Working Paper 237

Quality Perceptions and Dynamic Brand Choice

By

Srinivas Prakhya

June 2005

Please address all correspondence to:

Prof. Srinivas Prakhya
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 560 076, India
Phone: +91-80-2699 3186, Fax: +91-80-2658 4050,
E-mail: sprakhya@iimb.ernet.in

QUALITY PERCEPTIONS AND DYNAMIC BRAND CHOICE

Srinivas Prakhya
Indian Institute of Management Bangalore

June 2005

Abstract

A simple model of dynamic brand choice that takes into account the uncertainty faced by the consumer and the consequent information value associated with choice is developed in this paper. The model is developed using the Neumann-Morgenstern expected utility framework. The consumer, on each occasion, chooses from the brands available so as to maximize his expected utility over a finite horizon based on preferences that are dynamically updated with consumption experience. The dependence of current choice on past choices is hypothesized as being due to learning effects. Brands are allowed to vary in their ability to influence long term preferences with consumption experience. The dynamic model is estimated on scanner panel data using the conditional choice probability (CCP) estimation procedure, proposed by Hotz and Miller [1993], that is computationally simple compared to the extant methods of estimating dynamic models. Results support the hypothesis that the consumer is not perfectly informed and is forward looking. Models that explicitly take such effects into account fit the data better than reduced form approaches.

KEY WORDS: *Dynamic Choice; Quality Perceptions; CCP Estimator; Consumer Learning; Expected Utility; Risk-Aversion; Valuation Function*

1 INTRODUCTION

Choice models in the marketing literature have mostly been based on the assumption that consumers are perfectly informed about the attributes of the brands in the product category (Guadagni and Little [1983]; Kamakura and Russell [1989]; Gonul and Srinivasan [1995]). The consumer's intrinsic preferences for the alternatives in the choice set are assumed to be static and the choice on each occasion is made based on these preferences and the state of the marketing mix variables. Utility specifications follow the compensatory multi-attribute model where the intrinsic utility that can be derived from each brand and the marketing mix variables such as price, display etc. are additively represented for a representative consumer up to a vector of unknown parameters that are to be estimated.

Apart from brand specific intrinsic preferences and marketing mix variables, researchers typically include a loyalty term (Guadagni and Little [1983]) in the consumer's utility function and thus incorporate dynamics in the choice process. The Guadagni and Little [1983] loyalty measure is constructed from the consumer's choice history using an exponential weighting scheme with choices in the recent past being given more importance relative to those in the distant past. Loyalty variables are known to have significant coefficients and provide stability to the model. It is therefore not surprising that loyalty is always included in brand choice models and researchers have been motivated to develop other operationalizations of loyalty. Krishnamurthi and Raj [1988], for example, operationalize loyalty as simply the share of past purchases for each alternative. The above measures explain variation in choice behavior by capturing both household heterogeneity and purchase to purchase preference updations. Fader and Lattin [1993] develop a measure that disentangles the two effects of heterogeneity and non-stationarity. Their loyalty measure is derived from a non-stationary Dirichlet multinomial choice model.

These measures are constructed such that the consumer's intrinsic preference for a brand increases each time the brand is chosen. Usually, these measures are not brand specific. The updating of preferences implied by these loyalty measures is simplistic. First, patterns of evolution of preferences are assumed that may not hold in all markets for all consumers. In markets for frequently bought consumer goods that are characterized by frequent new product introductions and re-positioning, strong learning effects driven by uncertain quality perceptions could exist and these effects could vary across brands. Hence, choice behavior could radically change with consumption experiences of alternatives that the consumer is not fully informed about. Such effects could be particularly significant for experiential goods whose attributes become clear upon consumption. Second and more importantly, preference updation implies that the consumer is not sure about his evaluation of the alternatives before the consumption experience and is making his choice in an uncertain environment. Choice, made in an uncertain environment, typically leads to rewards in the form of information that would be useful in future periods over and above the utility derived from immediate consumption. The consumer would have to be myopic if he were to make his choice to maximize the current period utility while operating in an uncertain environment. Consumers, in uninformed states, have an extra incentive (over and above immediate inducements such as price promotions) to switch in the form of better information for future choice situations; and not incorporating this information factor would lead to biased estimates of marketing mix variables. In uncertain environments, it would be more appropriate to assume that consumers with imperfect information would make choices on each occasion to maximize an objective function that comprises of current utility and the utility that would be derived in the future.

In this paper, a structural choice model that would be applicable for low involvement products is derived from a learning framework . The objectives of this study are two fold;

first, to build a structural model of consumer choice that takes into account the implications of imperfect information; and second, to formulate the model such that it is easily estimable without compromising these implications.

The basic premise here is that consumers are operating with imperfect information. We assume that consumers are risk averse and that they are partially informed about each alternative. The consumer's evaluation of each brand is based on his information state which includes the number of past consumption experiences with each of the alternatives. Each time he samples, the consumer gains information. In such a situation, the consumer can be expected to take into account the future value of current choices. Sampling of an alternative that the consumer is not perfectly informed about entails a cost of experimenting which reduces current expected utility but leads to a reward in the form of resolution of uncertainty thereby increasing future expected utility from this alternative. Hence, the consumer is assumed to maximize a value function which has two components; utility in the current period and the utility expected in the future. Consequent to the gain in information from each consumption experience the consumer has an updated intrinsic preference and a reduced cost of experimenting. The first objective of building a model of consumer choice that takes into account the implications of imperfect information is thus addressed by 1) modeling the consumer as an expected utility maximizer and 2) incorporating the forward looking nature of consumers in such situations into the model by allowing the objective function to have current and future utility components.

The traditional approach to modeling choice under uncertainty is to specify prior distributions that get modified with consumption experiences into posteriors. The consumer's choice would be such that he maximizes utility given his expectation of the attributes and the risk associated with the variance of the attributes. In such a formulation, the consumer's perceptions are serially correlated and unobserved. These have to be integrated out to arrive at choice probabilities and the order of integration can be high; significantly increasing computational costs. A structural model that allows for learning using Bayesian updating schemes and is characterized by such computational complexity has been proposed by Erdem and Keane [1996]. Their estimation methodology is based on approximations to solve the dynamic discrete choice problem using simulation and interpolation following Keane and Wolpin [1994]. Overcoming this problem of computational complexity is the second objective in this paper. This is achieved by considering the implications of imperfect information; the cost of experimenting that is incurred in the current period and the reward that is realized in the future consequent to current sampling. Thus, we avoid constructing choice probabilities in terms of unobserved variables that have to be integrated out. Another element of dynamic discrete choice problems that contribute to computational costs is the lack of a closed form solution to the value function that is being maximized. The standard approach is to exploit Bellman's condition and use backward induction methods to compute the value function for a given set of parameter values. In this paper a methodology that is based on an alternative representation of the value functions (due to Hotz and Miller [1993]) is used to solve the consumer's dynamic problem. This methodology considerably reduces the computational cost associated with the maximum likelihood estimation of dynamic structural models. Here, conditional (on the state) choice probabilities (CCP) are used to compute the value functions and the choice probabilities implied by these value functions are used to derive estimators for the underlying structural parameters. A large class of stochastic dynamic models can be estimated with relative computational ease using this approach.

It is assumed that consumers derive utility from quality based on an elementary utility function characterized by constant risk-aversion. Quality perceptions are uncertain and described by distinct normal distributions foe each alternative. Scanner data is used to calibrate the model and the implications of such a model for choice behavior are examined.

The hypothesis that consumers are not perfectly informed about the alternatives in the choice set is supported by the empirical results. Consumers appear to be forward looking; they take into account the impact of current choice on the future. Static models, based on current utility maximization, that do not take the effects of uncertainty and learning with experience could generate biased marketing mix parameters.

The paper is organized as follows : a general model of consumer behavior in an uncertain environment is formulated in Section 2. In Section 3, a general dynamic discrete choice model is described and estimation issues are discussed. A simple model of experimenting with choice under uncertainty, assuming complete learning with one experience and a value function with a two period horizon, is then formulated. Identification issues are then considered and reformulations of the learning models are proposed accordingly. In Section 4, the data used to calibrate the model is described and preliminary statistics are reported. The results are then presented and the implications are discussed. In conclusion, the contributions of this paper are summarized and directions for future research are outlined.The CCP methodology (Hotz and Miller [1993]) that we use in this paper is outlined in an appendix along with an illustration of the methodology using a simple model of choice with a two period value function.

2 A MODEL OF CONSUMER BEHAVIOR

In this section, consumer behavior in an uncertain environment is examined; assumptions are described, and a model of sequential consumer decision making is proposed. Consider a consumer making choices in an uncertain environment. Specifically, the uncertainty is with respect to the quality that the consumer would derive from consuming each of the available alternatives. This uncertainty could be driven by the consumer having not tried alternatives in the recent past (leading to some decay in information possessed by the consumer) or claims from the marketers about new or improved attributes in the alternatives.

Let $\tilde{q}_j$ represent the consumer's uncertain quality perceptions about brand j, $\mathrm{E}(\tilde{q}_j) = \mu_j$ and $\mathrm{Var}(\tilde{q}_j) = \sigma_j^2$. In dealing with choices made in uncertain environments, it is important to distinguish between consequences and actions. In this context,consequences are quality realizations arising out of the actions of brand choice. Accordingly, $v(q)$ is the preference scaling function defined over consequences and $U(x)$ is the utility ordering function defined over the alternative actions of brand choice $U(x)$ is constructed using the expected utility rule proposed by Neumann and Morgenstern. In dealing with certainty choices, utility is treated as ordinal but cardinal utilities are required to ensure that the expected utility rule determines preferences over actions. Such a preference scaling function can be constructed using reference lotteries (see Keeney and Raiffa (1976) and Hirschleifer and Riley (1992) for an elaborate treatment of decision making under uncertainty).

Consumer attitude to risk is critical in determining behavior when faced with uncertainty. We assume that the consumer is risk-averse and not perfectly informed about the available alternatives. While exceptions have been reported, risk-averse behavior can be considered to be the normal case for it is observed that individuals typically hold diversified portfolios and do not invest in a single asset that offers the highest expected value (Hirschleifer and

Riley, 1992) In the context of brand choice, risk-aversion is reported by Erdem and Keane (1996). Risk-aversion and learning consequent to sampling offers a alternative explanation to loyalty, for the observed dependence of current choice on past choices. A consumer is said to be risk-averse if he prefers the expected value of a lottery for certain to the lottery. For a risk-averse consumer, $Ev(q) < v[E(q)]$, and this inequality holds only if $v''(q) < 0$ i.e., the preference scaling function is concave implying diminishing marginal utility. Since the expected utility is computed as the mathematical expectation of the utilities associated with possible consequences, concavity is necessary to ensure that the value associated with a certain consequence is higher than the value associated with an uncertain consequence whose expectation is equal to the certain consequence. A concave preference scaling function represents an interaction of diminishing marginal utility and risk-aversion. The ratio of the second and first derivatives, $-\frac{v''(q)}{v'(q)}$, is used a measure of local absolute risk aversion and preference functions that have equal absolute risk-aversion can be considered to be strategically equivalent.

Following Hirschleifer and Riley (1992), $v(\tilde{q})$ can be expanded in a Taylor's series about its expected value μ

$$v(\tilde{q}) = v(\mu) + \frac{v'(\mu)}{1!}(\tilde{q} - \mu) + \frac{v''(\mu)}{2!}(\tilde{q} - \mu)^2 + \frac{v'''(\mu)}{3!}(\tilde{q} - \mu)^3 + \ldots\ldots\ldots\ldots$$

and the expected value is

$$U = Ev(\tilde{q}) = v(\mu) + \frac{v''(\mu)}{2!}\sigma^2 + \frac{v'''(\mu)}{3!}E(q - \mu)^3 + \ldots\ldots\ldots$$

The curvature of the preference-scaling function and the properties of the probability distribution associated with the action will determine the expected utility. Simplification of the expected utility can be achieved by assumptions regarding the nature of the preference scaling function or the probability distribution of $\tilde{q}$. For instance, if v is quadratic, third and higher order derivatives are zero and the expected utility simplifies to a function of the mean and variance associated with $\tilde{q}$. (However, a quadratic specification would be appropriate in a range since the marginal utility is negative for large values of q. Also, the absolute risk aversion for a quadratic specification is increasing in q). For a normal distribution all higher odd moments are zero and even moments are functions of the mean and standard deviation. Hence the Taylor's series expansion simplifies to two terms.

We assume that the first derivative of $v(q)$ is positive (more is better) and that its second derivative is negative (diminishing marginal utility). Further, we assume that the absolute risk-aversion is constant with respect to quality levels. A specification that is consistent with these assumptions is $v(q) = -e^{-Kq}$, where $K = -\frac{v''(\mu)}{v'(\mu)}$ is a local measure the degree of absolute risk-aversion. This specification has been adopted by Roberts and Urban (1988) for modeling risk and belief dynamics in durable brand choice. We also assume that the consumer's quality perceptions are normally distributed; $\bar{q} \sim N(\mu, \sigma^2)$. Under these assumptions,

$$Ev(q) = -\int_{-\infty}^{\infty} e^{-Kq} \frac{1}{\sqrt{2\pi}} \exp\{-\tfrac{1}{2}\left(\tfrac{q-\mu}{\sigma}\right)^2\} dq = -e^{-K(\mu - \frac{1}{2}K\sigma^2)}$$

It follows that preference,when absolute risk-aversion is constant and beliefs are normally distributed, can be represented by the indirect utility function

$$U(\mu, \sigma) = \mu - \tfrac{1}{2}K\sigma^2$$

The consumer's valuation of an alternative is at the mean quality perception offset by the effect of uncertainty. The valuation is decreasing in the variance associated with the quality perception and K is a factor that scales the effect of uncertain perceptions. In a sequential decision making context, each consumption experience would provide information leading to updations to the mean and variance. A forward looking consumer would associate a future reward with every choice over and above the immediate consumption utility and choose to maximize utility expected over a horizon. Valuations for periods ahead of the current would be characterized by a reduction in the offsetting due to uncertainty in the perception.

Indexing brands with $j \in (1, 2, ..., J)$, the consumer's expected utilities for the alternatives are

$$U_j(\mu_j, \sigma_j) = \mu_j - \frac{1}{2}K\sigma_j^2, (j \in (1, 2, ..., J) \tag{1}$$

We index time periods using $t \in (1, 2, ..., T)$ and define a state variable s_{jt} with $s_{jt} = n$ if brand j has been sampled n times by time period t, so that the vector $s_t = (s_{1t}, s_{2t}, ..., s_{Jt})$ characterizes the consumers information set with regard to J brands at time period t.The expected utilities in these states for each brand are

$$u_j^*(s_{jt}, t) = \mu_{js_{jt}} - \frac{1}{2}K\sigma_{js_{jt}}^2 \tag{2}$$

and a forward looking consumer's future evaluation of the alternative j after it has been sampled n times is described by

$$u_j^*(s_{jt}, t+n) = \mu_{js_{jt}} - g_j^n\left(\frac{1}{2}K\sigma_{js_{jt}}^2\right) \tag{3}$$

where $g_j \in (0, 1)$ for all j. The future evaluation takes into account the effect of learning from consumption experience with g_j capturing the reduction, due to sampling the alternative, in the offsetting (from the mean perception) due to the uncertainty in quality perception. This expected reduction is the information value associated with current choice. In equation 2.3, $u_j^*(s_{jt}, t+n)$ represents the future expected utility associated with alternative j after it has been sampled n times. This is the consumer's evaluation of utility to be derived in the future conditional on information now (s_{jt}) and n samplings; the evaluation takes into account the information value of sampling. The consumer expects the effect of uncertainty to reduce by a multiplicative factor g_j with every consumption experience.

We adopt an additive multiattribute specification (Lancaster, 1966) of net utility with a state specific Neumann-Morgenstern expected utility and price as the two attributes. Incorporating a stochastic component ϵ_{jt},

$$u_{jt} = u_j^* (s_{jt}, t) + \beta_p P_{jt} + \epsilon_{jt} \quad (4)$$

The stochastic component represents what is unknown to the econometrician and is assumed to be independent and identically distributed across brands and occasions. Since it is not associated with the evolving consumer quality perceptions, the issue of serial correlation does not arise With this preference structure, the consumer chooses among the alternatives in each period to maximize the expected value over a horizon of an objective function of the form

$$\left(\sum_{t=0}^{T} \sum_{j=1}^{J} d_{jt} \delta^t \left[u_j^* (s_{jt}, t) + \beta_p P_{jt} + \epsilon_{jt} \right] \right)$$

where δ is a discount factor; $\delta \in (0, 1)$. There are no a priori considerations that would determine the appropriate horizon over which the objective function is computed. In this model, this would depend on the parameters g_j, since this determines the future value of current choice. The data could provide an answer to the question of an appropriate horizon. This is a stochastic dynamic programming problem that is typically solved using backward recursion. In the next section, an overview of the dynamic discrete choice framework is provided and estimation issues associated with this model are discussed.

3 ESTIMATION OF DYNAMIC DISCRETE CHOICE MODELS

The dynamic programming framework has been extensively used in the engineering. mathematics (stochastic control problems), economics fields and is now gaining ground in the marketing literature. This extensive use stems from the fact the framework is rich enough to model a variety of situations involving choices made over time and in uncertain environments. Recent applications in marketing where the dynamic programming framework has been used include Erdem and Keane (1995) and Gonul and Srinivasan (1996). In the former study, the authors hypothesize that consumers learn with choices made in each period since the decisions are made under uncertainty and model choice probabilities as functions of past choices in a Bayesian learning framework. Gonul and Srinivasan (1996) hypothesize that consumers form expectations about the availability of coupons in future periods and their decisions to buy or not to buy in the current period are influenced by these expectations. In these and other similar situations, the dynamic programming framework can provide a good empirical model of how decision makers in the real world actually behave, in additional to providing a normative theory of rational behavior. The dynamic programming framework is detailed in this section.

3.1 The Framework

Consider a typical dynamic discrete choice model which characterizes a consumer making choices over time in an uncertain environment. We assume that the consumer chooses one of J alternatives in each period with the objective of maximizing the expected value of a sum of period specific utilities over a horizon of T periods. Let $d_{tj} = 1$ if alternative j is chosen in period t and $d_{tj} = 0$ if any other alternative is chosen.

$$d_{tj} \in (0, 1) \text{ for all } t, j \in T \times J \text{ and} \sum_{j=1}^{J} d_{tj} = 1 \text{ for all } t \in T$$

The choice made in period t affects the outcome which arrives at the end of the period i.e., if H_t is the history at the beginning of period t, H_{t+1} is either fully determined by the choice made in period t or is determined according to some transition probabilities. In this application,choice fully determines the outcome; H_t is a vector whose elements are the number of observed consumption experiences for each of the J alternatives i.e.,

$$H_t = \{H_{1t}, H_{2t}, ..., H_{Jt}\}$$

and if alternative j is chosen,

$$H_{jt+1} = H_{jt} + 1 \text{ and } H_{it+1} = H_{it} \text{ for all } i \in J, i \neq j$$

In other applications, outcomes are associated with choices in terms of vectors of transition probabilities $F_j(H_t)$ that define the chances of the different outcomes being realized when alternative j is chosen. In such cases, H_{jt+1} evolves stochastically from H_t.

Let u^*_{jt} be the expected utility to be obtained in period t if alternative j is chosen. This expected utility depends on the history H_t and is usually specified in terms of a vector of structural parameters θ . Adding a stochastic component to this expected utility we have,

$$u_{jt} = u^*_{jt}(H_t) + \epsilon_{jt} \tag{5}$$

where ϵ_{jt} follows a well defined joint probability distribution function $G(\epsilon_{1t,}\epsilon_{2t}, ..., \epsilon_{Jt})$.

In this setting, the consumer sequentially chooses d_t to maximize expected utility over a horizon viz.,

$$E_0\left(\sum_{t=0}^{T}\sum_{j=1}^{J} d_{jt}\delta^t\left[u^*_{jt}(H_t) + \epsilon_{jt}\right]\right) \tag{6}$$

where E_0 is the expectation, conditional on information at $t = 0$, with respect to the controlled stochastic process induced by the decision rule. Let the conditional valuation function associated with choosing alternative j in period t be defined as

$$v_{jt} = v_j\left(H_t\right) = E\left(\sum_{s=t+1}^{T}\sum_{j=1}^{J} d_{js}^{o}\delta^{t}\left[u_{js}^{*} + \epsilon_{jt}\right]|d_{jt} = 1\right), \tag{7}$$

where d_{js}^{o} denotes the consumer's optimal choice in period s.

Optimal decision making implies that $d_{kt}^{o} = 1$ if and only if

$$k = \arg\max_{j\in J}\left[u_{jt}^{*} + \epsilon_{jt} + v_{jt}\right] \tag{8}$$

If we assume that ϵ_{jt} is an IID extreme value process, we obtain a dynamic version of the multinomial logit model.

$$\Pr\left(d_{kt} = 1\right) = \frac{\exp\left(u_{kt}^{*} + v_{kt}\right)}{\sum_{j=1}^{J}\left(\exp\left(u_{jt}^{*} + v_{jt}\right)\right)} \tag{9}$$

The only way in which the above is different from the static logit model is that here we have the sum of a one period utility u_{jt} and the expected discounted utility in all future periods over the horizon being considered v_{jt} while in the static logit model only the current utility u_{jt} is considered. It is this difference that makes dynamic discrete decision processes computationally complex. The functional form of the conditional value function is generally not known and its values must be computed numerically for any particular value of the parameter vector θ. Hence, algorithms that are used to find maximum likelihood estimates of the vector of structural parameters θ in such dynamic discrete choice models are constructed in two levels. An outer numerical search algorithm is used to search over the parameter space for values that maximize the likelihood function, and an inner algorithm solves the dynamic programming problem and computes choice probabilities for the current set of parameter values. The dynamic programming problem is usually solved using recursive backward induction methods in finite horizon cases. This recursive computation of the value function has to be repeated every time the parameter values are changed in the outer search algorithm. Clearly, this is a computationally expensive methodology that leads to researchers going to great lengths to simplify the problem and the state space.

The Conditional Choice Probability (CCP) methodology developed by Hotz and Miller (1993) is one way of simplifying the estimation problem. The main advantage of the CCP methodology is that it leads to a significant reduction in the computational cost associated with the estimation of dynamic discrete choice problems. A nested numerical solution method using the Bellman equation is avoided. The CCP methodology is detailed in an appendix and illustrated in an application of the methodology in estimating a dynamic brand choice model incorporating learning.The methodology involves the inversion of conditional (on state) choice probability estimates to arrive at estimates of normalized value functions. The CCPs are also used as weights for different states in computing the expected utility over a horizon.

3.2 Models

Different versions of the model of consumer behavior described in the second section, based on the number of updations to the mean, variance and the horizon over which the objective function is computed are considered for estimation. In the first version a two period horizon is assumed and a single updation to the mean and variance of the quality perception are considered. In terms of the notation adopted earlier in this section

$$u_j^*(0,t) = \mu_{j0} - \frac{1}{2}K\sigma_{j0}^2 \tag{10}$$

and

$$u_j^*(0,t+1) = \mu_{j0} - g_j \frac{1}{2} K \sigma_{j0}^2 \tag{11}$$

Sampling of an alternative leads to updation of the mean and the variance. Therefore, when $s_j = 1$,

$$u_j^*(1,t) = \mu_{j1} - \frac{1}{2} K \sigma_{j1}^2 \tag{12}$$

and

$$u_j^*(1,t+1) = \mu_{j1} - g_j \frac{1}{2} K \sigma_{j1}^2 \tag{13}$$

In the proposed model, the consumer makes his choice on each occasion to maximize a value function comprising of the utility in the current period and expected utility in the future. The value function for each alternative is assumed to be the sum of current utility for the alternative and the expected utility in the next period.

$$v_{jt} = u_{jt} + E\left[u_{t+1}/d_{jt} = 1\right] \tag{14}$$

Next, a three period horizon is considered with a single updation to the mean and variance. The specifications of the expected utility in the current period and next period conditional on sampling remain as above. The expected utility after two periods is

$$u_j^*(0,t+2) = \mu_{j0} - g_j^2 \frac{1}{2} K \sigma_{j0}^2 \tag{15}$$

and

$$u_j^*(1,t+2) = \mu_{j1} - g_j^2 \frac{1}{2} K \sigma_{j1}^2 \tag{16}$$

The consumer chooses to maximize an objective function with a three period horizon. The conditional valuation functions for each alternative are of the form

$$v_{jt} = u_{jt} + E\left(\sum_{s=1}^{2}[u_{t+s}/d_{jt} = 1]\right) \tag{17}$$

The purpose of considering both these versions is to obtain an empirical answer to the question of how far ahead the consumer looks. In this application, there is no a priori basis for fixing the horizon over which the consumer computes his value function. This would depend on the future rewards that the consumer expects from current choice and these rewards are tied to g_j. However, if $g_j \to 0$ for all j ; there is little difference between the expected utility in the next period and subsequent periods. In such a situation, a two period horizon for the value function would seem to be appropriate. These two models are referred to as single updation models in later sections.

Next, a model with an enlarged state space is considered. Learning effects for each of three consecutive consumption experiences are considered; i.e., μ_{s_j} and $\sigma^2_{s_j}$ are allowed to vary for $s_j = 0, 1, 2$, and 3. Thus,

$$\mu_{s_j} = \sum_{i=0}^{s_j}(\mu_{ji}) \text{ for } s_j = 0, 1, 2 \text{ and } 3$$

and

$$\frac{1}{2}K\sigma^2_{s_j} = \frac{1}{2}K\left(\sigma^2_{j0} - \sum_{i=1}^{s_j}\sigma^2_{ji}\right) \text{ for } s_j = 1, 2 \text{ and } 3$$

In later sections, this is referred to as the multiple updation model. The three versions of the learning model that are described above are estimated.

3.3 Identification Issues

The forward looking learning models proposed earlier are now reformulated in terms of parameters that are identified. In our model, expected utility is expressed in terms of the mean quality peception μ_j, the effect of risk-aversion $\frac{1}{2}K\sigma_j^2$, and the factor g_j which captures the consumer's evaluation of future benefits from current choice in terms of expected reduction in variance. In discrete choice models, choice probabilities are constructed from differences in the utitlities associated with the alternatives. Since $\mu_j$and $\frac{1}{2}K\sigma_j^2$ are both associated with current choice, to enable identification, we consider the expected utility in the current period and the expected reduction in the effect of risk-aversion in the future periods as a result of sampling. Let the utility in the current period for alternative j be γ_j,

$$\gamma_j = \mu_j - \frac{1}{2}K\sigma_j^2 \tag{18}$$

The consumer expects sampling to lead to more information and evaluates the future effect of uncertain perceptions conditional on sampling at $g_j\frac{1}{2}K\sigma_j^2$. Let the increase in utility that the consumer expects to gain in the next period from sampling alternative j now be c_j,

$$c_j = (1 - g_j)\frac{1}{2}K\sigma_j^2 \tag{19}$$

Incorporating definitions 18 and 19 in 10 , 11, 12 and 13; we can express current and future expected utilities for alternative j in states $s_{jt} = 0$ and $s_{jt} = 1$ as

$$u_j^*(0,t) = \gamma_{j0}, \tag{20}$$

$$u_j^*(0,t+1) = \gamma_{j0} + c_{j0} \tag{21}$$

and

$$u_j^*(1,t) = \gamma_{j0} + \gamma_{j1} \tag{22}$$

$$u_j^*(1,t+1) = \gamma_{j0} + \gamma_{j1} + c_{j0} - c_{j1} \tag{23}$$

where γ_{j1} and c_{j1} are updations to the expected utility and the effect of risk-aversion after one consumption experience. Thus, the expected utility (γ_{j0}) and the expected reduction in the effect of risk-aversion due to sampling once (c_{j0}) are estimated along with updations to these initial values after one consumption (γ_{j1} and c_{j1}) for the single updation model with the value function computed over a two period horizon. In this reformulation in terms of expected utility and expected reduction in the effect of risk-aversion, the reduction factor g_j is not estimated.

In the single updation model with the value function being computed over three periods, the reduction factor g_j is estimated along with the expected current utilities and expected future reward from sampling. In equation 19, the expected reward in the next period of sampling alternative j now is defined as $c_j = (1 - g_j)\frac{1}{2}K\sigma_j^2$. Continuing in this vein, it can be seen that the expected reward to the consumer two periods ahead from sampling alternative j now and in the next period is $\left(1 - g_j^2\right)\frac{1}{2}K\sigma_j^2$. Now,

$$\left(1 - g_j^2\right)\frac{1}{2}K\sigma_j^2 = (1 + g_j)\,c_j$$

so that the consumer's expected utility from sampling alternative j twice in each of the two states is

$$u_j^*(0,t+2) = \gamma_{j0} + c_{j0} + g_j c_{j0} \tag{24}$$

and

$$u_j^*(1,t+2) = \gamma_{j0} + \gamma_{j1} + c_{j0} - c_{j1} + g_j c_{j0} - g_j c_{j1} \tag{25}$$

respectively. Thus, in the single updation model with a horizon of three periods (including the current period) the expected reward from sampling an alternative twice is modeled as in 24 and 25. Again, initial expected utilities γ_{j0} and expected future rewards from sampling c_{j0} for each alternative and updations to these initial values after one sampling are estimated.

Finally, in the multiple updation model with value functions computed over two periods, updations to expected utilities and the expected reduction in the effect of risk-aversion for each of the first three experiences are estimated i.e.,

$$\gamma_j = \gamma_{j0} + \sum_{i=1}^{s_j} (\gamma_{ji}) \text{ for } s_j = 1, 2 \text{ and } 3$$

and

$$c_j = c_{j0} - \sum_{i=1}^{s_j} (c_{ji}) \text{ for } s_j = 1, 2 \text{ and } 3$$

In this model, we have assumed that initial conditions are unknown to the econometrician. Initialization of initial conditions is not appropriate if there is information decay and frequent repositioning. Structural parameters of interest (mean quality perception μ_j and the effect of risk-aversion c_j) are not seperately identified. However, the preference structure of the representative consumer changes with changes in learning states. This enables identification of updations to structural parameters. Inclusion of future expected utility in the value function enables a close look at the influence of risk-aversion. Identification of the expected reduction in the effect of risk-aversion is facilitated by including future expected utility in the consumer's value function. Since reduction in the effect of risk-aversion is modeled multiplicatively, the reduction factor is identified in the two-period ahead expected utility component of the value function. If the value function is computed over three periods, the reduction in cost of experimenting is identified.

4 DATA, RESULTS AND DISCUSSION

The forward looking models proposed in this study and myopic models (included for comparison) are calibrated using scanner panel data on purchases of ketchup and tuna. The analysis is restricted to households that have purchased one out of the top four (in terms of market share) brand size combinations so as to keep the state space small. In addition, households that have less than five observations are eliminated to enable capture of dynamics in choice behavior. In this study, marketing mix variables other than price are not included. Price is incorporated as an additive component in the utility function and not as a state variable. We assume that the consumer evaluates future choices on the basis of current prices. This is done to simplify the state space and facilitate focus on learning effects. We are interested in examining if incorporation of learning effects leads to a change in the price coefficeint. Low correlations between price and the omitted marketing mix variables (display and feature) in these data sets suggest that such interpretations are not confounded by this omission. Descriptive statistics for the ketchup and tuna datasets are reported in tables 1 and 2 respectively.

First, a baseline static logit model (Model 1) with brand specific intercepts and a price parameter that is common across alternatives is estimated. Here, the utility specification is

$$u_{jt} = \mu_j + \beta_p P_{jt} + \epsilon_{jt}$$

Next, a model (Model 2) that includes the loyalty measure introduced by Guadagni and Little [1983] apart from brand specific intercepts and a price parameter that is common across alternatives is estimated. The loyalty measure is included in the utility specification thereby increasing the probability of an alternative being chosen in the future with each consumption. It has been observed by researchers that this reduced form specification does very well in terms of fit and prediction and is therefore included in studies as a benchmark model. Here, the utility specification is

$$u_{jt} = \mu_j + \beta_p P_{jt} + \beta_l L_{jt} + \epsilon_{jt}$$

The Guadagni and Little [1983] loyalty measure implies a rigid updation pattern that is common across alternatives with recent past choices being weighted more than those in the distant past. In Model 3, free updations to the intercepts upon consumption experience are estimated for each alternative . Starting with an initial intrinsic preference μ_{j0} for alternative j three consumption experiences are each allowed to add to this initial preference, so that

$$\mu_j = \mu_{j0} + \sum_{i=1}^{s_j} (\mu_{ji}) \text{ for } s_j = 1, 2 \text{ and } 3$$

In this model, consumers are assumed to be myopic and do not consider the future value of current choice. Consumers choose on every occasion to maximize current utility. This model is included for purposes of comparison since in the forward looking model that is proposed as many updations are allowed for as in this model.

Model 4 is the single updation model, introduced in earlier. This is a forward looking model incorporating the future value of current choice with the consumer making choices to maximize a value function over a horizon of two periods. In ths model, the expected utility in the current period (γ_{j0}) and the expected reduction in the effect of risk-aversion as a result of sampling (c_{j0}) are estimated. Updations to these parameters from one consumption experience (γ_{j1} and c_{j1}) are also estimated.

The next model (Model 5) is the single updation forward looking model with the objective function being computed over a three period horizon. In this model, the parameters γ_{j0}, c_{j0}, γ_{j1}, c_{j1} and the reduction factor g_j are estimated for each alternative j. As mentioned earlier, the reduction factor g_j is identified only in a model where the value function is computed over a three period horizon.

Brand	Market Share	Mean Price (std dev)
Heinz (HE)	0.44	1.17 (0.13)
Hunt (HU)	0.24	1.15 (0.16)
Del Monte (DM)	0.17	1.06 (0.19)
Generic (GE)	0.16	0.77 (0.04)

Table 1: Descriptive Statistics - Ketchup

Brand	Market Share	Mean Price (std dev)
Star Kist Water (SKW)	0.40	0.672 (0.095)
Star Kist Oil (SKO)	0.13	0.673 (0.096)
Ckn of Sea Water (CSW)	0.37	0.665 (0.110)
Ckn of Sea Oil (CSO)	0.10	0.666 (0.110)

Table 2: Descriptive Statistics - Tuna

Finally, the multiple updation model with a two period horizon for the value function is estimated (Model 6). In this model, initial expected utilities and expected reduction in risk-aversion along with updations after each of three consumption experiences are estimated.

The value functions in the forward looking models are represented as weighted sums of period specific utilities with the conditional choice probabilities being used as weights. Hence, all choice probabilities take the simple multinomial logit forms and maximum likelihood estimation is undertaken to obtain parameter estimates. Sample frequencies are used as estimates of the conditional choice probabilities used in computing the value functions in models 4 and 5. The state space is larger in model 6 where each of four consumption experiences are considered for updating consumer beliefs for each of the four alternatives. Kernel estimates of the conditional choice probabilities are computed using sample frequen-

Variable	Model 1	Model 2	Model 3
HE-I	2.97 (30.0)	2.59 (28.8)	3.44 (27.4)
HU-I	1.65 (23.6)	1.46 (20.9)	0.37 (3.7)
DM-I	1.31 (18.7)	1.26 (18.9)	0.46 (3.7)
GE-I	0	0	0
Price	-5.03 (-27.9)	-4.82 (-25.4)	-5.23 (-26.8)
Loyalty		3.33 (14.5)	
HE-U 1, 2, 3			0.45 (3.7), 1.91(18.8), 0.83 (7.5)
HU-U 1, 2, 3			0.25 (1.8), 0.55 (3.9), 1.68 (15.8)
DM-U 1, 3, 3			0.87 (7.1), 0.44 (2.5), 0.56 (3.2)
GE-U 1, 2, 3			1.19 (9.4), 0.99 (5.5), 1.04 (5.6)
Log Likelihood	-4565.85	-4043.63	-3987.93

Table 3: Myopic models - Ketchup (t stats in parantheses)

cies for each state and these estimates are used to compute the value functions in model 6.The estimation procedure is detailed in the appendix.

The estimates for the six models outlined above for the ketchup and tuna data sets are presented in Tables 1 to 6. Models 1, 2 and 3 are comparable (since consumers are represented as being myopic) and are grouped together. Tables 3 and 4 contain the estimates for these myopic models 1, 2 and 3 for the ketchup and tuna data sets respectively. Abbreviations are used to represent parameters in the tables: HE (Heinz), HU (Hunt), DM (Del Monte) and GE (Generic) are the ketchup brands; SKW (Star Kist Water), SKO (Star Kist Oil), CSW (Chicken of Sea Water) and CSO (Chicken of Sea Oil) are the tuna brands. EU, FR and RF denote expected utility, future reward and reduction factor respectively; I and U denote initial values and updations. For example, DM-EU 3 denotes the third updation in expected utility for the Del Monte brand.

The addition of a loyalty variable improves the log-likelihood significantly in both data sets. For ketchup, the price parameter decreases with the inclusion of loyalty but increases in the case of tuna. All the updations in model 3 are significant and the log-likelihood improves

Variable	Model 1	Model 2	Model 3
SKW-I	1.54 (38.5)	0.90 (22.5)	1.28 (14.4)
SKO-I	0.43 (10.8)	0.29 (5.8)	1.05 (10.7)
CSW-I	1.34 (43.5)	0.76 (19.0)	0.64 (5.7)
CSO-I	0	0	0
Price	-12.6 (-7.0)	-15.4 (-4.9)	-14.2 (-26.8)
Loyalty		4.27 (47.5)	
SKW-U 1, 2, 3			0.95 (10.7), 0.56 (6.1), 1.12 (10.7)
SKO-U 1, 2, 3			0.60 (4.4), 1.18(10.0), 0.93 (10.9)
CSW-U 1, 2, 3			1.43 (15.5), 0.42 (4.0), 1.16 (13.7)
CSO-U 1, 2, 3			1.52 (12.6), 0.53 (3.5), 1.39 (10.6)
Log Likelihood	-12441.5	-8869.0	-9096.0

Table 4: Myopic models - Tuna (t stats in parantheses)

for both data sets. Model 3 performs better than Model 2 for ketchup, indicating that free updations capture better the change in intrinsic preferences than the exponential weighting scheme used in creating the loyalty term. In the case of tuna, however, Model 2 has a better log-likelihood than Model 3. This could be because the average number of observations per household is higher for tuna (25) than for ketchup (10).

The estimates for models 4, 5 and 6 for the two data sets in tables 5 and 6. These models represent the consumer as being forward looking. On each purchase occasion, the consumer chooses among the alternatives to maximize an objective function that includes current utility and expected utility in future periods. In model 5 the horizon is two periods into the future; models 4 and 6 are based on a one period ahead value function. Model 6 is comparable to the myopic model 3 in terms of the number of updations that are estimated. The intertemporal discount factor was set at 0.9 for all the forward looking models. The expected utility (γ) and the future reward (c) are set to 0 and e respectively, for the generic and Ckn of Sea Water brands while estimating the forward looking models using the ketchup and tuna data. The forward looking models fit the data better than the myopic models in

	Model 4	Model 5	Model 6
HE-EU I (γ)	-0.19 (-0.5)	-0.61 (-1.1)	-0.51 (-1.7)
HE-EU U 1, 2, 3	0.64 (1.7)	0.58 (1.1)	0.14 (0.7), 0.01 (0.1), -0.36 (-2.7)
HU-EU I (γ)	-0.23 (-1.0)	-0.35 (-0.8)	-0.62 (-1.0)
HU-EU U 1, 2, 3	0.58 (2.8)	0.63 (1.6)	0.58 (0.5), -0.07 (-0.2), 0.17 (1.1)
DM-EU I (γ)	-0.65 (-3.1)	-0.66 (-2.0)	-0.65 (-0.5)
DM-EU U 1, 2, 3	0.86 (3.6)	0.98 (2.7)	0.41 (0.1), -0.24 (-0.7), -0.07 (-0.1)
GE-EU I (γ)	0.00	0.00	0.00
GE-EU U 1, 2, 3	0.73 (7.9)	0.47 (1.3)	0.53 (4.6), 0.21 (1.3), 0.01 (0.1)
HE-FR I $c_j'(c_j = \exp(c_j'))$	2.01 (15.6)	1.81 (9.1)	2.08 (23.9)
HE-FR U 1, 2, 3	0.96 (2.4)	0.46 (0.8)	-0.01 (-0.1), -0.42 (0.4), -2.10 (-2.4)
HU-FR I (c_j')	1.92 (2.4)	1.65 (6.1)	2.07 (7.1)
HU-FR U 1, 2, 3	-0.43 (-0.4)	-0.47 (-0.3)	-0.55 (-0.1), -0.64 (-0.3), -1.10 (-2.8)
DM-FR I (c_j')	2.40 (27.9)	1.99 (11.1)	2.34 (3.2)
DM-FR U 1, 2, 3	1.42 (5.5)	1.04 (3.0)	0.09 (0.1), -0.01 (-0.1), -0.16 (-0.1)
GE-FR I (c_j')	1.00	1.00	1.00
GE-FR U 1, 2, 3	-4.32 (-26.5)	-4.30 (-0.1)	-4.22 (-36.7), -2.51 (-14.2), -1.74 (-5.8)
HE-RF $g'(g = \exp(g')/1 + \exp(g'))$		-4.56 (-0.1)	
HU-RF		-4.71 (-0.1)	
DM-RF		-4.64 (-0.1)	
GE-RF		-4.69 (-0.1)	
Price	-1.85 (-8.3)	-1.81 (-7.9)	-1.49 (-5.3)
Log-likelihood	-3690.6	-3688.1	-3668.3

Table 5: Forward Looking Models:Ketchup (t stats in parantheses)

both data sets. In the single updation model (4), most of the expected utility parameters (γ) and the future rewards (c) parameters are significant in both data sets. However, in the three updation model (6), the initial expected utilities and future rewards are significant but many of the updations to these initial values are insignificant in the ketchup estimates. The tuna data set with longer purchase histories yields better estimates.

The magnitude of the initial parameters and the changes after a few experiences in the forward looking models suggest that consumers have differing perceptions across alternatives.

	Model 4	Model 5	Model 6
SKW-EU I (γ)	0.47 (2.6)	0.44 (2.1)	0.59 (3.3)
SKW-EU U 1, 2, 3	1.14 (6.3)	1.11 (5.0)	0.74 (3.0), 0.29 (0.9), 0.74 (2.6)
SKO-EU I (γ)	0.15 (1.2)	0.18 (1.4)	0.08 (0.7)
SKO-EU U 1, 2, 3	0.51 (4.1)	0.50 (3.6)	0.53 (3.1), 0.20 (0.9), 0.70 (3.1)
CSW-EU I (γ)	0.26 (1.5)	0.28 (1.4)	0.50 (2.9)
CSW-EU U 1, 2, 3	0.73 (4.1)	0.72 (3.4)	0.88 (3.9), -0.14 (-0.4), 0.58 (2.0)
CSO-EU I (γ)	0.00	0.00	0.00
CSW-EU U 1, 2, 3	1.37 (11.2)	1.19 (7.1)	1.06 (5.6), 0.12 (0.4), 0.68 (2.3)
SKW-FR I c_j'	1.30 (10.9)	1.16 (7.1)	1.19 (8.9)
SKW-FR U 1, 2, 3	-0.47 (-0.7)	-0.46 (-0.8)	-1.28 (-0.6), -0.40 (-0.4), 0.27 (0.4)
SKO-FR I (c_j')	1.73 (15.7)	1.48 (13.3)	1.78 (16.4)
SKO-FR U 1, 2, 3	-1.37 (-0.6)	-1.36 (-0.8)	-1.64 (-0.4), -0.08 (-0.1), -0.03 (-0.1)
CSW-FR I (c_j')	1.40 (12.1)	1.18 (7.9)	1.14 (7.5)
CSW-FR U 1, 2, 3	-2.84 (-0.4)	-2.84 (-0.5)	-2.33 (-0.4), -3.36 (-0.2), -4.89 (-0.1)
CSO-FR I (c_j')	1.00	1.00	1.00
CSO-FR U 1. 2, 3	-9.3 (-0.1)	-5.3 (-0.1)	-7.58 (-0.1), -4.97 (-0.1), -3.32 (-0.2)
SKW-RF g'		-2.99 (-0.6)	
SKO-RF		-2.80 (-0.8)	
CSW-RF		-2.88 (-0.7)	
CSO-RF		-2.97 (-0.5)	
Price	-8.77 (-27.9)	-8.77 (-26.8)	-10.06 (-28.5)
Log-likelihood	-8831.6	-8792.8	-8723.2

Table 6: Forward Looking Models:Tuna (t stats in parantheses)

Model 5 estimates enable a close look at the uncertainty around quality perceptions for each of the alternatives. If the parameter g is close to zero, the consumer expects to totally resolve the uncertainty around quality perceptions and the information value of sampling (the future reward c) is close to the offsetting due to risk-aversion in the current period. The estimates of the reduction factors (g) are close to zero and insignificant; consumers do not seem to look beyond the immediate future for these low value products. The expected future reward can then be added back to the expected utility γ to arrive at the mean quality perception μ. The parameters estimated in Model 5 are used to arrive at descriptions of

Brand	Mean	Std Dev	Coeff of Var
Heinz	5.11	3.20	0.63
Hunt	4.21	2.90	0.69
Del Monte	4.15	3.00	0.72
Generic	1.72	1.85	1.08

Table 7: Quality Perceptions - Ketchup

the uncertainty associated with quality perceptions. Estimates of the mean and variance of quality perceptions for the ketchup brands after updation are provided in Table 7. The estimates have been computed assuming that the absolute risk aversion constant K is equal to 1. K scales the effect of uncertainty in the utility function and the model captures the product of K and σ^2. Ascertaining the value of K would help in rescaling the variances appropriately although the effect of risk-aversion would remain unchanged. The probability distributions along with the preference scaling function are plotted in Figure 1. The preference scaling function has been vertically shifted for ease of viewing. Heinz and the Generic brand have the highest and lowest mean quality levels. These brands are also characterized by the lowest and highest coefficients of variation respectively. Hunt and Del Monte are similar to each other in this regard.

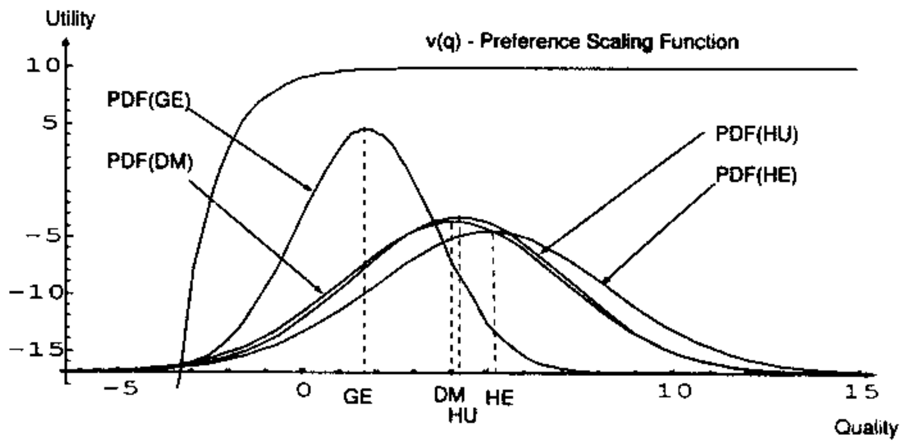

Figure 1: "Quality Preference and Uncertain Perceptions - Ketchup"

Model 6 provides estimates of initial perceptions and three updations. Heinz, Hunt and Del Monte have mean perceptions and expected future rewards that are similar in magnitude to start with. Subsequent to three consumption experiences, future rewards associated with Heinz, Hunt and Del Monte decrease while the generic brand remains a costly experiment. However, in terms expected utility in the current period, the generic brands fare well subsequent to sampling. Among the tuna brands, the initial expected future reward is highest for Star Kist Oil and lowest for Ckn of Sea Oil; the other two brands are similar in this respect. The magnitude of reduction in the expected future reward after sampling is high for the two Star Kist brands relative to the Ckn of Sea brands. Expected utilities improve for all brands suggesting that experiences are positive and inducing trial can lead to retention. As hypothesized, the price parameter in the forward looking models is significantly different from that in the myopic models. Even a two updation forward looking model (4) performs better than all the myopic models.

5 CONCLUSIONS

In this paper, a simple structural model of dynamic brand choice that takes into account the uncertainty faced by the consumers and the consequent information value associated with choice is proposed. The model is formulated with ease of estimation as an objective. The consumer, while making brand choice, maximizes his expected utility over a finite horizon based on uncertain preferences that are dynamically updated with consumption experience. The consumer is assumed to be risk-averse with the degree of risk-aversion remaining constant across quality levels. Uncertainty in quality perceptions is incorporated by assuming that perceptions are normally distributed. Brands are allowed to vary in their ability to influence long term preferences with consumption experience since quality means and variances are assumed to be brand specific.

The model suggests that estimates of the impact of marketing-mix variables will be biased if the information factor is not taken into account. The proposed dynamic, forward looking model is estimated using the conditional choice probability (CCP) estimation procedure, developed by Hotz and Miller [1993]. This methodology is characterized by considerably lower computation costs compared to the extant methods of estimating dynamic models. Results are consistent with the proposed model of consumer behavior. Models that take future value of current choice into account perform significantly better than myopic models. The hypothesis that consumers associate significant costs of experimenting with alternatives that have not been sampled in the recent past is supported.

Further research in this area is necessary to isolate the various elements that contribute to the consumer's perception of the cost of sampling. The quality signals that emanate from prices need to be looked at. Rich structural models that incorporate price expectations, stockpiling effects, etc along with learning effects are interesting options for further research.

Incorporating consumer heterogeneity in such dynamic models is a challenging avenue for further research. While the differing states that are the basis for the dynamics in the proposed models are also likely to capture consumer heterogeneity, differing initial conditions are not allowed for. Future research could attempt to develop models that capture the essence of differences in initial conditions and disentangle these from choice dynamics.

6 REFERENCES

1. Aguirregabria, V. [1994]: "Dynamic Models of Price and Inventory Decisions with Fixed Adjustment Costs," Ph.D dissertation, Universidad Complutense de Madrid.

2. Allenby, G. and P. Rossi [1991] : "Quality Perceptions and Asymmetric Switching between Brands," Marketing Science, Summer, 3, 185-204.

3. Altug, S. and R. Miller [1995]: "The Effect of Work Experience on Female Wages and Labor Supply," Working Paper, Carnegie Mellon University.

4. Blattberg, R. and K. Wisniewski [1989]: "Price Induced Patterns of Competition," Marketing Science, Fall, 4, 291-309.

5. Erdem, T. and Keane [1996]: "Decision-making Under Uncertainty : Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," Marketing Science, Fall, 15, 1-20.

6. Fader, P. and J. Lattin [1993]: "Accounting for Heterogeneity and Non Stationarity in a Cross Sectional Model of Consumer Purchase Behavior," Marketing Science, Fall, 12. 304-316.

7. Gonul. F and K. Srinivasan [1996]: "Estimating the Impact of consumer Expectations of Coupons on Purchase Behavior: A Dynamic Structural Model, Marketing Science," 13, 262-279

8. Guadagni, P. and J. Little [1983]: "A Logit Model of Brand Choice Calibrated on Scanner Data, Marketing Science," Summer, 2, 203-238.

9. Hardie, B., E. Johnson and P. Fader [1993]: "Modeling Loss Aversion and Reference Dependence Effects on Brand Choice," Marketing Science, Fall, 4, 378-394.

10. Hirschleifer, J and J.G. Riley [1992]: "The Analysis of Uncertainty and Information", Cambridge University Press, 1992.

11. Hotz, V. and R. Miller [1993]: "Conditional Choice Probabilities and the Estimation of Dynamic Models," Review of Economic Studies, 60, 497-529.

12. Hotz, V., R.Miller, S. Sanders and G. Smith [1994]: "A Simulation Estimator for Dynamic Models of Discrete Choice," Review of Economic Studies, 61, 265-289.

13. Kamakura, W. and G. Russel [1989]: "A Probabilistic Choice Model for Market Segmentation and Elasticity Structure," Journal of Marketing Research, 26, 379-390.

14. Keane and Wolpin [1994]: "Solution and Estimation of Dynamic Programming Models by Simulation," Review of Economics and Statistics, 76, 4 (November), 648-672.

15. Krishnamurthi, L. and S.P. Raj [1988]: "A Model of Brand Choice and Purchase Quantity Sensitivities," Marketing Science, Summer, 7, 1-20.

16. Lancaster, Kevin [1966], "A New Approach to Consumer Theory," Journal of Political Economy, 74, 132-157.

17. Roberts, J.H. and Urban, G.L. [1988]: "Modeling Multiattribute Utility, Risk, and Belief Dynamics for New consumer Durable Brand Choice," Management Science, 34, 2, February.

Appendix: The Conditional Choice Probability Methodology

The conditional choice probability (CCP) estimator, developed by Hotz and Miller [1993], adopts an approach which reduces much of the computational burden associated with estimating dynamic discrete choice models. In this approach, a nested numerical solution method using the Bellman equation is avoided. Instead the value function is expressed in terms of the utilities, choice probabilities and the transition probabilities of choices and outcomes that remain feasible in future periods. The basic intuition underlying the CCP estimator is that the normalized value functions (difference between an alternative's value function and that of a base alternative) for each of the $J-1$ alternatives are related to the probabilities of the alternatives being optimal choices conditional on the history. If the consumer were behaving optimally, his choice in each state would be a function of the value functions, given the state, for each of the alternatives. In discrete choice models, the probability of choice is defined in terms of inequalities; the probability of an alternative being chosen is the probability of the alternative's value function being greater than the value functions of all the other alternatives. Hence, the choice probability of an alternative in any state would be a function of the differences (normalized value functions) between this alternative's value function and the value functions of the other alternatives in this state. Given a set of normalized value functions, we can compute choice probabilities by specifying the distribution of the stochastic components. In a similar vein, if the choice probabilities were invertible in the normalized value functions, we could work backward to obtain the latter from the former. [1] Thus, if we are able to obtain consistent estimates of the CCPs, these could be inverted to arrive at estimates of the normalized value functions which could be used estimate the structural parameters θ.

From equation 8, it can be seen that the conditional probability of choosing alternative k in

[1] For a formal proof that the conditional choice probabilities are indeed invertible in the normalized value functions, see Hotz and Miller [1993].

period $t+1$

$$p_{kt+1} = \Pr\left(d_{kt+1} = 1 | H_{t+1}\right)$$

$$= \Pr\left\{\bigcap_{j=1}^{J} \epsilon_{jt+1} \leq \epsilon_{kt+1} + v_{kt+1} - v_{jt+1} + u^*_{kt+1} - u^*_{jt+1} | H_t\right\} \tag{26}$$

equation 26 can be expressed as

$$p_{kt+1} = \int_{-\infty}^{\infty} G_k \left(\begin{array}{c} \epsilon_{kt+1} + u^*_{kt+1} - u^*_{1t+1} + v_{kt+1} - v_{1t+1}, ..., \epsilon_{kt+1} \cdots \\ \cdot, \epsilon_{kt+1} + u^*_{kt+1} - u^*_{Jt+1} + v_{kt+1} - v_{Jt+1} \end{array} \right) d\epsilon_{kt+1} \tag{27}$$

where G_k is the probability density function for ϵ_{kt}, given history H_t when the kth choice is optimal (G_k is the derivative of the joint probability distribution function $G(\epsilon_{1t}, \epsilon_{2t}, ..., \epsilon_{Jt})$ with respect to ϵ_{kt}.

For each k belonging to the choice set, the choice probability is a function of the differences, between the kth alternative and the others, in current utilities and future utilities. Since equation 27 is indeed invertible in the normalized value functions; estimates of the conditional choice probabilities can be used to arrive at estimates of the normalized value functions.

Further, this result allows us to find the expected values of the stochastic components ϵ_t, drawn from a sample that is subject to choice based censoring. Let w_k be the expected value of ϵ_k conditional on the kth alternative being the optimal choice. In any period s,

$$w_k(p_{ks}) = E(\epsilon_{ks} | d_{ks} = 1)$$

$$= \int_{-\infty}^{\infty} \epsilon_{ks} G_k \left(\epsilon_{ks} + u^*_{ks} - u^*_{1s} + v_{ks} - v_{1s}, ..., \epsilon_{ks}, ..., \epsilon_{ks} + u^*_{ks} - u^*_{Js} + v_{ks} - v_{Js}\right) d\epsilon_{ks} \tag{28}$$

To compute the integral defined in equation 28, we need to know the differentials $v_{ks} - v_{1s}$ ... . Let $q_{kt} - q_{1t}$.... be the estimates of the normalized value functions obtained by inverting 0.2, i.e., $q_{kt} - q_{1t} = u^*_{kt} - u^*_{1t} + v_{kt} - v_{1t}$. The expected value in equation 28 can now be computed by using the estimates $q_{kt} - q_{1t}$ obtained by inverting equation 27; i.e.,

$$w_k(p_{ks}) = \int_{-\infty}^{\infty} \epsilon_{ks} G_k(\epsilon_{ks} + q_{ks} - q_{1s}, ..., \epsilon_{ks}, ..., \epsilon_{ks} + q_{ks} - q_{Js})\, d\epsilon_{ks} \tag{29}$$

Optimal behavior by the consumer would mean that the utility he expects to receive in any period in the future in any given state is

$$\sum_{j=1}^{J} p_j \left[u_j^* + w_j(p_j)\right] \tag{30}$$

Here. p_j is the conditional choice probability of alternative j being chosen; and w_j is the expected value of the stochastic component when alternative j is optimal. The conditional choice probabilities are used as the weights to compute the expected utility in each future period and w_j corrects for what is unobserved by the econometrician but known to the consumer.

Hence, an approximation of the valuation function can be obtained by summing the expected utilities from equation 30 over future periods and states (if the evolution of states were stochastic). This approximation can then be used to estimate the structural parameters at a relatively low computational cost instead of solving the dynamic programming problem using recursion. Although this representation reduces the computational burden relative to recursive solutions, it can be cumbersome to implement in applications characterized by relatively large state spaces and long horizons. However, in some special cases this alternate representation of the conditional valuation functions can be used to further simplify the problem. These special cases include problems with terminal actions, stationary markov problems and problems with finite dependence (Hotz and Miller [1993], Aguirregabria [1994] and Altug and Miller [1995]).

Thus, the CCP methodology allows for simplified representation of the conditional value functions which usually do not have closed forms and need to be computed recursively for any set of parameter values. The simplification results from the fact that estimates of the normalized value functions can be obtained by inverting estimates (using sample frequencies) of the conditional choice probabilities. In certain situations, as described above, it is an easy matter to recover the value functions for each alternative from the normalized value functions. In other situations, the value functions will have to be computed by integrating over future paths and states using the conditional choice probabilities as weights. The CCP methodology can again be used to develop a conditional choice simulation technique (Hotz, Miller, Sanders and Smith [1994]) in which integrating over future paths and states is avoided. The conditional choice simulation (CCS) estimator uses estimates of the conditional choice probabilities and the transition probabilities of outcomes resulting from each of the choices to arrive at a simulated value function. The computational ease of the CCS estimator follows from the fact that recursive computation of the value function and integration over future paths are avoided.

In summary, the main ideas that constitute the CCP methodology are 1) estimates of conditional choice probabilities can be inverted to obtain estimates of normalized value functions; 2) estimates of the normalized value functions can be used to estimate expected values of the stochastic components in future periods; 3) the conditional choice probabilities can be used as weights to compute expected utilities in future periods; and 4) estimates of the conditional choice probabilities and transition probabilities relating outcomes to choices can be used to simulate choices and outcomes; and these simulated sequences can be used to construct a simulation estimator of the structural parameters.

Application of the CCP methodology to a single updation model of learning

In this section, the single updation model with the assumption of complete learning from a single consumption experience is described and the implied choice probabilities are derived. Consumers are assumed to make choices that maximize utility in the current period and expected utility in the next period. This model is used to illustrate the CCP estimation methodology.

As before, brands are indexed using $j \in (1, 2, ..., J)$; and time periods as $t \in (1, 2, ..., T)$. The consumers information with respect to the utility of each brand is characterized by the state variable s_{jt}; with $s_{jt} = 1$ if brand j has been sampled by time period t and $s_{jt} = 0$ otherwise,so that the vector $s_t = (s_{1t}, s_{2t}, ..., s_{Jt})$ characterizes the consumers information set at t.

Let $u_j^* (s_{jt}, t)$represent the expected utility to the consumer from choosing brand j at time t. When the consumer is partially informed about a brand, he evaluates the brand on the basis of his priors and takes into account the experimental cost associated with choosing the brand. Thus, we assume

$$u_j^* (o, t) = \mu_{j0} - c_j \tag{31}$$

where μ_{j0} represents the expected value to the consumer of the intrinsic utility associated with brand and c_j is the cost associated with experimenting with brand j; and,

$$u_j^* (o. t + 1) = \mu_{j0} \tag{32}$$

since the uncertainty with respect to the brand gets resolved once it is sampled and there is no longer a cost associated with experimenting. This follows from our assumption of complete learning with one experience. The above evaluations are made by the consumer when he has not sampled alternative j i.e.,when $s_{jt} = 0$. Given imperfect information, the consumer associates an experimental cost c_j with current choice for this alternative leading to a net utility (net of the cost of experimenting) of $\mu_{j0} - c_j$ in the current period. The consumer also assumes that incurring this experimental cost c_j now leads to a better future in that this cost would no longer exist; hence the evaluation of the utility to be derived in the next period of μ_{j0}.

It is also assumed that the consumer updates his beliefs and his intrinsic preference for the brand upon consuming it changes. Thus

$$u_j^* (1, t) = \mu_{j1} \tag{33}$$

for all t. Now, the consumer is in state $s_{jt} = 1$ and no longer has any uncertainty regarding the utility to be derived from the brand since he has sampled it. Hence, there is no longer any cost of experimenting associated with this alternative and this evaluation holds for the current and future periods.

The consumer's indirect utility for brand j in time period t, adjusted for price, is

$$u_{jt} = u_j^*(s_{jt}, t) + \beta_p P_{jt} + \epsilon_{jt} \tag{34}$$

where P_{jt} is the price of brand j at time period t, β_p is the price coefficient and ϵ_{jt} is a stochastic term that is independent across alternatives and time periods.

The consumer is assumed to maximize a value function of the form

$$v_t = u_t + E_t(u_{t+1}) \tag{35}$$

where E_t is the expectation operator conditional on information available at t .

The conditional value function associated with each brand j can be written as

$$v_{jt} = u_{jt} + E[u_{t+1}/d_{jt} = 1] \tag{36}$$

where d_{jt} is an indicator variable that is equal to 1 if brand j is chosen at time period t and is zero otherwise. In equation 35, the first term is u_{jt}, the utility to be derived from alternative j in the current period: this is known upto a vector of parameters (other than the stochastic component ϵ_{jt}). The second term is the conditional (on choosing alternative j) expectation of utility to be derived in the next period. This term can be written as

$$E[u_{t+1}/d_{jt} = 1] = E\left[\sum_{k=1}^{J} d_{kt+1}^{o}\left[u_k^*(s_{kt}, t+1) + \beta_p P_{kt+1} + \epsilon_{kt+1}\right] | s_t, d_{jt} = 1\right] \tag{37}$$

Now, the expected value of alternative k being optimal in period $t+1$ conditional on s_t and $d_{jt}=1$ is the conditional choice probability $\Pr_k(s_{t+1})$, where s_{t+1} is defined by s_t and the alternative chosen in period t. Hence, equation 35 can be expanded out as

$$v_{jt} = u_j^*(s_{jt},t) + \beta_p P_{jt} + \epsilon_{jt} + \sum_{k=1}^{J} \Pr_k(s_{t+1}) \left[u_k^*(s_{kt},t+1) + \beta_p P_{kt+1} + E\left(\epsilon_{kt+1}|s_t, d_{kt+1}=1\right)\right] \quad (38)$$

In the preceding section, the general solution to computing the expected values of the stochastic components was outlined. The estimates of the normalized value functions arrived at by inverting conditional choice probability estimates can be used to compute the expected value of the stochastic component $\epsilon_{kt+1}$when k is the optimal choice given the state. These stochastic components are drawn from a sample that is subject to choice based censoring. If ϵ_{jt} follows an IID extreme value process, it follows that the expectation of ϵ_{kt+1}, conditional on $d_{kt+1}=1$ has the following simple form.

$$E\left(\epsilon_{kt+1}|s_t, d_{kt+1}=1\right) = \eta - \ln\left(\Pr_k(s_{t+1})\right) \quad (39)$$

where η is Euler's constant. It can be seen that the conditional expectation of ϵ_{kt+1} is declining in the conditional choice probability $\Pr_k(s_{t+1})$. The conditional choice probability $\Pr_k(s_{t+1})$ is a function of the normalized value function $v_{kt+1} - v_{1t+1}$, if alternative 1 is treated as the base alternative. As this difference $v_{kt+1} - v_{1t+1}$ increases, so does $\Pr_k(s_{t+1})$, and the threshold value of the difference between the stochastic components $\epsilon_{kt+1} - \epsilon_{1t+1}$ required to induce choice k becomes lower.

Hence, equation 38 simplifies to

$$v_{jt} = u_j^*(s_{jt},t) + \beta_p P_{jt} + \epsilon_{jt} + \sum_{k=1}^{J} \Pr_k(s_{t+1}) \left[u_k^*(s_{kt},t+1) + \beta_p P_{kt+1} + \eta - \ln\left(\Pr_k(s_{t+1})\right)\right] \quad (40)$$

Let $v_{jt}^* = v_{jt} - \epsilon_{jt}$. The probability of choosing alternative k in period t is

$$\Pr_k(s_t) = \Pr\left(k = \arg\max_{j\epsilon J}(v_{jt})\right)$$

and given the distributional assumptions, this probability takes a simple form

$$\Pr_k(s_t) = \frac{\exp\left(u^*_{kt} + v^*_{kt}\right)}{\sum_{j=1}^{J}\left(\exp\left(u^*_{jt} + v^*_{jt}\right)\right)} \tag{41}$$

Implementation of the above representation of the conditional value functions requires consistent estimates of the conditional choice probabilities which are necessary to arrive at the conditional expectations of the stochastic components and the conditional value functions. One possibility is to use sample frequencies as estimates of the conditional choice probabilities. This would require a require a reasonable number of observations in the data set for each of the elements in the state space. If some of the elements in the state space have very few or no observations, sample frequencies are non-operational as estimates of the conditional choice probabilities . In such situations, smoothing techniques can be used to arrive at estimates of the conditional choice probabilities for each element in the state space. Specifically. non-parametric techniques using kernels can be used to consistently estimate the conditional choice probabilities and thus overcome this problem which is likely to occur in most applications with large state spaces.

Consider the general multiple updation model and the single updation models of learning formulated in Section 3. For estimating the single updation model, sample frequencies were directly used as estimates of the conditional choice probabilities for computing the value functions. Since each of the four alternatives could be in one of two states (sampled - 0 and unsampled - 1), the state space had sixteen elements all of which were sufficiently visited in the data sets. In the case of the mulitple updation model. each of the four alternatives is allowed to be in one of 5 states (0 to 4 consumption experiences). Here, not all of the states are sufficiently visited in the data sets. Hence, non-parametric smoothing techniques are used to estimate conditional choice probabilities from the sample frequencies for these two models. The kernel function that was used was

$$K(s, s_d) = \prod_{i=1}^{4} \phi(\frac{s_{di} - s_i}{h})$$

where $\phi(.)$ is the standard normal density function, h is the bandwidth, s_d is the state pertaining to the observation and s is the state that the kernel weight is being computed for. The subscript i refers to the element in the state; there are four elements in this application since there are four alternatives in the choice set. Using these kernel weights, the conditional choice probability estimates are computed for each of the possible states s in the state space as

$$p(s) = \frac{\sum_{n=1}^{N} K(s, s_d)\, d_n}{\sum_{n=1}^{N} K(s, s_d)}$$

The subscript n refers to the observation in the data set, $p(s)$ is the vector of conditional choice probabilities in state s and d_n is the vector of $(0, 1)$ variables describing the observed choice in the nth observation. We use a bandwidth of 1 in computing kernel estimates of the conditional choice probabilities.

In the above representation of the state space, a variable that describes the number of times an alternative has been sampled is included for each alternative. Choice probabilities will no doubt also be a function of the price that the consumer is faced with on each occasion. The consumer's choice probabilities would also be impacted by other marketing mix variables such as display, feature advertising etc. In reduced form approaches to modeling choice behavior, these variables are usually included in the utility function and the same could be done here. However, in using the CCP methodology for short horizons as described above, the state space representation is not affected by inclusion of marketing mix variables in the utility function. If the horizon over which the value function is computed is relatively long, implementation of the CCP methodology as described above would be cumbersome. In such situations, a more appropriate estimation methodology within the CCP framework would be the CCS (conditional choice simulation) estimator. In the CCS methodology, future choices and realizations of the marketing mix variables would be simulated from appropriate distributions and these would be used to arrive at value functions as functions of the parameters to be estimated. In this situation, all variables that impact on choice behavior would have to be considered when defining the state space.

The reasoning behind conditional choice probabilities being estimated for each learning state only and other variables such as price etc., not being included in the description of the state is as follows. Sample frequencies (within each learning state) that are used to estimate the conditional choice probabilities reflect variation in the environment and the implicit assumption in using these sample frequencies is that the consumer anticipates the variation in the environment correctly. The two options available to the econometrician are to make the above assumption (to use sample frequencies in each learning state) or to simulate future realizations of price etc., and include these variables in the state description while estimating conditional choice probabilities. Another issue that arises is that if price promotions are usually conducted in tandem with displays and features, there is the possibility that the learning effects being captured are actually due to these variables not being included in the utility specifications. This possibility is checked for by looking at the simple correlations between prices and the omitted marketing mix variables. Correlations between price and other marketing mix variables, such as display and feature, are very low in these data sets. Hence, it is unlikely that the learning effects that we are capturing are actually driven by the marketing mix variables that have not been included in this study and hence the result of mis-specification.