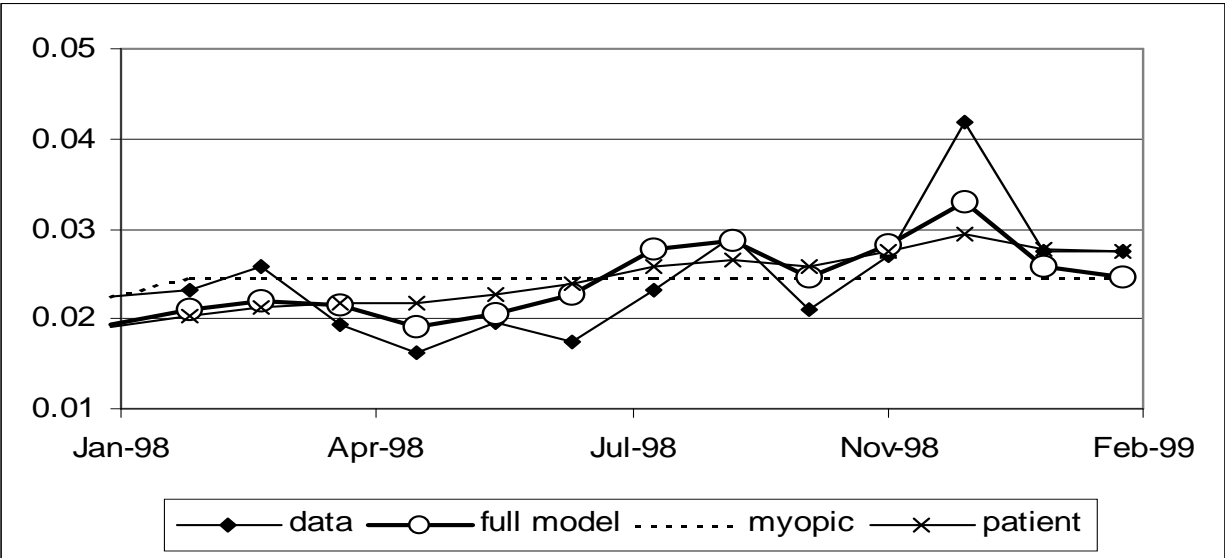


Figure 5.3: Aggregate Sales Levels: Actual vs. Fitted Values

aggregate sales. The unconstrained estimator strikes a proper balance between smoothing random demand fluctuations and the data fit.

Given the estimates for dynamic parameters, one can compute out-of-sample forecasts to compare model predictions with the data. Unfortunately, monthly sales data are difficult to obtain even on the industry level. PC Data quotes that in 1999, printer sales growth rate has fallen to 11.2%, despite the decline in nominal prices by 18%. As can be seen from Figure 5.3, the myopic model basically predicts flat sales for 1999. In contrast, projected sales growth rate for 1999 for the full model is 11.8%. Since I have very limited data, I hesitate to make any assessments of the model's forecasting power. Nevertheless, it seems that the model is capable of explaining rather complicated sales patterns.

6 Conclusions and Future Research

Summary and Applications. This paper considers demand implications of consumers' forward-looking behavior in a dynamic market for durable goods. An empirical technique for estimating durable products demand system is proposed. Our method incorporates the following market features:

- product durability;

- product differentiation;
- intertemporal demand substitution.

The advantages of the dynamic structural approach over static models are illustrated using Monte-Carlo simulations. Simulation results suggest that static demand systems perform poorly in the above setting. Our method produces more accurate estimates of the demand side parameters. Demand elasticities play a crucial role in many policy issues and hence the model developed here might facilitate their empirical analysis. In addition, for empirical applications such as marketing and welfare implications of new product introduction it is important to distinguish new demand from the replacement demand, and our technique sheds some light on this issue.

Static models are limited to providing a conditional, in-sample analysis of counterfactual experiments. The dynamic model can also generate out-of-sample predictions that might be useful for obtaining a more realistic account of the industry evolution. In particular, our approach permits to analyze short-run and long-run effects of aggregate macro shocks on the consumers' demand for durables.

The empirical part of the paper applies the model to analyze the demand structure in the U.S. computer printer market. The estimates suggest that consumers' intertemporal optimization is an important feature of this market. It seems that the dynamic model fits industry trends better than static models and provides more realistic out-of-sample forecasts.

Extensions. Several important aspects of markets for durable products are not addressed in this paper. The focus of this study is on the demand side. The supply side of the durable products industry is an important and controversial topic with far-reaching policy implications, in particular for anti-trust concerns. Coasian conjecture implies that durable goods manufacturers are not able to exercise any market power even in the absence of competitors, an argument that was recently brought up to public attention during the Microsoft trial. In contrast, numerous empirical studies find an evidence of the market power in various durable goods industries. It appears that in random utility models manufacturers may have a sufficiently large future demand to sustain positive markups, especially after accounting for the replacement demand. The technique proposed in this paper might simplify the analysis of dynamic multi-product pricing policies by reducing dimensionality of the problem. Endogenizing manufacturers' pricing and product entry decisions is a challenging topic that I intend to pursue in the future research.

Another important issue is an equilibrium in the secondary markets. While for our empirical application the market for used durables seems to be negligible, for many industries endogenous

replacement decisions and the distribution of consumer holdings are important and may play a crucial role in explaining procyclical behavior of the demand for durable products. This topic has not received an appropriate attention in the empirical industrial organization literature, indicating the need for further research.

7 Appendix

7.1 Appendix A: Proof of Proposition 1.

Here we derive the distribution function of $v_t = \max_{j \in \mathfrak{J}_t} u_{jt}$.

$$\begin{aligned}
F_v(z) &= P\{v_t \leq z\} = P\{u_{1t} \leq z, \dots, u_{Jt} \leq z\} \\
&= P\{\varepsilon_{1t} \leq z - \delta_{1t}, \dots, \varepsilon_{Jt} \leq z - \delta_{Jt}\} \\
&= \exp\left(-G(e^{-(z-\delta_{1t})}, \dots, e^{-(z-\delta_{Jt,t})})\right) \\
&= \exp\left(-e^{-z}G(e^{\delta_{1t}}, \dots, e^{\delta_{Jt,t}})\right) \text{ (by homogeneity of } G(\cdot)\text{)} \\
&= \exp(-R_t \exp(-z)) \text{ where } R_t = G(e^{\delta_{1t}}, \dots, e^{\delta_{Jt,t}}) \\
&= \exp(-\exp(-(z - \ln R_t))) = \exp(-\exp(-(z - r_t))),
\end{aligned}$$

which implies that v_t is distributed Type I extreme value with the mode of r_t . Density of v_t obtains trivially by differentiating above expression. Usually we consider the distribution of v_t to be parametrized by r_t , but occasionally it is more convenient to treat R_t as a distribution parameter. Both forms of density function are given below:

$$\begin{aligned}
f_v(z; r_t) &= e^{-(z-r_t)} F_v(z; r_t) = \exp(-e^{-(z-r_t)} - z + r_t) \\
&= R_t \exp(-R_t e^{-z} - z) = f_v(z; R_t).
\end{aligned}$$

Let $\varepsilon[r]$ denote an extreme value random variable with the mode of r . One useful property of extreme value deviates that we will need for some results is that $\varepsilon[r] \stackrel{d}{=} r + \varepsilon[0]$. This follows immediately from the functional form of the distribution:

$$P\{r + \varepsilon[0] \leq z\} = P\{\varepsilon[0] \leq z - r\} = \exp(-e^{-(z-r)}) = F_{\varepsilon[r]}(z). \blacksquare$$

7.2 Appendix B: Asymptotic Distribution for the Autoregressive Data Generating Process

This appendix derives an asymptotic distribution for the data generating process used in Monte Carlo experiments. It was assumed that the number of products available on the market does not

change over time ($J_t = J$) and the quality index of individual products evolves over time according to

$$x_{j,t+1} = x_{j,t} + l + \sigma\nu_{j,t+1}$$

where ν_{jt} are i.i.d. $N(0, 1)$ deviates. Let

$$\bar{R}_t = \frac{1}{J} \sum_j \exp(\xi_j + \alpha x_{jt}).$$

Then

$$\begin{aligned} E [\bar{R}_{t+1} | \bar{R}_t] &= E \left[J^{-1} \sum_j \exp(\xi_j + \alpha x_{jt} + \alpha l + \alpha \sigma \nu_{j,t+1}) | R_t \right] \\ &= \exp(\alpha l) J^{-1} \sum_j \exp(\xi_j + \alpha x_{jt}) E \exp(\alpha \sigma \nu_{j,t+1}) = \exp(\alpha l + \alpha^2 \sigma^2 / 2) \bar{R}_t. \end{aligned}$$

Consider asymptotic behavior of \bar{R}_t as $J \rightarrow \infty$. From above formula and by consistency of a sample mean

$$\text{plim } \bar{R}_{t+1} = \exp(\alpha l + \alpha^2 \sigma^2 / 2) \bar{R}_t.$$

Then by Slutsky theorem

$$\text{plim } \ln(\bar{R}_{t+1}) = \ln(\text{plim } \bar{R}_{t+1}) = \alpha l + \alpha^2 \sigma^2 / 2 + \ln \bar{R}_t$$

and

$$\text{plim } [\ln(\bar{R}_{t+1}) - \ln(\bar{R}_t)] = \text{plim } [\ln(R_{t+1}) - \ln(R_t)] = \alpha l + \alpha^2 \sigma^2 / 2.$$

References

- BAGNOLI, M., S. SALANT, AND J. SWIERZBINSKI (1989): “Durable Goods Monopoly with Discrete Demand,” *Journal of Political Economy*, 97(6), 1459–1479.
- BERRY, S. (1994): “Estimating Discrete Choice Models of Product Differentiation,” *RAND Journal of Economics*, 25(2), 242–262.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- BERTSEKAS, D. (1995): *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA.
- BRESNAHAN, T., S. STERN, AND M. TRAJTENBERG (1997): “Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s,” *RAND Journal of Economics*, 28(0), S17–S44.

- COASE, R. (1972): “Durability and Monopoly,” *Journal of Law and Economics*, 15, 143–149.
- DAVIS, P. (2000): “Spatial Competition in Retail Markets: Movie Theaters,” MIT, Working Paper.
- DEATON, A., AND J. MUELLBAUER (1980): “An Almost Ideal Demand System,” *American Economic Review*, 70(3), 312–326.
- DENECKERE, R., AND A. DE PALMA (1998): “The Diffusion of Consumer Durables in a Vertically Differentiated Oligopoly,” *RAND Journal of Economics*, 29(4), 750–771.
- GOLDBERG, P. (1995): “Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry,” *Econometrica*, 63(4), 891–951.
- HAUSMAN, J. (1997): “Valuation of New Goods Under Perfect and Imperfect Competition,” in *The Economics of New Goods*, ed. by T. Bresnahan, and R. Gordon. Chicago: University of Chicago Press.
- HIMMELBERG, C., AND S. OLLEY (1998): “New Products and the Dynamics of Market Power in the Hard Disk Industry,” Columbia University, Working Paper.
- MAHAJAN, V., E. MULLER, AND F. BASS (1990): “New Product Diffusion Models in Marketing: A Review and Directions for Research,” *Journal of Marketing*, 54, 1–26.
- MCCALL, J. (1970): “Economics of Information and Job Search,” *Quarterly Journal of Economics*, 84(1), 113–126.
- McFADDEN, D. (1973): “Conditional Logit Analysis of Qualitative Choice Behavior,” in *Frontiers of Econometrics*, ed. by P. Zarembka. New York: Academic Press.
- (1978): “Modelling the Choice of Residential Location,” in *Spatial Interaction Theory and Planning Models*, ed. by A. Karlqvist. Amsterdam: North-Holland.
- (1981): “Econometric Models of Probabilistic Choice,” in *Structural Analysis of Discrete Choice with Econometric Implications*, ed. by D. McFadden, and C. Manski. MIT Press.
- McFADDEN, D., AND K. TRAIN (200?): “Mixed MNL Models for Discrete Response,” *Journal of Applied Econometrics*, forthcoming.
- NEVO, A. (200?): “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, forthcoming.

- RUST, J. (1985): “Stationary Equilibrium in a Market for Durable Assets,” *Econometrica*, 53(4), 783–805.
- (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55(5), 999–1033.
- SOBEL, J., AND I. TAKAHASHI (1983): “A Multi-stage Model of Bargaining,” *Review of Economic Studies*, 50, 411–426.
- STIGLER, G. (1961): “The Economics of Information,” *Journal of Political Economy*, 69, 213–225.
- STOKEY, N. (1981): “Rational Expectations and Durable Goods Pricing,” *Bell Journal of Economics*, 12, 112–128.
- STOKEY, N., AND R. LUCAS (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA.
- STOLYAROV, D. (2000): “Turnover of Used Durables in a Stationary Equilibrium: Are Older Goods Traded More?,” University of Michigan, Working Paper.