

Learning Models: An Assessment of Progress, Challenges and New Developments*

by

Andrew T. Ching
Rotman School of Management
University of Toronto
Andrew.Ching@rotman.utoronto.ca

Tülin Erdem
Stern School of Business
New York University
terdem@stern.nyu.edu

Michael P. Keane
University of New South Wales, and
Arizona State University
m.keane@unsw.edu.au

This draft: August 25, 2012

* Ching's work on this project is supported by SSHRC. Keane's work on this project was supported by Australian Research Council grant FF0561843. We thank Masakazu Ishihara for helpful discussion.

Abstract

Learning models extend the traditional discrete choice framework by postulating that consumers may have incomplete information about product attributes, and that they learn about these attributes over time. In this survey we describe the literature on learning models that has developed over the past 20 years, using the model of Erdem and Keane (1996) as a unifying framework. We described how subsequent work has extended their modeling framework, and applied learning models to a wide range of different products and markets. We argue that learning models have contributed greatly to our understanding of consumer behavior, in particular in enhancing our understanding of brand loyalty and long run advertising effects. We also discuss the limitations of existing learning models and discuss potential extensions. One key challenge is to integrate learning models with other key mechanisms that may generate choice dynamics (inventories, habit persistence, etc.).

Keywords: Learning Models, Choice modeling, Dynamic Programming, Structural models, Brand equity

1. Introduction

In the field of discrete choice, the most widely used models are clearly the multinomial logit and probit.¹ Of course, there has been substantial effort over the past 20 years to generalize these workhorse models to allow richer structures of consumer taste heterogeneity, serial correlation in preferences, dynamics, endogenous regressors, etc. However, with few exceptions, work within the traditional random utility framework maintains the strong assumption that *consumers know the attributes of their choice options perfectly*.

Learning models extend the traditional discrete choice framework by postulating that consumers may have incomplete information about product attributes. Thus, they make choices based on *perceived* attributes. Over time, consumers receive information signals that enable them to learn more about products. It is this inherent temporal aspect of learning models that distinguishes them from static choice under uncertainty models.

Within this general framework, different types of learning models can be distinguished along four key dimensions. One is whether consumers behave in a forward-looking manner. If attributes are uncertain, and consumers are myopic, they choose the alternative with highest expected current utility. But forward-looking consumers may (i) make trial purchases to enhance their information sets, or (ii) actively seek out information about products (“active learning”).

A second key distinction is whether utility is linear in attributes or whether consumers exhibit risk aversion. In the linear case, forward-looking consumers are willing to pay a premium for unfamiliar products, as they receive not only the expected utility of consumption but also the value of the information acquired by trial.² But with risk aversion, consumers are willing to pay a premium for a more familiar product. This can generate “brand equity” for well-known brands.

A third distinction involves sources of information. In the simplest learning models trial is the only information source. In more sophisticated models consumers can learn from a range of sources, such as advertising, word-of-mouth, price signals, salespeople, product ratings, social networks, newspapers, etc.. A consumer must decide how much to use each available source.

¹ Until fairly recently, the logit was much more popular than probit, largely due to computational advantages. But advances in simulation methods, such as the GHK algorithm and Gibbs sampling (see Geweke and Keane (2001), McCulloch, Polson and Rossi (2000)) have greatly increased the popularity of probit, particularly among Bayesians.

² Thus, learning models play havoc with traditional welfare analysis, as consumer surplus is no longer the area under the demand curve. Parameters of the demand curve are no longer structural parameters of preferences, but depend on the information set (e.g., they can be shifted by advertising). See Erdem, Keane and Sun (2008) for a discussion.

A fourth distinction is how consumers learn. They may be Bayesians, or they may update perceptions in some other way. For instance, consumers may over/under weight new information relative to an optimal Bayesian rule, or forget information that was received too far in the past.

Learning models were first applied to marketing problems in pioneering work by Roberts and Urban (1988) and Eckstein, Horsky and Raban (1988).³ But, due to the technical limitations of the time (i.e., both in computer speed and estimation algorithms), their models had to be quite simple. Roberts and Urban (1988) study how (Bayesian) consumers learn about a new product from word-of-mouth signals. Consumers in their model are risk averse, but they are myopic so there is no active search. In contrast, in Eckstein et al (1988) consumers are forward-looking, but utility is linear, and trial is the only source of information. So their model exhibits the “value of information” phenomenon, but not the “brand loyalty” phenomenon created by risk aversion. For Roberts and Urban (1988) the converse is true. The strong simplifying assumptions of these early models, plus the difficulty of estimating even simple learning models 20 year ago, probably explains why no further learning models appeared in the literature for several years after 1988.

The paper by Erdem and Keane (1996) represented a significant methodological advance, because it greatly expanded the class of learning models that are feasible to estimate.⁴ Their approach could handle forward-looking consumers, risk aversion, multiple information sources, and both active and passive search in one model. Thus, their model exhibited both the “value of information” and “brand loyalty” phenomena. In their empirical application, consumers had uncertainty about the quality of brands, and learned both through use experience (active learning) and exogenously arriving advertising signals (passive learning).

Erdem and Keane (1996) assumed that consumers processed quality signals as Bayesians. This assumption imposes a very special structure on how past choice history affects current choice probabilities.⁵ A striking result in their paper was that this structurally motivated functional form for choice probabilities actually fit the data better than commonly used reduced form specifications, such as Guadagni and Little (1983)’s exponential weighted average of past

³ Pioneering papers that first applied learning models in labor economics were Jovanovic (1979) and Miller (1984). Both were concerned with workers and firms learning about the quality of job matches.

⁴ Their computational approach involved (1) using the method of Keane and Wolpin (1994) to obtain a fast but accurate approximate solution to the dynamic optimization problem of forward looking agents, and (2) using (then) recently developed simulation estimation methods (see, e.g., Keane (1994)) to approximate the likelihood function.

⁵ Specifically, only the total number of use experiences should matter, not their timing. The same is true for ad exposures. This structure changes with forgetting, and issue we will discuss later.

purchases specification (the so called “loyalty variable”).⁶ Erdem and Keane (1996) also found strong evidence that advertising has important long run effects on demand (via the total stock of advertising), but that short run effects of recent advertising were negligible.

The Erdem and Keane (1996) paper was influential because: (1) it provided a practical method for estimating complex learning models, (2) it showed that, far from imposing a “straight jacket” on the data, the Bayesian learning structure led to insights about the functional form for state dependence that improved model fit, and (3) it generated interesting results about long vs. short-run effects of advertising, and (4) it gave an economic rationale for the “brand loyalty” observed in scanner panel data.⁷ These results generated new interest in structural learning models (and dynamic structural models more generally) within the field of marketing.

Nevertheless, there was a time lag of roughly five years from Erdem and Keane (1996) to the publication of many additional papers on learning models. But, starting in the early 2000s, there has been an explosion of new work in both marketing and economics applying learning models to brand choice and many other problems. Other interesting applications include: (i) demand for new products, (ii) choice of TV shows and movies, (iii) prescription drugs, (iv) durable goods, (v) insurance products, (vi) choice of tariffs (i.e., price/ usage plans) , (vii) fishing locations, (viii) career options, (ix) service quality, (x) childcare options, and (xi) medical procedures. Some of these applications are based rather closely on the Erdem and Keane (1996) framework with forward-looking Bayesian consumers, while other papers depart from or extend that framework in important ways (often along one of the four key dimensions noted above).

The outline of the survey is as follows: In Section 2 we describe the learning model of Erdem and Keane (1996) in some detail. We will treat their model as a unifying framework to discuss the rest of the literature. In general, later developments can be viewed as extending the

⁶ In most applications, imposing structure involves sacrificing fit to some extent (i.e., not surprisingly, structural models usually fit worse than flexible reduced form or statistical/descriptive models). The payoff of imposing the structure is (1) greater interpretability of parameter estimates and (2) to do policy experiments. Erdem and Keane (1996) was a rare instance where a structural model actually fit better than popular competing reduced form models.

⁷ A key insight of Erdem and Keane (1996) was that uncertainty about quality combined with risk aversion could lead to brand loyal behavior (i.e., persistence in brand choice over time). Loyalty emerges as consumers stick with familiar products (whose attributes are precisely known) to avoid risk. Given equal prices, a familiar brand may be chosen over a less familiar brand even if it has lower expected quality, provided consumers are sufficiently risk averse. In this framework, “loyalty” is the price premium that consumers are willing to pay for greater familiarity (lower risk). Keller (2002) refers to the general framework laid out in Erdem and Keane (1996), and elucidated further in Erdem and Swait (1998), as “the canonical *economic* model of brand equity.” (Of course there are also a number of psychology-based models – see Keller (2002) for an overview).

Erdem-Keane model along certain dimensions (while typically restricting it on others to make those extensions feasible), or applying it in different contexts. Section 3 reviews the subsequent literature on learning models. This is divided in subsections that cover (i) more sophisticated models of learning, (ii) correlated learning, (iii) new/novel applications of learning models, (iv) non-standard Bayesian learning models, and (v) equilibrium models. Some papers span these boundaries, so we can only roughly divide them. Section 4 describes what we consider the key challenges for future research. In Section 5 we summarize and conclude.

2. The General Structure of the Erdem and Keane (1996) Model

As we noted in the introduction, the papers by Roberts and Urban (1988) and Eckstein, Horsky and Raban (1988) were the first applications of learning model to marketing problems. The model of Erdem and Keane (1996), henceforth “EK,” nests those models in a more general framework. Thus, in this section we describe the EK model in some detail. Readers interested in more detail about the earlier models can refer to a detailed description in [online Appendix A](#).

2.1. A Simple Dynamic Learning Model with Gains from Trial Information

Of course, the key feature of learning models is that consumers do not know the attributes of brands with certainty. While this may be true of many attributes, most papers, including EK, have focused on learning about brand quality. In their model, consumers receive signals about quality through both use experience and ad signals. But prior to receiving any information, consumers have a normal prior on brand quality:

$$(1) \quad Q_j \sim N(Q_{j0}, \sigma_{j0}^2) \quad j = 1, \dots, J$$

This says that, prior to receiving any information, consumers perceive that the true quality of brand j , denoted Q_j , is distributed normally with a mean of Q_{j0} and a variance of σ_{j0}^2 . The values of Q_{j0} and σ_{j0}^2 may be influenced by many factors, such as reputation of the manufacturer, pre-launch advertising, etc.

Use experience does not fully reveal quality because of “inherent product variability.” This has multiple interpretations. First, the quality of different units of a product may vary. Second, a consumer’s experience of a product may vary across use occasions. For instance, a cleaning product may be effective at removing the type of stains one faces on most occasions,

but be ineffective on other occasions. Alternatively, there may be inherent randomness in psychophysical perception. E.g., the same cereal tastes better to me on some days than others.

Given inherent product variability, there is a distinction between “experienced quality” for brand j on purchase occasion t , which we denote Q_{jt}^E , and true quality Q_j . Let us assume the “experienced quality” delivered by use experience is a noisy signal of true quality, as in:

$$(2) \quad Q_{jt}^E = Q_j + \varepsilon_{jt} \quad \text{where} \quad \varepsilon_{jt} \sim N(0, \sigma_\varepsilon^2)$$

Here σ_ε^2 is the variance of inherent product variability, which we often refer to as “experience variability.”

Note that we have conjugate priors and signals, as both the prior on quality (1) and the noise in the quality signals (2) are assumed to be normal. This structure gives simple formulas for updating perceptions as new information arrives, as we will see below. This is precisely why we assume priors and signals are normal. Few other reasonable distributions would give simple expressions. Also, as signals are typically unobserved by the researcher, it is not clear that more flexible distributions would be identified.

Thus, the posterior for perceived quality, given a single use experience signal, is given by the simple updating formulas:

$$(3) \quad Q_{j1} = \frac{\sigma_{j0}^2}{\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j1}^E + \frac{\sigma_\varepsilon^2}{\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j0}$$

$$(4) \quad \sigma_{j1}^2 = \frac{1}{(1/\sigma_{j0}^2) + (1/\sigma_\varepsilon^2)}$$

Equation (3) describes how a consumer’s prior on quality of brand j is updated as a result of the experience signal Q_{j1}^E . Note that the extent of updating is greater the more accurate is the signal (i.e., the smaller is σ_ε^2). Equation (4) describes how a consumer’s uncertainty declines as he/she receives more signals. σ_{j1}^2 is often referred to as the “perception error variance.”

Equations (3) and (4) generalize to any number of signals. Let $N_j(t)$ denote the total number of use experience signals received through time t . Then we have that:

$$(5) \quad Q_{jt} = \frac{\sigma_{j0}^2}{N_j(t)\sigma_{j0}^2 + \sigma_\varepsilon^2} \sum_{s=1}^t Q_{js}^E d_{js} + \frac{\sigma_\varepsilon^2}{N_j(t)\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j0}$$

$$(6) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{j0}^2) + N_j(t)(1/\sigma_{\varepsilon}^2)}$$

where d_{jt} is an indicator for whether brand j is bought/consumed at time t .

In equation (5), the perceived quality of brand j at time t , Q_{jt} , is a weighted average of the prior and all quality signals received up through time t , $\sum_{s=1}^t Q_{js}^E d_{js}$. Crucially, this is a random variable across consumers, as some will, by chance, receive better quality signals than others. Thus, the learning model endogenously generates heterogeneity across consumers in perceived quality of products (even starting from identical priors). This aspect of the model is appealing. It seems unlikely that people are born with brand preferences (as standard models of heterogeneity implicitly assume), but rather that they arrive at their views through heterogeneous experience.

Of course, as equation (6) indicates, the variance of perceived quality around true quality declines as more signals are received, and in the limit perceived quality converges to true quality. Still, heterogeneity in perceptions will persist over time, for several reasons: (i) both brands and consumers are finitely lived, (ii) there is a flow of new brands and consumers entering a market, and (iii) as people gather more information the value of trial diminishes, and the incentive to learn about unfamiliar products will become small. Intuitively, once a consumer is familiar with a substantial subset of brands, there is rarely much marginal benefit to learning about all the rest.

In general, learning models must be solved by dynamic programming, because today's purchase affects tomorrow's information set, which affects future utility. Thus, choosing the brand with the highest *perceived* quality in the *current* period is not necessarily optimal.

To see this, it is useful to consider the special case where the choice is between an old familiar brand (whose attributes are known with certainty) and a new brand. Denote these by $j=o,n$ (for old and new). The information set is $I_t = \{Q_{nt}, \sigma_{nt}^2\}$, where we suppress the values for the old brand which are simply Q_o and is $\sigma_{ot}^2 = 0$. The prices are given by P_{jt} for $j=o,n$. Then the values (denoted V) of choosing each option in the current period are:

$$(7) \quad V(n, t|I_t) = E[U(Q_{nt}^E, P_{nt})|I_t] + \beta EV(I_{t+1}) \quad \text{where} \quad I_{t+1} = \{Q_{nt+1}, \sigma_{nt+1}^2\}$$

$$(8) \quad V(o, t|I_t) = E[U(Q_{ot}^E, P_{ot})] + \beta EV(I_t) \quad \text{as} \quad I_{t+1} = I_t$$

At time t a consumer will choose the brand with the highest V , which is not necessarily the brand

with the highest expected utility. It is important to understand why. Purchase of the familiar brand gives expected utility $E[U(Q_{ot}^E, P_{ot})|I_t]$. This is increasing in true quality Q_o , which is known. It is decreasing in experience variability σ_ε^2 , assuming consumers are risk averse with respect to product variability. On the other hand, purchase of the new brand delivers expected utility $E[U(Q_{nt}^E, P_{nt})]$, which is increasing in Q_{nt} and decreasing in both experience variability and the perception error variance σ_{nt}^2 . Purchase of the new brand also increases next period's expected value function. That is, $EV(I_{t+1}) > EV(I_t)$ because I_{t+1} contains better information. As a result, it may be optimal to try the new brand even if $E[U(Q_{nt}^E, P_{nt})|I_t] < E[U(Q_{ot}^E, P_{ot})]$.

To gain further insight, it is useful to consider the special case where utility is linear in experienced quality Q_{jt}^E , as in Eckstein et al (1988), thus abstracting from risk aversion, and also linear in price. In that case (7) and (8) simplify to:

$$(9) \quad V(n, t|I_t) = Q_{nt} - P_{nt} + e_{nt} + \beta EV(I_{t+1}) \quad \text{where} \quad I_{t+1} = \{Q_{nt+1}, \sigma_{nt+1}^2\}$$

$$(10) \quad V(o, t|I_t) = Q_o - P_{ot} + e_{ot} + \beta EV(I_t) \quad \text{where} \quad I_{t+1} = I_t$$

Here the e_{jt} for $j=o,n$ are stochastic terms in the utility function that represent purely idiosyncratic tastes for the two brands. These play the same role as the brand specific stochastic terms in traditional discrete choice models like logit and probit.⁸

Now, a consumer will choose the new brand over the familiar brand if the value function in equation (9) exceeds that in (10). This means that $V_{nt}^* > 0$, where:

$$(11) \quad V_{nt}^* \equiv V(n, t|I_t) - V(o, t|I_t) = (Q_{nt} - P_{nt}) - (Q_o - P_{ot}) + (e_{nt} - e_{ot}) + G_t$$

$$(12) \quad G_t \equiv \beta[EV(I_{t+1}) - EV(I_t)]$$

We will refer to G_t as the “gain from trial.” It is the increase in expected present value of utility from $t+1$ until the terminal period T arising from the enhanced information the consumer obtains by trying the new brand at time t .

⁸ Without these terms, choice would be deterministic (conditional on Q_{nt} , Q_o , P_{nt} and P_{ot}). But in contrast to standard discrete choice models, it is not strictly necessary to introduce the $\{e_{jt}\}$ terms to generate choice probabilities. This is because perceived quality Q_{nt} is random from the perspective of the econometrician. However, we feel it is advisable to include the $\{e_{jt}\}$ terms regardless. This is because, in their absence, choice becomes deterministic conditional on price as experience with the new brand grows large. Yet, both introspection and simple data analysis suggest consumers do switch brands for purely idiosyncratic reasons even under full information.

Intuitively, the gain from trial comes from two sources. Most obviously, the consumer may learn that the new brand is better than the old brand. More subtly, suppose the evidence indicates the new brand is inferior to the familiar brand. There is nevertheless a price differential large enough that the consumer would prefer the new brand. More precise information about the quality of the new brand enables the consumer to set this price differential more accurately.

Here, we give a sketch of a proof that $EV(I_{t+1}) > EV(I_t)$, and hence that G_t is positive. This is a very general result of information economics, but it is easiest to show in the linear case. It is also easiest to consider a finite horizon problem with terminal period T . As there is no future, the consumer at time T will simply choose the brand with the highest expected utility. Thus, the utility a consumer with incomplete information (i.e., $Q_{nT} \neq Q_n$) receives at T is simply:

$$V_T(I_T = \{Q_{nT}, \sigma_{nT}^2\}) = \max\{Q_{nT} - P_{nT} + e_{nT}, Q_o - P_{oT} + e_{oT}\}$$

On the other hand, a consumer with complete information would receive utility:

$$V_T(I_T = \{Q_n, 0\}) = \max\{Q_n - P_{nt} + e_{nt}, Q_o - P_{ot} + e_{ot}\}.$$

This depends on true quality Q_n , not on perceived quality Q_{nT} . Thus, a consumer with incomplete information is in effect making decisions at T using the “wrong” decision rule, so in general he/she will make suboptimal decisions. More formally, letting a^* be a noisy measure of a , we have $E\{V|V = \max(a, b)\} > E\{V|V = a \text{ if } a^* = \max(a^*, b), V = b \text{ otherwise}\}$. This is the key intuition for why information is valuable. A complete proof involves two more steps. First, one needs to show that as Q_{nT} becomes more accurate the consumer’s decisions become closer to optimal, so that $EV(I_T)$ is decreasing in the perception error variance σ_{nT}^2 . Second, by backwards induction it can be shown that this is true back to any period t .

Although $G_t > 0$, i.e., more information is better, it is notable that, G_t is smaller if (i) the consumer has more information (σ_{nt}^2 smaller) or (ii) use experience signals are less accurate (σ_ε^2 larger). Both lower the value of trial. Notice that (11) can be rewritten as “Choose brand n if:”

$$(13) \quad (Q_{nt} - P_{nt}) + G_t + e_{nt} > (Q_o - P_{ot}) + e_{ot}$$

This shows that the trial value G_t *augments* the perceived value of the new brand ($Q_{nt} - P_{nt}$). Thus, *ceteris paribus*, the new brand can command a price premium over the old brand because

it delivers valuable trial information. So the model with linear utility (i.e., no risk aversion) generates a “value of information” effect that is *opposite* to the conventional brand loyalty phenomenon. [In [online Appendix A](#) we give some details on estimation of this model.]

2.2. Introducing Risk Aversion and Exogenous Signals

Next, we introduce two key features of the Erdem and Keane (1996) model that generalize the simple setting described above. First, we introduce exogenous signals of quality (e.g., advertising) as an additional source of information besides use experience. Second, we consider utility functions that exhibit risk aversion with respect to variation in brand attributes (focusing again on quality). We should note that both these features were already present in Roberts and Urban (1988), but in a static choice context.

There are numerous ways one can obtain information about a brand other than trial purchase.⁹ Examples are advertising, word-of-mouth, magazine articles, dealer visits, etc. For simplicity we will often refer to these as “exogenous” signals, as we may think of them as arriving randomly from the outside environment. (Of course, a consumer may actively seek out such signals, an extension we discuss below). For frequently purchased goods the most important source of information is probably advertising, and this is the source that EK consider.

Let A_{jt} denote an exogenous signal (advertising, word of mouth, etc.) that a consumer receives about brand j at time t . We further assume that:

$$(14) \quad A_{jt} = Q_j + \varepsilon_{jt}^A \quad \text{where} \quad \varepsilon_{jt}^A \sim N(0, \sigma_A^2)$$

This says the signals A_{jt} provide unbiased but noisy information about brand quality, where the noise has variance σ_A^2 . The noise is assumed normal, to maintain conjugacy with the prior in (1).

It is important to compare (14) with (2). The noise in trial experience is from inherent product variability, which is largely a feature of the product itself. The noise in a signal like advertising or word-of-mouth is, in contrast, largely a function of the medium. Presumably some media convey information more accurately than others, and no medium is as accurate as direct use experience. We also stress that the noise in (14) differs fundamentally from that in (2), as inherent product variability affects a consumer’s *experienced* utility from consuming the product,

⁹ Indeed, when considering a durable good, as opposed to a frequently purchased good, trial purchase is no longer even a relevant consideration (assuming that one cannot return a purchase to get a refund).

while exogenous quality signals do not. Nevertheless, both types of signal enter the consumer's learning process in the same way. Given the exogenous signal A_{jt} , we can rewrite (5)-(6) as:

$$(15) \quad Q_{jt} = \frac{\sigma_{j0}^2}{N_j^A(t)\sigma_{j0}^2 + \sigma_A^2} \sum_{s=1}^t A_{js} d_{js}^A + \frac{\sigma_A^2}{N_j^A\sigma_{j0}^2 + \sigma_A^2} Q_{j0}$$

$$(16) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{j0}^2) + N_j^A(1/\sigma_A^2)}$$

where d_{jt}^A is an indicator for whether a signal for brand j is received at time t , and $N_j^A(t)$ is the total number of signals received for brand j up through time t .

It is simple to extend the Bayesian updating rules in (5)-(6) and (15)-(16) to allow for two types of signals – i.e., both use experience and exogenous signals. Then we obtain the formulas:

$$(17) \quad Q_{jt} = \frac{(1/\sigma_\varepsilon^2)}{S_j(t)} \sum_{s=1}^t Q_{js}^E d_{js} + \frac{(1/\sigma_A^2)}{S_j(t)} \sum_{s=1}^t A_{js} d_{js}^A + \frac{(1/\sigma_{j0}^2)}{S_j(t)} Q_{j0}$$

$$(18) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{j0}^2) + N_j(t)(1/\sigma_\varepsilon^2) + N_j^A(t)(1/\sigma_A^2)}$$

where $S_j(t) \equiv 1/\sigma_{jt}^2$. Note that the timing of signals does not matter. In (17)-(18) only the total stock of signals determines a consumer's state. Furthermore, receiving N signals with variance σ^2 affects the perception variance in the same way as receiving one signal with variance σ^2/N .

As we will see, these properties are important for simplifying the solution to consumers' dynamic optimization problem. This is because the consumer's level of uncertainty, as captured by the $\{\sigma_{jt}^2\}_{j=1}^J$ depends only on the number of signals received, not the order or timing with which they were received. One could imagine scenarios where more recent signals are more salient, or, conversely, where first impressions are most important. These are important potential extensions of the model, but they would make computation more difficult.

Next, we must map perceived quality and quality uncertainty into choice probabilities. To do this within a structural framework one must assume a specific utility function. Of course, many functions are possible. Erdem and Keane (1996) assumed a utility function of the form:

$$(19) \quad U(Q_{jt}^E, P_{jt}) = w_Q Q_{jt}^E - w_Q r (Q_{jt}^E)^2 + w_P (X - P_{jt}) + e_{jt}$$

Here utility is quadratic in the experienced quality of brand j at time t , and linear in consumption of the composite outside good $C_t = X - P_{jt}$, where X is income. The parameter w_Q is the weight on quality, r is the risk coefficient, w_P is the marginal utility of the outside good, and e_{jt} is an idiosyncratic brand and time specific error term.¹⁰ Note that, as choices only depend on utility differences, and as income is the same regardless of which brand is chosen, income drops out of the model. So we can simply think of w_P as the price coefficient.

Given (19), combined with (2) and (18), expected utility is given by:

$$(20) \quad E[U(Q_{jt}^E, P_{jt})|I_t] = w_Q Q_{jt} - w_Q r Q_{jt}^2 - w_Q r (\sigma_{jt}^2 + \sigma_\varepsilon^2) - w_P P_{jt} + e_{jt} \quad j = 1, \dots, J$$

Also, as the Erdem-Keane model was meant to be applied to weekly data, and as consumers may not buy in every week, a utility of the no purchase option was also specified. This was written simply as $E[U_0|I_t] = \Phi_0 + \Phi_1 t + e_{0t}$. The time trend captures in a simple way the possibility of changing value of substitutes for the category in question.

Agents are forward-looking, making choices to maximize value functions of the form:

$$(21) \quad V(j, t|I_t) = E[U(Q_{jt}^E, P_{jt})|I_t] + \beta EV(I_{t+1}|I_t, j) \quad \text{for } j=0, \dots, J$$

where the consumer's information set is given by:

$$(22) \quad I_{t+1} = \{Q_{1,t+1}, \dots, Q_{J,t+1}, \sigma_{1,t+1}^2, \dots, \sigma_{J,t+1}^2\}$$

A key point is that a consumer's information about a brand may be updated between t and $t+1$ for two reasons: (i) the consumer buys the brand, or (ii) the consumer receives an exogenous signal about the brand. Henceforth we simply refer to these as "ad signals." In forming $EV(I_{t+1}|I_t, j)$ we allow for both sources of information. We describe the process in detail in the next section.

With the introduction of risk aversion, the EK model can capture both gains from trial and brand loyalty phenomena. As we discussed earlier, the $EV(I_{t+1}|I_t, j)$ terms in a dynamic learning model capture the gain from trial information. These are greater for less familiar brands,

¹⁰ It is common to assume utility is linear in consumption of the outside good, so there are no income effects, when dealing with inexpensive items like frequently purchased consumer goods. This also means the marginal utility of consumption w_P is constant within the range of outside good consumption levels spanned by different brand choices. However, it is likely that the marginal utility of consumption w_P would be lower for households at higher wealth levels. And the assumption of no income effects would not be tenable for expensive durable goods.

where the gain from trial is greater.¹¹ At the same time, the risk terms σ_{jt}^2 are also greater for such brands. These two forces work against each other, and which dominates determines whether a consumer is more or less likely to try a new unfamiliar brand vs. a familiar brand. In categories that exhibit a high degree of brand loyalty (i.e., persistence in choice behavior) we would expect the risk aversion effect to dominate. In categories that exhibit a high degree of brand switching (due to experimentation and variety seeking) we would expect the gains from trial to dominate.

2.3. Solving the Dynamic Optimization (DP) Problem

The expected value functions in (21) have the form:

$$(23) \quad EV(I_{t+1}|I_t, j) = E_t \max\{V(0, t|I_{t+1}), \dots, V(J, t|I_{t+1})\}$$

That is, the consumer at time t knows that at time $t+1$ he/she will choose from among the J options the one with the highest value function. The consumer can form the expected maximum over these value functions, because his/her information and decision at time t (i.e., the (I_t, j)) generate a distribution of I_{t+1} , in the manner described earlier.

However, it is not clear that writing equation (23) helps us to solve the consumers' optimization problem, because the V s on the right hand side of (23) themselves contain expected value functions dated at $t+2$, that is, functions of the form $EV(I_{t+2}|I_{t+1}, j)$. So it seems we have only pushed the problem one period ahead. One key insight for solving a dynamic programming problem is to assume there exists a terminal period T beyond which a consumer does not plan. At T , the consumer will simply choose the option with highest expected utility. Thus, we have that:

$$(24) \quad V(j, T|I_T) = E[U(Q_{jT}^E, P_{jT})|I_T] \quad \text{for} \quad j=0, \dots, J$$

$$(25) \quad EV(I_T|I_{T-1}, j) = E_{T-1} \max\{E[U_0|I_T], E[U(Q_{1T}^E, P_{1T})|I_T], \dots, E[U(Q_{jT}^E, P_{jT})|I_T]\}$$

The integral in (25) is straightforward to evaluate. Indeed, if the I_T and P_T were known at $T-1$, then the only unknowns appearing in (25) would be the logistic errors $\{e_{0T}, \dots, e_{jT}\}$. In that case (25) would have a simple closed form given by the well-known nested logit "inclusive value" formula. This illustrates the point that estimating a finite-horizon dynamic model is very much

¹¹ It is worth noting that a consumer may have incomplete information about all J brands in the market. Alternatively, he/she may have very accurate knowledge of a few favorite brands.

like estimating a nested logit model – if one thinks of moving down the nesting structure as a process that plays out over time.

Of course, evaluating (25) in the Erdem-Keane model is more difficult, because, as we noted earlier, both use experience signals and ad signals may arrive between $T-1$ and T , causing the consumer to update his/her information set. The expectation in (25) must be taken over the possible I_T that may arise as a result of these signals. Specifically, we must: (i) update σ_{jt}^2 to $\sigma_{j,t+1}^2$ using (18) to account for additional use experience, (ii) integrate over possible values of the use experience signal in (2) to take the expectation over possible realizations of $Q_{j,t+1}$, (iii) integrate over possible ad exposures that may arrive between t and $t+1$ (i.e., over realizations of $d_{j,t+1}^A$ for $j=1, \dots, J$) to account for ad induced changes in the $\{\sigma_{j,t+1}^2\}$, and (iv) integrate over possible values of the ad signals in (14), as these will lead to different values of the $\{Q_{j,t+1}\}$.¹²

Clearly the integrals described here are high dimensional, and simulation methods will be needed. That is, one must integrate by simulation over draws from the distributions of the signal processes. Obviously, if consumers learn about multiple brands, and have more than one source of information, it greatly complicates the nature of the integrals that must be evaluated to form the expected value functions in (25).

Once we have calculated the value $EV(I_T|I_{T-1}, j)$ for every possible (I_{T-1}, j) – a point we return to below – it is possible to form the version of equation (21) dated at $T-1$:

$$(26) \quad V(j, T-1|I_{T-1}) = E[U(Q_{jT-1}^E, P_{jT-1})|I_{T-1}] + \beta EV(I_T|I_{T-1}, j)$$

Note that (26) is just like (24), except for the $EV(I_T|I_{T-1}, j)$ terms that are appended. But we have already solved for these and saved them in memory, so they are just numbers. Given that we can now construct the $V(j, T-1|I_{T-1})$, we can proceed backwards to calculate the time $T-1$ version of (25), and obtain the $EV(I_{T-1}|I_{T-2}, j)$. Then we can work back again and obtain the $V(j, T-2|I_{T-2})$. This backwards induction process is repeated until we have solved the entire dynamic programming problem back to $t=1$. Detailed descriptions of this process, known as “backsolving” are contained in many sources. See, for instance, Keane et al (2011).

¹² It is also necessary to integrate over future price realizations for all brands. To make this integration as simple as possible, Erdem and Keane (1996) assumed that the price of each brand is distributed normally around a brand specific mean, with no serial correlation other than that induced by these mean differences.

In practice, T is generally chosen to be some period beyond the end of the sample period. This can be chosen far enough out so that results are not sensitive to the exact value of T .

Unfortunately, the above description is oversimplified as it assumes it is feasible to calculate the value $EV(I_T|I_{T-1}, j)$ for every possible (I_{T-1}, j) . But note that the number of variables that characterize the state of an agent in equation (22) is $2 \cdot J$. Solving a dynamic programming problem exactly requires that one solve the expected value function integrals at every point in the state space, and this is clearly not feasible here, because there are too many state variables. Of course, as the state variables in (22) are continuous, it would be literally impossible to solve for the expected value functions at every state point (as the number of points is infinite). A common approach is to discretize continuous state variables using a fairly fine grid. Say we use G grid points for each state variable.¹³ As we have $2 \cdot J$ state variables, this gives $G^{2 \cdot J}$ grid points, which is impractically large even for modest G and J . This is known as the “curse of dimensionality.” A number of ways to deal with this problem have been proposed:

To solve the optimization problem in their model (that is, to construct the $EV(I_{t+1}|I_t, j)$ in (21)), Erdem and Keane (1996) used an approximate solution method developed in Keane and Wolpin (1994). The idea is to evaluate the expected value function integrals at a randomly selected subset of state points (where this set is relatively small compared to the size of the total state space).¹⁴ The expected value functions are then constructed at other points via interpolation. For instance, one can run a regression of the value functions on the state variables (at the random subset of state points), and use the regression to predict the value functions at other points. We give more detail on how to apply this method to estimate learning models in online Appendix B.

The Keane-Wolpin approximation method, or variants on its basic idea, has become widely used in both economics and marketing in the past 15 years to solve many types of dynamic models. This has greatly increased the richness and complexity of the dynamic models that it is feasible to estimate. We will not give details of these computational methods here, but refer the reader to surveys by Keane, Todd and Wolpin (2011), Aguirregebirria and Mira (2010) and Geweke and Keane (2001) among others.

¹³ Note that range of the discretization needs to be big enough to cover the true Q_j 's (or $\sigma_{j,0}^2$'s), which are unknown to researchers a priori. In [online appendix B](#), we outline a procedure to determine the bounds.

¹⁴ Also, in most applications the expected value function integrals are simulated using Monte Carlo methods rather than evaluated numerically. This makes it practical to deal with the three aspects of integration described below equation (31) – integration over content of use experience signals, over exposure to ads, and over the content of ads.

Finally, a common question is how we can solve the DP problem when we do not know the true parameter values, either for the utility function or the stochastic processes that generate signals. The answer is that the DP problem must be solved at each trial parameter value that is considered during the search process for the maximum of the likelihood function. In other words, the DP solution is nested within the likelihood evaluation. We consider the construction of the likelihood function in the next section.

2.4. Evaluating the Likelihood Function

In this section we discuss how to form the likelihood function for the EK learning model. Let $\theta = \{w_Q, w_P, r, \{Q_{j0}, \sigma_{j0}^2\}, \sigma_\varepsilon^2, \sigma_A^2\}$ denote the entire vector of model parameters. Combining (20) and (21) we have the choice specific value functions:

$$(27a) \quad V(j, t|I_t) = w_Q Q_{jt} - w_Q r Q_{jt}^2 - w_Q r (\sigma_{jt}^2 + \sigma_\varepsilon^2) - w_P P_{jt} + e_{jt} + \beta EV(I_{t+1}|I_t, j) \quad j = 1, J$$

$$(27b) \quad V(0, t|I_t) \equiv \Phi_0 + \Phi_1 t + e_{0t} + \beta EV(I_{t+1}|I_t, 0)$$

Erdem and Keane assume that the idiosyncratic brand and time specific error terms e_{jt} in (27) are *iid* extreme value. In this case, the choice probabilities have a simple multinomial logit form:

$$(28) \quad P(j|I_t) = \frac{\exp(V_j(\theta))}{\exp(V_0(\theta)) + \sum_{k=1}^J \exp(V_k(\theta))}$$

where:

$$(29a) \quad V_j(\theta) \equiv w_Q Q_{jt} - w_Q r Q_{jt}^2 - w_Q r (\sigma_{jt}^2 + \sigma_\varepsilon^2) - w_P P_{jt} + \beta EV(I_{t+1}|I_t, j) \quad j = 1, \dots, J$$

$$(29b) \quad V_0(\theta) \equiv \Phi_0 + \Phi_1 t + \beta EV(I_{t+1}|I_t, 0)$$

Equations (28)-(29) illustrate the point, emphasized in Keane et al (2011), that choice probabilities in dynamic discrete choice models look exactly like those in static discrete choice models (multinomial logit in the present case), *except* that the $V_j(\theta)$ terms in the dynamic model include the extra $EV(I_{t+1}|I_t, j)$ terms. However, once one has solved the dynamic programming problem, these extra terms are merely numbers that one can look up in a table that is saved in computer memory – i.e., a table that lists expected value functions at every point in the state space. Alternatively, if one has used interpolating method rather than saving every value, the

appropriate $EV(I_{t+1}|I_t, j)$ may be constructed as needed using the interpolating function. Erdem and Keane (1996) use the latter procedure.

Some key points about identification become apparent from examining (27)-(29). First, suppose that consumers have complete information about all brands. Then we have $EV(I_{t+1} | I_t, j) = EV(I_t) = k$ for $j=1, \dots, J$, where k is a constant. That is, there is no updating of information sets based on choice, and so the $\beta EV(I_{t+1}|I_t, j)$ terms drop out of the model (just like any term that is constant across choices in a discrete choice model). What remains is a static model where:

$$(30) \quad V(j, t | I_t = \{Q_1, \dots, Q_J, 0, \dots, 0\}) \equiv w_Q Q_j - w_Q r Q_j^2 - w_Q r \sigma_\varepsilon^2 - w_P P_{jt} + e_{jt}$$

Here, we have set $\sigma_{jt}^2 = 0$ and $Q_{jt} = Q_j$ because there is no uncertainty about quality.

Obviously we cannot identify β , σ_A^2 or the priors $\{Q_{j0}, \sigma_{j0}^2\}$ as they drop out of the model. More subtly, it follows from (30) that we cannot identify σ_ε^2 either, as it is constant across alternatives $j=1, \dots, J$. Although σ_ε^2 does not enter the value of the no purchase option, any shift in σ_ε^2 can be undone by a shift in Φ_0 , leaving utility differences unchanged. Thus, β , σ_A^2 , σ_ε^2 and the priors $\{Q_{j0}, \sigma_{j0}^2\}$ only affect choice probabilities through the $EV(I_{t+1} | I_t, j)$ terms. Finally, careful inspection of (30) reveals that r is not identified either, as it cannot be disentangled from the scaling of Q_j . (Obviously, if Q_j had a known scale this would not be a problem). Comparing (27a) and (30) we see that r is identified only when σ_{jt}^2 is positive and varies across consumers.

Given complete information, all that can be identified are the price coefficient w_p , the products $w_Q Q_j$, and the terms Φ_0 and Φ_1 in the value of the no purchase option.¹⁵ Furthermore, as only utility differences matter for choice, we need a normalization to establish a reference alternative. A standard procedure would be to set $Q_j = 0$ for one brand. However, this would not let one disentangle the $w_Q Q_j$ products. Thus, EK instead set $Q_j = 1$ for one brand, so quality of all other brands are measured relative to brand j .¹⁶ Alternatively, one could fix w_Q instead.

Thus, by observing consumers with essentially complete information (i.e., those with a great deal of experience with all brands), we can identify w_Q and the $\{Q_j\}$, given normalization, as well as w_p , Φ_0 and Φ_1 . It is interesting that identification of the parameters β , σ_A^2 , σ_ε^2 and r , as

¹⁵ The scale normalization on utility is imposed by assuming the scale parameter of the extreme value errors is one.

¹⁶ With quadratic utility, it is desirable to constrain the largest Q_j to fall in the region of increasing utility. EK impose this constraint in estimation by updating the level of Q_j at each step (while keeping relative Q values fixed).

well as the priors $\{Q_{j0}, \sigma_{j0}^2\}$, requires that incomplete information actually exist. In that case we get variation in $EV(I_{t+1} | I_t, j)$ and σ_{jt}^2 across consumers generated by variation in their information sets I_{it} . Intuitively, the parameters $\beta, \sigma_A^2, \sigma_\varepsilon^2, r$ and $\{Q_{j0}, \sigma_{j0}^2\}$ are identified by the extent to which, *ceteris paribus*, consumers with different information sets are observed to have different choice probabilities. For instance, we see clearly from (27a) that variation in σ_{jt}^2 identifies r (given w_Q). The learning parameters $\sigma_A^2, \sigma_\varepsilon^2$ and $\{Q_{j0}, \sigma_{j0}^2\}$ are pinned down by how much the arrival of signals alters behavior. And we will only see strategic trial purchases to the extent that $\beta > 0$.

Of course, this variation in choice behavior across consumers due to (i) uncertainty and (ii) variation in information sets cannot be distinguished from variation in behavior due to a completely general form of state dependence. The form of state dependence must be constrained for the learning model to be identified. From the perspective of a structural econometrician this is not a limitation – a model that simply specifies a very general form of state dependence is merely a statistical model that has no structural/behavioral interpretation.

To construct the likelihood function for the EK model, we need some definitions. Let $j(t)$ denote the choice actually made at time t (we continue to suppress the i subscripts to conserve on notation). Let $D_{t-1} \equiv \{j(1), \dots, j(t-1)\}$ denote the history of purchases made up through time $t-1$. Similarly, let $D_{t-1}^A \equiv \{d_1^A, \dots, d_{t-1}^A\}$ denote the history of ads received up through time $t-1$, where $d_t^A \equiv \{d_{jt}^A\}_{j=1}^J$. Also, let $\{Q_{t-1}^E, A_{t-1}\} \equiv \{Q_{j(t),t}^E, A_t\}_{t=1}^{T-1}$ denote the actual content of experience and ad signals received up through $t-1$. Finally, we define $P(j(t) | I_0, D_{t-1}, D_{t-1}^A, Q_{t-1}^E, A_{t-1})$ as the probability of a person's choice at time t given his/her history of use experience and ad signal exposures up through time $t-1$, as well as the content of those signals.

Unfortunately, we cannot observe the actual content of ad and experience signals, so we must integrate over that content to obtain unconditional probabilities $P(j(t) | I_0, D_{t-1}, D_{t-1}^A)$. Thus, the probability of a choice history for an individual takes the form:

$$(31) \quad P(\{j(t)\}_{t=1,T}) = \prod_{t=1}^T P(j(t) | I_t) = \prod_{t=1}^T P(j(t) | I_0, D_{t-1}, D_{t-1}^A) \\ = \int_{\{Q_{j(t),t}^E, A_t\}_{t=1}^{T-1}} \prod_{t=1}^T P(j(t) | I_0, D_{t-1}, D_{t-1}^A, Q_{t-1}^E, A_{t-1}) d\mu\{Q_{j(t),t}^E, A_t\}_{t=1}^T$$

In (31) we integrate over all experience and advertising signals that the consumer may have

received from $t=1, \dots, T-1$. That is, we integrate over the distribution of $\{Q_{j(t),t}^E, A_t\}_{t=1}^{T-1}$. Clearly, the required order of integration is substantial.¹⁷

To deal with this problem, Erdem-Keane used simulated maximum likelihood (see, e.g., Keane (1993, 1994)). Specifically, draw D sets of signals $S^d = \{Q_{j(t),t}^d, A_t^d\}_{t=1}^T$ for $d=1, \dots, D$, using the distributions defined in (2) and (16). Then form the simulated probability:

$$(32) \quad P(\{j(t)\}_{t=1,T}) = \frac{1}{D} \sum_{d=1}^D \prod_{t=1}^T P(j(t)|I_0, D_{t-1}, D_{t-1}^A, S^d)$$

Finally, sum the logs of these probabilities across individuals $i=1, \dots, N$.

A key complication is that consumer purchase histories and ad exposures are not usually observed prior to the start of the sample period. This creates an “initial conditions problem.” A consumer who likes a particular brand will have bought it often *before* the sample period starts. Thus, brand preference is correlated with the information set at the start of the period (I_0). The usual consequence is to exaggerate the impact of lagged purchases on current choices. An exact solution to the initial conditions problem requires integrating over all possible initial conditions when forming the likelihood, but in most cases this is not computationally feasible. Thus, a number of approximate (ad hoc) approaches have been proposed. For example, EK had scanner data for three years, but they used the first two years to estimate the initial conditions for each consumer at the start of the third year, and then used only the third year in estimation.

2.5. Key Substantive Results of Erdem-Keane (1996)

Erdem and Keane (1996) estimated their model on Nielsen scanner panel data on liquid detergent from Sioux Falls, SD. The sample period was 1986-88, but only the last 51 weeks were used; at that time, telemeters were attached to panelists’ TVs, to measure household specific ad exposures. The data include 7 brands. A nice feature is that three brands were introduced during the period, generating variability in consumers’ familiarity with the brands. The estimation sample contained 167 households who met various criteria, like having a working telemeter.

Parameters to be estimated are: (i) true quality levels of the brands, (ii) the prior mean and variance of quality, (iii) variances of experience and ad signals, (iv) utility parameters – i.e.,

¹⁷ It is worth emphasizing that this high-order integration problem arises even in a static learning model (i.e., with myopic agents), as long as the contents of signals is not observed.

quality and price coefficients and the risk parameter, (v) the intercept and trend for no-purchase (and small brands), and (vi) the discount factor. Some key issues that arose in estimation are worth discussing, as they are common across many applications of dynamic learning models:

First, EK had difficulty obtaining a precise estimate of the weekly discount factor, and so pegged it at 0.995.¹⁸ The identification of the discount factor is often problematic in dynamic models, as we discuss in Section 4.2. Second, EK found it can be difficult to pin down the prior mean of quality along with all other model parameters. Hence, they constrained the prior mean of quality to equal the average of the true quality level across all the brands. This implies that peoples' priors are correct *on average*. They also constrained the prior uncertainty to be equal across brands $\sigma_{j0} = \sigma_0$, as allowing it to differ did not significantly improve the likelihood.

Along with the dynamic learning model, EK estimated two other models for comparison purposes. These were a myopic learning model ($\beta = 0$), and a simple reduced form model similar to Guadagni and Little (1983), henceforth GL. This was a multinomial logit that included an exponentially smoothed weighted average of past purchases (the “loyalty” variable), a similar variable for ad exposures, a price coefficient, brand intercepts, and trends for values of no purchase and small brands.

Strikingly, EK found that both learning models fit substantially better than the GL model.¹⁹ This is surprising, as GL specifies flexible (albeit ad hoc) functional forms for effects of past usage and ad exposures on current choice probabilities, while the Bayes learning models impose a very special structure. Specifically, as we saw in (17) and (18), only the sum of past experience or ad exposures matter in the Bayesian model, not the timing of signals.

Another striking result is that advertising is not significant in the GL model, implying advertising has no effect on brand choice. The EK model does not have one coefficient to capture the effect of advertising. The parameter r is significant and positive, so consumers are risk averse with respect to quality, while the σ_0 , σ_ε and σ_A imply: (i) consumers have rather precise priors about new brands in the detergent category, and (ii) experience signals are much more accurate than ad signals. But the effect of advertising can only be assessed via simulations.

¹⁸ In trying to estimate the weekly discount factor, they obtained 1.001 with a standard error of 0.02. This standard error implies a large range of annual discount factors. It is also worth noting that Erdem and Keane set the terminal period for the DP problem at $T=100$, which is 50 weeks past the end of the data set.

¹⁹ The dynamic learning model had 16 parameters while the other two models both had 15. EK obtained BIC values of 7531, 7384 and 7378 for GL and the myopic and forward-looking learning models, respectively.

EK used their model to simulate the impact of an increase in ad frequency for Surf from 23% to 70%.²⁰ The simulation is also done for a hypothetical new brand with the characteristics of Surf. The results indicate that an increase in advertising has little effect on market share for about 4 months, but the impact is substantial after about 7 or 8 months. Thus, the model implies that advertising has little impact in the short run, but sustained advertising is important in the long run. And, as expected, the impact of advertising is much greater for a new brand (as there is more scope for learning).²¹

The advertising simulation results are not surprising in light of the parameter estimates. As consumers have rather precise priors about brands in the detergent category, and as ad signals are imprecise, it takes sustained advertising over a long period to move priors and/or reduce perceived risk of a brand to a significant degree.²² A clear prediction is that the higher is prior uncertainty, and the more precise are ad signals, the larger will be advertising effects and the quicker they will become noticeable. Thus, an important agenda for the literature on learning models is to catalogue the magnitudes of prior uncertainty and signal variances across categories.

Finally, it is notable that, while the forward-looking model provides a statistically better fit than the myopic model, the improvement is modest (only 6 points in BIC). This is because subjective prior uncertainty is fairly low for detergent, so perceived gains from trial are small. An important agenda item for the literature on learning models is to compare the degree of prior uncertainty across categories, to determine when forward-looking behavior is most important.

3. A Review of the Recent Literature on Learning Models

Here we review developments in learning models subsequent to the foundational work discussed in Section 2. Almost all of this work is post-2000, but it already forms a large literature. We divide the review into (i) more sophisticated models of learning, (ii) correlated learning, (iii) new applications of learning models (beyond brand choice), (iv) non-standard Bayesian learning models and (v) equilibrium models.

²⁰ “Ad frequency” is weekly probability of a household seeing an ad for a brand. In the data this was 23% for Surf.

²¹ We have noticed that Figure 1 in Erdem and Keane (1996) contains a typo. The scale on the y-axis in Figure 1, which reports results for the new brand with the myopic model, is incorrectly labeled. It should be labeled in the same way as Figure 5. This doesn’t affect any of the results we discuss here.

²² The simulation results also clarify why the GL model fails to find significant advertising effects. The “loyalty” variable tends to put more weight on recent advertising, and discounts advertising from several months in the past. But the simulations show that advertising in the past few months does little to move market shares.

3.1. More Sophisticated Models of Consumer Learning

3.1.1. Models with Forward-Looking Consumers and More Sophisticated Learning

Following Erdem and Keane (1996), several papers have made significant contributions in the area of learning models with forward-looking consumers. We discuss these in turn:

Ackerberg (2003) deviates from Erdem-Keane in several dimensions. Most notably, he models both informative and persuasive effects of advertising. The persuasive effect is modeled as advertising intensity shifting consumer utility directly. The informational effect is modeled by allowing consumers to draw inferences about brand quality based on advertising intensity.²³ This is quite different from the information mechanism in Erdem-Keane, where ad content provides noisy signals of quality. Other differences are: (i) he is primarily interested in learning about a new product, and his model allows for heterogeneity in consumers' match value with the new product, and (ii) he assumes it takes only one trial for consumers to learn the true match value. Estimating the model on scanner data for yogurt, Ackerberg (2003) finds a strong, positive informational effect of advertising. But the persuasive effect is not significant.

The key innovation of Crawford and Shum (2005) is to allow for multi-attribute learning. In their application to prescription drugs, they argue that panel data can allow them to separately identify two effects: (i) symptomatic, which impact a patient's per period utility via symptom relief, and (ii) curative, which alter the probability of recovery. They allow physicians/patients to have uncertainty along both dimensions. They also endogenize how long patients need treatment by allowing for the possibility they can recover (and leave the market). The main goal of their study is to examine the welfare cost of uncertainty relative to the "first-best" environment with no uncertainty. This research question would be hard to answer using a reduced form framework because it is hard to find a control environment where consumers do not face uncertainty about prescription drugs. However, after estimating their structural model, Crawford and Shum can simulate removal of uncertainty by setting the initial prior variance to be zero, and setting each consumer's prior match value to be the true match value. By conducting this counterfactual experiment, they find that consumer learning allows consumers to dramatically reduce the costs of uncertainty (as there is substantial heterogeneity in drug efficacy across patients).

Erdem, Keane and Sun (2008) was the first paper to model the quality signaling role of

²³ That is, ad frequency itself signals brand quality, as in the theoretical literature on "advertising as burning money" (which only high quality brands can afford to do) (Kihlstrom and Riordan 1984).

price in the context of frequently purchased goods. Also, they allow both advertising frequency and advertising content to signal quality (combining features of Akerberg (2003) and Erdem and Keane (1996)). And they allow use experience to signal quality, so that consumers may engage in strategic sampling. Thus, this is the only paper that allows for these four key sources of information simultaneously. In the ketchup category they find that use experience provides the most precise information, followed by price, then advertising. The direct information provided by ad signals is found to be more precise than the indirect information provided by ad frequency.

The main finding of Erdem, Keane and Sun (2008), obtained via simulation of their model, is that, when price signals quality, frequent price promotions can erode brand equity in the long run. As they note, there is a striking similarity between the effect of price cuts in their model and in an inventory model. In each case, frequent price cuts reduce consumer willingness to pay for a product; in the signaling case by reducing perceived quality, in the inventory case by making it optimal to wait for discounts. We return to this issue in Sections 4 and 5.

Osborne (2011) is the first paper to allow for both learning and switching costs as sources of state dependence in a forward-looking learning model. This is important because learning is the only source of brand loyalty in Erdem and Keane (1996). So it is possible they only found learning to be important due to omitted switching costs. However, Osborne finds evidence that both learning and switching costs are present in the laundry detergent category. When learning is ignored, cross elasticities are underestimated by up to 45%.²⁴

Erdem, Keane, Öncü and Strebel (2005) represents a significant extension of previous learning models, as it is the first paper where consumers actively decide how much effort to devote to search before buying a durable. This contrasts with Roberts and Urban (1988) where word-of-mouth (WOM) signals are assumed to arrive exogenously, or Erdem and Keane (1996) where ad signals arrive exogenously. Another novel feature of the paper is that there are several information sources to choose from (WOM, advertisements, magazine articles, etc.) and in each period consumers decide how many of these sources to utilize.²⁵

In their application, Erdem et al. (2005) consider technology adoption (Apple/Mac vs.

²⁴ Osborne (2011) allows for a continuous distribution of consumer types. Of course, it is literally impossible to solve the DP problem for each type (which is why the DP literature usually assumes discrete types). Thus, some approximation is necessary here. Osborne is able to estimate his model by adapting the MCMC algorithm developed by Imai, Jain and Ching (2009), and extended by Norets (2009) to accommodate serially correlated errors.

²⁵ Chintagunta et al. (2012) model how physicians learn about drugs using both patients' experiences and detailing.

Windows) in personal computer markets where there is both quality and price uncertainty. As in the brand choice problem, consumers are not perfectly informed about competing technologies. But a unique aspect of high-tech durables is rapid technical progress. This causes the price of PCs to fall rapidly over time. This creates two incentives to delay purchase: (i) to get a better price, and (ii) to search for more information about the competing technologies. Waiting, however, implies a forgone utility of consumption. Erdem et. al. estimate their model on survey panel data collected from consumers who are in the market for a PC. Their findings indicate that consumers defer purchases both to gather more information and to get a better price. Simulations of their model imply that learning is the more important reason for purchase delay.

3.1.2. Models with Myopic Agents and More Complex Learning Mechanisms

Another stream of literature has focused on extending learning models by allowing for more complex learning mechanisms or more complicated consumer heterogeneity. To make such extensions feasible, these papers have assumed that consumers are myopic:

For instance, besides accounting for informative and persuasive effects of advertising, Mehta, Chen and Narashiman (2008) model the “transformative” effect of advertising, where advertising affects consumer assessment of the consumption experience. They find the strength of advertising’s different effects change over time, informative advertising being more important early in the product life-cycle while transformative advertising is more important in later.

With the availability of rich pharmaceutical data, several researchers have focused on detailing as an information channel. A good example is Narayanan, Manchanda and Chintagunta (2005), who model how physicians learn about drugs. They consider a set of brand-name drugs before patent expiration and investigate the informative and persuasive roles of detailing. They assume firms know the true quality of their products. Thus, they follow Erdem and Keane (1996) and assume that detailing provides noisy but unbiased signals about drug efficacy. This crucial assumption allows them to separate informative and persuasive roles of detailing. It implies that, in the long-run, the impact of informative detailing on demand is negligible. Therefore, any long-run effect of detailing on sales must come from its persuasive role. Their framework also allows them to quantify the temporal evolution of responsiveness of physicians to detailing.

Narayanan and Manchanda (2009) also study the informative and persuasive roles of detailing. They extend Narayanan et al. (2005) by allowing for physician heterogeneity in

learning rates. They find that physicians who are more responsive to detailing in early periods are less responsive later on, and vice versa. They conclude that firms could increase their revenue if they took this heterogeneity into account while deciding on allocations of detailing.

Ching and Ishihara (2012a) use a new approach to separate persuasive and informative roles of detailing. They focus on drugs that have signed a co-marketing agreement, under which two firms can market the same chemical under different brand-names. A key assumption is that the informative component of detailing is chemical specific while the persuasive component is brand specific. Then, the persuasive effect of detailing is identified by variation in the relative market share of two brands. The informative effect of detailing is identified by variation in the market share of chemicals (given detailing effort summed across brands of the same chemical).

Using data for ACE-inhibitors, Ching and Ishihara (2012a) find that the informative role of detailing is mainly responsible for the diffusion patterns of chemicals. But the persuasive role of detailing plays a crucial role in determining the demand for brands which co-market the same chemical. They also find that patients could be worse off if the government bans detailing.

Another way to extend basic learning models is to focus on mechanisms by which consumers learn from others. As we've noted Roberts and Urban (1988) focused on word-of-mouth, but other mechanisms are possible. For instance, Yang, Zhao, Erdem and Koh (2010) model dyadic learning where a spouse updates his/her belief of product quality based on both own and spouse's past experience with a product. Looking at cell-phone providers, they find that as communication between spouses becomes more effective their choice shares converge.

Another paper that studies both across and within consumer learning is by Chintagunta, Jing and Jin (2009). They study doctor's prescribing decisions for Cox-2 inhibitors, a new class of pain killers. They make use of patient diary data that records actual use experience. Thus, like Roberts and Urban (1988), they do not require simulation methods to integrate latent experience signals out of the likelihood. Having measures of perceived quality also enables them to separate across and within consumer learning.²⁶ They find that both types of learning are important. In a counterfactual experiment, they find that creating a nationwide database of patient feedback would speed up doctors' decisions to switch from traditional pain killers to Cox-2 inhibitors.

Three papers that also study learning from others are Ching (2010a, b) and Ching and

²⁶ A limitation is that they treat the discrete signals from patients' diaries as a continuous variable. Chernew et al (2008) show how to use this type of data by estimating a model with a discrete learning process.

Ishihara (2010). These papers model the supply side as well, so we discuss them in Section 3.5.

3.2. Models of Correlated Learning

By “correlated learning” we mean learning about a brand in one category through use experience of the same brand in another category, and/or learning about one attribute (e.g., drug potency) from another (e.g., side effects). This will occur if priors and/or signals are correlated across products or attributes. Several papers have considered such effects:

Erdem (1998) considers a model where priors are correlated across umbrella brands. She finds evidence that consumers learn via experience across categories for umbrella brands in the toothpaste and toothbrush categories. She shows that brand dilutions can occur if a brand in the “parent” category (toothbrush) is extended to a new product in a different category (toothpaste) and the new product is not well-received. This framework has been extended to study decisions about fishing locations (Marcoul and Weninger (2008)), and adoption of organic food products (Sridhar, Bezawada and Trivdei (2012)). These papers assume a single quality dimension.

In a multi-attribute setting, consumers may make inferences about one attribute based on their experience of another.²⁷ Prescription drugs are a good example: Coscelli and Shum (2004) estimate a diffusion model for the anti-ulcer drug Omeprazole. An anti-ulcer drug can treat: (i) heartburn, (ii) hypersecretory conditions, and (iii) peptic ulcer, and provide (iv) maintenance therapy. In their model, patients’ experience signals on these four dimensions are correlated. Physicians know the correlation structure, and use of their patients’ experience signals to update their multi-dimensional prior beliefs. The estimates imply signals for the first three attributes are positively correlated, but they are negatively correlated with the fourth attribute.²⁸

Chan, Narasimhan and Xie (2012) also consider prescription drugs. They focus on side-effects and effectiveness, and assume experience signals are correlated on these dimensions. They achieve identification by combining revealed and stated preference data (patients report the reason they switch: side-effects or ineffectiveness). Perhaps not surprisingly, they find detailing visits are much more effective in reducing uncertainty about effectiveness than side effects.

Lim and Ching (2012) apply a multi-dimensional learning model with correlated beliefs to the major class of anti-cholesterol drugs, statins. They focus on two quality dimensions,

²⁷ The model in Erdem (1998) can, in principle, be extended to this case by assuming consumers use a vector of signals to learn about a vector of attributes, and by allowing attributes and/or signals to be correlated.

²⁸ Both Coscelli and Shum (2004) and Crawford and Shum (2005) study the anti-ulcer drug market. But in order to study forward-looking physicians’ behavior, Crawford and Shum (2005) abstract away from correlated learning.

effectiveness in lowering cholesterol (q^c) and heart disease risk (q^h). High cholesterol is a risk factor for heart disease. While statins may reduce heart disease risk by reducing cholesterol, the exact relationship (q^h/q^c) may differ across drugs. Lim and Ching argue that physicians are uncertain about the relationship for a particular drug, but they have correlated prior beliefs about q^c and q^h . Since heart attack is a rare event, physicians cannot readily use patient's experiences to update their priors. Therefore, Lim and Ching argue that clinical studies, informative detailing and media coverage play crucial roles in physicians' updating process. To estimate their model, they combine product level market share data with data on these four information signals. They find that physicians' priors about effectiveness are correlated across drugs. This implies that clinical studies have spillover effects. Interestingly, the drug that is most effective in lowering cholesterol is best able to "free-ride" on clinical studies conducted by other firms.

Dickstein (2012) considers a model in which forward-looking physicians are uncertain about patients' intrinsic preferences for drug characteristics. The physician uses patients' utility of consuming a drug at time t to update his/her belief about their preferences. Physicians do not observe patients' preferences for each drug characteristic. But patient utility is assumed to be a linear function of characteristics. Thus, the Bayesian updating procedure for physicians is similar to Bayesian inference in a linear regression model. This updating procedure generates correlated posterior beliefs about patient preferences, which leads to a complex DP problem.²⁹ This is the first model with forward-looking agents that allows for information spillover.

Finally, Che, Erdem and Öncü (2011) develop a forward-looking demand model with spillover effects in learning and changing consumer needs over time. They estimate their model using scanner data for the disposable diapers category, where consumers have to switch to the next bigger size periodically as babies grow older. This leads to an increase in strategic trial around the time of needing to change size. Their results imply that consumer experience of a particular size of a brand provides a quality signal for other sizes, and consumer quality sensitivities are lower and price sensitivities higher for larger sizes than smaller sizes.

3.3. Other Applications of Learning Models – Services, Insurance, Media, Tariffs, Etc.

Learning models have been applied to many problems other than choice among different

²⁹ Dickstein uses an extension of the Gittin's index method proposed by Lai (1987). The original method only applies to multi-armed bandit problems with independent signals. Lai's approach allows for correlated signals. But, as in Eckstein et al (1988), one must assume consumers are risk neutral. Ferreyra and Kosenok (2010) use a method similar to Gittin's index to estimate a simpler dynamic learning problem.

brands of a product. In particular, Bayesian learning models have been applied to choice among services, insurance plans, media, tariffs, etc. Here we discuss these types of applications.

Israel (2005) uses a learning model to study the customer-firm relationship in the auto insurance market. This environment is well suited for studying learning because opportunities to learn arrive exogenously when an accident happens. Presumably, when consumers file a claim, they learn something about the customer service of the insurance company. If a consumer leaves the company after filing a claim, it may indicate that they had a negative experience.

Bayesian learning models have also been used to infer the value of certification systems that are meant to assist consumers in making well-informed choices. Chernew et al (2008) study the value of the health plan reports that GM provides to their employees, and Xiao (2010) studies the value of quality accreditation for the childcare market. Interestingly, both studies find that the certification system provides a very limited amount of information to consumers.

Anand and Shachar (2011) apply a Bayesian learning model to choice of TV programs. They study the role of advertising in matching consumers to programs when they are uncertain about observed and unobserved program attributes. They also allow for the persuasive effect of advertising. The results indicate that an exposure to a single advertisement reduces a consumers' chance of choosing a sub-optimal alternative by approximately 10%.

Chan and Hamilton (2006) is a surprising application of an Erdem-Keane style structural learning model to study clinical drug trials. The “natural experiment” school of econometrics views clinical trials as the “gold standard” to which economics should aspire to avoid the need for structural modeling – see Keane (2010a). But Chan and Hamilton show how a structural learning model can help evaluate clinical trial outcomes by correcting for endogenous attrition. In their model, patients in a trial are uncertain about effectiveness of the treatment, but they learn from experience signals. In each period, they need to decide whether to quit based on costs (side-effects) vs. expected benefits of continuing the trial. They find that a structural interpretation of clinical trial results can be very different from the standard approach. For example, a treatment that is less effective in reducing CD4 count may be preferable because it has fewer side-effects.³⁰

Bayesian learning models have also been applied to tariff choice. It is widely believed that consumers are irrational when they choose between the flat-rate and per-use plans. This is

³⁰ CD4 is part of the immune system. The CD4 count is used to decide when to begin treatment for HIV infection.

because several studies have found that many consumers could save by choosing per-use option instead of the flat-rate option. But these papers tend to look at behavior over a short period. Miravete (2003) finds strong evidence to contradict the “irrational consumer” view. Using data from the 1986 Kentucky tariff experiment, he provides evidence that consumers learn their actual usage rates over time, and switch plans in order to minimize their monthly bills.

Narayanan, Chintagunta and Miravete (2007) interpret the same data using a Bayesian learning model with myopic consumers. To explain why consumers make mistakes in choosing an initial plan, they assume they are uncertain about their actual usage. The structural approach allows them to quantify changes in consumer welfare under different counterfactual experiments. Iyengar, Ansari and Gupta (2007) develop another myopic learning model that is closely related. But they also allow consumers to be uncertain about the quality of service.

Goettler and Clay (2011) use a Bayesian learning model to infer switching costs for tariff plans. They do not observe consumers switching plans in their data. Identification of switching costs is achieved by assuming consumers are forward-looking, have rational expectation about their own match value, and make plan choice decisions every period after their initial enrollment. The implied cost to rationalize no switching is quite high (\$208 per month).

Grubb and Osborne (2012) argue an alternative way to explain infrequent plan switching is that consumers do not consider plan choice every period (as consideration is costly and/or time consuming). They formulate the consideration decision using the “Price Consideration Model” of Ching, Erdem and Keane (2009), and the switching decision as Bayesian learning model. Their rich data set allows them to investigate prior mean bias, projection bias and overconfidence.

Finally, learning models have also been extended to study fertility decisions (Mira 2007), bargaining problems (Watanabe 2010), a manager’s job assignment problem (Pastorino 2012), voters’ decision problems (Knight and Schiff 2010), and even the U.S. kidney market (Zhang 2010).

3.4. Non-standard Learning Models

Bayesian learning models make some strong assumptions about how consumers process information – they never forget past signals, they use optimal Bayesian updating rules (and do not over or underreact to new information), they do not suffer from various cognitive biases like confirmatory bias, etc. There have been some attempts to relax these assumptions.

Mehta, Rajiv and Srinivasan (2004) extend the Bayesian model to account for forgetting. Consumers imperfectly recall prior brand experiences, and the extent of forgetting increases with time. Then, a consumer's state depends on the timing of signals, not just the total number (as in equation (6)). Thus, it is necessary to assume myopia to make modeling forgetting feasible.

Camacho, Donkers and Stremersch (2011) argue that some types of experience may be more salient in certain contexts. For example, a physician may pay special attention to feedback from patients who have just switched treatment. They modify the standard Bayesian model by introducing a salience parameter to capture the extra weight physicians may attach signals in that case. Using data on asthma drugs, they find evidence that feedback from switching patients receives 7-10 times more weight in physician learning than feedback from other patients.

Zhao, Zhao and Helsen (2011) consider whether consumers who receive a very negative experience signal may change their perception of signal variance. In their model consumers are uncertain about the precision of quality signals, and they update their perception of this precision over time. They estimate the model using scanner data that spans the period of a product-harm crisis affecting Kraft Australia's peanut butter division in June 1996. Their model fits the data better than a standard consumer model, which assumes consumers know the true signal variance.

3.5. Equilibrium Models

The demand system generated by a consumer learning model is always dynamic in the sense that current sales affect future demand.³¹ This is true regardless of whether consumers are forward-looking or myopic. Thus, in a consumer learning environment, we expect firms to be forward-looking when choosing their marketing-mix. But it is very difficult to estimate this type of model because the firm's state, and hence its optimal pricing policy, is so complex. That is, it depends on the whole *distribution* of perceived quality of its own brand and all competing brands across all consumers. Obviously the exact solution of such a complex problem is infeasible.

Ching (2010a) developed a market level demand model that captures consumer learning in a parsimonious way, so when combined with an oligopolistic supply-side the size of the state space is manageable. He modifies the learning model of Erdem and Keane (1996) so consumers learn from each others' experience via social networks ("*social learning*"). Once information on

³¹ This is why Erdem and Keane (1996) adopted the cumbersome terminology of "forward-looking dynamic model" vs. "immediate-utility-maximization dynamic model" to denote the difference between learning models with forward-looking vs. myopic agents. Here we instead refer to "forward-looking" vs. "myopic" learning models.

quality is socially aggregated, consumers share common beliefs. This assumption eliminates the distribution of consumers across different state as the state variables for firms.

Another contribution of Ching (2010a) is to show how to estimate this social learning model using product level market share data. Even with this simplified state space, estimating the demand and supply-side models jointly is very difficult. So, instead of solving the firm's problem exactly, he posits an approximation to the optimal pricing policy function. It depends on a relatively small set of state variables that approximate the state of the system (but that is parsimonious compared to the complete state space). To control for endogeneity of price in the demand side model, Ching estimates the pseudo-policy function jointly with the learning model.

Ching applies this approach to study demand for brand-name drugs and their generic counterparts in the 80s. Social learning is plausible in this market because each physician sees many patients, and physicians interact with each other via their own network. His results suggest that physicians/patients have pessimistic initial priors about new drugs, and that social learning plays a role in explaining the slow diffusion of generics in the 80s.³²

Ching (2010b) combines the social learning demand model with a dynamic oligopolistic supply side model. As far as we know, this is the first empirical paper to combine a dynamic demand system with forward-looking firms. Also, in his model, *both* consumers and firms are uncertain about the quality of generic drugs. The model is tailored to study competition between brand-name and generic drugs. Generic firms' entry decisions are endogenous, but entry timing depends on an exogenous random approval process by the FDA. Equilibrium is Markov-perfect Nash, as in Maskin and Tirole (1988). Ching applies the model to the market for clonidine. By simulating the model, he finds it can rationalize two stylized facts: (i) the slow diffusion of generic drugs, and (ii) the fact that brand-name firms raise their prices after generic entry.³³

Ching also uses the model to simulate a reduction in expected approval time for generics. Surprisingly, while this brings generics to the market sooner, it may reduce the total number of generic entrants. This is because it increases a potential entrant's chance of entering a crowded

³² Gu and Yang (2011) also use this pseudo-policy policy function approach to estimate an attribute-based learning model for the breakfast cereal market. Their main goal is to correct the price endogeneity problem that may arise in equilibrium. Chen et al. (2009) extend Ching's empirical demand framework to study the impact of word-of-mouth on movie demand over time for limited release movies. But they did not incorporate the supply side.

³³ Ching's model also allows for heterogeneity in consumer price-sensitivity. His model generates price increases by name-brand firms because, as learning takes place, an increasing proportion of the price-sensitive consumers switch to generics. Hence, the demand faced by the brand-name firms becomes more inelastic over time.

market. The overall impact on consumer welfare is ambiguous.

The demand system in Ching (2010a) was extended by Ching and Ishihara (2010) to study the effects of detailing from a new perspective. A stylized fact in the drug market is that effectiveness of detailing changes when new information arrives – e.g., when results of a new clinical trial are positive for a drug, the effectiveness of detailing increases. But the standard Bayesian learning model implies the marginal impact of information signals must fall over time.

Similar to Ching (2010b), Ching and Ishihara (2010) assume both consumers and firms are uncertain about the quality of the product. Social learning takes place via an intermediary (opinion leader, consumer watch group, etc.). As firms are uncertain about quality of their products, ads are not modeled as noisy signals of quality. Instead, the purpose of advertising or detailing is to build up a stock of physicians who are familiar with the most recent information of the intermediary. This leads to heterogeneity in the population, as the proportion of informed physicians is a function of detailing.

Ching and Ishihara (2010) estimate the demand model jointly with a pseudo-policy function to control for potential endogeneity of detailing.³⁴ They estimate the model using data on ACE-inhibitor with diuretic in Canada. The results are consistent with the stylized fact that the effectiveness of detailing changes when a new clinical trial outcome is released. The demand model developed by Ching and Ishihara is quite parsimonious and it could, in principle, be combined with an oligopolistic supply side. But this has not yet been done. In general, there is scope for much more research using equilibrium models to study consumer learning. **For instance, so far there is no research that combines dynamic oligopolistic supply side model with correlated learning and non-standard learning demand models that we discussed earlier.**

Two-sided learning makes estimation of equilibrium models computationally difficult. An interesting paper by Hitsch (2006) focuses on a firm learning about the demand for a new product. His approach is to abstract from the consumer learning problem by simply using a reduced form demand model. A firm needs to learn about the true demand parameters. In other words, in contrast to all the papers that we describe above, consumers are assumed to have no uncertainty about their preferences for the new product. Hence, unlike Ching (2010b), he considers a one-side learning equilibrium model. He also abstracts from competition. These

³⁴ Price endogeneity is not a concern because they apply their model to the Canadian prescription drug market, where prices are subject to regulation and cannot be changed freely.

simplifications significantly reduced the computational burden of the estimation, and yet the model delivers important new insights to the product launch and exit problem.

4. Limitations of the Existing Literature and Directions for Future Work

Here we discuss limitations of existing learning models and directions for future research. In our view, the main limitations are: (1) identification can be difficult in complex models, especially in the case of behaviorally rich specifications, (2) there is no clear consensus on the evidence for forward looking vs. myopic consumers as often the discount factor is difficult to identify, and (3) it may be difficult to disentangle different sources of dynamics (for example, there is no clear consensus on the importance of learning vs. inventories as a source of dynamics).

4.1. Behaviorally Richer Specifications Can Make Identification More Difficult

In Section 2 we discussed the formal identification of learning models, and this topic is also addressed in a number of other papers such as Erdem, Keane and Sun (2008). By formal identification we refer to a proof that a parameter is identified, given the structure of the model, as well as a discussion of any normalizations that are needed to achieve identification. An important point, however, is that it is possible for a parameter be formally identified and yet (i) the intuition for what patterns in the data actually pin it down are not clear, and/or (ii) the likelihood is so close to flat in the parameter that it is not practical to estimate it in practice (what Keane (1992) called "fragile identification"). These problems are not unique to dynamic learning models but they deserve further attention in this context.³⁵

As we discussed in Section 2, the taste parameters in learning models (e.g., the utility weight on quality and price) and the quality levels of brands are identified from the choice

³⁵ Structural models in general, and learning models in particular, are sometimes criticized on the grounds that they make a large number of assumptions about the learning process, the functional form of utility, etc.. Identification of these models relies on these functional form assumptions. The criticism of structural models often takes the form that they rely on "too many assumptions" and that we should prefer "simple methods" and/or "model free" evidence. The debate on this topic is extensive, and beyond the scope of this survey. For further discussion we refer the reader to the articles such as Heckman (1997), Keane (2010a,b) and Rust (2010). These authors argue that drawing inferences from data always relies on some set of maintained assumptions. They argue that "simple" reduced form or statistical models typically rely on just as many assumptions as structural models, the main difference being that the "simple" models leave many assumptions implicit.

Worth particular mention is the literature on "non-parametric identification," which has apparently been misinterpreted by many researchers as suggesting that it may be possible to obtain "model free evidence" in some cases. In fact, the standard approach of the non-parametric identification literature is to make *a priori* assumptions about certain parts of a model, and then show that some other part of the model (e.g., the functional form of utility or an error distribution) is identified without further assumptions. Thus, what is being non-parametrically identified is a part of the model, not all of it. Furthermore, a proof that a distribution or functional form is non-parametrically identified does not say anything about whether identification is feasible in practice.

behavior of individuals with sufficient experience with all brands so that their learning has effectively ceased, so their choice behavior is stationary. In contrast, the learning parameters (risk coefficient, initial perception variance and signal variabilities) are pinned down by the dynamics of choice behavior amongst households with little prior experience (and how their behavior differs from more experienced households). However, when one extends the basic set-up of the EK learning model discussed in Section 2, it becomes more difficult to understand what data patterns help identify the structural parameters.

As saw in Section 3, there is a healthy trend towards specifying behaviorally richer, and hence more complex, learning models. Examples are models where consumers have multiple sources of information or learn about multiple objects (correlated learning), or models that incorporate findings from psychology/behavioral economics. But this added complexity creates three problems: First, showing formal identifications for such complex models may be difficult. Second, even if formally identified, such complex models may be difficult to estimate in practice (e.g., there may be practical problems with “fragile” identification). Third, even if one succeeds in estimating the parameters of a complex learning model, the intuition for what data patterns or sources of variation pin down certain parameters may be hard to understand.

One particular issue is that, given only data on purchase decisions and signal exposures, it is very hard to distinguish alternative learning mechanisms, as multiple mechanisms may fit the data about equally well. Thus, given only revealed preference (RP) data (e.g., scanner data), identification of models with complex learning mechanisms may be difficult or impossible.

One solution to this problem is to combine multiple data sources – e.g., using both RP data and survey data on expectations/perceptions, or supplementing scanner data with publicly available information like product reviews or media coverage. For example, consider the paper Erdem, Keane, Oncu and Strebel (2005) on how consumers learn about computers. In addition to RP data, they also had stated preference (SP) data on how people rated each brand in each period leading up to purchase. They treated the SP data as providing noisy measures of consumer’s perceptions. This enabled them to identify variances of different information sources. Intuitively, if peoples’ ratings tend to move a lot after seeing an information source (and their perceived uncertainty tended to fall a lot), it implies that information source is perceived as accurate. Other papers that combine RP and SP data to aid in identification are Chintagunta et al. (2009),

Chan et al. (2012), Shin, Misra and Horsky (2012). The latter is an attempt to disentangle state dependence from learning.

In another approach, some papers combine choice data with direct measures of information signals. Ching and Ishihara (2010) and Lim and Ching (2012) use results of clinical trials to pin down the content of signals received by physicians, and incorporate them into their structural learning models. In a reduced form study, Ching et al. (2011) use data on media coverage of prescription drugs and find evidence that when patients learn that anti-cholesterol drugs can reduce heart disease risk, they become more likely to adopt them. Kalra et al. (2011) attempt to pin down content of information signals by examining news articles. Thus, there has been some work in this area but there is obviously much room for further progress.

4.2. Forward-Looking vs. Myopic Consumers

As we discussed in Section 2, the key distinction between forward-looking and myopic models is whether consumers engage in strategic trial. But the evidence on whether consumers are forward-looking is mixed. Indeed, in many applications, researchers have found that it is difficult to identify the discount factor in practice, because the likelihood is rather flat in this parameter. For instance, Erdem and Keane (1996) found that increasing the discount factor from 0 (a myopic model) to 0.995 improved the likelihood by only 6 points. That was significant, but if the likelihood is so flat in the discount factor, it is hard to discern forward-looking behavior.

As forward-looking models often have not provided substantial fit improvements, and as they are hard to estimate, it is not surprising that many researchers have adopted myopic models. But before taking this path, it is important to emphasize that trial purchase is the distinguishing feature of forward-looking models. In a mature category, consumers may have nearly complete information, leaving little gain from strategic trial. Then a forward-looking consumer will behave much like a myopic one – it is impossible to tell the two types apart, and the discount factor is not identified. In contrast, in a market with significant uncertainty about product attributes, the rate of trial identifies the discount factor – the higher the discount factor, the more consumers try new brands. Forward-looking models are more likely to provide a superior fit in such markets.

Given this background, we think it would be a mistake to infer from results on relatively mature categories that forward-looking models are unnecessary. This decision should be made on a case-by-case basis given the characteristics of the category under consideration.

Finally, we should note that if a dynamic model has exclusion restrictions (i.e., there are some variables that do not affect current utility, but that affect expected future payoffs via their impact on the transition density of the state variables), it can help identify the discount factor (see Magnac and Thesmar (2002) or Keane et al. (2011)). For example, Ching, Imai, Ishihara and Jain (2012) and Ching and Ishihara (2012b) use reward points for a brand with a frequent-buyer program as an example of an exclusion restriction (i.e., earning points does not affect current utility until total points reach a certain level). Fang and Wang (2010) show one can identify the parameters of quasi-hyperbolic discounting using at least three periods of data, provided one has exclusion restrictions. Chung, Steenburgh and Sudhir (2011) estimate the parameters of a quasi-hyperbolic discounting function in a dynamic sales force model where the accumulated sales made by a sales person acts as an exclusion restriction – one has to accumulate enough sales to reach the cutoff in order to earn a bonus. But before that cutoff is reached, accumulated sales has no impact on the current utility, and can only affect choice behavior via expected future payoffs.

4.3. Distinguishing Among Different Sources of Dynamics

Learning is one of many mechanisms that may cause structural state dependence. Other potential sources of state dependence include inertia, switching costs, habit persistence and inventories. In this section we discuss attempts to distinguish among these sources of dynamics. We particularly emphasize the problem of distinguishing between learning and inventories, because most dynamic structural models have assumed one of these mechanisms as the source of dynamics. Furthermore, and perhaps surprisingly, the behavioral patterns generated by learning can be quite similar to those generated by inventory behavior. Thus it can be very difficult to identify which mechanism generates the state dependence we see in the data.

Learning and inventory models generate dynamics in very different ways. Learning models generate persistence in choices (i.e., brand loyalty) as risk aversion leads consumers to stay with “familiar” brands. This familiarity arises endogenously, via information signals that cause consumers to gravitate toward particular brands early in the choice process. In contrast, inventory models do not generate persistence in brand choice endogenously. Rather they must postulate the existence of a priori consumer heterogeneity in tastes for particular brands. Obviously, a great appeal of learning models is that they provide a behavioral explanation for the emergence of brand preferences.

However, once we introduce unobserved taste heterogeneity, the dynamics generated by learning and inventory models are rather hard to distinguish empirically. The similarity of the two models is discussed extensively by Erdem, Keane and Sun (2008). They fit a learning model to essentially the same data used in the inventory model of Erdem, Imai and Keane (2003), and find that both models fit the data about equally well, and make very similar predictions about choice dynamics. For instance, both models predict that, in response to a price cut, much of the increase in a brand's sales is due to purchase acceleration rather than brand switching.

The similarity of the two models is even greater if we allow for price as a signal of quality. Then, both models predict that frequent price promotion will reduce consumer willingness to pay for a product; in the signaling case by reducing perceived quality, in the inventory case by changing price expectations and making it optimal to wait for discounts.

Obviously an important avenue for future research is to determine whether learning or inventory effects are of primary importance for explaining consumer choice behavior, or, indeed, whether both mechanisms are important. But unfortunately, computational limitations make it infeasible to estimate models with both learning and inventories. There are simply too many state variables – levels of perceived quality and uncertainty for all brands, inventory of all brands, current and lagged prices of all brands – to make solution and estimation feasible. This makes it impossible to nest learning and inventory models and test which mechanism is more important (or, alternatively, whether both mechanisms are important). Presumably, advances in computation will remove this barrier in the future.

In the meantime, some authors have proposed simpler approaches for testing whether learning or inventories or other factors generate choice dynamics. Ching, Erdem and Keane (2011) present a new “quasi-structural” approach that allows one to estimate models with both learning and inventory effects, while also testing whether consumers engage in strategic trial. Their idea is to approximate the *E_{max}* functions in (27) using simple functions of the state variables. They show how the learning model generates a natural exclusion restriction, in that the *E_{max}* function associated with choice of brand j contains the updated perception variance for brand j , while the *E_{max}* functions for all brands $k \neq j$ contain only the current perception variance for brand j . They present an application to the diaper category, which is ideal for studying learning because an exogenous event, birth of a first child, triggers entry into the market. Their

results suggest that learning and strategic trial are quite important, while inventories are a much less important source of dynamics in the diaper category.

Erdem, Katz and Sun (2010) developed a reduced form test of the relative importance of inventories vs. learning. They consider the learning mechanism where consumers use price as a signal of quality. They also exploit the fact that inventory models generate “reference” price effects (i.e., purchase decisions are based on the current price *relative* to the reference price of a brand). Their test relies on the interaction between a use experience term and the reference price (operationalized as an average of past prices). In a learning model, higher use experience should be associated with less use of price as a quality signal. Based on this test, they find evidence for both learning and inventory (i.e., reference price) effects for several frequently purchased goods.

Dubé, Hitsch and Rossi (2010) test for evidence of learning using a reduced form discrete choice model of the form:

$$(33) \quad V_{jt} = \alpha_j - w_P P_{jt} + \gamma_1 d_{j,t-1} + \gamma_2 d_{j,t-1} \cdot N_j(t) + \gamma_3 N_j(t) + e_{jt}$$

Here, given adequate controls for taste heterogeneity (α_j), the lagged choice variable $d_{j,t-1}$ can matter for several reasons, including inertia, switching costs, habit persistence or learning (i.e., it increases experience with the brand). However, Dubé et al argue that lagged choice should not be a significant predictor of current choice in a learning model for experienced consumers who have complete information. This is because the lagged choice does not augment their information set. The more general implication is that the interaction coefficient $\gamma_2 < 0$, meaning that more use experience $N_j(t)$ reduces the effect of lagged purchase. But, using data on margarine and frozen orange juice, they find that more experience with a product does not reduce the lagged choice effect ($\gamma_2 \approx 0$), contradicting the learning model prediction.

In our view there are some problems with interpreting this result. First, if we look back at equation (27a), we see that, in a static learning model ($\beta=0$) the use experience term $N_j(t)$ would enter expected utility through the perception variance σ_{jt}^2 . But σ_{jt}^2 is not linear in $N_j(t)$. Rather, it is given by equation (6), where $\sigma_{jt}^2 = [(1/\sigma_{j0}^2) + N_j(t)(1/\sigma_\varepsilon^2)]^{-1}$. Thus, even if a static learning model were true, the lagged choice variable may be significant in (33) because the functional form of the proxy for σ_{jt}^2 is mis-specified.

Second, as we have discussed earlier, there is no reason to presume that learning is the

only source of dynamics. A finding that lagged choice is significant for consumers with complete information would not necessarily mean that learning had not been important when a product was new (or when consumers were young). It may simply mean that some other source of dynamics (like habit persistence) is also present.

Third, not all types of learning model imply that choice behavior becomes stationary given sufficient experience. For instance, as we discussed in Section 3.4, the basic model may be extended to include forgetting (see Mehta, Rajiv and Srinivasan (2004), or Ching and Ishihara (2010)), or non-Bayesian learning more generally. It is also possible that product attributes change over time, **or the quality signal variance becomes more precise over time**.³⁶ Conceptually, it is straightforward to construct learning models where recent experience is more salient for a variety of reasons. But, as we discussed earlier, structural models of this variety are very difficult to implement because a consumer's state then depends on the timing of signals, not just the total number of signals received.

5. Summary and Conclusion

In this survey we laid out the basic Bayesian learning model of brand choice, pioneered by Eckstein et al (1988), Roberts and Urban (1988) and Erdem and Keane (1996). We described how subsequent work has extended the model in important ways. For instance, we now have models where consumers learn about multiple product attributes, and/or use multiple information sources, and even learn from others via social networks. And the model has also been applied to many interesting topics well-beyond the case of brand choice, such as how consumers learn about different services, tariffs, forms of entertainment, medical treatments and drugs.

We also identified some limitations of the existing literature. Clearly an important avenue for future research is to develop richer models of learning behavior. For instance, it would be desirable to develop models that allow for consumer forgetting, changes in product attributes over time, a greater variety of information sources, and so on. But such extensions present both computational problems and problems of identification. We suggest it would be desirable to

³⁶ **Allowing use experience quality signal variance to change over time is also another research topic for future research. The only related paper is Zhao, Zhao and Helsen (2011). But they only allow for a one time change in quality variance after the product harm crisis.**

augment RP data with direct measures of consumer perceptions and direct measures of signal content to help resolve these identification problems.

One clear limitation of the existing literature has been the difficulty of precisely estimating the discount factor in dynamic learning models. This makes it difficult to distinguish forward-looking and myopic behavior. We discussed the search for exclusion restrictions (i.e., variables that affect future but not current payoffs) to help resolve this issue.

Finally, another key challenge for future research is to develop models that combine learning with other potentially important sources of dynamics, such as inventories or habit persistence. We noted it has not been possible to build inventories into dynamic learning models due to computational limitations. However, this course of research is important, because the dynamics generated by inventories can be quite similar to those generated by learning. Thus, it is important to try to distinguish between the two mechanisms. The identification of different sources of dynamics is also a challenge, and we again conclude that progress would be aided by the combination of RP and SP data.

In summary, it is clear that learning models have contributed greatly to our understanding of consumer behavior over the past 20 years. Two of the best examples still come from the original Erdem and Keane (1996) paper: First, that when viewed through the lens of a simple Bayesian learning model the data are consistent with strong long-run advertising affects. Second, that a Bayesian learning model can do an excellent job of capturing observed patterns of brand loyalty. Future work will reveal if such key findings are robust to the extension of these models to include multiple sources of dynamics and behaviorally richer models of learning behavior.

References

- Akerberg, D. (2003) Advertising, learning, and consumer choice in experience good markets: A structural empirical examination. *International Economic Review*, 44(3): 1007-1040.
- Aguirregeburria, V. and P. Mira (2010) Dynamic Discrete Choice Structural Models: A Survey. *Journal of Econometrics*, 156(1): 38-67.
- Anand, B. and R. Shachar (2011) Advertising, the Matchmaker. *RAND Journal of Economics*, 42(2): 205-245.
- Camacho, N., B. Donkers and S. Stremersch (2011) Predictably non-bayesian: quantifying sailence effects in physician learning about drug quality. *Marketing Science*, 30(2): 305-320.
- Chan, T. and B. Hamilton (2006) Learning, Private Information, and the Economic Evaluation of Randomized Experiments. *Journal of Political Economy*, 114(6): 997-1040.
- Chan T., C. Narasimhan and Y. Xie (2012) Treatment Effectiveness and Side-effects: A Model of Physician Learning. Working paper, Washington University in St. Louis.
- Che, H. T. Erdem and S. Öncü (2011) Periodic Consumer Learning and Evolution of Consumer Brand Preferences. Working paper, New York University.
- Chen, X., Y. Chen, C. Weinberg (2009) Learning about movies: The impact of movie release types on the nationwide box office. Working paper, University of British Columbia.
- Chernew, M., G. Gowrisankaran and D. Scanlon (2008) Learning and the value of information: Evidence from health plan report cards. *Journal of Econometrics*, 144(1): 156-174.
- Ching, A. (2010a) Consumer learning and heterogeneity: dynamics of demand for prescription drugs after patent expiration. *International Journal of Industrial Organization*, 28(6): 619-638.
- Ching, A. (2010b) A dynamic oligopoly structural model for the prescription drug market after patent expiration. *International Economic Review*, 51(4): 1175-1207.
- Ching, A. and M. Ishihara (2010) The effects of detailing on prescribing decisions under quality uncertainty. *Quantitative Marketing and Economics*, 8(2): 123-165.
- Ching, A. and M. Ishihara (2012a) Measuring the informative and persuasive roles of detailing on prescribing decisions. *Management Science*, 58(7):1374-1387.
- Ching, A. and M. Ishihara (2012b) Identification of dynamic models of rewards program. Working paper, Rotman School of Management, University of Toronto.
- Ching, A., S. Imai, M. Ishihara and N. Jain (2012) A Practitioner's Guide to Bayesian Estimation of Discrete Choice Dynamic Programming Models. *Quantitative Marketing and Economics*,

10(2): 151-196.

Ching, A., T. Erdem and M. Keane (2009) The Price Consideration Model of Brand Choice. *Journal of Applied Econometrics*, 24(3): 393-420.

Ching, A., T. Erdem and M. Keane (2011) Learning Models: An Assessment of Progress, Challenges and New Developments. Working paper, Rotman School of Management, University of Toronto. Available at SSRN: <http://ssrn.com/abstract=1910839>.

Ching, A., R. Clark, I. Horstmann and H. Lim (2011) The Effects of Publicity on Demand: The Case of Anti-cholesterol Drugs. Working paper, Rotman School of Management, University of Toronto. Available at SSRN: <http://ssrn.com/abstract=1782055>

Chintagunta, Jing and Jin (2009) Information, Learning, and Drug Diffusion: The Case of Cox-2 Inhibitors. *Quantitative Marketing and Economics*, vol.7(4): 399-443.

Chintagunta, P., R. Goettler and M. Kim (2012) New Drug Diffusion when Forward-looking Physicians Learn from Patient Feedback and Detailing. Forthcoming in *Journal of Marketing Research*.

Chung, Doug, Thomas Steenburgh and K. Sudhir (2011) Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans. Working Paper, Yale School of Management.

Coscelli, A. and Shum M. (2004) An empirical model of learning and patient spillover in new drug entry. *Journal of Econometrics*, 122(2): 213-246.

Crawford, G. and M. Shum (2005). Uncertainty and learning in pharmaceutical demand. *Econometrica*, 73(4): 1137-1173.

Dickstein, M. (2012) Efficient provision of experience goods: evidence from antidepressant choice. Working paper, Stanford University.

Dubé, J.P., G. Hitsch and P. Rossi (2010) State Dependence and Alternative Explanations for Consumer Inertia. *RAND Journal of Economics*, 41(1): 417-445.

Eckstein, Z., D. Horsky, Y. Raban (1988) An empirical dynamic model of brand choice. Working paper 88, University of Rochester.

Erdem, T. (1998) An empirical analysis of umbrella branding. *Journal of Marketing Research*, 35(3): 339-351.

Erdem, T., M. Katz and B. Sun (2010) A Simple Test for Distinguishing between Internal Reference Price Theories. *Quantitative Marketing and Economics*, 8(3): 303-332.

- Erdem, T., M. Keane (1996) Decision-making under uncertainty: capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1): 1-20.
- Erdem, T. and J. Swait (1998) Brand Equity as a Signaling Phenomenon. *Journal of Consumer Psychology*, 7(2): 131-157.
- Erdem, T., S. Imai, M. Keane (2003) Brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and Economics*, 1(1): 5-64.
- Erdem, T. M. Keane, S. Öncü and J. Strebel (2005) Learning About Computers: An Analysis of Information Search and Technology Choice. *Quantitative Marketing and Economics*, 3(3): 207-246.
- Erdem, T., M. Keane and B. Sun (2008) A Dynamic Model of Brand Choice when Price and Advertising Signal Product Quality. *Marketing Science*, 27(6): 1111-25.
- Fang, H. and Y. Wang (2010) Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions. NBER working paper no. 16438.
- Ferreira, M.M. and G. Kosenok (2011) Learning about New Products: An Empirical Study of Physicians' Behavior. *Economic Inquiry*, 49(3): 876-898.
- Geweke, J. and M. Keane (2001) Computationally Intensive Methods for Integration in Econometrics. *Handbook of Econometrics*: Vol. 5, Heckman and Leamer (eds.), Elsevier Science B.V., 3463-3568.
- Goettler, R. and K. Clay (2011) Tariff Choice with Consumer Learning and Switching Costs. *Journal of Marketing Research*, 48(4): 633-652.
- Grubb, M. and M. Osborne (2012) Cellular service demand: Tariff choice, usage uncertainty, biased beliefs, learning and bill shock. Working paper, MIT Sloan School of Management.
- Gu, Z. (Jane) and S. Yang (2011) Consumer attribute-based learning and retailer category management strategies. Working paper, University of Albany.
- Guadagni, P. and J. Little (1983) A Logit Model of Brand Choice Calibrated on Scanner Data. *Marketing Science*, 2(3): 203-238.
- Heckman, J. (1997) Instrumental Variables: A Study of Implicit Behavioral Assumptions Used in Making Program Evaluations. *Journal of Human Resources*, 32(3): 441-462.
- Hitsch, G. (2006) An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty. *Marketing Science*, 25(1): 25-50.

- Imai, S., N. Jain and A. Ching (2009) Bayesian Estimation of Dynamic Discrete Choice Models. *Econometrica*, 77(6): 1865-1899.
- Israel, M. (2005) Services as experience goods: an empirical examination of consumer learning in automobile insurance. *American Economic Review*, 95(5):1444-1463.
- Iyengar, R., A. Ansari and S. Gupta (2007) A model of consumer learning for service quality and usage. *Journal of Marketing Research*, 44(4): 529-544.
- Jovanovic, B. (1979) Job Matching and the Theory of Turnover. *Journal of Political Economy*, 87(5): 972-990.
- Kalra, A., S. Li and W. Zhang. (2011) Understanding Responses to Contradictory Information About Products. *Marketing Science*, 30(6) 1098-1114.
- Keane, M., P. Todd and K. Wolpin (2011) The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications. *Handbook of Labor Economics*, Volume 4A, O. Ashenfelter and D. Card (eds.), 331-461.
- Keane, M. (1993) Simulation Estimation for Panel Data Models with Limited Dependent Variables. *The Handbook of Statistics*, G.S.Maddala, C.R. Rao and H.D. Vinod (eds), North Holland publisher, 545-571.
- Keane, M. (1994) A Computationally Practical Simulation Estimator for Panel Data. *Econometrica*, 62(1): 95-116.
- Keane, M. and K. Wolpin (1994) The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation: Monte Carlo Evidence. *Review of Economics and Statistics*, 76(4): 648-72.
- Keane, M. (1992) A Note on Identification in the Multinomial Probit Model. *Journal of Business and Economic Statistics*, 10(2): 193-200.
- Keane, M. (2010a) Structural vs. atheoretic approaches to econometrics. *Journal of Econometrics*, 156(1): 3-20.
- Keane, M. (2010b). A Structural Perspective on the Experimentalist School. *Journal of Economic Perspectives*, 24(2): 47-58.
- Keller, K. (2002) Branding and Brand Equity. *Handbook of Marketing*, B. Weitz and R. Wensley (eds.), Sage Publications, London, 151-178.
- Kihlstrom R. and M. Riordan (1984), "Advertising as a Signal," *Journal of Political Economy*, 92(3): 427-450.

- Knight, B and N. Schiff (2010) Momentum and Social Learning in Presidential Primaries. *Journal of Political Economy*, 118(6):1110-1150.
- Lai, T. L. (1987) Adaptive treatment allocation and the multi-armed bandit problem. *The Annals of Statistics*, 15(3): 1091-1114.
- Lim, H. and A. Ching (2012) A Structural Analysis of Promotional mix, Publicity and Correlated Learning: The Case of Statins. Working paper, University of Toronto.
- Magnac, T. and D. Thesmar (2002) Identifying Dynamic Discrete Decision Processes. *Econometrica*, 20(2): 801-816.
- Marcoul, P. and Q. Weninger (2008) Search and active learning with correlated information: empirical evidence from mid-Atlantic clam fisherman. *Journal of Economic Dynamics and Control*, 32(6): 1921-1948.
- Maskin, E. and J. Tirole (1988) A Theory of Dynamic Oligopoly, I and II. *Econometrica*, 56(3): 549-570.
- McCulloch, R. N. Polson and P. Rossi (2000) A Bayesian analysis of the multinomial probit model with fully identified parameters. *Journal of Econometrics*, 99(1): 173-193.
- Mehta N., S. Rajiv and K. Srinivasan (2004) The Role of Forgetting in Memory-based Choice Decisions. *Quantitative Marketing and Economics*, 2(2): 107-140.
- Mehta, N., X. Chen and O. Narasimhan (2008) Informing, transforming, and persuading: disentangling the multiple effects of advertising on brand choice decisions. *Marketing Science*, 27(3): 334-355.
- Miller, R. (1984) Job matching and occupational choice. *Journal of Political Economy*, 92(6): 1086-1120.
- Mira, P. (2007) Uncertain infant mortality, learning, and life-cycle fertility. *International Economic Review*, 48(3): 809-846.
- Miravete, E. (2003) Choosing the wrong calling plan? Ignorance and learning. *American Economic Review*, 93(1): 297-310.
- Narayanan, S., P. Chintagunta and E. Miravete (2007) The role of self selection, usage uncertainty and learning in the demand for local telephone service. *Quantitative Marketing and Economics*, 5(1): 1-34.
- Narayanan, S. and P. Manchanda (2009) Heterogeneous learning and the targeting of marketing communication for new products. *Marketing Science*, 28(3): 424-441.

- Narayanan, S., P. Manchanda and P. Chintagunta (2005) Temporal differences in the role of marketing communication in new product categories. *Journal of Marketing Research*, 42(3): 278-290.
- Norets, A. (2009) Inference in Dynamic Discrete Choice Models with Serially Correlated Unobserved State Variables. *Econometrica*, 77(5): 1665-1682.
- Osborne, M. (2011) Consumer learning, switching costs, and heterogeneity: A structural Examination. *Quantitative Marketing and Economics*, 9(1): 25-46.
- Pastorino, E. (2012) Careers in Firms: Estimating a Model of Learning, Job Assignment and Human Capital Acquisition. Research Department Staff Report 469, Federal Reserve Bank of Minneapolis.
- Roberts, J. and G. Urban (1988) Modeling Multiattribute Utility, Risk, and Belief Dynamics for New Consumer Durable Brand Choice. *Management Science*, 34 (2): 167-185.
- Rust, J. (2010), "Comments on: "Structural vs. atheoretic approaches to econometrics" by Michael Keane," *Journal of Econometrics*, 156(1), 21-24.
- Shin, S., S. Misra and D. Horsky (2012) Disentangling preferences and learning in brand choice models. *Marketing Science*, 31(1): 115-137.
- Sridhar, K., R. Bezawada and M. Trivedi (2012) Investigating the Drivers of Consumer Cross-Category Learning for New Products Using Multiple Data Sets. *Marketing Science*, 31(4): 668-688.
- Watanabe, Y. (2010) Learning and bargaining in dispute resolution: theory and evidence for medical malpractice litigation. Working paper, Northwestern University.
- Xiao, M. (2010) Is quality accreditation effective? Evidence from the childcare market. *International Journal of Industrial Organization*, 28(6): 708-721.
- Yang, S., Y. Zhao, T. Erdem and D. Koh (2010) Modeling Consumer Choice with Dyadic Learning and Information Sharing: An Intra-household Analysis. Working Paper, Stern School of Business, New York University.
- Zhang, J. (2010) The sound of silence: observational learning in the U.S. kidney market. *Marketing Science*, 29(2): 315-335.
- Zhao, Y., Y. Zhao and K. Helsen (2011), "Consumer learning in a turbulent Market environment: Modeling Consumer Choice Dynamics after a Product-harm Crisis," *Journal of Marketing Research*, 48(2): 255-267.

**Online Appendices to Accompany:
Learning Models: An Assessment of Progress,
Challenges and New Developments***

by

Andrew T. Ching
Rotman School of Management
University of Toronto
Andrew.Ching@rotman.utoronto.ca

Tülin Erdem
Stern School of Business
New York University
terdem@stern.nyu.edu

Michael P. Keane
University of New South Wales, and
Arizona State University
m.keane@unsw.edu.au

This draft: August 25, 2012

* Ching's work on this project is supported by SSHRC. Keane's work on this project was supported by Australian Research Council grant FF0561843.

Appendix A: Details of Eckstein, Horsky, Raban (1988) and Roberts and Urban (1988)

A-1. Eckstein, Horsky and Raban (1988)

The basic setup of the Eckstein, Horsky and Raban (1988) model is that a consumer has full information about the attributes of an existing brand. In each period, he/she decides whether to buy the familiar brand, or a newly introduced brand whose attributes are uncertain. For simplicity, assume quality is the only attribute. Let Q_j denote the quality of brand j , where in this case $j=o,n$ for old and new.

Of course, the key feature of learning models is that consumers do not know the attributes of brands with certainty. Before receiving any information via use experience, assume consumers have a normal prior on brand quality:

$$(A1) \quad Q_j \sim N(Q_{j0}, \sigma_{j0}^2)$$

This says consumers perceive that the true quality of the new brand (Q_n) is distributed normally with a mean of Q_{n0} and a variance of σ_{n0}^2 . The values of Q_{n0} and σ_{n0}^2 may be influenced by many factors, such as reputation of the manufacturer, pre-launch advertising, etc. For the familiar brand, enough information has already been received so its quality (Q_o) is known with certainty.

Use experience does not fully reveal quality because of “inherent product variability.” This has multiple interpretations. First, the quality of different units of a product may vary. Second, a consumer’s experience of a product may vary across use occasions. For instance, a cleaning product may be effective at removing the type of stains one faces on most occasions, but be ineffective on other occasions. Alternatively, there may be inherent randomness in psychophysical perception. E.g., the same cereal tastes better to me on some days than others.

Given inherent product variability, there is a distinction between “experienced quality” for brand j on purchase occasion t , which we denote Q_{jt}^E , and true quality Q_j . Let us assume the “experienced quality” delivered by use experience is a noisy signal of true quality, as in:

$$(A2) \quad Q_{jt}^E = Q_j + \varepsilon_{jt} \quad \text{where} \quad \varepsilon_{jt} \sim N(0, \sigma_\varepsilon^2)$$

Here σ_ε^2 is the variance of inherent product variability, which we often refer to as “experience variability.” Both new and familiar brands have experience variability, so (A2) holds for $j=o,n$.

Note that we have conjugate priors and signals, as both the prior on quality (A1) and the noise in the quality signals (A2) are assumed to be normal. Thus, the posterior for perceived quality, given a single use experience signal, is given by the simple updating formulas:

$$(A3) \quad Q_{j1} = \frac{\sigma_{j0}^2}{\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j1}^E + \frac{\sigma_\varepsilon^2}{\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j0}$$

$$(A4) \quad \sigma_{j1}^2 = \frac{1}{(1/\sigma_{j0}^2) + (1/\sigma_\varepsilon^2)}$$

Equation (A3) describes how a consumer's prior on quality of brand j is updated as a result of the experience signal Q_{j1}^E . Note that the extent of updating is greater the more accurate is the signal (i.e., the smaller is σ_ε^2). Equation (A4) describes how a consumer's uncertainty declines as he/she receives more signals. σ_{j1}^2 is often referred to as the "perception error variance."

Equations (A3) and (A4) generalize to any number of signals. Let $N_j(t)$ denote the total number of use experience signals received through time t . Then we have that:

$$(A5) \quad Q_{jt} = \frac{\sigma_{j0}^2}{N_j(t)\sigma_{j0}^2 + \sigma_\varepsilon^2} \sum_{s=1}^t Q_{js}^E d_{js} + \frac{\sigma_\varepsilon^2}{N_j(t)\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j0}$$

$$(A6) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{j0}^2) + N_j(t)(1/\sigma_\varepsilon^2)}$$

where d_{jt} is an indicator for whether brand j is bought/consumed at time t .

In equation (A5), the perceived quality of brand j at time t , Q_{jt} , is a weighted average of the prior and all quality signals received up through time t , $\sum_{s=1}^t Q_{js}^E d_{js}$. Crucially, this is a random variable across consumers, as some will, by chance, receive better quality signals than others. Thus, the learning model endogenously generates heterogeneity across consumers in perceived quality of products (even starting from identical priors). This aspect of the model is appealing. It seems unlikely that people are born with brand preferences (as standard models of heterogeneity implicitly assume), but rather that they arrive at their views through heterogeneous experience.

Of course, as equation (A6) indicates, the variance of perceived quality around true quality declines as more signals are received, and in the limit perceived quality converges to true quality. Still, heterogeneity in perceptions will persist over time, because: (i) both brands and consumers are finitely lived, (ii) there is a flow of new brands and new consumers entering a market, and (iii) as people gather more information the value of trial purchases diminishes, and so eventually learning about unfamiliar products will become slow.

In general, learning models must be solved by dynamic programming, because today's purchase affects tomorrow's information set, which affects future utility. Thus, choosing the

brand with the highest *perceived* quality in the *current* period is not necessarily optimal. Take the case of a new vs. familiar brand, where the information set is $I_t = \{Q_{nt}, \sigma_{nt}^2\}$ and prices are given by P_{jt} for $j=o,n$. Then the values (denoted V) of choosing each option in the current period are:

$$(A7) \quad V(n, t|I_t) = E[U(Q_{nt}^E, P_{nt})|I_t] + \beta EV(I_{t+1}) \quad \text{where} \quad I_{t+1} = \{Q_{nt+1}, \sigma_{nt+1}^2\}$$

$$(A8) \quad V(o, t|I_t) = E[U(Q_{ot}^E, P_{ot})] + \beta EV(I_t) \quad \text{as} \quad I_{t+1} = I_t$$

Purchase of the familiar brand gives expected utility $E[U(Q_{ot}^E, P_{ot})]$. This is increasing in true quality Q_o , which is known. It is decreasing in experience variability σ_ε^2 , assuming consumers are risk averse with respect to product variability. Purchase of the new brand delivers expected utility $E[U(Q_{nt}^E, P_{nt})]$, which is increasing in Q_{nt} and decreasing in both experience variability and the perception error variance σ_{nt}^2 . Purchase of the new brand also increases next period's expected value function. That is, $EV(I_{t+1}) > EV(I_t)$ because I_{t+1} contains better information.

Eckstein et al (1988) assume that utility is linear in experienced quality Q_{jt}^E , thus abstracting from risk aversion, and also linear in price. In that case (A7) and (A8) simplify to:

$$(A9) \quad V(n, t|I_t) = Q_{nt} - P_{nt} + e_{nt} + \beta EV(I_{t+1}) \quad \text{where} \quad I_{t+1} = \{Q_{nt+1}, \sigma_{nt+1}^2\}$$

$$(A10) \quad V(o, t|I_t) = Q_o - P_{ot} + e_{ot} + \beta EV(I_t) \quad \text{where} \quad I_{t+1} = I_t$$

Here the e_{jt} for $j=o,n$ are stochastic terms in the utility function that represent purely idiosyncratic tastes for the two brands. These play the same role as the brand specific stochastic terms in traditional discrete choice models like logit and probit.¹ Eckstein et al (1988) assume they are normally distributed in order to obtain probit-like brand choice probabilities, and without loss of generality (as only utility differences matter in discrete choice models) they set $e_{nt} = 0$.

Now, a consumer will choose the new brand over the familiar brand if the value function in equation (A9) exceeds that in (A10). This means that $V_{nt}^* > 0$, where:

$$(A11) \quad V_{nt}^* \equiv V(n, t|I_t) - V(o, t|I_t) = (Q_{nt} - P_{nt}) - (Q_o - P_{ot}) + (e_{nt} - e_{ot}) + G_t$$

¹ Without these terms, choice would be deterministic (conditional on Q_{nt} , Q_o , P_{nt} and P_{ot}). But in contrast to standard discrete choice models, it is not strictly necessary to introduce the $\{e_{jt}\}$ terms to generate choice probabilities. This is because perceived quality Q_{nt} is random from the perspective of the econometrician. However, we feel it is advisable to include the $\{e_{jt}\}$ terms regardless. This is because, in their absence, choice becomes deterministic conditional on price as experience with the new brand grows large. Yet, both introspection and simple data analysis suggest consumers do switch brands for purely idiosyncratic reasons even under full information.

$$(A12) \quad G_t \equiv \beta[EV(I_{t+1}) - EV(I_t)]$$

Eckstein et al (1988) call G_t the “information gain,” although we prefer to call it the “gain from trial.” It is the increase in expected present value of utility from $t+1$ until the terminal period T arising from the enhanced information the consumer obtains by trying the new brand at time t .

The gain from trial comes from two sources. Most obviously, the consumer may learn that the new brand is better than the old brand. More subtly, suppose the evidence indicates the new brand is inferior to the familiar brand. There is nevertheless a price differential large enough that the consumer would prefer the new brand. More precise information about the quality of the new brand enables the consumer to set this price differential more accurately.

Note that $G_t > 0$, i.e., more information is always better. However, G_t is smaller if (i) the consumer has more information (σ_{nt}^2 smaller) or (ii) use experience signals are less accurate (σ_ε^2 larger). Both lower the value of trial. It is notable that (A11) can be rewritten “Choose brand n if:”

$$(A13) \quad (Q_{nt} - P_{nt}) + G_t + e_{nt} > (Q_o - P_{ot}) + e_{ot}$$

This shows that the trial value G_t *augments* the perceived value of the new brand ($Q_{nt} - P_{nt}$). Thus, *ceteris paribus*, the new brand can command a price premium over the old brand because it delivers valuable trial information. The Eckstein et al (1988) model generates a “value of information” effect that is *opposite* to the conventional brand loyalty phenomenon.

Gittins and Jones (1974) noted that, under certain assumptions, one can construct G_t without having to solve the full DP problem. Eckstein et al (1988) use fairly simple closed form (approximate) solutions for G_t that were developed by Whittle (1983), Bather (1983) and Gittins (1983). We emphasize that the special case of linear utility is important here.² Then, given a normality assumption on the idiosyncratic errors, equation (A11) implies simple binomial probit probabilities for the choice probabilities *conditional* on Q_{nt} .

However, Q_{nt} is not observed by the econometrician. And, as we see in (A5), it is serially correlated, creating dependence in choice probabilities across periods. This makes estimation difficult because we have a panel probit model with serially correlated errors (see Keane (1994)).

Say we have household level data on purchases and prices of both brands for periods from $t=0, \dots, T$. And let the periods be long enough that we can ignore no-purchase. Then the

² The method used by Eckstein et. al. (1988), which is from Gittins and Jones (1974), is much more difficult to implement without linear utility. Furthermore, it cannot allow for exogenously arriving signals (such as advertising or word-of-mouth). It can only accommodate learning by trial.

likelihood *conditional* on a hypothetical $\{Q_{nt}\}$, $t=0, \dots, T$ sequence takes the form:

$$(A14) \quad P(\{d_{nt}\}_{t=0}^T | \{Q_{nt}\}_{t=0}^T) = \prod_{t=0}^T [d_{nt} P(V_{nt}^*(Q_{nt}) > 0 | D_t) + d_{ot} P(V_{nt}^*(Q_{nt}) < 0 | D_t)]$$

where D_t represents the history of past purchases up through t . To obtain the *unconditional* likelihood we must integrate out the $\{Q_{nt}\}_{t=0}^T$ history, as in:

$$(A15) \quad P(\{d_{nt}\}_{t=0}^T) = \int_{\{Q_{nt}\}_{t=0}^T} P(\{d_{nt}\}_{t=0}^T | \{Q_{nt}\}_{t=0}^T)$$

Note that Q_{nt} stays fixed except for time periods when the new brand is purchased. So the order of integration in (A15) is equal to $\sum_{t=1}^T d_{nt}$, the number of times the consumer buys the new brand and receives a trial signal.³ Thus, using 1988 technology (i.e., without simulation techniques), estimation is not feasible unless the number of purchases of the new brand is quite small.⁴

Eckstein et al (1988) estimate their model using data from NPD Research on a new cereal called Corn Bran introduced by Quaker Mills in January 1979. The data follow 100 households over 20 months after the brand introduction. All existing brands are grouped together as an “old” brand. This grouping is facilitated by the fact that the authors actually ignore price information. They view the entire quantity $f_{nt} \equiv (Q_{nt} - P_{nt})$ as the “premium” of the new brand and assume it is f_{nt} that consumers learn about. A similar assumption is made for the “old” brand(s). Thus, price fluctuations around mean price form part of the normally distributed error term (e_{ot}).

The parameters of the model are the ratio of the signal to the prior variance (they are not identified separately here), the discount rate, and the prior mean, which depends on household demographics (income, education, family size). The variance ratio is greater than one, implying fairly tight priors, but the hypothesis of no learning is strongly rejected. The monthly discount factor is only 0.02%, implying a high degree of forward-looking behavior. Demographics are not significant for the prior mean, which is negative. Eckstein et al (1988) assume that true quality and the prior mean on quality are equal.⁵ The subsequent learning literature has not followed this assumption, as it rules out the possibility that a brand may be better or worse than initially

³ This contrasts with the general case of a panel probit with correlated errors, in which the order of integration is T.

⁴ It is important to stress the difference in the econometric treatment of Q_{jt} and σ_{jt}^2 . The fact that a consumer’s perceived quality of brand j at time t , Q_{jt} , is a random variable from the point of view of the econometrician means we must integrate it out of the likelihood function. In contrast, the perception variance σ_{jt}^2 is a deterministic function of the number of past purchases, so it can be treated as known (given the model parameters). Thus, the evolving Q_{jt} are a source of difficulty in forming the likelihood, while the evolving σ_{jt}^2 cause no special problems.

⁵ They argue that: “It is natural to assume that the prior distribution of the brand premium is equal to the true distribution of the quality of the brand as it is perceived by the population of consumers.”

expected, and thus imposes strong restrictions on the evolution of market shares. Still, the model predicts a long run market share of 4.5% for the new brand compared to 4.9% in the data, and does a reasonably good job of matching the time path of market share in this case.

A-2. Roberts and Urban (1988)

Roberts and Urban (1988) use experimental data to calibrate a Bayesian learning model. The goal is to predict market share of a new durable good (i.e., a car). In considering a durable, as opposed to a frequently purchased good, trial purchase is no longer relevant. Instead, the relevant information sources are advertising, word-of-mouth, dealer visits, etc. For simplicity we will often refer to these as “exogenous” signals, as we may think of them as arriving randomly from the outside environment. (Of course, a consumer may actively seek out such signals, an extension we discuss below).

Let A_{jt} to denote an exogenous signal (e.g., advertising, word-of-mouth) that a consumer receives about brand j at time t . We further assume that:

$$(A16) \quad A_{jt} = Q_j + \varepsilon_{jt}^A \quad \text{where} \quad \varepsilon_{jt}^A \sim N(0, \sigma_A^2)$$

This says the signals A_{jt} provide unbiased but noisy information about brand quality, where the noise has variance σ_A^2 . The noise is assumed normal, to maintain conjugacy with the prior in (A1).

It is important to compare (A16) with (A2). The noise in trial experience is from inherent product variability, which is largely a feature of the product itself. The noise in a signal like advertising or word-of-mouth is, in contrast, largely a function of the medium. Presumably some media convey information more accurately than others, and no medium is as accurate as direct use experience. We also stress that the noise in (A16) differs fundamentally from that in (A2), as inherent product variability affects a consumer’s *experienced* utility from the product, while the noise in exogenous signals does not. Nevertheless, both types of signal enter the learning process of a consumer in the same way. Given the exogenous signal A_{jt} , we can rewrite (A5)-(A6) as:

$$(A17) \quad Q_{jt} = \frac{\sigma_{j0}^2}{N_j^A(t)\sigma_{j0}^2 + \sigma_A^2} \sum_{s=1}^t A_{js} d_{js}^A + \frac{\sigma_A^2}{N_j^A\sigma_{j0}^2 + \sigma_A^2} Q_{j0}$$

$$(A18) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{j0}^2) + N_j^A(1/\sigma_A^2)}$$

where d_{jt}^A is an indicator for whether a signal for brand j is received at time t , and $N_j^A(t)$ is the total number of signals received for brand j up through time t .

Roberts and Urban (1988) studied the introduction of a new car by General Motors. They interviewed 336 consumers in 1983, and attempted to measure actual perceptions of the new car. First, they used a rating scale to measure Q_{j0} for each consumer. Then, they exposed consumers to a video meant to simulate word-of-mouth exposure. They measured each consumer's post exposure perception of the quality of the car, Q_{j1} , as well as their assessment of the quality signal conveyed by the video, A_{j1} . From (A17) we see that, with these three pieces of information, we can back out σ_A^2/σ_{j0}^2 , the ratio of the signal to prior variance. The larger the ratio, the more weight a consumer puts on his/her prior, and the less he/she updates perceived quality on seeing the video. Next, by taking a measure of the perception error variance post-video, one can use (A18), along with the σ_A^2/σ_{j0}^2 estimate, to back out σ_A^2 and σ_{j0}^2 separately. This is done for each consumer.

Next, Roberts and Urban (1988) map perceived quality and quality uncertainty into choice probabilities. They assume utility is negative exponential over car quality Q , to introduce risk aversion with respect to quality variation. There are a few ways one might proceed from here, of which we consider two: First, one could assume all goods other than cars are grouped into a single composite commodity (or "outside good"), with consumption level $C=I-P_j$, where I is income and P_j is the car price. Uncertainty about the outside good is not of interest, so assume that utility is quasi-linear. Then, utility conditional on purchase of car j of quality level Q_j is:

$$(A19) \quad U(Q_j, P_j) = -\exp(-rQ_j) + \alpha_P(I - P_j) + e_j$$

Here $r > 0$ is the risk aversion parameter. Parameter α_P may be interpreted as the marginal utility of consumption of the outside good (since utility is quasi-linear, α_P is a constant). Finally, e_j is a purely idiosyncratic taste for the car, introduced to make choice probabilistic (as in (A9)-(A10)).

As the quality of the new car is uncertain, the consumer makes decisions based on expected utility. From (A19) we have:

$$(A20) \quad E[U(Q_j^E, P_j)|I_t] = -\exp\left(-r\left(Q_{jt} - \frac{r}{2}(\sigma_{jt}^2 + \sigma_\varepsilon^2)\right)\right) + \alpha_P(I - P_j) + e_j$$

Notice that here – in contrast to Eckstein et al (1988), equations (A9)-(A10) – both the perception error variance and experience variability appear in current expected utility. And, assuming agents are risk averse with respect to quality variation, they both reduce current expected utility.

The second approach, which Roberts and Urban (1988) adopt, is to “linearly discount preference by price,” as in Hauser and Urban (1986). This means replacing (A19) and (A20) by:

$$(A21) \quad U(Q_j, P_j) = -\exp\left(-r(Q_j - \alpha_P P_j + e_j)\right)$$

$$(A22) \quad E[U(Q_j^E, P_j)|I_t] = -\exp\left(-r\left(Q_{jt} - \frac{r}{2}(\sigma_{jt}^2 + \sigma_\varepsilon^2) - \alpha_P P_j + e_j\right)\right)$$

Note that for $r > 0$ expected utility in (A22) is monotonically increasing in the quantity:

$$(A23) \quad \chi_j = Q_{jt} - \frac{r}{2}(\sigma_{jt}^2 + \sigma_\varepsilon^2) - \alpha_P P_j + e_j$$

Roberts and Urban (1988) call this the “risk adjusted preference” for car j , and choice behavior can be represented as choosing the car that maximizes χ_j from the set $j=1, \dots, J$. If the error term e_j is extreme value with scale parameter $(1/\beta)$, the probability of purchase of new car j is:

$$(A24) \quad P(j|I_t) = \exp\left(\beta\left(Q_{jt} - \frac{r}{2}(\sigma_{jt}^2 + \sigma_\varepsilon^2) - \alpha_P P_j\right)\right) / \sum_{j=1}^J \exp\left(\beta\left(Q_{jt} - \frac{r}{2}(\sigma_{jt}^2 + \sigma_\varepsilon^2) - \alpha_P P_j\right)\right)$$

Next, Roberts and Urban obtain purchase probabilities on an 11 point scale, for both the new car and an existing reference car. Letting o denote the reference car, we manipulate (24) to obtain:

$$(A25) \quad \ln P(j|I_t) / \ln P(o|I_t) = \beta(Q_{jt} - Q_{ot}) - \frac{1}{2}\beta r(\sigma_{jt}^2 - \sigma_{ot}^2) - \beta\alpha_P(P_j - P_o)$$

The estimates of the three coefficients in this equation, $(\beta, -\frac{1}{2}\beta r, -\beta\alpha_P)$, were 2.6, -2.3 and -3.6 respectively, implying that consumers are risk averse with respect to quality uncertainty (the implied estimate of r is 1.77).

In the final step, Roberts and Urban use a previous new car introduction to calibrate the diffusion of work-of-mouth signals over time as a function of market share. Then, using an assumption about true quality of the new car (informed by focus groups, etc.), they use (A17)-(A18) and (A25) to predict how the demand for the car will evolve over time. Unfortunately, they could not validate their market share predictions due to supply problems that occurred after launch.⁶

As we have seen, Roberts and Urban (1988) study how Bayesian consumers learn about a new product from word-of-mouth signals. Consumers in their model are risk averse, but myopic.

⁶ One interesting aspect of Roberts and Urban (1988) is that, if consumers are shown a negative video about the new car, their perceived quality goes down and perceived variance goes up. This contradicts the prediction of Bayesian learning models that additional information (positive or negative) should always reduce uncertainty.

Thus, their model exhibits “brand loyalty” due to risk aversion, but it does not generate a “value of trial information” (i.e., there is no active search). In contrast, Eckstein et al (1988) allow for forward-looking Bayesian consumers, but they assume risk neutrality. Thus, their model exhibits the “value of trial” phenomenon, but not “brand loyalty” – which requires risk aversion.

A-3. Erdem and Keane (1996) Interpreted as Unifying the Earlier Models

The paper by Erdem and Keane (1996) can be viewed as integrating the key features of Roberts and Urban (1988) and Eckstein et al (1988) into a single framework. That is, consumers in their model are both forward-looking and risk averse with respect to product uncertainty. Thus their model exhibits both “brand loyalty” (due to risk aversion) and a value of trial information.

Finally, Erdem and Keane (1996) differed from earlier work in that it was a true brand choice model. That is, the model did not focus only on the choice between a new product and all existing products, or on the evolution of market share for a single new product. Their model could handle a portfolio of products, about which consumers had differing degrees of familiarity.

Finally, as Erdem and Keane included both use experience and advertising signals as information sources, it was the first paper to capture learning about brand quality through both strategic sampling (or trial) and advertising (i.e., both active and passive search).⁷ It is simple to extend the Bayesian updating rules in (A5)-(A6) and (A17)-(A18) to accommodate two types signals:

$$(A26) \quad Q_{jt} = \frac{(1/\sigma_{\epsilon}^2)}{S_j(t)} \sum_{s=1}^t Q_{js}^E d_{js} + \frac{(1/\sigma_A^2)}{S_j(t)} \sum_{s=1}^t A_{js} d_{js}^A + \frac{(1/\sigma_{j0}^2)}{S_j(t)} Q_{j0}$$

$$(A27) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{j0}^2) + N_j(t)(1/\sigma_{\epsilon}^2) + N_j^A(t)(1/\sigma_A^2)}$$

where $S_j(t) \equiv 1/\sigma_{jt}^2$. Note that in (A26)-(A27) the timing of signals does not matter. Only the total stock of signals determines a consumer’s current state. Also, receiving N signals with variance σ^2 has the same impact on the perception variance as receiving one signal with variance σ^2/N .

Erdem and Keane (1996) assumed a utility function of the form:

$$(A28) \quad U(Q_{jt}^E, P_{jt}) = w_Q Q_{jt}^E - w_Q r (Q_{jt}^E)^2 + w_P (I - P_{jt}) + e_{jt}$$

⁷ Given that the Erdem and Keane (1996) model was designed with an application to frequently purchased goods in mind, it was natural to include both use experience and advertising as sources of product information.

Here utility is quadratic in the experienced quality of brand j at time t . The parameter w_Q is the weight on quality, r is the risk coefficient, w_P is the marginal utility of the outside good, and e_{jt} is an idiosyncratic brand and time specific error term. Note that, as choices only depend on utility differences, and as income, I , is the same regardless of which brand is chosen, income drops out of the model. So we can simply think of w_P as the price coefficient. Expected utility is given by:

$$(A29) \quad E[U(Q_{jt}^E, P_{jt})|I_t] = w_Q Q_{jt} - w_Q r Q_{jt}^2 - w_Q r (\sigma_{jt}^2 + \sigma_\varepsilon^2) - w_P P_{jt} + e_{jt}$$

Note that (A29) is similar to (A23) in Roberts and Urban (1988), except for the quadratic in quality.

Finally, as the Erdem-Keane model was meant to be applied to weekly data, and as consumers may not buy in every week, a utility of the no purchase option was also specified. This was written simply as $E[U_0|I_t] = \Phi_0 + \Phi_{0t}t + e_{0t}$. The time trend captures in a simple way the possibility of changing value of substitutes for the category in question.

Agents are forward-looking, making choices to maximize value functions of the form:

$$(A30) \quad V(j, t|I_t) = E[U(Q_{jt}^E, P_{jt})|I_t] + \beta EV(I_{t+1}|I_t, j) \quad \text{for } j=0, \dots, J$$

Equation (A30) looks similar to equations (A7)-(A8) in Eckstein et al (1988). But there are important differences in terms of how the information set evolves. Recall that in (A7)-(A8) there were only two brands, new and old, and consumers only learned about the new brand. Thus the information set was updated to $I_{t+1} = \{Q_{n,t+1}, \sigma_{n,t+1}^2\}$ if a consumer bought the new brand at t , while staying fixed at $\{Q_{nt}, \sigma_{nt}^2\}$ if not. Here there are two main differences: First, the consumer may have incomplete information about all J brands in the market (although, of course, he/she may have very accurate knowledge of a few favorite brands). Second, a consumer might learn about any brand through ad exposures, even if he/she does not purchase it, as ads arrive exogenously.

These two differences lead to several complications. Now the information set is:

$$(A31) \quad I_{t+1} = \{Q_{1,t+1}, \dots, Q_{J,t+1}, \sigma_{1,t+1}^2, \dots, \sigma_{J,t+1}^2\}$$

In forming $EV(I_{t+1}|I_t, j)$, we must: (i) update σ_{jt}^2 to $\sigma_{j,t+1}^2$ using (A27) to account for additional use experience, (ii) integrate over possible values of the use experience signal in (A2) to take the expectation over possible realizations of $Q_{j,t+1}$, (iii) integrate over possible ad exposures that may arrive between t and $t+1$ (i.e., over realizations of $d_{j,t+1}^A$ for $j=1, \dots, J$) to account for ad induced

changes in the $\{\sigma_{j,t+1}^2\}$, and (iv) integrate over possible values of the ad signals in (A16), as these will lead to different values of the $\{Q_{j,t+1}\}$.⁸ One can clearly see that if consumers learn about multiple brands, and have more than one source of information, it greatly complicates the nature of the integrals that must be evaluated to form the expected value functions in (A30).

Furthermore, the number of variables that characterize the state of an agent goes from two when $I_t = \{Q_{nt}, \sigma_{nt}^2\}$ as in Eckstein et al (1988) to $2 \cdot J$ in equation (A31). Solving a dynamic program requires that one solve the expected value function integrals at every point in the state space. This is clearly not feasible here. Of course, as the state variables in (A31) are continuous, it would be literally impossible to solve for the expected value functions at every state point (as the number of points is infinite). A common approach is to discretize continuous state variables using a fairly fine grid. But, as we have $2 \cdot J$ state variables, this gives $G^{2 \cdot J}$ grid points, which is impractically large even for modest G and J . This is known as the ‘‘curse of dimensionality.’’

To solve the optimization problem of the agents in their model (that is, to construct the $EV(I_{t+1}|I_t, j)$ in (A30)), Erdem and Keane (1996) used an approximate solution method developed in Keane and Wolpin (1994). We provides details on how this method works in Appendix B.

Assuming the idiosyncratic brand and time specific error terms e_{jt} in (A29) are iid extreme value, the choice probabilities in the Erdem and Keane (1996) model take the form:

$$(A32) \quad P(j|I_t) = \frac{\exp(w_Q Q_{jt} - w_Q r Q_{jt}^2 - w_Q r (\sigma_{jt}^2 + \sigma_\varepsilon^2) - w_P P_{jt} + \beta EV(I_{t+1}|I_t, j))}{\exp(\Phi_0 + \Phi_0 t + \beta EV(I_{t+1}|I_t, 0)) + \sum_{k=1}^J \exp(w_Q Q_{kt} - w_Q r Q_{kt}^2 - w_Q r (\sigma_{kt}^2 + \sigma_\varepsilon^2) - w_P P_{kt} + \beta EV(I_{t+1}|I_t, k))}$$

Aside from the slight difference in the utility function, this equation bears a striking resemblance to equation (A24) from Roberts and Urban (1988). That is, choice probabilities in the dynamic learning model look just like those from a static model, *except* for the additional $\beta EV(I_{t+1}|I_t, j)$ terms that capture the gain from trial information. These are greater for less familiar brands, where the gain from trial is greater. At the same time, the risk terms σ_{jt}^2 are also greater for such brands. These two forces work against each other, and which dominates determines whether a consumer is more or less likely to try a new unfamiliar brand vs. a familiar brand. In categories that exhibit a high degree of brand loyalty (i.e., persistence in choice behavior) we would expect

⁸ It is also necessary to integrate over future price realizations for all brands. To make this integration as simple as possible, Erdem and Keane (1996) assumed that the price of each brand is distributed normally around a brand specific mean, with no serial correlation other than that induced by these mean differences.

the risk aversion effect to dominate. In categories that exhibit a high degree of brand switching (due to experimentation and variety seeking) we would expect the gains from trial to dominate.

We give details on how to simulate the likelihood function in the Erdem and Keane (1996) model in the main text.

Appendix B: Details of the Keane-Wolpin (1994) Algorithm Applied to Learning Models

In this appendix, we describe how to apply Keane-Wolpin (1994) to approximate the solution to the dynamic learning model described in section 2. Let us return to the terminal period problem. At T , the consumer will simply choose the option with highest expected utility. Recall equations (24) and (25) from the main text:

$$(24) \quad V(j, T | I_T) = E[U(Q_{jT}^E, P_{jT}) | I_T] \quad \text{for} \quad j=0, \dots, J$$

$$(25) \quad EV(I_T | I_{T-1}, j) = E_{T-1} \max \{ E[U_0 | I_T], E[U(Q_{1T}^E, P_{1T}) | I_T], \dots, E[U(Q_{jT}^E, P_{jT}) | I_T] \}$$

As we discussed earlier, because both use experience signals and ad signals may arrive between $T-1$ and T , causing the consumer to update his/her information set, the expectation in (25) must be taken over the possible I_T that may arise as a result of these signals. Specifically, we must: (i) update σ_{jt}^2 to $\sigma_{j,t+1}^2$ using (18) to account for additional use experience, (ii) integrate over possible values of the use experience signal in (2) to take the expectation over possible realizations of $Q_{j,t+1}$, (iii) integrate over possible ad exposures that may arrive between t and $t+1$ (i.e., over realizations of $d_{j,t+1}^A$ for $j=1, \dots, J$) to account for ad induced changes in the $\{\sigma_{j,t+1}^2\}$, and (iv) integrate over possible values of the ad signals in (14), as these will lead to different values of the $\{Q_{j,t+1}\}$.

For each state point in I_{T-1} , the expectation can be calculated using Monte Carlo integration. As we mentioned in section 2, the state variables in $I_{T-1} = \{Q_{1,T-1}, \dots, Q_{J,T-1}, \sigma_{1,T-1}^2, \dots, \sigma_{J,T-1}^2\}$ are continuous. Therefore, it is not possible for us to solve for (25) at all points in I_{T-1} in the computer. This is why Keane-Woplin approximation method is particularly useful in this setting. When applying their interpolation method, we only need to solve for (25) at $K_{T-1} < \infty$ randomly drawn state points.

How do we randomly draw K_{T-1} state points? If I_{T-1} is bounded, this is straightforward to do. One can draw state points from a multi-dimensional uniform distribution. However, it should be highlighted that the support for $Q_{j,T-1}$ is unbounded, for all j . To see this, note that according to (2) and (14), both $Q_{j,T-1}^E$ and A_{T-1} have unbounded support because they are just realizations of normal distributions. Therefore, it follows from (17) that $Q_{j,T-1}$ is unbounded for all j . This adds a complication when applying the interpolation method – no matter how many state points we randomly sample, we cannot sample all regions of an unbounded state space. As a result, it would be questionable how good this method can predict the E_{\max} in regions where

we do not obtain state points to estimate the interpolation regression. To get around this problem, it is important to note that for a given parameter vector, the distribution of the state variables should concentrate around the path that connect between the initial prior mean quality, Q_{j_0} , and the true mean quality, Q_j . Therefore, in practice, if we choose a range that include Q_{j_0} and Q_j in the “center,” we will be just cutting the tail of the distribution. We will discuss more details about how to pick the max and min values of the range later. For now, suppose we have already set the range for each unbounded state variable.

Let I_{T-1}^- denotes the restricted state space defined by the pre-set range for each state variable (in general, we will use I_t^- to denote the restricted state space at time t). We randomly draw K_{T-1} state points from I_{T-1}^- . Using Monte Carlo integration to approximate the expectation in (25), we obtain $\{\hat{E}V(I_T|I_{T-1,k}^-, j)\}_{k=1}^{K_{T-1}}$, where \hat{E} denotes that the expectation is evaluated based on a simulation method. We can now treat $\{\hat{E}V(I_T|I_{T-1,k}^-, j)\}_{k=1}^{K_{T-1}}$ as a vector of dependent variables, and the K_{T-1} vectors of state variables $\{I_{T-1,k}^-\}_{k=1}^{K_{T-1}}$ as independent variables in an interpolating regression:

$$(B1) \quad \hat{E}V(I_T|I_{T-1,k}^-, j) = g_{T-1}(I_{T-1,k}^-, j; \gamma_{T-1}) + v_{T-1,k},$$

where γ_{T-1} is a time $T-1$ vector of regression coefficients and $g_{T-1}(\cdot; \cdot)$ is a flexible function of state variables. The idea is to pick a parsimonious functional form that gives us very good fit (i.e., high R^2). Experience shows that highly parameterized functional forms (e.g., high order polynomials) can often perform poorly when one extrapolates to state points not included in the regression. Thus, it is important for a researcher to engage in considerable experimentation in order to find a functional form that is accurate yet at the same time parsimonious. Once a form is chosen, the estimates of γ_{T-1} can be quickly obtained based on OLS.

With this interpolating function in hand, estimates of $EV(I_T|I_{T-1}, j)$ can be obtained at any state point in the set I_{T-1}^- . In fact, even for state points that are outside I_{T-1}^- , we can use the interpolating regression to predict their *Emax* values. But the caveat is that the accuracy of prediction would likely deteriorate as we move farther away from I_{T-1}^- .

Given $\hat{E}V(I_T|I_{T-1}, j)$, we can now obtain the estimates of $V(j, T-1|I_{T-1})$ (see (21)). We can then apply Monte Carlo integration again to obtain $\{\hat{E}V(I_{T-1}|I_{T-2,k}^-, j)\}_{k=1}^{K_{T-2}}$ for K_{T-2} randomly drawn state space elements from I_{T-2}^- . Similarly, we use $\{\hat{E}V(I_{T-1}|I_{T-2,k}^-, j)\}_{k=1}^{K_{T-2}}$ as a vector of dependent variables, and $\{I_{T-2,k}^-\}_{k=1}^{K_{T-2}}$ as independent variables in estimating *another*

interpolating regression:

$$(B2) \quad \hat{E}V(I_{T-1}|I_{T-2,k}^-, j) = g_{T-2}(I_{T-2,k}^-, j; \gamma_{T-2}) + v_{T-2,k},$$

which can provide estimates for $EV(I_{T-1}|I_{T-2}, j)$ at any state point in I_{T-2} . It should be highlighted that the functional form for $g_{T-2}(\cdot, \cdot; \cdot)$ could be different from that for $g_{T-1}(\cdot, \cdot; \cdot)$, although they are usually very similar in practice. Even though we use the same functional forms for these two interpolating regressions, we need to run another OLS regression to obtain the estimates of γ_{T-2} . This is why we use a different subscript here. Continuing this procedure backward, we can obtain the interpolating functions for all t , i.e., $g_{T-1}, g_{T-1}, \dots, g_1$.

Now we return to discuss how to set the boundaries of $I_t^- \subset I_t$. Let's focus on how to set the max and min for the range of $Q_{j,t}$ first. Choosing a range for $\sigma_{j,t}^2$ is a simpler task and we will discuss it later. Just like any other numerical methods, setting bounds for $Q_{j,t}$ will inevitably introduce approximation errors. But a careful choice of bounds can minimize their impacts.⁹ Recall that theory tells us that the stochastic evolution of $Q_{j,t}$ will distribute around a path that connects Q_{j_0} and Q_j . Therefore, to reduce the approximation errors, we would like to set the bounds such that Q_j and Q_{j_0} are at the center of the pre-set range. By setting the range sufficiently large, we will essentially just truncate the tail of the distribution with very little density. Therefore, even though we are not sampling state points outside I_{t-1}^- when estimating the interpolation regression, which may result in poor out-of-sample predictions for $\hat{E}V(I_t|I_{t-1}^+, j)$, where $I_{t-1}^+ \equiv I_t \setminus I_{t-1}^-$, we do not expect they will have much impact on the estimate of $V(j, T-1|I_{T-1})$. By locating Q_j and Q_{j_0} at the center of the range, and set the range large enough, we can reduce the probability density that $Q_{j,t}$ moves near the bounds. This is how we can reduce the approximation errors. We now outline a general procedure to achieve this goal.

1. Estimate a static learning model first. Without forward-looking behavior, there is no need to do discretization of state variables because we do not need to solve for the dynamic programming problem. Let $\hat{\theta}^s$ denote the parameter estimates obtained in this step.
2. Take the estimates from step 1 as the reference, and then choose the bounds of I_t^- so

⁹ The discussion here extends Ching (2010).

that \hat{Q}_j^s 's and $\hat{Q}_{j_0}^s$'s are located at the center.

For σ_{jt} , it is easier to choose the range because it is a non-increasing function of t , and it must be greater than 0. With the estimate of σ_{j_0} from step 1, we can set the range of the possible values for σ_{jt} to be $[0, \hat{\sigma}_{j_0}^s]$. To be safe, one can pick an upper bound to be a slightly larger than $\hat{\sigma}_{j_0}^s$.

3. Estimate the dynamic model. Let $\hat{\theta}^d$ denote the parameter estimates obtained in this step. Keep in mind that for each trial of parameter vector, we need to solve for the dynamic programming problem. As described above, we suggest using Keane-Wolpin's interpolation method to solve for the approximated value functions at each t . In particular, when estimating the interpolation regressions, we randomly sample the state points from I_t^- .

If forward-looking incentive plays an important role, the parameter estimates will likely be different from those obtained from estimating the static version of the model. If it turns out that \hat{Q}_j^s 's and $\hat{Q}_{j_0}^s$'s move towards the boundary, we should expand I_t^- , and re-estimate the model.

4. In general, how large should we set I_t^- ? This question is similar to how to choose the terminal period, T , for solving the dynamic programming problem. The trade-off is that the larger I_t^- , the more accurate the approximation. However, the larger I_t^- , the more simulated state points are required to do the interpolating regressions in order to achieve certain goodness-of-fit. In general, we should experiment expanding I_t^- . If the estimates are not sensitive to the expansion, then we stop.

Recently, several research papers have also applied the modified Bayesian MCMC algorithm first proposed by Imai, Jain and Ching (2009) to estimate learning models with forward-looking consumers: Ching et al. (2009), Osborne (2011), Roos et al. (2012). Readers may refer to these papers for details about how to implement this new estimation method. Narayanan and Manchanda (2009) has shown how to extend the standard MCMC and data augmentation procedure to estimate learning models with myopic consumers.

Appendix References

- Bather, J. (1983) Optimal Stopping of Brownian Motion: A Comparison Technique. Recent Advances in Statistics, Papers in honor of Herman Chernoff on his sixtieth birthday, M.H. Rizvi et al. (eds), Academic Press, 19-49.
- Ching, A. (2010) A dynamic oligopoly structural model for the prescription drug market after patent expiration. *International Economic Review*, 51(4): 1175-1207.
- Ching, A., S. Imai, M. Ishihara and N. Jain (2009) A dynamic model of consumer learning and forgetting. Working paper, Rotman School of Management, University of Toronto.
- Eckstein, Z., D. Horsky, Y. Raban (1988) An empirical dynamic model of brand choice. Working paper 88, University of Rochester.
- Erdem, T., M. Keane (1996), "Decision-making under uncertainty: capturing dynamic brand choice processes in turbulent consumer goods markets," *Marketing Science* **15** 1-20.
- Gittins, J. (1983) Dynamic Allocation Indices for Bayesian Bandits. Lectures Notes in Statistics, vol.20, Mathematical Learning Models, Theory and Algorithms, V. Kalin and W. Wogel (eds), Springer-Verlag.
- Gittins, J. and D. Jones (1974) A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, 66(3): 561-565.
- Hauser, J. and G. Urban (1986) Value Priority Hypotheses for Consumer Budget Plans. *Journal of Consumer Research*, vol.12(4): 446-462.
- Imai, S., N. Jain and A. Ching (2009) Bayesian Estimation of Dynamic Discrete Choice Models. *Econometrica*, vol.77(6): 1865-1899.
- Keane, M. (1994) A Computationally Practical Simulation Estimator for Panel Data. *Econometrica*, 62(1): 95-116.
- Keane, M. and K. Wolpin (1994) The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation: Monte Carlo Evidence. *Review of Economics and Statistics*, 76(4): 648-72.
- Narayanan, S. and P. Manchanda (2009) Heterogeneous learning and the targeting of marketing communication for new products. *Marketing Science*, 28(3): 424-441.
- Osborne, M. (2011) Consumer learning, switching costs, and heterogeneity: A structural Examination. *Quantitative Marketing and Economics*, 9(1): 25-46.
- Roberts, J. and G. Urban (1988) Modeling Multiattribute Utility, Risk, and Belief Dynamics for New Consumer Durable Brand Choice. *Management Science*, 34(2): 167-185.
- Roos, Jason M.T., Carl F. Mela and Roni Shachar (2011) Hyper Media Search and Consumption. Working paper, Fuqua School of Business, Duke University.

Whittle, P. (1983) Enlightened Approximations in Optimal Stopping. Recent Advances in Statistics, Papers in honor of Herman Chernoff on his sixtieth birthday, H.H. Rizvi et al. (eds), Academic Press, 19-49.

A Simple Approach to Estimate the Roles of Learning, Inventory and Experimentation in Consumer Choice*

by

Andrew T. Ching
Rotman School of Management
University of Toronto
Andrew.Ching@rotman.utoronto.ca

Tülin Erdem
Stern School of Business
New York University
terdem@stern.nyu.edu

Michael P. Keane
University of Oxford
michael.keane@nuffield.ox.ac.uk

This draft: October 23, 2012

* Ching's work on this project is supported by SSHRC. Keane's work on this project was supported by Australian Research Council grant FF0561843. We thank Masakazu Ishihara for helpful discussion.

1. Introduction

In the past two decades, the marketing and empirical industrial organization literatures have been focusing on two sources of state dependence in choice dynamics: (i) consumer learning (e.g., Erdem and Keane, 1996; Akerberg, 2003; Crawford and Shum, 2006) and (ii) consumer stockpiling and inventory (e.g., Erdem, Imai and Keane, 2003; Hendel and Nevo, 2006; Ching, Erdem and Keane, 2009). These models are still difficult to estimate – requiring the solution of a dynamic programming (DP) problem if we allow consumers to be forward-looking. A key open unanswered question is whether learning or inventories provide a better explanation of choice dynamics (Ching, Erdem and Keane, 2012). Or, put it differently, how does accounting for consumer inventory affects the inference of learning and experimentation behavior based on panel data of consumer choice?

Obviously, to answer this question, one avenue is to develop a structural model that incorporates these two sources of dynamics explicitly. But unfortunately, computational limitations make it infeasible to estimate models with both learning and inventories where consumers are forward-looking. There are simply too many state variables – levels of perceived quality and uncertainty for all brands, inventory of all brands, current and lagged prices of all brands – which make the computation of optimal solution too time-consuming, and hence make the estimation infeasible. This makes it impossible to nest learning and inventory models and test which mechanism is more important for explaining choice dynamics (or, alternatively, whether both mechanisms are important).

In this paper, we take a step towards addressing this issue. Specifically, we present a new approach that allows one to estimate models with both learning and inventory effects, while also testing whether consumers engage in strategic trial (i.e., experimentation). Here, we extend a method developed by Geweke and Keane (2000) to make it possible to estimate these models without solving the DP problem. By using the Geweke and Keane (2000) method, we can, for the first time, estimate a model with both learning and inventories, and shed light on the role of each. Our example focuses on demand for diapers.

The outline of the paper is as follows: **Section 2** describes our new modeling framework. We will first present a standard learning model, and then discuss how our new “quasi-structural” approach allows us to (a) nest learning and inventory models and (b) test for forward looking

behavior. **Section 3** presents the results. **Section 4** concludes.

2. The Modeling Framework

2.1 A Standard Learning Model

We will start describing a standard learning model in this section. Then we will discuss how to modify the model to account for consumer's inventory. The key feature of learning models is that consumers do not know the attributes of brands with certainty. Before receiving any information via use experience, assume consumers have a normal prior on brand quality:

$$(1) \quad Q_j \sim N(Q_{j0}, \sigma_{j0}^2)$$

This says consumers perceive that the true quality of brand j (Q_j) is distributed normally with a mean of Q_{j0} and a variance of σ_{j0}^2 . The values of Q_{j0} and σ_{j0}^2 may be influenced by many factors, such as reputation of the manufacturer, pre-launch advertising, etc. We will later show how to specify it.

Usage experience does not fully reveal quality because of “inherent product variability.” This has multiple interpretations. First, the quality of different units of a product may vary. Second, a consumer's experience of a product may vary across use occasions. For instance, a drug may be more effective in relieving symptoms in some occasions, but be ineffective in other occasions, depending on the exact virus/bacteria that causes the symptoms. Alternatively, there may be inherent randomness in psychophysical perception. E.g., the same cereal tastes better to me on some days than others.

Given inherent product variability, there is a distinction between “experienced quality” for brand j on purchase occasion t , which we denote Q_{jt}^E , and true quality Q_j . Let us assume the “experienced quality” delivered by use experience is a noisy signal of true quality, as in:

$$(2) \quad Q_{jt}^E = Q_j + \varepsilon_{jt} \quad \text{where} \quad \varepsilon_{jt} \sim N(0, \sigma_\varepsilon^2)$$

Here σ_ε^2 is the variance of inherent product variability, which we often refer to as “experience variability.” It should be noted that all brands have experience variability, so (2) holds for all j .

Note that we have conjugate priors and signals, as both the prior on quality (1) and the noise in the quality signals (2) are assumed to be normal. Thus, the posterior for perceived

quality, given a single use experience signal, is given by the simple updating formulas:

$$(3) \quad Q_{j1} = \frac{\sigma_{j0}^2}{\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j1}^E + \frac{\sigma_\varepsilon^2}{\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j0}$$

$$(4) \quad \sigma_{j1}^2 = \frac{1}{(1/\sigma_{j0}^2) + (1/\sigma_\varepsilon^2)}$$

Equation (3) describes how a consumer's prior on quality of brand j is updated as a result of the experience signal Q_{j1}^E . Note that the extent of updating is greater the more accurate is the signal (i.e., the smaller is σ_ε^2). Equation (4) describes how a consumer's uncertainty declines as he/she receives more signals. σ_{j1}^2 is often referred to as the "perception error variance."

Equations (3) and (4) generalize to any number of signals. Let $N_j(t)$ denote the total number of use experience signals received through time t . Then we have that:

$$(5) \quad Q_{jt} = \frac{\sigma_{j0}^2}{N_j(t)\sigma_{j0}^2 + \sigma_\varepsilon^2} \sum_{s=1}^t Q_{js}^E d_{js} + \frac{\sigma_\varepsilon^2}{N_j(t)\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j0}$$

$$(6) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{j0}^2) + N_j(t)(1/\sigma_\varepsilon^2)}$$

where d_{jt} is an indicator for whether brand j is bought/consumed at time t .

In equation (5), the perceived quality of brand j at time t , Q_{jt} , is a weighted average of the prior and all quality signals received up through time t , $\sum_{s=1}^t Q_{js}^E d_{js}$. Crucially, this is a random variable across consumers, as some will, by chance, receive better quality signals than others. Thus, the learning model endogenously generates heterogeneity across consumers in perceived quality of products (even starting from identical priors).

As equation (6) indicates, the variance of perceived quality around true quality declines as more signals are received, and in the limit perceived quality converges to true quality. Still, heterogeneity in perceptions may persist over time, because: (i) both brands and consumers are finitely lived, as people gather more information the value of trial purchases diminishes, and so eventually learning about unfamiliar products will become slow; (ii) there is a flow of new brands and new consumers entering a market.

We assume consumer i 's utility of consuming brand j is:

$$(7) \quad U_i(Q_{jt}^E, P_{jt}) = f(Q_{jt}^E) - w_P P_{jt} + e_{ijt}$$

where w_P is the marginal utility of income (which captures the price sensitivity), and e_{jt} is an idiosyncratic brand and time specific error term. If we assume that $f(Q_{jt}^E)$ takes the constant absolute risk aversion form (CARA), then the expected utility is given by:

$$(8) \quad E[U_i(Q_{jt}^E, P_{jt})|I_t] = -\exp\left(-r\left(Q_{jt} - \frac{r}{2}(\sigma_{jt}^2 + \sigma_\varepsilon^2)\right)\right) - w_P P_{jt} + e_{ijt}.$$

We also introduce the utility of the no purchase option. Typically, it is modeled using a reduced form approach, e.g., $E[U_{i0}|I_t] = \Phi_0 + e_{i0t}$.

In general, learning models are solved by dynamic programming, because today's purchase affects tomorrow's information set, which affects future utility. As a result, consumers choose a brand that gives them the highest expected current utility plus the expected future payoffs:

$$(9) \quad V(j, t|I_t) = E[U(Q_{jt}^E, P_{jt})|I_t] + \beta EV(I_{t+1}|I_t, j) \quad \text{for } j=0, \dots, J$$

where $I_{t+1} = \{Q_{1,t+1}, \dots, Q_{J,t+1}, \sigma_{1,t+1}^2, \dots, \sigma_{J,t+1}^2\}$. Note that choosing the brand with the highest *expected current* utility is not necessarily optimal. Eq(9) is commonly called alternative specific value function in the discrete choice dynamic programming literature.

2.2 Our Approach to Incorporate Inventory and Forward-looking Incentives

A key limitation of existing learning models is that they ignore inventory effects, which are another key source of dynamics in consumer choice behavior. As we noted earlier, it is not feasible to estimate a model of forward-looking agents with both learning and inventory effects, as the state space is intractably large. Here we propose a potential solution to this problem.

Geweke and Keane (2000) present a method to estimate parameters of dynamic models without the need to solve agents' DP problem. The basic idea is to replace equation (9) by:

$$(10) \quad V(j, t|I_t) = E[U(Q_{jt}^E, P_{jt})|I_t] + F[I_{t+1}(I_t, j)|\pi_t] \quad \text{for } j=0, \dots, J$$

Here $F[I_{t+1}(I_t, j)|\pi_t] \approx \beta EV(I_{t+1}|I_t, j)$ is a polynomial in the state variables that approximates

the “future component” of the value function. And π_t is a vector of reduced form parameters that characterize the future component. The idea of the Geweke-Keane (GK) method is to estimate the π_t jointly with the structural parameters that enter the current period expected utility function.

The GK method has been applied to life-cycle labor supply (Houser (2003)), and to model behavior in games (Houser, Keane and McCabe (2004)). In those applications, current payoffs are at least partially observed. This allows one to separate the π_t from the utility function parameters. Otherwise, one could not identify if a variable affects choices by shifting current utility or the future component. But in the typical learning model this approach is not useful, as we do not observe the current payoffs (i.e., we observe only choices, not utility).

Here, we show that an alternative approach to identification is possible if the structure of the problem implies that some state variables enter the future component but not current payoffs. As we’ll see, this is true of the (updated) perception error variances in the Bayesian learning model. If the π_t parameters characterizing the future component are identified, one can estimate them without imposing all the structural restrictions implied by the Bayesian learning model. Then, various modeling assumptions about how agents process information become testable. In particular, we provide a way to test whether strategic trial is important.

We first show how to estimate the Erdem and Keane (1996) learning model using the Geweke-Keane method. Then later we discuss inclusion of inventories. To proceed, write the future component of the value function in the Bayesian learning model as follows:

$$(11) \quad \beta E_t V[I(t+1)] = F(Q_{1t}, \dots, Q_{jt}, \sigma_{1,t+1}^2, \dots, \sigma_{j,t+1}^2)$$

Note the crucial distinction between Q_{jt} and σ_{jt}^2 . The σ_{jt}^2 evolve deterministically, so a consumer knows their values at $t+1$, given any current action. Thus, the $\{\sigma_{j,t+1}^2\}$ can be conditioned on when forming the expected value at time t . In contrast, the $\{Q_{jt}\}$ are rational forecasts of brand quality, given all current information. A consumer cannot anticipate any change in the $\{Q_{jt}\}$, and the expected value only depends on their current values.¹

Now, consider the expected value function conditional on brand j being chosen at time t :

¹ Informally, a rational consumer cannot *expect* to discover (through further information gathering) that a brand is better or worse than he/she thinks it is (i.e., a consumer cannot *expect* to be positively or negatively surprised). Formally, the law of iterated expectations applies to rational forecasts: $E_t [Q_{j,t+1} | I_{t+1}] = E_t [Q_{j,t+1} | I_t] = Q_{jt}$.

$$(12) \quad \beta E_t V[I(t+1) | d_{jt} = 1] = F(Q_{1t}, \dots, Q_{jt}, \sigma_{1t}^2, \dots, \sigma_{j,t+1}^2, \dots, \sigma_{jt}^2)$$

Note that all state variables upon which $F(\cdot)$ depends are dated at time t except for $\sigma_{j,t+1}^2$. This is because (i) the chosen brand j is the only one whose perception error variance is expected to drop due to trial, and (ii) as we just discussed, expected qualities at $t+1$ are the same as those at time t .

A key point is that only utility differences matter for choices in any discrete choice model (including dynamic models). Thus, it is standard to define a base alternative and measure utilities (or value functions) relative to that alternative. In the present case, we choose the no-purchase option as the base alternative. The future component associated with no-purchase has the form:

$$(13) \quad \beta E_t V[I(t+1) | d_{jt} = 0 \forall j] = F(Q_{1t}, \dots, Q_{jt}, \sigma_{1t}^2, \dots, \sigma_{jt}^2, \dots, \sigma_{jt}^2)$$

If we compare (12) and (13), we see the only source of difference between the two functions is the difference between $\sigma_{j,t+1}^2$ and σ_{jt}^2 . This is useful, as it means the differenced value functions (between choice of j and choice of the no purchase option) take on a rather simple form.

For example, assume the future component is approximated as a polynomial in the state variables. Then, if we take the difference between (12) and (13), any terms that do not involve the change in σ_{jt}^2 will drop out.² For instance, suppose we have:

$$(14) \quad F(Q_{1t}, \dots, Q_{jt}, \sigma_{1t}^2, \dots, \sigma_{jt}^2) = P_0(Q_t) + P_1(\sigma_{-jt}^2) + P_2(Q_t, \sigma_{-jt}^2) + P_3(Q_t, \sigma_{jt}^2)$$

Here, the $P_k(\cdot)$ denote polynomials in the indicated arguments. For instance, $P_2(Q_t, \sigma_{-jt}^2)$ is a polynomial that includes interactions between perceived qualities and perception error variances *other than* that for brand j . When we take differences with respect to the no purchase option, all the polynomials except $P_3(Q_t, \sigma_{jt}^2)$ will drop out. To be concrete, let P_3 take the simple form:

$$(15) \quad P_3(Q_t, \sigma_{jt}^2) = (\pi_{0j} + \pi_{11}Q_{1t} + \dots + \pi_{1j}Q_{jt}) \cdot \sigma_{jt}^2$$

Then, conditional on brand j being chosen at time t , P_3 takes the value:

² A key point is that the $F(\cdot)$ function does not differ across alternatives in its parameters, only in its arguments. This is clear from comparison of (30) and (34), and is reflected in the notation in (36)-(37), i.e., F has no j subscript.

$$(16) \quad P_3(Q_t, \sigma_{j,t+1}^2) = (\pi_{0j} + \pi_{11}Q_{1t} + \dots + \pi_{j1}Q_{jt}) \cdot \sigma_{j,t+1}^2$$

If we difference the future components associated with choice of brand j vs. no purchase we get:

$$(17) \quad F(\cdot | d_{jt} = 1) - F(\cdot | d_{0t} = 1) = (\pi_{0j} + \pi_{11}Q_{1t} + \dots + \pi_{j1}Q_{jt}) \cdot \Delta\sigma_{j,t+1}^2$$

where $\Delta\sigma_{j,t+1}^2 \equiv \sigma_{j,t+1}^2 - \sigma_{jt}^2$. Then, we have that:

$$(18) \quad V_{jt}^s \equiv V_{jt} - V_{0t} = E[U(Q_{jt}^E, P_{jt}) | I_t] - E[U_{0t} | I_t] + (\pi_{0j} + \pi_{11}Q_{1t} + \dots + \pi_{j1}Q_{jt}) \cdot \Delta\sigma_{j,t+1}^2$$

Note that we could complicate (16) by specifying that $P_3(\cdot)$ involves higher order terms in Q_{jt} and/or σ_{jt}^2 . We could also simplify (16) by assuming that $P_3(\cdot)$ does not involve quality levels Q_{kt} for brands other than j . This would give us $P_3(Q_t, \sigma_{jt}^2) = (\pi_{0j} + \pi_{1j}Q_{jt}) \cdot \sigma_{jt}^2$ and hence:

$$(19) \quad V_{jt}^s \equiv V_{jt} - V_{0t} = E[U(Q_{jt}^E, P_{jt}) | I_t] - E[U_{0t} | I_t] + (\pi_{0j} + \pi_{1j}Q_{jt}) \cdot \Delta\sigma_{j,t+1}^2$$

In our empirical work we adopt the simple specification in (19) for three main reasons:

First, the GK method has never been applied to learning models, so it seems best to start with a parsimonious specification (to avoid any risk of identification problems). Second, we find that even the simple model in (19) leads to an improvement in fit over a myopic model.

Third, equation (19) is intuitive and easily interpretable: Parameter π_{0j} captures the value of information gained through a trial purchase of brand j . Of course, this value may differ across brands, a point we examine in the empirical work.³ Parameter π_{1j} captures that this trial value may differ depending on the perceived quality of a brand. The sign of π_{1j} is not clear *a priori*.⁴

Estimation of a learning model using the GK method is of considerable interest in itself. It enables us test whether strategic trial is empirically important without needing to fully impose the strong assumptions about learning behavior implied by the Bayesian learning model.

³ There are many reasons why the value of trial information may differ across brands. For example, say brand 2 has a much lower market share than brand 1 (perhaps because it is priced high relative to its quality). In that case, information about brand 2 may be almost irrelevant for the value function compared to information about brand 1.

⁴ The question is whether trial information is more valuable for a brand perceived to be high or low quality? CARA utility predicts $\pi_{1j} < 0$. If $r=1$, the expectation of CARA is $-\exp(-Q+(1/2)\sigma^2)$. This implies that, for any given σ^2 , the lower is Q the more expected utility one gains by reducing σ^2 by a fixed amount.

Next we proceed to include inventory behavior in the model. Ideally, we would like to add inventories to the state space in (9). Unfortunate, as is typical, we do not observe inventory directly. So we proxy for inventories using weeks since last purchase (henceforth $pgap$). The exclusion restrictions relevant for the state variable $pgap_t$ are obvious: If one buys any brand at t then $pgap_{t+1} = 1$. If one chooses no purchase, then $pgap_{t+1} = pgap_t + 1$. Thus, the value function difference, (19), is simply augmented to include $pgap_t$ itself:

$$(20) \quad V_{jt}^* \equiv V_{jt} - V_{0t} = E[U(Q_{jt}^E, P_{jt}) | I_t] - E[U_{0t} | I_t] + (\pi_{0j} + \pi_1 Q_{jt}) \cdot \Delta \sigma_{j,t+1}^2 + \pi_{pg} pgap_t$$

We also incorporate the “Price Consideration” (PC) mechanism of Ching, Erdem and Keane (2009). The PC model generalizes standard inventory models by positing that consumers may not consider a category in every period. The idea is that, shortly after buying in a category, consumers may not look at it again for a few periods (so as to conserve on mental effort). During this interlude, consumers may not respond even to a deep price discount by their favorite brand, simply because they do not see it. Ching et al (2009) find the PC model gives a much better fit to observed purchase hazard rates – particularly for short inter-purchase spells – than traditional models where consumers consider a category every period.⁵

The PC model consists of two stages. In the first stage, consumers decide whether to consider a category in period t . If they consider, then, in stage two, choice is made using the model in (20). Note that no-purchase is still a possible outcome given consideration. We model the consideration decision using a simple logit model (suppressing household i subscripts):

$$(21) \quad U_{ct} = \gamma_0 + \log(1 + (\gamma_{pg} + \gamma_q Q_t^{lp}) pgap_t) + \gamma_{t\#} t_t^\# + \gamma_x X + \gamma_{ft} \sum_j ft_{jt} + \gamma_{dp} \sum_j dp_{jt}$$

Note that $pgap$ is entered in log form, and it is interacted with Q_t^{lp} , quantity purchased on the last purchase occasion. The variable $t_t^\#$ is time since the household first entered the diaper market,⁶ while X contains indicators for whether the household is young, high income or low education. Marketing activity may induce a household to consider the category; ft_{jt} and dp_{jt} are indicators for whether brand j is on feature or display at time t , respectively. We sum these over brands, so

⁵ Of course, the PC model nests standard inventory models as the consideration probability approaches one.

⁶ Time in the market may affect the consideration probability for several reasons. One is that familiarity with the price process (timing of discounts may increase). Another is that the usage rate for diapers falls as a child ages.

marketing activity for the whole category shifts the consideration probability.

To complete the model, we need to select a functional form for the expected utility function in (21).⁷ We assume an augmented version of the CARA form in equation (20):

$$(21) \quad E[U(Q_j^E, P_j)|I_t] = -\exp\left(-r\left(Q_{jt} - \frac{r}{2}(\sigma_{jt}^2 + \sigma_\varepsilon^2)\right)\right) + \alpha_P P_{jt} \\ + \alpha_f f_{jt} + \alpha_d dp_{jt} + \alpha_a ad_{jt} + e_{jt}$$

Here, allow for persuasive effects of advertising by letting feature, display and advertising expenditure (ad_{ijt}) affect utility directly.

Next, we specify that expected utility from no purchase is simply $E[U_{0t}|I_t] = e_{0t}$. Note that no-purchase contains no intercept. This is an identification assumption, as all other utilities are measured relative to no-purchase. Also, the data contains two very small brands (labeled 5 and 6). We model utility from these as just a constant plus an extreme value error.

As noted earlier, the learning aspect of our model is based on Erdem and Keane (1996), where advertising and use experience signal brand quality (see Section 3 for details). However, unlike Erdem and Keane, we do not observe ad exposure at the individual level. Instead, we let ad expenditure be a proxy for the volume of ads received by consumers. Specifically, we modify (16) by assuming that $\varepsilon_{jt}^A \sim N(0, \sigma_A^2(i))$ and $\sigma_A^2(i) = \sigma_A^2 / (1 + \eta \cdot ad_{ijt})^2$ where ad_{ijt} denotes advertising expenditure on brand j to consumer i at time t .⁸ Note that consumers from two different cities will have different exposure to advertising. Here the parameter η is used to map ad spending into an “effective” number of ads seen.

Next, consider the prior on quality. Most papers assume this is common across brands and consumers.⁹ We relax this by letting prior means be consumer and brand specific, as in:

$$(22) \quad Q_{ijt} = Q_{j0} + (\lambda_1 + \lambda_2 low_ed_i + \lambda_3 young_i) \cdot \overline{ad}_{ijt}$$

⁷ Note that, for simplicity, the timing of consumption is not modeled. The use experience signal is assumed to arrive at the time of purchase.

⁸ This approach is similar to Ching (2010a, 2010b) who uses purchase volume as a proxy for the number of experience signals shared by the market in a social learning environment.

⁹ Also, the prior mean on quality for each brand is often assumed to be equal to the true mean quality level across all brands in the category. This is a type of rational expectations assumption, as it means prior beliefs are consistent with the population distribution of true mean qualities (see Erdem and Keane (1996)).

Here $\overline{ad}_{ijt} = (ad_{ijt} + ad_{ij,t-1} + ad_{ij,t-2} + ad_{ij,t-3})/4$. Thus, the prior mean of person i for brand j depends on the brand's ad spending (in the prior month), and the person's education and age. We continue to assume the prior variance (σ_0^2) is the same across households and brands.

Finally, it is clear from (21) that one cannot identify r , the quality levels, and the signal variances separately. Hence, we fix $r=1$ for identification (see Ching and Ishihara (2010a)).

3. Empirical Results

We apply our model to weekly scanner panel data for diapers provided by AC Nielsen. Diapers are a good category for studying learning, because an exogenous event, birth of a first child, triggers entry into the market. Before that, most people presumably know little about diapers. Unfortunately, the Nielsen data do not record births. Thus, we use selection criteria that make it *likely* our households entered the market due to a first birth during our sample period. Loosely speaking, we take households who took a long time before making their first in-sample purchase, and who make more frequent purchases thereafter.¹⁰ Such households *probably* did not enter the diaper market until after the start of the sample.¹¹

Our sample consists of 91 households and 4588 observations (50.4 weeks per household). We estimate the model using simulated maximum likelihood (Keane, 1993). Tables 1 and 2 give summary statistics on households and brands, respectively. There are four major brands: The leading name brands are Pampers (34.1%) and Huggies (30.7%), followed by LUVS (10.9%). The market share of Store Brands is 20.8%. Households rarely buy small brands, which we label "Other" or multiple brands in a shopping trip, which we label "multiple brands."

To investigate the importance of different types of dynamics, we also estimate models without the PC stage and without strategic trial. The latter shut down the $(\pi_{0j} + \pi_1 Q_{jt}) \cdot \Delta \sigma_{j,t+1}^2$

¹⁰ We include households who satisfy the following criteria: (a) the first inter-purchase spell must be longer than any subsequent inter-purchase spell, (b) the child must be under six, (c) age of the primary shopper is under 55, (d) they cannot purchase over 500 diapers in any week, (e) they must purchase at least 5 times and less than 60 times during the sample period. To model consumer heterogeneity in the initial prior on quality, we also need to observe three weeks of advertising expenditures before households enter the diaper market.

¹¹ Of course, there are other explanations. Most obviously, a household may have had one or more children several years earlier, followed by a long gap before the present child. In that case they would have had prior experience of diapers with the older children. On the other hand, that experience may be too old to be relevant to the current market.

terms in (20), leaving a myopic learning model. However, we always include the $\pi_{pg}pgap_t$ term in (20), so consumers are always forward-looking with regard to inventories.¹²

First consider the myopic learning model with no PC stage (our most basic model). In Table 3 we see that the prior mean quality levels for the two large name brands are 11.7, while those for LUVS and the Store brand are 11.3 and 11.5, respectively. Consumers appear to be over-optimistic, as the estimated true mean quality levels range from 7.9 to 9.6. The λ parameters (see equation (22)) imply that lagged advertising has no significant effect on priors for quality.

Table 3 also reports the key learning parameters. The prior variance is 8.58 (=2.93²) and the experience signal variance is 17.39 (=4.17²). As noted earlier, we map ad expenditure into the ad signal variance using the equation $\sigma_A^2/(1 + \eta \cdot ad_{ijt})^2$. As $\sigma_A^2 = 370.1$ and $\eta = 1.42 \times 10^{-5}$, an expenditure of \$6300 (mean for Huggies) implies an ad signal variance of 312 (=17.7²). Note that the relative size of the standard deviations for the three signals, 2.9, 4.2, 17.7, are similar to those found by Erdem and Keane (1996), i.e., 0.20, 0.57 and 1.75.¹³ Thus, just as Erdem-Keane found, category priors are rather tight, and experience signals are much less noisy than ads.

The parameters of the expected utility function (21) are reported in Table 4. The price coefficient is -.022 with a t-statistic of -4.4. As expected, display ads increase the probability of purchase (i.e., the coefficient on the display indicator is a substantial 2.66). Interestingly, feature and advertising have no *direct* (i.e., persuasive) effect on the probability of purchase.

Finally, the $\pi_{pg}pgap_t$ term is reported in Table 5. Strikingly, it is small and insignificant. This suggests inventory is not a significant determinant of purchase probability for diapers. This is not as surprising as it may appear. Inventory does not *necessarily* affect purchase timing in an inventory model.¹⁴ The results imply that diapers are so expensive and high involvement, with such a high usage rate and long lifespan, that consumers are always interested in buying on deal, even if they have bought recently and inventory is high.

Next we turn to the forward-looking learning model that accommodates strategic trial. Normally this would require solving a DP problem. But in the GK approach all we need do is add the $(\pi_{0j} + \pi_1 Q_{jt}) \cdot \Delta \sigma_{j,t+1}^2$ terms that approximate the future component in equation (20).

¹² It makes no sense to consider a myopic inventory model, as myopic agents would not care about inventories.

¹³ It doesn't make sense to compare the levels, as the scale normalizations are different in the two studies.

¹⁴ The main prediction of any inventory model is that consumers time purchases for when prices are relatively low. Only if carrying costs are high enough to constrain this behavior will inventories affect purchase probability.

Table 5 reports the estimates of the π_0 parameters. They are highly significant for all four major brands. Thus, consumers act as if they value strategic trial. The interaction terms imply that trial is less valuable for brands with higher perceived quality (as predicted by CARA utility).

Table 6, columns 1 and 2, compare the fit of the myopic and forward-looking learning models. As we see, the log-likelihood improves from -4041 to -3978 (or 63 points) when the future component terms are included. This adds 8 parameters, but BIC improves by 55 points. This is clear evidence that strategic trial is an important factor in consumer choice dynamics.

In Table 4, we see the parameters of the expected utility function are not greatly affected by including the future component (although the price coefficient is 50% greater). But in Table 3 we see that both prior means and true values of quality are scaled up considerably, as are prior and signal variances. We conjecture this occurs because, *ceteris paribus*, adding an incentive for strategic trial leads to more brand switching. To control this effect (and maintain a plausible level of brand loyalty) the model needs to create more perceived differentiation among the brands.¹⁵

Next we consider the forward-looking learning model with the PC stage. Estimates of the logit for category consideration are reported in Table 5. Several variables are significant determinants of consideration: display, time since first purchase, age, income and education. But time since last purchase is not significant. This is consistent with our earlier discussion of why inventory is not an important determinant of purchase probability.

Adding the category consideration stage has several notable effects on the estimates. As we see in Table 3, it causes the prior and true quality parameters to be scaled down somewhat (but they are still much larger than in the myopic model). Also, estimates of true quality now exceed the priors, implying consumers are pessimistic. And, unlike the two previous models, we now get the plausible result that true quality is highest for the two major brands. Next, as we see in Table 4, the price coefficient more than doubles (-0.08), and display is no longer significant. This is intuitive, as display now matters only through the consideration stage (Table 5). And finally, as we see in Table 5, the future component parameters now imply trial is more valuable for brands with higher perceived quality.

In Table 6, we compare the fit of the alternative models. Adding the PC stage improves the log-likelihood of the forward-looking learning model from -3978 to -3858, or 120 points. We

¹⁵ Another way to see this is that, given our identification normalization ($r=1$), scaling up all the Q s is equivalent to scaling up r , which means more risk aversion and hence more loyalty to familiar brands.

have added 9 parameters, so the BIC still improves by 159 points. This is clear evidence that accounting for the consideration stage helps improve the fit to consumer choice dynamics.

If we compare the full model (forward-looking agents and consideration) with the most basic model, the BIC improvement is 213 points. An interesting question is whether most of this improvement could have been achieved by adding the PC stage alone. Is strategic trial really important? To answer this question, we estimate a learning model with myopic consumers and a PC stage. The fit is reported in Table 6 column 3 (“myopic”). The log-likelihood of this myopic model is 37 points worse than the full model, and the AIC is 58 points worse, but the BIC is only 2 points worse. This is because the forward-looking model adds 8 parameters (the future components), and BIC imposes a harsher penalty than AIC for these extra parameters. Thus, we have evidence that strategic trial is important, but this evidence is somewhat weak. Most of the improvement over the “basic” model can be achieved by including the PC stage alone.

In summary, our results suggest that learning, category consideration, and, to a lesser extent, strategic trial, are all key factors driving choice dynamics in the diaper category. However, we find no evidence that inventories have a significant effect on choice probabilities.

4. Conclusion

In this paper, we have (1) developed a new and simpler approach to estimating dynamic models, and (2) used this new approach to estimate a model that includes learning and inventory effects simultaneously. In an application to the diaper category, we found that learning, category consideration, and, to a lesser extent, strategic trial, are all key factors driving choice dynamics. However, we find no evidence that inventories have a significant effect on choice probabilities.

We stress that these findings are category specific, and caution against generalizing to other contexts. Our main contribution is to provide a *framework* within which it is feasible to: (1) test for forward-looking behavior and (2) investigate the extent to which learning vs. inventories explain the observed dynamics in consumer choice behavior.

Obviously, more work remains to sort out which structural explanations are mainly responsible for the observed dynamics in consumer choice data. Erdem, Keane and Sun (2008) point out that learning models and consumer inventory models can give similar empirical implications. The basic idea of inventory models is that consumers wait for a good time to buy – i.e., a time when the price of their favorite brand is relatively low. This generates a similarity

between learning and inventory models that is especially apparent when price and quality are correlated. In that case, a learning model predicts frequent price cuts will dilute brand equity by lowering perceived quality. An inventory model says frequent deals lower the reservation price by changing price expectations. These two mechanisms are fundamentally different, but observationally they seem to be hard to distinguish.

More precisely, EKS model the quality signaling role of price in the context of frequently purchased goods. Also, they allow both advertising frequency and advertising content to signal quality (combining features of Akerberg (2003) and Erdem and Keane (1996)). And they allow use experience to signal quality, so that consumers may engage in strategic sampling. Thus, this is the only paper that allows for these four key sources of information simultaneously. In the ketchup category they find that use experience provides the most precise information, followed by price, then advertising. The direct information provided by ad signals is found to be more precise than the indirect information provided by ad frequency. They fit a learning model to essentially the same data used in the inventory model of Erdem, Imai and Keane (2003), and find that both models fit the data about equally well, and make very similar predictions about choice dynamics. For instance, both models predict that, in response to a price cut, much of the increase in a brand's sales is due to purchase acceleration rather than brand switching.

Often time, we would not be able to fully understand the empirical implications of incorporating different sources of dynamics in a structural model until we explicitly model them and estimate the structural parameters. The computational burden of solving for the optimal solution using dynamic programming has been the main hurdle for such a research agenda. However, our proposed approach should allow us to specify more sources of learning, and allowing for strategic trials simultaneously. This would allow us to take a step further to understand choice dynamics.

However, we should also highlight the drawbacks of our proposed approach. Because our approach does not impose all the structures implied by a dynamic model, it is more data demanding when we estimate the model (as we need to have more parameters that capture the expected future payoffs). Moreover, it is not suitable for conducting counterfactual experiments which could result in changing the functional form of the expected future payoffs. Therefore, for researchers who are interested in predicting counterfactual outcomes, they probably would still

need to fully specify the dynamic model and estimate it. However, since our approach potentially allows us to recover the expected future payoffs as a flexible function of the state variables, we believe that it should provide useful guidance about how to specify consumer expectation about the future. [Consider mentioning Yang and Ching, 2010].

References

- Ackerberg, D. (2003) "Advertising, learning, and consumer choice in experience good markets: A structural empirical examination," *International Economic Review*, 44: 1007-1040.
- Che, H. T. Erdem and S. Öncü (2011) "Periodic Consumer Learning and Evolution of Consumer Brand Preferences," working paper, New York University.
- Ching, A. (2010a), "Consumer learning and heterogeneity: dynamics of demand for prescription drugs after patent expiration," *International Journal of Industrial Organization*, 28(6), 619-638.
- Ching, A. (2010b), "A dynamic oligopoly structural model for the prescription drug market after patent expiration," *International Economic Review*, 51(4), 1175-1207.
- Ching, A. and M. Ishihara (2010), "The effects of detailing on prescribing decisions under quality uncertainty," *Quantitative Marketing and Economics*, 8(2), 123-165.
- Ching, A. T. Erdem and M. Keane (2009) "The Price Consideration Model of Brand Choice," *Journal of Applied Econometrics*, 24(3), 393-420.
- Crawford, G. and M. Shum (2005). Uncertainty and learning in pharmaceutical demand, *Econometrica*, 73(4), 1137-1173.
- Eckstein, Z., D. Horsky, Y. Raban (1988), "An empirical dynamic model of brand choice," Working paper 88, University of Rochester.
- Erdem, T., M. Keane (1996), "Decision-making under uncertainty: capturing dynamic brand choice processes in turbulent consumer goods markets," *Marketing Science* 15 1-20.
- Erdem, T., S. Imai, M. Keane (2003), "Brand and quantity choice dynamics under price uncertainty," *Quantitative Marketing and Economics* 1 5-64.
- Erdem, T. M. Keane, S. Öncü and J. Strebel (2005), Learning About Computers: An Analysis of Information Search and Technology Choice, *Quantitative Marketing and Economics* 3, 207-246.
- Erdem, T., M. Keane and B. Sun (2008). "A Dynamic Model of Brand Choice when Price and Advertising Signal Product Quality," *Marketing Science*, 27:6, 1111-25.
- Geweke, J. and M. Keane (2000). Bayesian Inference for Dynamic Discrete Choice Models without the Need for Dynamic Programming. In Mariano, Schuermann, and Weeks, editors, *Simulation Based Inference and Econometrics: Methods and Applications*, 100-131. Cambridge University Press.

Geweke, J. and M. Keane (2001). Computationally Intensive Methods for Integration in Econometrics. *Handbook of Econometrics*: Vol. 5, Heckman and Leamer (eds.), Elsevier Science B.V., 3463-3568.

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

Hendel, I. and A. Nevo (2006), "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 74(6), 1637-1673.

Houser, D. (2003), "Bayesian analysis of a dynamic stochastic model of labor supply and saving," *Journal of Econometrics*, vol.113, 289-335.

Houser, D., M. Keane and K. McCabe (2004), Behavior in a dynamic decision problem: an analysis of experimental evidence using a Bayesian type classification algorithm, *Econometrica*, 72(3), 781-822.

Imai, S., N. Jain and A. Ching (2009), "Bayesian Estimation of Dynamic Discrete Choice Models," *Econometrica*, vol. 77(6), 1865-1899.

Keane, M., P. Todd and K. Wolpin (2011), The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications, in *Handbook of Labor Economics*, Volume 4A, O. Ashenfelter and D. Card (eds.), 331-461.

Keane, M. (1993). Simulation Estimation for Panel Data Models with Limited Dependent Variables. *The Handbook of Statistics*, G.S.Maddala, C.R. Rao and H.D. Vinod editors, North Holland publisher.

Keane, M. (1994). A Computationally Practical Simulation Estimator for Panel Data. *Econometrica*, 62:1, 95-116.

Keane, M. and K. Wolpin (1994). The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation: Monte Carlo Evidence. *Review of Economics and Statistics*, 76:4, 648-72.

Sun, B., S. Neslin and K. Srinivasan (2003), "Measuring the impact of promotions on brand switching under rational consumer behavior," *Journal of Marketing Research*, 40:4, 389-405.

Van Heerde, S. Gupta and D. Wittink (2003). "Is 75% of the sales promotion bump due to brand switching? No, only 33% is," *Journal of Marketing Research*, 40:4, 481-491.

Van Heerde, H., P. Leeflang and D. Wittink (2000), "The Estimation of Pre- and Postpromotion Dips with Store-Level Scanner Data," *Journal of Marketing Research*, 383-95.

Van Heerde, H., P. Leeflang and D. Wittink (2004), "Decomposing the sales promotion bump with store data," *Marketing Science*, 23:3, 317-334.

Table 1. Summary statistics of household characteristics

#Households	91
Average household income*	19.98
Average household size	3.46
Percentage of households with female head < 30	45.1%
Percentage of households with female head below college education	83.5%
Total number of observations	4588
Total number of purchases	1166

Table 2. Summary statistics of product characteristics

	No purchase	Huggies	Pampers	LUVS	Store Brand	Others	Multiple Brands
Share(%)	74.6%	7.80%	8.74%	2.77%	5.32%	0.15%	0.63%
mean(p_jt)	n.a.	26.44	27.48	22.03	19.58	n.a.	n.a.
mean(ad_jt)	n.a.	0.021	0.025	0.007	0.011	n.a.	n.a.
mean(display_jt)	n.a.	0.001	0.002	0.001	0.003	n.a.	n.a.
mean(ad_exp_jt)	n.a.	6.30	9.78	3.29	0.255	n.a.	n.a.
mean(inter-purch spell)*	n.a.	4.04	3.84	2.95	4.28	8.38	4.03
mean(normalized spell)**	n.a.	0.060	0.056	0.057	0.074	0.189	0.033

Table 3. Estimates of learning related parameters

parameter	Without category consideration				With category consideration			
	Myopic		Forward-looking		Myopic		Forward-looking	
	estimate	s.e.	estimate	s.e.	estimate	s.e.	estimate	s.e.
Q_{10}	11.7*	2.11	57.3*	1.69	12.5*	2.56	20.1*	0.725
Q_{20}	11.7*	2.11	60.1*	1.84	12.4*	2.56	20.4*	0.735
Q_{30}	11.3*	2.08	54.5*	1.72	12.1*	2.55	19.3*	0.715
Q_{40}	11.5*	2.09	57.5*	1.78	12.3*	2.56	20.0*	0.721
Q_1	8.17*	1.15	39.1*	1.41	10.2*	2.33	23.0*	0.84
Q_2	8.44*	1.11	40.7*	1.43	10.8*	2.42	24.7*	1.09
Q_3	9.60*	1.26	40.0*	1.43	11.5*	2.57	22.9*	0.807
Q_4	7.87*	1.23	38.6*	1.38	10.1*	2.32	22.8*	0.811
λ_1	0.006	0.008	0.026	0.033	0.003	0.006	0.003	0.002
λ_2	-0.008	0.008	0.015	0.035	-0.005	0.007	-0.004	0.002
λ_3	3.56E-04	0.013	0.192***	0.116	0.003	0.011	-0.006	0.005
σ_0	2.93*	0.421	5.26*	0.078	2.66*	0.199	1.16*	0.013
σ_δ	4.17*	0.271	9.10*	0.166	4.52*	0.491	6.59*	0.111
σ_τ	370.1*	2.3	658.0*	23.1	338.6*	1.67	464.8*	1.51
η	1.42E-05	3.16E-02	4.22E-07	3.30E-03	7.54E-04	6.25E-02	2.00E-02	4.33E-02

*1% level of significance; **5% level of significance; ***10% level of significance.

Table 4: Estimates of Expected Utility Function Parameters

parameter	Without category consideration				With category consideration			
	Myopic		Forward-looking		Myopic		Forward-looking	
	estimate	s.e.	estimate	s.e.	estimate	s.e.	estimate	s.e.
Utility parameters for control variables								
ϕ_p (price)	-0.022*	0.005	-0.031*	0.003	-0.005	0.005	-0.08*	0.009
ϕ_f (feature)	0.162	0.404	0.26	0.418	0.418	0.449	0.459	0.419
ϕ_d (display)	2.66*	1	2.75*	0.995	1.76***	0.923	0.987	0.984
ϕ_{ad} (ad exp)	-0.001	0.002	-1.76E-04	0.002	-7.47E-04	0.002	-0.003	0.003
Utility of other small brands (just brand intercepts)								
α_5	-6.98*	0.375	-7.01*	0.356	-6.35*	0.368	-5.84*	0.374
α_6	-5.56*	0.213	-5.60*	0.182	-4.93*	0.181	-4.42*	0.189

*1% level of significance; **5% level of significance; ***10% level of significance.

Table 5. Estimates of the Category Consideration Stage and Future Component Polynomial

	Without category consideration				With category consideration			
	Myopic		Forward-looking		Myopic		Forward-looking	
parameter	estimate	s.e.	estimate	s.e.	estimate	s.e.	estimate	s.e.
Category Consideration Stage								
γ_0 (intercept)					1.16*	0.362	0.086	0.237
γ_{et} (time since 1st buy)					-0.045*	0.005	-0.03*	0.003
γ_{age} (age)					1.75*	0.457	1.72*	0.398
γ_{inc} (income)					1.26*	0.32	0.611*	0.177
γ_{edu} (I_(fedu=low))					0.912*	0.317	0.617*	0.237
γ_f (sum_feature_j)					-0.641	0.495	-0.361	0.403
γ_d (sum_display_j)					160.1*	3.24	150.5*	1.47
γ_{pg} (pgap)					4.26E-11	9.15E-05	2.82E-09	2.43E-04
γ_{ql} (Q_lp)					2.87E-11	1.11E-05	4.91E-13	3.71E-06
Expected Future Components								
π_{01} ($\Delta\sigma_{i1}$)			3.20*	0.319			-66.5*	6.38
π_{02} ($\Delta\sigma_{i2}$)			1.07*	0.27			-41.5*	5.15
π_{03} ($\Delta\sigma_{i3}$)			10.2*	0.302			-133.5*	14.8
π_{03} ($\Delta\sigma_{i4}$)			2.74*	0.287			-78.5*	7.86
π_{11} ($Q_{i1}*\Delta\sigma_{i1}$)			-0.064*	0.005			14.2*	0.558
π_{12} ($Q_{i2}*\Delta\sigma_{i2}$)			-0.027*	0.004			9.86*	0.448
π_{13} ($Q_{i3}*\Delta\sigma_{i3}$)			-0.195*	0.011			31.5*	1.24
π_{14} ($Q_{i4}*\Delta\sigma_{i4}$)			-0.058*	0.006			14.6*	0.647
π_{pg} (pgap)	-5.18E-11	4.70E-05	-3.57E-10	8.84E-05	-5.22E-10	1.39E-04	-3.52E-09	2.25E-04

*1% level of significance; **5% level of significance; ***10% level of significance.

Table 6. Goodness-of-fit, AIC, BIC

	Without category consideration		With category consideration	
	Myopic	Forward-looking	Myopic	Forward-looking
log-likelihood	-4041.1	-3977.78	-3894.99	-3858.2
(-2*log-likelihood)	8082.20	7955.56	7789.98	7716.40
AIC	8126.20	8015.56	7851.98	7794.40
BIC	8279.28	8224.30	8067.68	8065.76
#parameters	22	30	31	39
#obs	7770	7770	7770	7770