

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

by

Andrew T. Ching
Rotman School of Management
University of Toronto

Tülin Erdem
Stern School of Business
New York University

Michael P. Keane
Nuffield College
University of Oxford

This draft: June 16, 2013

Abstract

Learning models extend the traditional discrete choice framework by postulating that consumers 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 disentangle learning as a source of dynamics from other key mechanisms that may generate choice dynamics (inventories, habit persistence, etc.). Another is to enhance identification of learning models by collecting and utilizing direct measures of signals, perceptions and expectations.

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

Acknowledgements: We thank Preyas Desai, Eric Bradlow, the AE and two anonymous referees, along with five Wharton graduate students, for very helpful comments. Keane's work on this project was supported by Australian Research Council grants FF0561843 and FL110100247.

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 search for information about products via other sources.

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. In particular, consumers may engage in “passive search,” using only information sources that arrive

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.

exogenously, or “active search,” exerting effort to gather information. Or they may do both.

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 technical limitations of the time (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 myopic, so there is no active search. In contrast, in Eckstein et al (1988) consumers are forward-looking, and trial purchase is the (only) source of information. But utility is linear, so their model exhibits the “value of information” phenomenon, but not the “brand equity” 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

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, an issue we will discuss later.

form specifications, such as Guadagni and Little (1983)'s exponential weighted average of past 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, (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 fields of marketing and economics.

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 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

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) the ability 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).

discuss the rest of the literature. In general, later developments can be viewed as extending the 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. It is divided in subsections that cover: (i) more sophisticated learning models with myopic consumers, (ii) more sophisticated learning models with forward-looking consumers, (iii) models for product level/market share data, and (iv) new or novel applications of learning models. 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 models 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:

$$(1) \quad Q_j \sim N(Q_{j1}, \sigma_{j1}^2) \quad j = 1, \dots, J.$$

This says that, prior to receiving any information, consumers perceive that the true quality of brand j , denoted Q_j , is distributed normally with mean Q_{j1} and variance σ_{j1}^2 . So in the first period, the consumer’s information set is just $I_1 = \{Q_{j1}, \sigma_{j1}^2 \forall j\}$. The values of Q_{j1} and σ_{j1}^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:

$$(2) \quad Q_{jt}^E = Q_j + \varepsilon_{jt} \quad \text{where} \quad \varepsilon_{jt} \sim N(0, \sigma_\varepsilon^2) \quad \text{for} \quad t=1, \dots, T.$$

Here σ_ε^2 is the variance of inherent product variability, which we often refer to as “experience variability.” Of course experience signals are consumer i specific. But here and in later equations we will suppress the i subscript whenever possible to save on notation.

Note that we have conjugate priors and signals, as both the prior on quality in (1) and the noise in the quality signals in (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 from choice data alone.

Thus, the posterior for perceived quality, given a single use experience signal (received after the first purchase of brand j), is given by the simple Bayesian updating formulas:

$$(3) \quad Q_{j2} = \frac{\sigma_{j1}^2}{\sigma_{j1}^2 + \sigma_\varepsilon^2} Q_{j1}^E + \frac{\sigma_\varepsilon^2}{\sigma_{j1}^2 + \sigma_\varepsilon^2} Q_{j1},$$

$$(4) \quad \sigma_{j2}^2 = \frac{1}{(1/\sigma_{j1}^2) + (1/\sigma_\varepsilon^2)}.$$

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 . The extent of updating is greater the more accurate the signal (i.e., the smaller is σ_ε^2). Equation (4) describes how a consumer’s uncertainty declines when he/she receives the signal. The quantity σ_{j2}^2 is often referred to as the “perception error variance.”

Equations (3) and (4) generalize to multiple signals. Let $N_j(t)$ denote the total number of use experience signals received up until the purchase occasion at time t . Then we have:

$$(5) \quad Q_{jt} = \frac{\sigma_{j1}^2}{N_{j(t)}\sigma_{j1}^2 + \sigma_{\varepsilon}^2} \sum_{s=1}^{t-1} Q_{js}^E d_{js} + \frac{\sigma_{\varepsilon}^2}{N_{j(t)}\sigma_{j1}^2 + \sigma_{\varepsilon}^2} Q_{j1} \quad \text{for } t=2, \dots, T.$$

$$(6) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{j1}^2) + N_{j(t)}(1/\sigma_{\varepsilon}^2)} \quad \text{for } t=2, \dots, T.$$

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

In (5), perceived quality of brand j at time t , Q_{jt} , is a weighted average of the prior and all quality signals received up until the beginning of time t , $\sum_{s=1}^{t-1} Q_{js}^E d_{js}$. 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 (DP), because today's purchase affects tomorrow's information set, which affects future utility. The key idea of DP is that, at each time t , the value of choosing option j consists of an immediate payoff, *plus* the expected present value of the future payoff stream which arises in period $t+1$ onward. This "forward looking" term is *conditional* on the option j chosen at time t , as the choice at j alters the consumer's information set, which in turn affects the choices that he/she makes in the future.

In our notation, the information set is I_t , the value of choosing alternative j at time t is $V(j, t|I_t)$, the current payoff will be a context-specific utility function, and the expected present value of future payoffs (conditional on I_t and j), or "future component," is $EV(I_{t+1}|I_t, j)$.

If choices convey not just utility but also information, it may not be optimal to choose the brand with the highest *perceived* quality in the *current* period. 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 just Q_o and $\sigma_{ot}^2 = 0$. Prices are given by P_{jt} for $j=o,n$. Then values of choosing each brand in the current period are:

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

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

At time t a consumer chooses 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 old 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. On the other hand, purchase of the new brand delivers expected utility $E[U(Q_{nt}^E, P_{nt})|I_t]$, 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, $EV(I_{t+1}|I_t, n) > EV(I_{t+1}|I_t, o) = EV(I_t|I_t, o)$ because I_{t+1} contains better information. As a result, it may be optimal to try the new brand even if $E[U(Q_{nt}^E, P_{nt})|I_t] < E[U(Q_{ot}^E, P_{ot})|I_t]$.

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:

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

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

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

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

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:

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

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

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 that arises because the consumer obtains information by trying the new brand at time t .

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 large enough price differential such that the consumer will choose the new brand over the familiar brand if the new brand is cheaper by at least that amount. More precise information about the quality of the new brand enables the consumer to set this reservation price differential more accurately.

Here, we give a sketch of a proof that $EV(I_{t+1}|I_t, n) > EV(I_{t+1}|I_t, o)$, 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 simply chooses the brand with 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,

consumers do switch brands for purely idiosyncratic reasons even under full information.

one needs to show that as Q_{nT} becomes more accurate the consumer's decisions become closer to optimal, so that $EV_T(I_T = \{Q_{nT}, \sigma_{nT}^2\})$ 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 (prior to the time t purchase decision). We further assume that:

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).

$$(14) \quad A_{jt} = Q_j + \varepsilon_{jt}^A \quad \text{where} \quad \varepsilon_{jt}^A \sim N(0, \sigma_A^2) \quad \text{for} \quad t=1, \dots, T$$

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, 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:

$$(15) \quad Q_{jt} = \frac{\sigma_{j1}^2}{N_j^A(t)\sigma_{j1}^2 + \sigma_A^2} \sum_{s=1}^t A_{js} d_{js}^A + \frac{\sigma_A^2}{N_j^A\sigma_{j1}^2 + \sigma_A^2} Q_{j1} \quad \text{for} \quad t=1, \dots, T$$

$$(16) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{j1}^2) + N_j^A(1/\sigma_A^2)} \quad \text{for} \quad t=1, \dots, T$$

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:

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

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

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 *much* more difficult.

In order to progress further and develop a model that can be taken to the data, one must specify a particular functional form for the 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 must also be specified. EK wrote this 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.

We have now specified the complete EK model, and we are in a position to formally state the consumer's problem. Consumers are assumed to be forward-looking, making choices to maximize value functions of the form:

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.

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

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, 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 where risk aversion dominates, we would expect to see a high degree of brand loyalty (i.e., persistence in choice behavior). In categories where the gains from trial dominate, we would expect to see a high degree of brand switching (due to experimentation).¹¹

2.3. Solving the Dynamic Optimization (DP) Problem

Here we show how to solve a consumer's dynamic optimization problem. Solving the DP problem is computationally difficult for two reasons: (i) The expected value functions in (21) are high dimensional integrals, and (ii) these integrals must be evaluated at many state points.

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

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

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 set and decision at time t (i.e., the (I_t, j))

¹¹ We discuss the identification of learning models in detail in Section 2.5. But here we note that it is not possible to determine from raw data patterns the magnitudes of risk aversion vs. gains from trial. Such an inference can only be made conditional on an assumed model structure.

generate a distribution of I_{t+1} , in the manner described earlier.

However, it is not immediately obvious how (23) helps us to solve the consumers' optimization problem. 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 feasible to evaluate. Suppose that, hypothetically, the I_T and P_T were known at $T-1$. Then, as we see from (20), 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 (see Rust (1994)). This illustrates the point that estimating a finite-horizon dynamic model is very much 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 EK model is more difficult, because the I_T and P_T are not known at $T-1$. Both 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 $\sigma_{j,T-1}^2$ to $\sigma_{j,T}^2$ 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_{jT} , (iii) integrate over possible ad exposures that may arrive between $T-1$ and T (i.e., over realizations of d_{jT}^A for $j=1, \dots, J$) to account for ad induced changes in the $\{\sigma_{jT}^2\}$, and (iv) integrate over possible values of the ad signals in (14), as these will lead to different values of the $\{Q_{jT}\}$.¹²

Clearly the integrals in (25) are high dimensional, and simulation methods are needed. That is, we *integrate by simulation* over draws from the distributions of the signal processes. The

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.

computational burden increases if consumers learn about multiple brands, and/or have more than one source of information. Memory is also an issue, as all the $EV(I_T|I_{T-1}, j)$ must be saved.

Having calculated the values of $EV(I_T|I_{T-1}, j)$ for every possible (I_{T-1}, j) and saved the results – a point we return to below – we can move back to time $T-1$, where (21) becomes:

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

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. So we can now construct the $V(j, T - 1|I_{T-1})$. This enables us to proceed backwards and 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 “back-solving” are contained in many sources. See, for instance, Keane et al (2011).

In practice, T is generally chosen to be some time 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 (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

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.

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) and Rust (1994), among others.

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 Eqs (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:

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.

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

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, stressed by 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 DP 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 an 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, because the number of possible states in their model is so large.

To proceed in constructing the likelihood 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 before time t . Similarly, let $D_t^A \equiv \{d_1^A, \dots, d_t^A\}$ denote the history of ads received up through time t , where $d_t^A \equiv \{d_{jt}^A\}_{j=1}^J$. Also, let $Q_{t-1}^E \equiv \{Q_{j(s),s}^E\}_{s=1}^{t-1}$ and $A_t \equiv \{A_s\}_{s=1}^t$ denote the actual content of experience and ad signals received up through t . (Recall that D_{t-1} and D_t^A are observed by the econometrician, while Q_{t-1}^E and A_t are not).

Finally, we define $P(j(t)|I_0, D_{t-1}, D_t^A, Q_{t-1}^E, A_t)$ as the probability of a person's choice at time t given his/her history of use experience up through time $t-1$, and advertising exposures up through time t , as well as the content of those signals. It is worth emphasizing the timing convention that the ads at time t are observed before the time t choice is made.

Unfortunately, we cannot observe the actual content of ad and experience signals. Thus, we must integrate over that content to obtain unconditional probabilities $P(j(t)|I_0, D_{t-1}, D_t^A)$.

Thus, the probability of a choice history for an individual takes the form:

$$\begin{aligned}
(30) \quad P(\{j(t)\}_{t=1,T}) &= \prod_{t=1}^T P(j(t)|I_t) = \prod_{t=1}^T P(j(t)|I_0, D_{t-1}, D_t^A) \\
&= \int_{\{Q_{j(t-1),t-1}^E, A_t\}_{t=1}^T} \prod_{t=1}^T P(j(t)|I_0, D_{t-1}, D_t^A, Q_{j(t-1),t-1}^E, A_t) d\mu\{Q_{j(t-1),t-1}^E, A_t\}_{t=1}^T
\end{aligned}$$

In (30) we integrate over all experience and advertising signals that the consumer may have received from $t=1, \dots, T$. That is, we integrate over the distribution of $\{Q_{j(t-1),t-1}^E, A_t\}_{t=1}^T$.

Clearly, the required order of integration is substantial.¹⁵

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

$$(31) \quad P(\{j(t)\}_{t=1,T}) = \frac{1}{D} \sum_{d=1}^D \prod_{t=1}^T P(j(t)|I_0, D_{t-1}, D_t^A, S^d)$$

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 $t=0$. 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. Identification

The discussion of identification can be confusing, as the word has multiple meanings. It can mean showing the parameters of a model are identified *given the assumed model structure*. This may involve both formal proof as well as intuitive discussion of what data patterns drive the estimates. We discuss identification in this “narrow” sense in section 2.5.A.

¹⁵ 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.

Identification can also mean analysis of what assumptions are *necessary* to estimate a model, or just convenient.¹⁶ For example, can assumptions like Bayesian updating or normal signals be relaxed? Even more generally, can one distinguish the learning model from other plausible models that also generate state dependence? How can we tell if consumers are forward looking? We discuss identification in this “broad” sense in section 2.5.B.

Finally, some parameters may be formally identified but difficult to pin down in finite samples. We discuss this issue in Section 2.6, when we discuss the estimates of the EK model.

2.5.A. Identification of Learning Model Parameters (Given the Model Structure)

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:

$$(32) \quad V(j, t | I_t = \{Q_1, \dots, Q_J, 0, \dots, 0\}) \equiv w_Q Q_j - w_Q r Q_j^2 - w_Q r \sigma_\varepsilon^2 - w_P P_{jt} + e_{jt}.$$

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 $\{Q_{j1}, \sigma_{j1}^2\}$ as they drop out of the model. We also cannot identify σ_ε^2 , as it is constant across alternatives $j=1, \dots, J$.¹⁷ And careful inspection of (32) 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). Thus, r , β , σ_A^2 , σ_ε^2 and the priors $\{Q_{j1}, \sigma_{j1}^2\}$ only affect choice probabilities through the $EV(I_{t+1} | I_t, j)$ terms.

So, in an environment of 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

¹⁶ This is known as “non-parametric” identification analysis. Unfortunately, this literature has been misinterpreted by many researchers as suggesting it may be possible to obtain “model free evidence” about behavior. In fact, the 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 (e.g., the functional form of utility or an error distribution) is identified without further assumptions. Thus, what is non-parametrically identified is just a part of the model, not all of it. For instance, Matzkin (2007) says the “ideal” of non-parametric estimation is to *start* with a structural model and then impose only restrictions implied by theory (e.g., continuity, monotonicity, homogeneity, equilibrium conditions). One then uses the data to identify functional forms and distributions that are not pinned down by theory. A related point is that observing data patterns that seem consistent or inconsistent with a model can make that model seem more or less plausible, given our priors. But they can never provide non-parametric evidence that a model is correct.

¹⁷ Note that σ_ε^2 does not enter the value of the no purchase option. However, any shift in σ_ε^2 can be undone by a shift in Φ_0 , leaving utility differences unchanged.

option.¹⁸ Furthermore, as only utility differences matter for choice, we need a normalization to establish a reference alternative. EK set $Q_j = 1$ for one brand, so quality of all other brands are measured relative to brand j .¹⁹ Alternatively, one could fix w_Q . To summarize, 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 .

The identification of the parameters β , σ_A^2 , σ_ε^2 and r , as well as the priors $\{Q_{j1}, \sigma_{j1}^2\}$, requires that incomplete information actually exist. In that case, variation in $EV(I_{t+1} | I_t, j)$ and σ_{jt}^2 across consumers is generated by variation in the information sets I_{it} . Intuitively, the parameters β , σ_A^2 , σ_ε^2 , r and $\{Q_{j1}, \sigma_{j1}^2\}$ are identified by the extent to which, *ceteris paribus*, consumers with different information sets are observed to have different choice probabilities. For instance, by comparing (27a) and (32) we can clearly see that variation in σ_{jt}^2 across consumers, arising from variation in use experience and ad exposures, enables us to identify r . (This is because w_Q is already identified from consumers with complete information, as we noted earlier).

Similarly, variation of I_{it} within consumers over time is also relevant. The learning parameters σ_A^2 , σ_ε^2 and $\{Q_{j1}, \sigma_{j1}^2\}$ determine how the arrival of ad and use experience signals change σ_{jt}^2 and the $EV(I_{t+1} | I_t, j)$. Thus, these parameters are pinned down by the extent to which the arrival of signals alters behavior over time. For instance, if behavior is greatly altered by arrival of one use experience signal, it implies that σ_{j1}^2 is large and σ_ε^2 is small.

It is worth stressing that this argument for identification based on comparing behavior of consumers with different amounts of information applies in both dynamic and static models. Indeed, this is the source of identification of the learning related parameters in the static Bayesian learning model of Roberts and Urban (1988). It is also worth stressing that variation in ad exposures and in prices are plausibly exogenous sources of variation in the I_{it} .

Now we turn to dynamic considerations. Recall from Section 2.1 that consumers will only engage in strategic trial if $\beta > 0$. But in the typical scanner data set we cannot observe if a purchase is a “trial.” Thus, aside from the functional forms of utility and the EV functions, the

¹⁸ The scale normalization on utility is imposed by assuming the scale parameter of the extreme value errors is one.

¹⁹ It is worth noting that an alternative normalization 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. Also, 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).

discount factor β is pinned down by variables that affect the $EV(I_{t+1} | I_t, j)$ but do not affect current utility. In the EK model there are two exogenous variables that play this role. These are the brand specific price variances and advertising frequencies. There is no reason for these variables to affect behavior in a static model – in a static model one only cares about the current price and the current stock of information, not the likelihood of future deals or future information arrival.

2.5.B. Identification in the “General” or “Non-Parametric” Sense

As EK discuss in some detail, the Bayesian learning model implies a particular form of state dependence and serial correlation in the errors. This can be seen from careful examination of Equations (17)-(18) and (27). A frequently asked question is how learning behavior can be distinguished from other forms of state dependence/serial dependence.

In his fundamentally important paper on panel data, Chamberlain (1984) defined the relationship between two variables y_t and x_t as “static” conditional on a latent variable c if (i) y_t is independent of lagged x conditional on x_t and c , and (ii) x_t is *strictly* exogenous with respect to y conditional on c (i.e., y_t does not cause future x). As Chamberlain shows, this “static” condition is actually *stronger* than the condition that there is no structural state dependence (see Heckman, 1981), as the latter does not require strict exogeneity.²⁰

Rather remarkably, Chamberlain shows that in nonlinear models (like discrete choice models), one can always find a distribution of the latent variable c such that the relationship between y_t and x_t is static. In simple terms, one can always find a sufficiently flexible/complex heterogeneity distribution such that state dependence is not needed to explain the data. The key implication is that one cannot construct a non-parametric test of whether state dependence exists.

Obviously, if one cannot form a non-parametric test of whether state dependence exists, then it is true *a fortiori* that one cannot form a non-parametric test of whether any *particular* form of state dependence (such as learning) exists (i.e., if heterogeneity can account for general state dependence it can obviously account for any particular form of state dependence). Nor can we form non-parametric tests to distinguish among competing forms of state dependence (e.g., learning vs. inventories vs. adjustment costs).

Chamberlain’s result is an instance of the Cowles Foundation view that one cannot

²⁰ Note that a dependence of y_t on lagged x is the defining characteristic of structural state dependence. This is ruled out by condition (i), but condition (ii) is additional. In the learning model the exogenous x variables correspond to advertising exposures and prices. These variables affect purchase decisions by shifting the I_{it} and budget constraints.

deduce interesting economic relationships from the data alone. One needs *a priori* identifying assumptions, regardless of what sort of idealized variation is present in the data.²¹ Thus, our interpretation of data will always be subjective, as it is contingent on our model. To be concrete, both the extent and nature of any state dependence we find in discrete choice data will depend on the assumed functional forms for state dependence and heterogeneity (see, e.g., Keane (1997)).

As we described in Section 2.5.A, in the parameterized EK learning model, we identify parameters that describe dynamics from variation in choice behavior *across* consumers with different information sets (I_{it}), and *within* consumers as their information sets change over time. This variation in I_{it} arises from different histories of use experience, ad exposures and prices. However, Chamberlain's results imply that differences in behavior due to differences in history (i.e., state dependence) cannot be distinguished non-parametrically from differences in behavior due to a completely general form of heterogeneity. Nor can learning behavior be distinguished non-parametrically from other mechanisms that may induce state dependence.

Thus, functional forms of both state dependence and heterogeneity must be constrained for the learning model to be identified. But this is true of *any* non-linear dynamic model. For a structural econometrician this is not a limitation – a model that simply specifies very general forms of state dependence and/or heterogeneity so as to obtain a good fit to the data is merely a statistical model with no structural/behavioral interpretation. Such a model cannot be used for policy experiments. Furthermore, Occam's razor suggests that we do not wish to work with such general models. What we seek are parsimonious models that fit well, that give useful insights into the data and that can be used for policy experiments.

Recognizing the impossibility of completely non-parametric identification of learning effects, we can still give some contingent answers to the questions we asked at the start of Section 2.5. First, note that the Bayesian updating and normal signaling assumptions can be relaxed. We discuss some papers that do this in Section 3.1.1.

Second, in principle one can distinguish learning from other plausible mechanisms that may generate state dependence (like inventories or switching costs), but only if one is willing to

²¹ As Koopmans, Rubin and Leipnik (1950) state: "Suppose ... B is faced with the problem of identifying ... the structural equations that alone reflect specified laws of economic behavior ... Statistical observation will in favorable circumstances permit him to estimate ... the probability distribution of the variables. Under no circumstances whatever will passive statistical observation permit him to distinguish between different mathematically equivalent ways of writing down that distribution ... The only way in which he can hope to identify and measure individual structural equations ... is with the help of *a priori* specifications of the form of each structural equation."

specify parametric forms for all the competing models. This is consistent with the view of Bayesian decision theory that “one needs a model to beat a model.” We will return to this point in section 4.

Third, the questions of whether we can identify the discount factor and whether we can test if consumers are forward-looking are obviously closely related. Interestingly, however, Ching, Erdem and Keane (2012) show that, in the learning model, one can identify whether consumers are forward-looking using only (i) the laws of motion of the state variables and (ii) the form of current utility. But identification of the discount rate requires assumptions about the full structure (i.e., expectation formation), so that one can construct the expected value functions.

To see this, suppose we adopt the Geweke and Keane (2000) method to estimate dynamic models without the need to solve agents’ DP problem, and without imposing the full structure of the model. To implement their method we take the value function in (21):

$$(21') \quad V(j, t|I_t) = E[U(Q_{jt}^E, P_{jt})|I_t] + \beta EV(I_{t+1}|I_t, j) \quad \text{for} \quad j=0, \dots, J$$

and replace it by the equation:

$$(33) \quad V(j, t|I_t) = E[U(Q_{jt}^E, P_{jt})|I_t] + F[I_{t+1}(I_t, j)|\pi_t] \quad \text{for} \quad j=0, \dots, J$$

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.

Notice that, as F is just a flexible function of the state variables, all that is assumed is that consumers understand the laws of motion of the state variables (i.e., how $(I_{t+1} | I_t, j)$ is formed). They need not form expectations based on the true model. The approach is also agnostic about whether consumers use Bayesian updating or some other method. In general, identification of π_t requires exclusion restrictions such that some variable enters F but not U .²²

²² Geweke and Keane (2000) point out that in the absence of exclusion restrictions, one must observe current payoffs (at least partially) in order to identify F . In labor economics, researchers may argue that wages capture much of the current payoff (e.g., Houser, 2003). Or, researchers can control current payoffs in a lab experiment (e.g., Houser, Keane and McCabe, 2004). Recently, Yao et al. (2012) proposed another strategy to identify the discount factor. They argue that if a data set consists of two regimes: a static environment and a dynamic environment, one can first estimate the parameters of the current payoff function using the static environment data, and then hold them fixed

We see from (33) that, when the full structure is not imposed, one cost is that we lose identification of the discount factor. The β is subsumed as a scaling factor for the parameters π_t of the F function.²³ On the other hand, we can test whether $\pi_t = 0$, which is a test for forward-looking behavior. Although the test makes weak assumptions about F , it is *not* non-parametric, as a functional form must be chosen for the current payoff function. As Ching, Erdem and Keane (2012) show, given the current payoff function, the π_t are identified in the learning model because different current choices lead to different values of next period's state variables.

2.6. 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.

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 0.995.²⁴ Identification of the discount factor is often a practical problem in dynamic models, even when it is formally identified. (We discuss this further in Section 4.3). Second, EK also found it difficult to pin down the prior mean of quality. Hence, they constrained it to equal the average true quality level across all brands. This implies peoples' priors are correct *on average*. They also constrained the prior uncertainty to be equal across brands, $\sigma_{j1} = \sigma_1$, as allowing it to differ did not significantly improve the fit.

when estimating the discount factor using the dynamic environment data. Their approach requires the assumption that the current payoff function remains unchanged across regimes.

²³ Recently, several papers have explored using exclusion restrictions to estimate the discount factor. Chevalier and Goolsbee (2009) and Ishihara and Ching (2012) use the resale value of a used good as an exclusion restriction in estimating dynamic demand models for new and used goods. In a dynamic store choice model, Ching, Imai, Ishihara and Jain (2012) use cumulative points earned via a reward program as an exclusion restriction. In a study of sales person productivity, Chung, Steenburgh and Sudhir (2013) use cumulative sales as an exclusion restriction. The ideas in Ching et al. (2012) and Chung et al. (2013) are similar: cumulative points (or sales) do not affect current payoffs until they reach certain cutoffs so that customers (sales reps) can receive a bonus. Fang and Wang (2010) show that even parameters of quasi-hyperbolic discounting can be identified if a dynamic model has exclusion restrictions and one has panel data with at least three periods.

²⁴ 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.

Aside from the dynamic learning model, EK estimated two other models for comparison. These are a myopic learning model ($\beta = 0$), and a reduced form model similar to Guadagni and Little (1983), henceforth GL. The latter is a multinomial logit with 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 structural learning models fit substantially better than the GL model.²⁵ 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 models, 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. In the EK model there is no 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.

EK used their model to simulate an increase in ad frequency for Surf from 23% to 70%.²⁶ The simulation was also done for a hypothetical new brand with the characteristics of Surf. The results imply 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 advertising has little impact in the short run, but sustained advertising is important in the long run. As expected, the impact of advertising is much greater for a new brand (as there is more scope for learning).²⁷

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

²⁵ 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.

²⁶ “Ad frequency” is weekly probability of a household seeing an ad for a brand. In the data this was 23% for Surf.

²⁷ 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.

perceived risk of a brand to a significant degree.²⁸ 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.

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 this work is post-2000, but it already forms a large literature. We divide the review into (i) more complex learning models with myopic agents, (ii) more complex learning models with forward-looking consumers; (iii) learning models for product level/market share data; (iv) new applications of learning models (beyond brand choice). We should note that our survey focuses on *empirical* structural learning models where agents are uncertain about product attributes. [There is a literature on dynamic games where agents learn how to play equilibrium strategies, or learn how to coordinate in multiple equilibria settings including social learning environments. This literature is beyond the scope of our survey.]

3.1. Models with Myopic Agents

One stream of literature has focused on extending learning models by allowing for more complex learning mechanisms. To make such extensions feasible, it is often necessary to assume that consumers are myopic. We consider such models in the next two sub-sections that cover: (i) models with more complex learning mechanisms and (ii) models with correlated learning.

3.1.1. More Complex Learning Mechanisms

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.

Deighton (1984) proposed that advertising has a “transformative” effect whereby it alters consumer assessment of the consumption experience. Mehta, Chen and Narashiman (2008) include this effect in a learning model. They allow information signals from advertising to be biased, and this bias can change how consumers interpret their consumption experience. The

²⁸ 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.

identification of such a model is very challenging. Mehta et al. (2008) can achieve identification because their data set includes consumers who have hardly watch TV commercials. Choices of these consumers allow one to identify true brand quality levels because their experience signals are not “contaminated” by the transformative effect of advertising. After controlling for the true mean brand qualities, the choices of the consumers who do watch TV commercials allows them to identify the bias of the advertising signals and the transformative effects of advertising.

Camacho, Donkers and Stremersch (2011) also model perception biases, but in a simpler framework. They 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 to 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) allow for consumer uncertainty about the precision of quality signals. They update their perception of this precision over time. In particular, consumers who receive a very negative experience signal may change their perception of signal variance. 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 is able to fit the data better than a standard learning model, which assumes consumers know the true signal variance.

3.1.2. Models of Correlated Learning

Another stream of literature models information spillover across brands, or “correlated learning.” By this we mean learning about a brand in one category by using the same brand in another category, and/or learning about one attribute (e.g., drug potency) from another (e.g., side effects). This occurs if priors and/or signals are correlated across products or attributes.

Erdem (1998) considers a model where priors are correlated across “umbrella brands” (i.e., a brand that operates in multiple categories). She finds evidence that consumers learn via experience across 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).²⁹

Other papers have extended learning models to multi-attribute settings where consumers use experience of one attribute to draw inferences about other attributes. Prescription drugs are a good example: Coscelli and Shum (2004) estimate a diffusion model for Omeprazole, an anti-ulcer drug. It can treat: (i) heartburn, (ii) hypersecretory conditions, (iii) peptic ulcer, and provide (iv) maintenance therapy. In the model, physicians know how signals are correlated across the four conditions. In each patient-physician encounter, a physician only observes a signal of the condition being treated, but he/she uses it to update his/her multi-dimensional prior belief.

Chan, Narasimhan and Xie (2012) also apply a multi-attribute learning model to the drug market. They assume experience signals are correlated on the two dimensions of side-effects and effectiveness. They achieve identification by supplementing revealed preference data with data on self-reported reasons for switching: side-effects or ineffectiveness. Interestingly, they find detailing visits are more effective in reducing uncertainty about effectiveness than side-effects.

3.2. More Sophisticated Learning Models with Forward-looking Consumers

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.³⁰ 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.³¹ This is quite different from the information mechanism in Erdem-Keane, where ad

²⁹ Hendricks and Sorensen (2009) use a similar idea of information spillover to explain the skewness of music CD sales. They find evidence that a successful new album release by an artist increases the likelihood that consumers purchase older albums of the same artist.

³⁰ The separate identification of informative and persuasive effects of advertising relies on this qualitative implication of learning models: As consumers gather more information over time, the marginal benefits of informative advertising must fall; therefore, if advertising has any impact on brand choice in the long run, it is due to persuasive advertising. To our knowledge, Leffler (1981) is the first paper that proposes this identification strategy. He implements it in a reduced form model using product level sales data for new and old prescription drugs. Narayanan et al. (2005) make use of the same identification argument when estimating their structural model using product level data. Recently, Ching and Ishihara (2012) propose a new identification strategy to attack this problem – they argue that informative advertising should affect all products that share the same features/ingredients equally, but persuasive advertising should be brand specific. Ching and Ishihara implement their identification strategy in a prescription drug market where some drugs are made of the same chemical, but with different brand-names.

³¹ 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).

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, Akerberg (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 an application to prescription drugs, they argue that panel data allows them to identify two effects: (i) symptomatic effects, 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 (although they abstract from correlated learning). They also endogenize length of treatment by allowing patients to recover. Their estimates imply substantial patient heterogeneity in drug efficacy. They go on to study the welfare cost of uncertainty relative to the first-best environment with no uncertainty. Welfare questions cannot be addressed without a structural model. However, after estimating their model, Crawford and Shum can simulate removal of uncertainty by setting the initial prior variance to zero, and setting each consumer's prior match value to be the true match value. By conducting this experiment, they find that consumer learning allows consumers to dramatically reduce the costs of uncertainty.

Erdem, Keane and Sun (2008) was the first paper to model the quality signaling role of price in the context of frequently purchased goods. They also 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 Section 4.

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%.³²

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.³³

In their application, Erdem et al. (2005) consider technology adoption (Apple/Mac vs. 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 special aspect of high-tech durables is rapid technical progress. This causes the price of PCs to fall rapidly over time. Thus, there are two incentives to delay purchase: (i) to get a better price, and (ii) to search for more information about the technologies. Delay, 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 results indicate that consumers defer purchases both to gather more information and to get a better price. But, perhaps surprisingly, simulations of their model imply that learning is the more important reason for purchase delay.

Another way for consumers to learn is by observing other consumers' choices, instead of their opinions, i.e., observational learning (Banerjee, 1992). To capture this idea, Zhang (2010)

³² 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. It is worth noting that Narayanan and Manchanda (2009) also estimate a learning model with continuous distribution of consumer types. But to estimate their model, they need to assume agents are myopic.

³³ Chintagunta et al. (2012) model how physicians learn about drugs using both patients' experiences and detailing.

develops a new product adoption model with observational learning. In her model, consumers are forward-looking and heterogeneous. One interesting implication of her model is that observational learning would lead to slower product adoption compared with full informational sharing (i.e., all experience signals are common knowledge). She estimates her model using data from the U.S. kidney market, and by using a counterfactual experiment, quantifies the extent of inefficiency caused by observational learning (compared with full-informational sharing).³⁴

Che, Erdem and Öncü (2011) develop a forward-looking consumer brand choice model with spillover effects in learning and changing consumer needs over time. This is the first paper to model correlated learning with forward-looking consumers. 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.

Finally, Dickstein (2012) considers a model in which forward-looking physicians are uncertain about patients' intrinsic preferences for *multiple* drug attributes.³⁵ This is the first model of forward-looking agents that allows for information spillover across *alternatives*. A physician uses patients' utility of consuming a drug at time t to update his/her belief about patients' preferences parameters. The Bayesian updating procedure for physicians is similar to Bayesian inference in a linear regression model. An interesting implication of this approach is that, after seeing a negative outcome of drug A, physicians may want to avoid other drugs that share some of the attributes of drug A.³⁶

3.3. Modeling Consumer Learning using Product Level Data

The estimation technique developed by Berry, Levinsohn and Pakes (1995) (BLP) led to a large body of demand analysis that applies static discrete choice models primarily to product

³⁴ Cai et al. (2009) provide interesting evidence for observational learning by studying a natural experiment where customers of a restaurant are given a ranking of some popular dishes.

³⁵ Learning about preferences may appear to be different from learning about attributes. But the two are equivalent as long as one assumes the utility function is linear in attributes and preference weights.

³⁶ To estimate his model, Dickstein uses the Gittin's index approach (Gittins and Jones, 1979). But, as in Eckstein et al (1988), who also use that approach, he needs to assume consumers are risk neutral. It is worth noting that Ferreyra and Kosenok (2010) use a method similar to Gittin's index to estimate a simpler dynamic learning problem.

level or market share data.³⁷ In general, however, BLP cannot be used to estimate the demand systems generated by consumer learning models. Learning models are always dynamic in that current sales affect future demand, regardless of whether consumers are forward-looking or myopic. Demand for one brand in a learning model depends on the whole distribution (across consumers) of perceived quality for all brands. We are skeptical about whether individual heterogeneity distributions can be *credibly* identified from aggregate (i.e., product level) data. Furthermore, it is difficult to combine such a complex demand system with a supply side model.

If one wants to estimate a dynamic demand system with consumer learning, and only market share data is available, it is clear that one has to abstract from the endogenous consumer heterogeneity generated by individual level purchase histories (see section 2.1, Equation (5)). To address this issue, Narayanan, Manchanda and Chintagunta (2005) and Ching (2000, 2010a) propose two related modifications of the EK framework. Narayanan et al. (2005) assume every agent has an identical purchase history, i.e., for each brand, the quantity purchased is equally distributed across agents in each period. Ching (2000, 2010a) assumes consumers can learn from each other's experiences via social networks or information gathering institutions (e.g., physician networks). As a result, consumers use the same set of experience signals to update their beliefs, and all consumers share a common belief at any point of time.³⁸ This assumption eliminates the distribution of consumers across different states as the state variables for firms. Both Narayanan et al. (2005) and Ching (2000, 2010a) capture consumer learning in a parsimonious way, so that, when combined with an oligopolistic supply-side, the size of the state space is manageable. Both papers apply their frameworks to study the demand for prescription drugs.^{39,40}

As a demand system with consumer learning is always dynamic, we would expect firms to be forward-looking when choosing their marketing-mix. As a first attempt to address this issue, Ching (2010b) extends Ching (2010a) by combining a social learning demand model with a dynamic oligopolistic supply side model. As far as we know, this is the first empirical paper to

³⁷ Akerberg et al. (2007) provides an excellent survey of this area. Note that product or market level data are more readily available than scanner data for many industries.

³⁸ An interesting paper that studies both across and within consumer learning is by Chintagunta, Jing and Jin (2009). They apply their model to doctors' prescribing decisions for Cox-2 inhibitors, a new class of pain killers, and find that both types of learning are important.

³⁹ Chen et al. (2012) and Moretti (2011) use a similar framework to study the impact of WOM on movie sales.

⁴⁰ The BLP estimation method can be applied to the demand model of Narayanan et al. (2005), but it cannot be applied to the model of Ching (2000, 2010a). Due to space constraints, we will not discuss the details here. Interested readers may refer to Ching (2010a) for a detailed discussion.

combine a dynamic demand system with forward-looking firms. In the model, *both* consumers and firms are uncertain about the quality of generic drugs, but they rely on the same information set to update their belief over time, and hence perceived quality and variance are common for consumers and firms. Equilibrium is Markov-perfect Nash, as in Maskin and Tirole (1988), and Ericson and Pakes (1995). The model is tailored to study the competition between brand-name and generic drugs. Ching applies his model to the market for clonidine. Simulations of the model show that it can rationalize two important stylized facts: (i) the slow diffusion of generic drugs, and (ii) the fact that brand-name firms *slowly* raise their prices after generic entry.⁴¹

Ching and Ishihara (2010) (CI) extend the model in Ching (2010a) in order to explain an important stylized fact about the drug market: the effectiveness of detailing changes when new information arrives (e.g., if a new clinical trial is positive, effectiveness of detailing increases). Standard Bayesian learning models like Erdem and Keane (1996) cannot generate this pattern as they imply the marginal impact of information signals falls over time (see Equations (17)-(18)). Thus, CI deviate from the EK framework by introducing three new features: (i) both consumers and firms are uncertain about the quality of the product; (ii) social learning takes place via an intermediary (opinion leader, consumer watch group, etc.), who updates the information set for each brand; (iii) the purpose of detailing is to build up a stock of physicians who are familiar with the most recent information set of the promoted brand.⁴² The model generates heterogeneity in information sets, as the fraction of physicians with the most up-to-date information about a brand is a function of its cumulative detailing. Using this framework, CI are able to quantify how effectiveness of detailing changes when a new clinical trial outcome is released.

Lim and Ching (2012) extend the CI framework to a multi-dimensional learning model with correlated beliefs, and apply it to study demand for the major class of anti-cholesterol drugs, statins. The CI model may be applicable to settings besides drugs where the interaction between news and informative advertising is of first-order importance. The model is parsimonious and it could, in principle, be combined with an oligopolistic supply side. But this has not yet been done.

An interesting paper by Hitsch (2006) studies firm learning about the demand for a new product. He abstracts from consumer learning by using a reduced form demand model. That is,

⁴¹ Ching's model includes heterogeneity in consumer price-sensitivity. 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. This is why they gradually increase price.

⁴² This is in contrast to standard models where advertising is modeled as a noisy signal of true mean brand quality.

unlike other papers we have discussed, consumers have no uncertainty about new products. But a firm needs to learn the true demand parameters. Hitsch considers a one-side learning equilibrium model, and he also abstracts from competition. These simplifications significantly reduce the computational burden of the estimation, and yet the model delivers important new insights to the product launch and exit problem. Finally, it is worth noting that, due to computational burden, no paper has yet estimated a model with both forward-looking consumers and firms.

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

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

Chan and Hamilton (2006) is a surprising application of an Erdem-Keane style learning model to study a clinical trial of HIV treatments. The natural experiment school of econometrics views clinical trials as the gold standard to which economics should aspire to avoid structural modeling – see Keane (2010a). But Chan and Hamilton show how a structural learning model helps to evaluate clinical trial outcomes by correcting for endogenous attrition. In their model, patients in a trial are uncertain about effectiveness and side-effects of 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, low CD4 implies a weaker immune system. A treatment that is less effective in reducing CD4 may still be preferable because it has fewer side-effects. Fernandez (2013) extends Chan and Hamilton by allowing patients to be uncertain about whether they are assigned to a treatment or control group.

Tariff choice is another area where Bayesian learning models have been useful. It is widely believed that consumers are irrational when they choose between flat-rate and per-use plans, as several studies have found that many consumers could save by switching to a per-use

option. But these papers tend to look at behavior over a short period. By using a longer 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 expectations 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 a 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 the value of certification systems (Chernew et al. 2008; Xiao 2010), TV program choice (Anand and Shachar 2011), fertility decisions (Mira 2007), spousal interactions (Yang, Zhao, Erdem and Koh 2010), bargaining problems (Watanabe 2009), a manager’s job assignment problem (Pastorino 2012), human capital investment problems (Stange 2012, Stinebrickner and Stinebrickner 2013, Hoffman and Burks 2013) and voters’ decision problems (Knight and Schiff 2010).

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

In this section we discuss limitations of existing learning models and directions for future research. In our view, the four main limitations are: (1) It is difficult to identify complex models with rich specifications of consumer behavior, (2) it is difficult to disentangle different sources

of dynamics, (3) there is no clear consensus on forward-looking vs. myopic consumers, and (4) more work is needed on how to estimate equilibrium models with consumer learning.

4.1. Identification in Behaviorally Rich Specifications

In Section 2.5.A we discussed the formal identification of learning models. 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. However, an important point (see Section 2.6), is that it is common in complex models 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 at all special to dynamic learning models, but they deserve further attention in this context.

As we noted in Section 2.5.A, in the EK model the true qualities of brands, utility weights and the price coefficient are identified from choices of households with sufficient experience of all brands so their learning has effectively ceased – so their choice behavior is stationary. In contrast, the learning parameters (risk aversion, prior uncertainty, signal variability) are pinned down by choice behavior of households with less experience (and how it differs from more experienced households). However, when one extends the basic set-up of the EK learning model, 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, or models that incorporate findings from psychology/behavioral economics. But this added complexity creates three problems: First, showing formal identification for complex models may be difficult. Second, even if formally identified, complex models may suffer from “fragile” identification in practice. Third, it may be difficult to understand what sources of variation in the data pin down certain parameters.

A particularly important issue is that, given only revealed preference (RP) data on purchase decisions and signal exposures, it may be hard to identify models with complex learning mechanisms, or to distinguish among alternative learning mechanisms (i.e., multiple mechanisms may fit the data about equally well).

A promising solution to this problem is to combine RP data with stated preference (SP) data that attempts to directly measure the learning process (also known as process data). 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 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 tends to fall a lot), it implies that information source is perceived as accurate. Another paper that combines RP and SP data to aid in identification is Shin, Misra and Horsky (2012), which is an attempt to disentangle preference heterogeneity from learning.

An alternative approach is to combine choice data with direct measures of information signals. Roberts and Urban (1988) did this in their original paper. Ching and Ishihara (2010) and Lim and Ching (2012) use results of clinical trials to measure the content of signals received by physicians, and incorporate this 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. Chintagunta et al. (2009), in a study of doctor's prescribing decisions for Cox-2 inhibitors, use patient diary data that record actual use experience.⁴³ In sum, there has been some work in this area but there is obviously much room for further progress.

4.2. 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 inventories. Thus it can be very difficult to identify which mechanism generates the state dependence we see in the data.

⁴³ One limitation of their paper 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.

Learning and inventory models generate dynamics in very different ways. Learning models generate persistence in choices (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. Inventory models, in contrast, do not generate persistence in brand choices. Rather they must assume the existence of *a priori* consumer taste heterogeneity to generate loyalty. Obviously, a great appeal of learning models is they provide a behavioral explanation for the emergence of brand loyalty.

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 if learning or inventory effects are of primary importance for explaining consumer choice behavior, or, indeed, if 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 assess the quantitative importance of each mechanism. Presumably, advances in computation will remove this barrier in the future.

Meanwhile, some authors have proposed simpler approaches to test whether learning or inventories (or both) generate choice dynamics. Ching, Erdem and Keane (2012) present a new quasi-structural approach that lets one estimate models with both learning and inventory effects, while also testing for forward-looking behavior (strategic trial). The idea is to approximate the *Emax* functions in (27) using simple functions of state variables. The learning model generates a

natural exclusion restriction: the *E*max function associated with choice of brand j contains the *updated* perception variance for brand j , while the current payoff and the *E*max functions for all brands $k \neq j$ contain the *current* perception variance for brand j . Ching et al. apply the method 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 for diapers.

Erdem, Katz and Sun (2010) propose a simple test of the relative importance of learning vs. inventories. 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., choices 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 two frequently purchased goods (ketchup and diapers).

As an alternative to nesting learning with other models of dynamics, a simple idea is to estimate a structural learning model and include a lagged choice variable in the payoff functions to capture any “left-over” state dependence in a reduced-form way. This model is identified, as lagged choice does not enter the EK learning model (only cumulative choices matter). However, it is difficult to interpret the lag coefficient. Osborne (2011) adopted this approach and called the lag coefficient “switching costs.” But there are many possible explanations, including inertia, inventories, habit persistence and recency effects in learning. Suppose the standard Bayesian model is not literally true, and consumers put extra weight on recent signals. Then a lagged purchase variable may just absorb the misspecification of the learning process. In general, we are skeptical that including non-structural elements in a learning model can be informative about the importance of learning vs. other mechanisms that generate dynamics. We believe that nesting of learning and other mechanisms, and incorporation of process data, are needed to make progress.

4.3. 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 it difficult to identify the discount factor, because the likelihood is rather flat in this parameter. For instance, in

the detergent category, 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.⁴⁴

As forward-looking models may not provide substantial fit improvements, and as they are much harder to estimate, it is not surprising that many researchers have adopted myopic models, as we saw in Section 3.1. But before taking this path, it is important to emphasize that strategic trial is the distinguishing feature of forward-looking models. In a mature category, consumers may have nearly complete information, leaving little to gain from trial purchase. 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.

Given this observation, the small likelihood improvement that EK found may not be surprising; subjective prior uncertainty is fairly low for detergent, so perceived gains from trial are small. In contrast, in a market with significant uncertainty about product attributes, the rate of trial would be higher.⁴⁵ Forward-looking models may provide a superior fit in such markets.

Thus, 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. A key 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.

Furthermore, we believe developing simple tests for forward-looking behavior (such as the test in the Ching, Erdem and Keane (2012) quasi-structural model) should be an important topic for future research. However, we also believe such tests must be derived explicitly from a theoretical model. Recently, there has been a trend whereby researchers seek to develop “model-free” tests for the assumptions of a structural model. We comment on this trend in section 4.5.

4.4. Integration of Learning Models with Supply Side Models

As we discussed in Section 3.3, there has been significant progress in developing dynamic demand systems with consumer learning that can be estimated using product level data. Examples are Ching (2000, 2010a), Ching and Ishihara (2010) and Narayanan et al. (2005). Nevertheless, there is clearly a large discontinuity between their models and those that are

⁴⁴ Other examples include Chintagunta et al. (2012), Dickstein (2012) and Yang and Ching (2010).

⁴⁵ Note that the higher the discount factor, the higher the rate of strategic trial.

applied to individual level data. All the models in Section 3.3 abstract from self-learning, and, for reasons of tractability, assume the existence of an information aggregator.⁴⁶ In our opinion, self-learning is still an important source of information for frequently purchased goods despite the advance of social networks. Moreover, none of the models in Section 3.3 allow for forward-looking consumers. Therefore, we believe that developing a richer aggregate dynamic demand system with learning remains a challenging and important area for future research.

Most of the demand analyses that use product level data are motivated by the ultimate goal of combining it with firms' problems, in order to build an equilibrium model to study the long-term outcomes of certain policy changes (e.g., advertising regulations, anti-competitive pricing regulation, merger analysis). The key challenge is how to model the firms' problem when facing such a complex dynamic demand system. The demand system generated by the EK framework (or similar models) is so complex that it is very difficult to analyze even in a monopoly situation. It is not clear if the fully rational approach to modeling firms' decisions is possible given computational costs of keeping track of such a complicated state space.

Thus, an important avenue for future research is to develop a dynamic demand system with learning that is not too costly for firms to use, yet can capture the potential forward-looking and strategic trial behavior of consumers. Hendel and Nevo (2011) take this research direction, but in the context of storable goods and not experience goods. Their demand model is motivated by the dynamic stockpiling models in Erdem, Imai and Keane (2003) and Hendel and Nevo (2006), but is much simpler to estimate, and tractable to combine with forward-looking firms. They show how their model can be used to study intertemporal price discrimination empirically.

4.5. Model-Free Evidence on the Validity of Structural Models?

Structural models in general, and learning models in particular, are often criticized on the grounds that they make a large number of assumptions (e.g., about how consumers learn and form expectations, the functional form of utility, etc.). Identification of these models relies on these functional form assumptions. Critics of structural models often argue 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 articles such as Heckman (1997), Keane (2010a,b) and Rust (2010). These authors argue that drawing inferences

⁴⁶ This is a good illustration of how structural modeling often involves a tradeoff between richness and tractability.

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. Here, instead of repeating their general arguments, we will discuss two example papers that use such “simple” approach to test learning behavior to illustrate our point.⁴⁷

Chintagunta, Goettler and Kim (2012) present reduced-form evidence of forward-looking behavior by physicians. More specifically, when a new drug is just introduced, they focus on the set of physicians who have not yet been exposed to detailing. They run a logit model to predict if a physician will prescribe the new drug to a patient. The key point is they include *future detailing* as a regressor. Say there is some risk involved in experimenting with the drug now, but future detailing is an opportunity to learn without risk. Hence, they argue, if physicians are forward-looking, then, the higher is future detailing, the less likely they are to prescribe the drug now. So a negative coefficient on future detailing suggests physicians are forward-looking.

However, this “model-free” test implicitly assumes there is no physician heterogeneity in receptivity to detailing. But it is plausible that some physicians are more skeptical about sales rep presentations, so they require more detailing to be convinced. This could cause sales reps to spend more time with less receptive physicians. Then, the coefficient on future detailing may be negative even if physicians are myopic. More generally, including a future variable in a regression is a Sims strict exogeneity test Sims (1972). It may just be that a current prescription reduces future detailing. Thus, while the test result is interesting, it is difficult to interpret.

We now turn to our second example. In an attempt to distinguish learning from other sources of state dependence such as switching costs, inertia or habit persistence, Dubé, Hitsch and Rossi (2010) estimate the following simple discrete choice model:

$$(34) \quad V_{jt} = \alpha_j - w_P P_{jt} + \gamma_1 d_{j,t-1} + \gamma_2 d_{j,t-1} \cdot N_j(t) + \gamma_3 N_j(t) + e_{jt}.$$

Note that (34) contains lagged choice, cumulative use experience, $N_j(t)$, and their interaction. So it could be viewed as a linear approximation to the more complex nonlinear form implied by the learning model (see Equations (5), (6) and (27)).

⁴⁷ Also see Ching (2013) for a critique on Moretti (2011).

Now, suppose the learning model is correct. Dubé et al argue that, for experienced consumers who have complete information, lagged choice should not be a predictor of current choice.⁴⁸ This is because, when cumulative experience is large, the additional impact of more experience on the perceived variance of a brand is trivial (see Equation (6)).⁴⁹ More generally, the fact that use experience $N_j(t)$ reduces the effect of lagged purchase implies the interaction coefficient γ_2 should be negative. However, using data on margarine and frozen orange juice, Dubé et al find that more experience does not reduce the lagged choice effect ($\gamma_2 \approx 0$). They interpret this as evidence against consumer learning.

It is tempting to treat this as a “model-free” test, as it does not impose the functional form assumptions required to estimate a fully specified learning model. But this interpretation is not correct. First, the test fails to account for a key feature of the Bayesian learning model: when $N_j(t)$ is large, so a consumer knows almost everything about brand j , any further increase in $N_j(t)$ has a negligible impact on utility. However, Equation (34) does not allow for this possibility, as $\partial V / \partial N = \gamma_3 + \gamma_2 d$, which is independent of N . Second, when $N_j(t)$ is small, the impact of $d_{j,t-1}$ can be positive or negative, depending on the realization of the experience signal relative to one’s prior. Therefore, the sign of the interaction term is ambiguous. Combining these two points, it is possible that γ_2 may be close to zero, even if there is consumer learning in the data. So again, the test result is interesting, but it is difficult to interpret.

We believe that searching for data patterns that are potentially consistent or inconsistent with a structural model is a useful exercise. It can often provide valuable insights, and can be a useful part of the process of building, validating and improving structural models. However, we do not believe that “simple models” and/or “model free” evidence can ever replace structural models or the key role of theory in empirical work more generally.

It is important to remember that truly “model free” evidence cannot exist. The “simple” empirical work that promises to deliver such evidence always relies on some assumptions. But these assumptions are often left implicit, due to failure to present an explicit model. Often these

⁴⁸ Given controls for taste heterogeneity (α_j), the lagged choice variable $d_{j,t-1}$ can matter for several reasons, such as inertia, switching costs, habit persistence, inventories or learning. So finding that lagged choice is significant for consumers with complete information may simply mean that sources of dynamics besides learning are also present.

⁴⁹ It is important to note that not all learning models imply that choice behavior becomes stationary given sufficient use experience. For instance, as we discussed in Section 3.1.1, Mehta, Rajiv and Srinivasan (2004) extend the basic model to allow forgetting. It is also possible that product attributes change over time. Thus, it is conceptually straightforward to construct learning models where recent experience is more salient for a variety of reasons.

implicit assumptions are (i) not obvious, (ii) hard to understand and (iii) very strong. One of the main contributions of structural learning models to marketing science has been to generate far more interest in the structural paradigm. We hope this will be a long term trend, regardless of future evaluations of the usefulness of the learning model *per se*.

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 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.

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.

Finally, we point out that integrating learning models of demand with supply side models remains under-explored and should be another important area for future research.

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.
- Akerberg, D., L. Benkard, S. Berry and A. Pakes (2007) Econometric tools for analyzing market outcomes. Chapter 63 in the *Handbook of Econometrics*, Vol. 6A, J.J. Heckman and E. Leamer (eds), North Holland Press.
- Aguirregebirria, 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.
- Banerjee, A.V. (1992) A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3): 797-817.
- Berry, S., J. Levinsohn, and A. Pakes (1995) Automobile prices in market equilibrium. *Econometrica*, 63(4): 841-890.
- 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.
- Cao, H., Y. Chen, H. Fang (2009) Observational learning: Evidence from a Randomized Natural Field Experiment. *American Economic Review*, 99(3): 864-882.
- 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. Forthcoming in *Management Science*.
- 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 (2012) Learning about movies: The impact of movie release types on the nationwide box office. *Journal of Cultural Economics*, published online in October 2012.
- 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.
- Chevalier, J. and A. Goolsbee (2009) Are Durable Goods Consumers Forward-Looking? Evidence from College Textbooks? *Quarterly Journal of Economics*, 124 (4), 1853-1884.

- Ching, A.T. (2000) Dynamic Equilibrium in the U.S. Prescription Drug Market after Patent Expiration. Ph.D. dissertation, University of Minnesota.
- Ching, A.T. (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.T. (2010b) A dynamic oligopoly structural model for the prescription drug market after patent expiration. *International Economic Review*, 51(4): 1175-1207.
- Ching, A.T. (2013) Comments on: "Social learning and peer effects in consumption: evidence from movie sales" by E. Moretti. Working paper, Rotman School of Management, University of Toronto.
- 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. and M. Ishihara (2012) Measuring the informative and persuasive roles of detailing on prescribing decisions. *Management Science*, 58(7):1374-1387.
- Ching, A.T., 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., T. Erdem and M. Keane (2012) A simple approach to estimate the roles of learning and inventories in consumer choice. Working paper, Rotman School of Management, University of Toronto.
- Ching, A.T., 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. *Journal of Marketing Research*, 49(6): 807-821.
- Chung, Doug, Thomas Steenburgh and K. Sudhir (2013) Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans. Working Paper, Harvard Business School.

- 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.
- Deighton, J. (1988) The interaction of advertising and evidence. *Journal of Consumer Research*, 15: 262-264.
- 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.
- Ericson R. and A. Pakes (1995) Markov-perfect industry dynamics: a framework for empirical work. *Review of Economic Studies*, 62: 53-82.
- 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.

Fernandez, J.M. (2013) An empirical model of learning under ambiguity: The case of clinical trials. *International Economic Review*, 54(2): 549-573.

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 (2000) Bayesian Inference for Dynamic Discrete Choice Models without the Need for Dynamic Programming. In Mariano, Schuermann, and Weeks (eds.), *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.

Gittins, J.C. and D.M. Jones (1979) A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, 66: 771-784.

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.

Guadagni, P. and J. Little (1983) A Logit Model of Brand Choice Calibrated on Scanner Data. *Marketing Science*, 2(3): 203-238.

Heckman, J.J. (1981) Heterogeneity and State Dependence. In S. Rosen (ed.), *Studies in Labor Markets*: 91-140.

Heckman, J.J. (1997) Instrumental Variables: A Study of Implicit Behavioral Assumptions Used in Making Program Evaluations. *Journal of Human Resources*, 32(3): 441-462.

Hendel, I. and A. Nevo (2006) Measuring the Implications of Sales and Consumer Inventory Behavior. *Econometrica*, 74(6): 1637-73.

Hendel, I. and A. Nevo (2011) Intertemporal Price Discrimination in Storable Goods Markets. Forthcoming in *American Economic Review*.

Hendricks, K. and A. Sorensen (2009) Information and the skewness of music sales. *Journal of Political Economy*, 117(2): 324-369.

Hitsch, G. (2006) An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty. *Marketing Science*, 25(1): 25-50.

Hoffman, M. and S. Burks (2013) Training contracts, worker overconfidence, and the provision of firm-sponsored general training. Working paper.

Houser, D. (2003) Bayesian analysis of a dynamic stochastic model of labor supply and saving. *Journal of Econometrics*, 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*, 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.
- Koopmans, T.C., H. Rubin and R.B. Leipnik (1950). Measuring the Equation Systems of Dynamic Economics. Cowles Commission Monograph No. 10: Statistical Inference in Dynamic Economic Models, T.C. Koopmans (ed.), John Wiley & Sons, New York.
- Leffler, K. (1981) Persuasion or Information? The Economics of Prescription Drug Advertising. *Journal of Law and Economics*, 24: 45-74.
- 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.
- 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.
- Matzkin, R.L. (2007) Nonparametric Identification. Chapter 73 in the *Handbook of Econometrics*, Vol. 6B, J.J. Heckman and E. Leamer (eds), North Holland Press: 5307-5368.
- 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.
- Moretti, E. (2011) Social learning and peer effects in consumption: evidence from movie sales. *Review of Economic Studies*, 78: 356-393.
- 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) Notes and Comments: 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. (1984) Structural Estimation of Markov Decision Processes. *Handbook of Econometrics*: Vol. 4, R.F. Engle and D.L. McFadden (eds.), Elsevier Science B.V., 3081-3143.
- 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.
- Sims, C. (1972) Money, Income, Causality. *American Economic Review*, 62(4): 540-552.

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.

Stange, K. (2012) An empirical investigation of the option value of college enrollment. *American Economic Journal: Applied Economics* 4: 49-84.

Stinebrickner, T. and R. Stinebrickner (2013) Academic performance and college dropout decision. Forthcoming in *Journal of Labor Economics*.

Watanabe, Y. (2009) 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, B. and A.T. Ching (2010) Dynamics of Consumer Adoption of Financial Innovation: The Case of ATM Cards. Working paper, University of Toronto.

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.

Yao, S., C. F. Mela, J. Chiang and Yuxin Chen (2012) Determining Consumers' Discount Rates with Field Studies. *Journal of Marketing Research*, 49 (6), 822-841.

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. Helsén (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.