

A Simple Method to Estimate the Roles of Learning, Inventories and Category Consideration in Consumer Choice*

by

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

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

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

This draft: June 10, 2014

Abstract: Models of consumer learning and inventory behavior have both proven to be valuable for explaining consumer choice dynamics. In their pure form these models assume consumers solve complex dynamic programming (DP) problems to determine optimal choices. For this reason, these models are best viewed as “as if” approximations to consumer behavior. In this paper we present an estimation method, based on Geweke and Keane (2000), which allows us to estimate dynamic models without solving a DP problem and without strong assumptions about how consumers form expectations about the future. The relatively low computational burden of this method allows us to nest the learning and inventory models. We also incorporate the “price consideration” mechanism of Ching, Erdem and Keane (2009), which essentially says that consumers may not pay attention to a category in every period. The resulting model may be viewed as providing a more “realistic” or “descriptive” account of consumer choice behavior.

* Ching’s work on this project is supported by the SSHRC. Keane’s work was supported by Australian Research Council grants FF0561843 and FL110100247. We thank Masakazu Ishihara for helpful discussion.

1. Introduction

The goal of this paper is develop and estimate a single choice model that incorporates learning, inventories and category consideration. Prior literature has considered these aspects of behavior separately but not jointly. Early papers on consumer learning were Erdem and Keane (1996), Ackerberg (2003) and Crawford and Shum (2006); for inventory behavior see Erdem, Imai and Keane (2003) and Hendel and Nevo (2006), and for category consideration see Ching, Erdem and Keane (2009).¹ Key open questions are whether learning or inventories provide a better explanation of consumer choice dynamics (see Erdem, Keane and Sun (2008)), and whether the category consideration mechanism is also an important aspect of choice behavior.

One approach to this question is to develop a structural model that incorporates all three mechanisms (learning inventories, consideration), and to test which are most important for explaining choice dynamics. Unfortunately, both learning and inventories lead to models that are difficult to estimate, as they require solution of dynamic programming (DP) problems.² This computational burden is greatly compounded if we combine learning and inventories into one model. There are simply too many state variables – levels of perceived quality and uncertainty for all brands, current and lagged prices of all brands, and inventories – making the DP problem very time-consuming. Thus, it has not been feasible to estimate dynamic structural models that include both mechanisms (let alone to also include category consideration).

In this paper, we take a step towards addressing this issue. We present a new approach that allows one to estimate models with both learning and inventory effects, while also including category consideration. Specifically, we extend a method developed by Geweke and Keane (2000) that makes it possible to estimate dynamic models without having to solve the DP problem. By using the Geweke-Keane method, we can, for the first time, estimate a model with both learning and inventories, and shed light on the role of each. Our approach also allows us to relax many of the strong assumptions about choice behavior that underlie the conventional DP

¹ Seiler (2013) extends Ching, Erdem and Keane (2009) by introducing category consideration in an inventory stockpiling model.

² Dynamic programming is needed to model behavior of forward-looking consumers. Consumers in an inventory model must be forward-looking, as *there is no reason that myopic consumers would hold inventories*. On the other hand, consumers in a learning model need not be forward-looking. If they are forward-looking they will engage in strategic trial or experimentation, as in Eckstein et al (1988). If they are myopic they still learn over time, updating their priors as more information is received, as in Roberts and Urban (1988). But they do so passively, with no strategic trial. But to nest learning and inventories within one model, we must allow for forward-looking consumers. Otherwise the inventory model would be ruled out *a priori*.

approach. We can also test whether consumers engage in strategic trial as the forward-looking learning model implies.

Our empirical application is to the demand for diapers. This category is ideal for studying learning and inventory behavior. This is because there is a well-defined point when a consumer enters the market – the birth of a child. In contrast, in most categories (e.g., detergent, cereal) consumers will usually have been in the market for years before we first observe them.

The outline of the paper is as follows: Section 2 describes the standard Bayesian learning model. Section 3 describes the Geweke-Keane method. Sections 4-7 shows how we apply it to a model that contains learning, inventories and category consideration, and how we test for forward-looking behavior. Section 8 presents the estimation results. Section 9 concludes.

2. The Basic Bayesian Learning Model

Here we describe a standard learning model, similar to Erdem and Keane (1996). The key feature of the model is that consumers do not know the attributes of brands with certainty. Before receiving any information via use experience, consumers have a normal prior on brand quality:

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

This says that, prior to any use experience, consumers perceive that the true quality of brand j (Q_j) is distributed normally with a mean of Q_{j0} and variance σ_{j0}^2 . The values of Q_{j0} and σ_{j0}^2 may be influenced by many factors, such as reputation of the manufacturer, advice from friends, etc.

Use experience does not fully reveal quality because of “inherent product variability.” This has multiple interpretations. First, quality of different units of a product may vary. Second, and more importantly, a consumer’s experience of a product will vary across use occasions. For instance, a diaper may hold all of a baby’s urine on some occasions but not on others (depending on how much milk the baby drank). And a baby would take some time to develop a diaper rash.

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

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

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

variability.” It should be noted that all brands have experience variability, so (2) holds for all j .

Note that we have conjugate priors and signals, as both the prior on quality (1) and the noise in the quality signals (2) are assumed to be normal. The posterior for perceived quality at $t=1$, after a single use experience signal is received, is given by the updating formulas:

$$(3) \quad Q_{ij1} = \frac{\sigma_{j0}^2}{\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{ij1}^E + \frac{\sigma_\varepsilon^2}{\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j0}$$

$$(4) \quad \sigma_{ij1}^2 = \frac{1}{(1/\sigma_{j0}^2) + (1/\sigma_\varepsilon^2)}$$

Equation (3) describes how a consumer’s prior on quality of brand j is updated as a result of the experience signal Q_{ij1}^E . Note that the extent of updating is greater the more accurate is the signal (i.e., the smaller is σ_ε^2). Equation (4) describes how a consumer’s uncertainty declines as he/she receives more signals. σ_{ijt}^2 is often referred to as the “perception error variance.”

Equations (3) and (4) generalize to any number of signals. Let $N_{ij}(t)$ denote the total number of use experience signals received by person i through time t . Then we have that:

$$(5) \quad Q_{ijt} = \frac{N_{ij}(t)\sigma_{j0}^2}{N_{ij}(t)\sigma_{j0}^2 + \sigma_\varepsilon^2} \sum_{s=1}^t Q_{ijs}^E d_{ijs} + \frac{\sigma_\varepsilon^2}{N_{ij}(t)\sigma_{j0}^2 + \sigma_\varepsilon^2} Q_{j0}$$

$$(6) \quad \sigma_{ijt}^2 = \frac{1}{(1/\sigma_{j0}^2) + N_{ij}(t)(1/\sigma_\varepsilon^2)}$$

where d_{ijt} is an indicator for whether brand j is bought/consumed at time t by person i .

In (5), the perceived quality of brand j at time t , Q_{ijt} , is a weighted average of the prior and all quality signals received up through time t , $\sum_{s=1}^t Q_{ijs}^E d_{ijs}$. Perceived quality is random across consumers, as some receive, by chance, better quality signals than others. So the learning model endogenously generates heterogeneity across consumers in their perceptions of products. We will also include advertising as a signal of quality. We discuss this extension in Section 6.

Let S_{it} denote consumer i ’s information set. 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 S_{it} may persist over time, because: (i) both brands and consumers are finitely lived, (ii) as people gather more information the value of trial purchases diminishes, and so eventually learning about unfamiliar products will

become slow; (iii) there is a flow of new brands and new consumers entering a market.

We assume consumer i 's (conditional indirect) utility of consuming brand j is:

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

Where P_{ijt} is price, w_P is marginal utility of income and e_{ijt} is an idiosyncratic brand, time and person specific error, distributed *iid* extreme value. If we assume that $f(Q_{ijt}^E)$ takes the constant absolute risk aversion form (CARA), then the expected utility is given by:

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

where $r > 0$ captures risk aversion with respect to variation in product quality. There is also a no purchase option, whose expected utility is simply $E[U_{i0}|S_{it}] = e_{i0t}$ (see Section 7 for details).

3. An Overview of the Geweke-Keane Method for Estimating Dynamic Models

In general, learning models like the one in Section 2 must be solved by dynamic programming (DP). This is because today's purchase affects tomorrow's information set, which affects future utility. Consumers choose a brand that gives them the highest expected *current* utility *plus* expected discounted *future* payoffs; that is, the highest value function:

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

where the next period information set, or state space, S_{t+1} is given by:

$$(9b) \quad S_{t+1} = \{Q_{1,t+1}, \dots, Q_{J,t+1}, \sigma_{1,t+1}^2, \dots, \sigma_{J,t+1}^2\}.$$

Equation (9a) is called the “alternative specific” value function, as it is the value of choosing option j . As we see in (9a), choosing the brand with the highest expected *current* utility is not necessarily optimal, because the “future component” $\beta EV(S_{t+1}|S_t, j)$ is also relevant.

Notice that in (9a)-(9b) we have suppressed the consumer subscript i to simplify notation. We will continue to do so in the remainder of this section, as well as in Sections 4-5.

A key limitation of existing learning models is they ignore inventories, which are another key source of dynamics in consumer choice behavior. It is not feasible to estimate a model of forward-looking agents with both learning and inventories, as the state space is even larger than

in (9b), making it intractably large. Here we propose a potential solution to this problem.

Geweke and Keane (2000, 2001) present a method to estimate dynamic models without the need to solve agents' DP problem. The basic idea is to replace equation (9) by:

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

Here $F[S_{t+1}(S_t, j)|\pi_t] \approx \beta EV(S_{t+1}|S_t, j)$ is a flexible function of the state variables designed to approximate the “future component” of the value function. And π_t is a vector of reduced form parameters that characterize the $F(\cdot)$ function. The idea of the Geweke-Keane (GK) method is to estimate the π_t jointly with the structural parameters that enter the current period utility function.

The Bayesian learning model implies a particular mapping from the structural parameters to the reduced form parameters π_t . However, if the π_t are identified, one can estimate them freely, without imposing the structural restrictions. In the GK method, they are estimated freely. Thus, the estimated π_t may approximate the optimal DP solution, but they may imply other types of behavior, like myopia ($\pi_t=0$), or failure to form expectations optimally. Thus we can use the estimated π_t to test assumptions about how consumers behave (as in Houser, Keane and McCabe (2004)). In particular, we will show how to test if strategic trial is important.

The intuition for the GK method is similar to Hotz and Miller (1993). Consumers solve their own problem, so we can infer the parameters of the solution (π_t) from their behavior, without having to solve the problem ourselves. The key difference is that the GK method assumes only that both the econometrician and the consumers have common knowledge about the laws of motion of the state variables. Unlike Hotz and Miller (1993), GK do not assume the econometrician knows how consumers form expectations (which would be necessary to form the mapping from the structural parameters to the π_t). Hence GK do not assume we know *a priori* exactly what problem the consumers are actually solving. Rather, we make inferences about their solution (i.e., their behavior) from the π_t estimates.

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

we do not observe the current payoffs (i.e., we observe only choices, not utility).

Here, we show that an alternative method of identification is possible if some state variables enter the future component of the value functions but not current payoffs. These state variables must also vary across the equations of the discrete choice model, thus generating a cross-equation exclusion restriction. As we'll see, this is true of the (updated) perception error variances in the Bayesian learning model.

4. Estimating the Bayesian Learning Model using the Geweke-Keane Method

In this section we show how to estimate the Erdem and Keane (1996) learning model using the GK method. (We add inventories in section 5). First, write the future component of the value function in the Bayesian learning model as follows:

$$(11) \quad \beta E_t V(I_{t+1}|I_t, j) = F(Q_{1t}, \dots, Q_{Jt}, \sigma_{1,t+1}^2, \dots, \sigma_{J,t+1}^2 | \pi_t)$$

We obtain (11) from (10) by specifying that the state variables in I_{t+1} are $\{Q_{jt}\}$ and $\{\sigma_{j,t+1}^2\}$. The following point is important: The σ_{jt}^2 evolve *deterministically*, so a consumer at time t knows what their values will be at $t+1$. Thus, the expectation at time t conditions on the $\{\sigma_{j,t+1}^2\}$. In contrast, the $\{Q_{jt}\}$ are rational forecasts of quality, given all current information. A consumer cannot anticipate changes in the $\{Q_{jt}\}$, so the expected value depends only on the current value.³

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

$$(12) \quad \beta E_t V(I_{t+1}|I_t, d_{jt} = 1) = F(Q_{1t}, \dots, Q_{Jt}, \sigma_{1t}^2, \dots, \sigma_{j,t+1}^2, \dots, \sigma_{Jt}^2 | \pi_t)$$

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

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

³ Informally, a rational consumer cannot *expect* to discover (through further information gathering) that a brand is better or worse than he/she thinks it is. Formally, the law of iterated expectations applies to rational forecasts.

option as the base alternative. The future component associated with no-purchase has the form:

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

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

For example, assume the future component is approximated as a polynomial in the state variables. Then, if we take the difference between (12) and (13), any terms that do not involve the change in σ_{jt}^2 will drop out.⁴ For instance, suppose we have:

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

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

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

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

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

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

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

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

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

⁴ A key point is that the $F(\cdot)$ function does not differ across alternatives in its parameters, only in its arguments. This is clear from comparison of (9) and (10), and is reflected in the notation in (11)-(13), i.e., F has no j subscript.

It is worth emphasizing that we only used the laws of motion of the state variable to derive (18), along with the assumed form of the approximating function $F(\cdot)$. We have not assumed anything about how consumers actually solve the DP problem (i.e. how they form expectations $E_t V(\cdot | \cdot)$).

Of course we may choose different forms for the function F in (14). We should choose a form based on fit to the data, as well as considerations of parsimony and interpretability. For instance, we could complicate (16) by specifying that $P_3(\cdot)$ involves higher order terms in Q_{jt} and/or σ_{jt}^2 . We could also simplify (16) by assuming that $P_3(\cdot)$ does not involve quality levels Q_{kt} for brands other than j . This would give us $P_3(Q_t, \sigma_{jt}^2) = (\pi_{0j} + \pi_{1j} Q_{jt}) \cdot \sigma_{jt}^2$ and hence:

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

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

First, the GK method has never been applied to learning models, so it seems best to start with a parsimonious specification (to avoid any risk of identification problems). Second, we find that even the simple model in (19) leads to an improvement in fit over a myopic model. Third, equation (19) is intuitive and easily interpretable: Parameter π_{0j} captures the value of information gained through a trial purchase of brand j . Of course, this value may differ across brands, a point we examine in the empirical work.⁵ Parameter π_{1j} captures that this trial value may differ depending on the perceived quality of a brand. The sign of π_{1j} is not clear *a priori*.

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

5. Incorporating Inventories in the Basic Learning Model

Next we include inventory behavior in the model in a simple way. Let the law of motion of inventories be $I_{t+1} = I_t + N_t - R_t$, where $N_t \geq 0$ is the number of units bought at time t , and R_t is the product usage rate. In the event of a stock out, $[I_t + N_t < R_t]$, we have $I_{t+1}=0$. Next, add stock out costs (c_s) and inventory carrying costs (cc) to the expected current payoff in equation (8). Specifically, add the terms $-cc(I_t + N_t) - c_s P[I_t + N_t < R_t]$, where $P[I_t + N_t < R_t]$ is the probability

⁵ There are many reasons why the value of trial information may differ across brands. For example, say brand 2 has a much lower market share than brand 1 (perhaps because it is priced high relative to its quality). In that case, information about brand 2 may be almost irrelevant for the value function compared to information about brand 1.

of a stock out. We assume the number of units purchased (N_t) is the same regardless of which brand is bought, which is a fairly common assumption in discrete/continuous choice modeling.

Consumers in an inventory model must be forward-looking, as *there is no reason that myopic consumers would hold inventories*. So next we add inventories I_{t+1} to the state space in (9b). We could also add variables that help predict future prices, as in Erdem, Imai and Keane (2003). But for the sake of simplicity we assume that prices are serially uncorrelated (aside from mean differences across brands). So we now have $S_{t+1} = \{Q_{1,t+1}, \dots, Q_{J,t+1}, \sigma_{1,t+1}^2, \dots, \sigma_{J,t+1}^2, I_{t+1}\}$.

The intuition of the inventory model is that: (i) the good is storable, so consumers try to time purchases for when price is low, and (ii) consumers hold inventory to avoid stock outs. For stock out risk to exist, consumption must occur *after* the purchase decision, and R_t must be stochastic. Then, reservation prices rise as inventory drops, because one tries to avoid stock outs.

An important point is that the expected value function $EV(S_{t+1}|S_t, j)$ in the inventory model is a function of $I_t + N_t$, not I_{t+1} . This is because R_t is not known at the time the purchase decision is made, so it must be integrated out when the expectation is formed. As we have seen, the quantity $I_t + N_t$ also determines the carrying cost and stock out risk that enter *current* payoffs.

Thus, in contrast to the learning model, the inventory model does not generate exclusion restrictions whereby some quantities appear in the future component of the value function but not in current payoffs. This means the structural parameters of the inventory model, c_s and cc , are not identified separately from the reduced form parameters π related to inventories in the future component approximation. We now describe how we add inventory to the future component.

Ideally, we would like to add inventory to the state space in (9b). Unfortunately, as is typical, we do not observe inventory directly. So we proxy for inventory using weeks since last purchase, henceforth $pgap_t$. The law of motion for $pgap_t$ is obvious: If one chooses no purchase, $pgap$ is increased to $pgap_{t+1} = pgap_t + 1$. But if one buys any brand at t it is reset to $pgap_{t+1} = 1$. Thus, the value function difference, (19), is simply augmented to include $pgap_t$ itself:

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

Note that π_{pg} is common across brands as $pgap_t$ shifts all their values relative to no purchase.

It worth emphasizing that $pgap_t$ proxies for the state variable $I_t + N_t$, which, as discussed earlier, shifts both the future component and current payoffs. Thus, the parameter π_{pg} combines

effects of inventories on current payoffs and $EV(S_{t+1}|S_t, j)$. So π_{pg} is a reduced form parameter, and its sign is ambiguous. Concern about stock-outs when the inventory is low will tend to make π_{pg} positive, but the storage/carrying costs associated with high inventory push it in a negative direction. If the former effect dominates, so that purchase probability increases with time since last purchase, we would have $\pi_{pg} > 0$. But as we'll see, once we introduce category consideration the situation is more complex: inventory may have different effects on the consideration probability and the probability of purchase *conditional* on consideration.

6. Adding the “Price Consideration” Mechanism to the Basic Model

Next, we incorporate the “Price Consideration” (PC) mechanism of Ching, Erdem and Keane (2009). This generalizes conventional choice models by assuming consumers may not consider a category in every period. For example, after buying and replenishing inventory, consumers may not look at a category again for a few weeks. During this interlude, they may not respond even to deep price discounts, simply because they do not see them. Or, some consumers may only consider a category when it is on display or promotion. Ching et al (2009) find the PC model gives a much better fit to observed purchase hazard rates – especially for short inter-purchase spells – than traditional models where consumers consider a category every period.

In each period t , the PC model has two stages. In stage one, consumers decide whether to consider a category. If they consider, then, in stage two, they choose an option $j=0, \dots, J$ using the model in (20). We model consideration using a simple logit (suppressing household i subscripts):

$$(21) \quad U_{ct} = \gamma_0 + \gamma_{pg}pgap_t + \gamma_{te}t_t^e + \gamma_x X + \gamma_{ft} \sum_j ft_{jt} + \gamma_{dp} \sum_j dp_{jt}$$

Factors that affect consideration include time since last purchase ($pgap$) and time since the household first entered the diaper market (t_t^e).⁶ Vector X contains factors that may shift the usage rate (i.e., indicators for a household being young, high income or low education). Marketing activity may also cause a household to consider the category; ft_{jt} and dp_{jt} are indicators for if brand j is on feature or display at time t , respectively. These are summed over brands, so marketing activity for the whole category shifts the consideration probability. Note that if $\gamma_{pg} > 0$ the interpretation is that low inventory makes consumers more likely to consider a category. Our

⁶ Time in the market may affect the consideration probability because the usage rate for diapers falls as a child ages.

estimates of γ_{pg} and π_{pg} will shed some light on whether inventories matter primarily through their effect the consideration probability or on the *conditional* purchase probability.

7. Details of the Specification

To complete the model, we need to select a functional form for the expected utility function in (9a).⁷ We assume an augmented version of the CARA form in equation (8):

$$(22) \quad E[U(Q_j^E, P_j)|I_t] = -\exp\left(-r\left(Q_{jt} - \frac{r}{2}(\sigma_{jt}^2 + \sigma_\varepsilon^2)\right)\right) + \alpha_P P_{jt} \\ + \alpha_f f_{jt} + \alpha_d dp_{jt} + \alpha_a ad_{ijt} + e_{jt}$$

Equation (22) generalizes (8) by allowing for direct persuasive effects of advertising. We do this by letting feature (f_{jt}), display (dp_{jt}) and ad expenditure (ad_{ijt}) affect utility directly. The literature on advertising typically views it as affecting demand either through information (e.g., learning) or persuasion; i.e., a direct effect on utility, as in Nerlove and Arrow (1962). We allow for direct persuasive effects of ads to allay any concern that we only find learning is important only because we omit the persuasion channel.⁸ Also, note that in (22) a brand's utility depends on its *own* advertising, while in (21) category consideration depends on *total* advertising activity of the category. This exclusion restriction helps distinguish consideration from promotion effects.

Recall that that expected utility from no-purchase is simply $E[U_{i0t}|I_{it}] = e_{i0t}$, where e_{i0t} is extreme value. No-purchase contains no intercept. This is an identification assumption, as all other utilities are measured relative to no-purchase. Also, the data contains two very small brands (labeled 5 and 6). We model utility from these as just a constant plus an extreme value error.

As we discussed in Section 2, the learning aspect of our model is based on Erdem and Keane (1996). In addition to use experience, they also let advertising exposures signal brand quality. However, unlike EK, we do not observe ad exposure at the *individual* level. So instead, we use ad expenditure to proxy for the volume of ads received by consumers.

Specifically, let A_{ijt} denote the advertising signal received by person i for brand j in period t . Analogous to equation (2), we assume:

⁷ For simplicity, timing of consumption is not modeled. The use experience signal arrives at the time of purchase.

⁸ As we discuss in Ching, Erdem and Keane (2013), we don't think this is the correct way to estimate the importance of the persuasion channel, as the feature, display and ad expenditure that we include in the utility function may be significant for other reasons (e.g., they might "sop up" misspecification of the inventory mechanism).

$$(23) \quad A_{ijt} = Q_j + \varepsilon_{ijt}^A \quad \text{where} \quad \varepsilon_{ijt}^A \sim N(0, \sigma_{A_{ijt}}^2) \quad \text{and} \quad \sigma_{A_{ijt}}^2 = \sigma_A^2 / ad_{ijt}$$

As in EK, ad signals A_{ijt} provide unbiased but noisy information about brand quality.⁹ The noise has variance $\sigma_{A_{ijt}}^2$. In our Bayesian updating framework, the information in N ads, each with variance σ_A^2 , is equivalent to that in a one ad with variance σ_A^2/N . We let $\sigma_{A_{ijt}}^2 = \sigma_A^2/ad_{ijt}$ where ad_{ijt} is advertising *expenditure* by brand j to consumer i at time t . This equation embeds an assumption about how ad spending maps into an effective number of ads seen (i.e., $N \propto ad_{ijt}$). Finally, note that consumers from two different cities will have different exposure to advertising.

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

$$(24) \quad Q_{jt} = Q_{jt-1} + d_{jt} \beta_{1jt} (Q_{jt}^E - Q_{jt-1}) + d_{jt}^A \beta_{2jt} (A_{jt} - Q_{jt-1}),$$

where $\beta_{1jt} = \sigma_{jt}^2 / (\sigma_{jt}^2 + \sigma_\varepsilon^2)$ and $\beta_{2jt} = \sigma_{jt}^2 / (\sigma_{jt}^2 + \sigma_{A_{ijt}}^2)$, and, for the perceived variance:

$$(25) \quad \sigma_{jt}^2 = \frac{1}{(1/\sigma_{jt-1}^2) + (1/\sigma_\varepsilon^2) + (1/\sigma_{A_{ijt}}^2)}$$

Note that, in contrast to Erdem and Keane (1996), the timing of ad signals does matter here (not just the total stock). This is because the $\sigma_{A_{ijt}}^2$ varies over time as ad spending varies over time.

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

$$(26) \quad Q_{ijt} = Q_{j0} + (\lambda_1 + \lambda_2 low_ed_i + \lambda_3 young_i) \cdot \overline{ad}_{ijt}$$

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

Finally, it is clear from (22) that one cannot identify r , the quality levels, and the signal variances separately. Hence, we fix $r=1$ for identification (see Ching and Ishihara (2010)).

⁹ Note that use experience variability affects a consumer's *experienced* utility from consuming the product, while advertising signals do not. But both types of signals are used in a similar way to learn about brands.

8. Empirical Results

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

Our estimation sample consists of 91 households and 4588 observations (on average 50.4 weeks per household). Tables 1 and 2 give summary statistics on households and brands, respectively. There are four major brands: The leading name brands are Pampers (34.4%) and Huggies (30.7%), followed by LUVS (10.9%). The market share of Store Brands is 20.9%. Households rarely buy small brands, which we label “Other” or multiple brands in a shopping trip, which we label “multiple brands.”

To investigate the importance of different types of dynamics, we estimate the full model of equations (20)-(21), as well as models (i) without the PC stage (i.e., without equation (21)), and (ii) without strategic trial. The latter shut down the $(\pi_{0j} + \pi_1 Q_{jt}) \cdot \Delta \sigma_{j,t+1}^2$ terms in (20), leaving a myopic learning model. However, we always include the $\pi_{pg} pgap_t$ term in (20), so consumers are always allowed to be forward-looking with regard to inventories. Myopic agents would not care about inventories, so it makes no sense to consider a myopic inventory model. We estimate all the models by simulated maximum likelihood (see e.g., Geweke and Keane, 2001). We now discuss the estimates of our models:

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

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

(i) Myopic Learning model with no PC stage. This is our most basic model, as it does not contain either strategic trial or category consideration decisions. In Table 3, column 1, we see that *prior* mean quality levels for Huggies and Pampers are 1.96 and 1.93, respectively, while those for LUVS and the Store brand are 1.59 and 1.80, respectively. The estimates of *true* mean quality levels range from 1.63 to 3.17. Priors appear to be pessimistic for Huggies and LUVS, and over-optimistic for Pampers and the store brand. Interestingly, the estimate of true quality is highest for the 3rd highest market share brand (LUVS). The λ parameters (see equation (26)) imply lagged advertising has no significant effect on priors for quality.

Table 3 also reports the key learning parameters. The prior variance is 2.64 ($=1.626^2$) and the experience signal variance is 3.77 ($=1.941^2$). As noted earlier, we map ad expenditure into the ad signal variance using the equation $\sigma_{Aijt}^2 = \sigma_A^2 / ad_{ijt}$. For example, as $\sigma_A = 784.6$, an expenditure of \$6,300, which is the mean for Huggies, implies an ad signal standard deviation of $784.6 / \sqrt{6300} = 9.89$.

Note the *relative* size of the standard deviations for the prior, experience and advertising signals, which are 1.63, 1.94 and 9.89, respectively, are rather different from those found by Erdem and Keane (1996) in their study of the detergent market; i.e., 0.20, 0.57 and 1.75. It would not make sense to compare the *levels* of the standard deviations between the two studies, as the scale normalizations are different. So we should instead compare their relative magnitudes. When we do this, we see that for diapers the experience signal standard deviation is only slightly larger than the prior standard deviation (i.e., 1.94 vs. 1.63). But in detergent the experience signal standard deviation is roughly three times larger (0.57 vs 0.20). Thus, in the diaper category, experience signals are more informative *relative* to the prior.

The parameters of the expected utility function (22) are reported in Table 4. The price coefficient is -0.010 with a t-statistic of -2. As expected, display increases the probability of purchase (i.e., the coefficient on the display indicator is a substantial 2.60). But, interestingly, feature and advertising have no *direct* (i.e., persuasive) effect on the probability of purchase.

Finally, the inventory term π_{pg} from equation (20) is reported in Table 5. Interestingly, it is small and insignificant. This suggests inventory is not an important determinant of purchase probability for diapers. This may appear surprising, but it is important to note that inventory does not *necessarily* have much effect on purchase timing in an inventory model. The main prediction

of inventory models is that consumers time purchases for when prices are low. If carrying costs for a good are low (relative to its value), then (i) consumers will always be interested in buying on deal, even if inventory is high, and (ii) consumers may buy in large quantities when price is low, so inventory ceases to be a concern for many periods. Given that diapers are (i) expensive, (ii) needed over a long horizon, (iii) have a high usage rate, and (iv) a long life span, it is intuitive that parents would always be interested in buying on deal, even if they have inventory.

(ii) Forward-looking Learning Model with no Consideration (PC) Stage. The estimation of a model with forward-looking consumers would normally require solving a DP problem. But in the GK approach all we need do is add the $(\pi_{0j} + \pi_1 Q_{jt}) \cdot \Delta\sigma_{j,t+1}^2$ terms that approximate the future component to equation (20). Table 5, column 2, reports the estimates of the π_0 parameters. They are highly significant for all four major brands. Thus, consumers act as if they are forward-looking and engage in strategic trial behavior. In all cases the linear term (π_{0j}) is negative while the interaction with quality (π_1) is positive. This means that trial is more valuable for brands with higher perceived quality.

Table 6, columns 1 and 2, compare the fit of the myopic and forward-looking learning models. As we see, the log-likelihood improves from -3600 to -3584 (or 16 points) when the future component terms are included. This adds 8 parameters. The AIC improves by 17 points, but BIC, which imposes a larger penalty for extra parameters, deteriorates by 34 points. Based on these results it is somewhat ambiguous whether strategic trial is really an important factor in consumer choice dynamics. But results discussed below will clarify this issue.

In terms of other results, in Table 3, column 2, we see that the estimates of the prior means on quality and the true mean quality levels all shift up by factors of 2 to 3. Consumers clearly have pessimistic priors on all four brands. Estimates of the forward-looking learning model also imply a lower level of prior uncertainty. The estimated standard deviations of the prior, experience and advertising signals, are now 1.03, 3.03 and 9.86, respectively. The relative magnitudes of these standard deviations are very close to those found in Erdem and Keane (1996) in their study of the detergent market (i.e., 0.20, 0.57 and 1.75). In Table 4, we see the price coefficient increases substantially (from -0.01 to -0.054) when we move from the myopic model to the forward-looking learning model, but the other parameters of the expected utility function (display, feature, ad exposure) are little affected.

(iii) Forward-looking Learning model with Consideration (PC) Stage. This is our most complete model, as it takes the myopic learning model and adds both forward-looking behavior (i.e., strategic trial) and a category consideration stage. Thus, consumers may not consider the category in every period, but, when they do, they may take into account both inventories and learning opportunities when making purchase decisions. Recall that we do not impose that consumers take these factors into account (as would be the case in a fully structural approach). Instead, the estimates of the π parameters determine the importance of learning and inventories.

The estimates of the category consideration equation are reported in Table 5, column 4, top panel. Four variables are significant determinants of consideration: display, time since first purchase, age and *pgap*. Most interestingly, the coefficient on time since last purchase (our proxy for inventories) is positive, substantial (.482), and highly significant ($t=4.4$). This estimate implies that consumers become more likely to consider the diaper category as their inventory gets low (i.e., time since last purchase gets large), presumably due to concern about stock-outs.

In Table 6, last column, we see that including the consideration stage leads to a substantial improvement in model fit. Adding the PC stage improves the BIC by 60 points over the forward-looking learning model, and the AIC by 111 points. This is clear evidence that accounting for the consideration stage helps improve the fit to consumer choice dynamics.¹²

These results are consistent with a multi-stage decision process with bounded rationality. In particular, if mental resources are scarce, consumers may rely on non-price factors (e.g., display) to trigger their purchase mindset. Once the purchase mindset is on, they devote more mental resources to purchase decisions, and a “rational” forward-looking consumer theory becomes relatively more applicable.

Adding the category consideration stage also has several notable effects on the other estimates. As we see in Table 3, it causes the prior and true quality parameters to be scaled up substantially. This is because, if probability of consideration is less than one, choice probabilities (*conditional on consideration*) must be proportionately higher to generate the same observed level of sales. As a result, the utilities from the purchase options must be scaled up. If we look at the relative size of the Q_{j0} and Q_j parameters, we see that estimates of true quality *still* exceed the priors (as in model#2), implying consumers are pessimistic. But, unlike the two previous models,

¹² Of course, the PC model nests standard inventory models as the consideration probability approaches one.

we now get the more plausible result that true quality for Pampers and Huggies (the two major brands) are higher than that for LUVS.

Next, as we see in Table 4, the price coefficient increases by 60% in absolute value (from -0.054 to -0.083), and the display coefficient drops from 2.76 to 2.04, compared to the forward-looking model without category consideration. These changes are as expected. First, as price is not relevant in the consideration stage, it appears to be more important in the purchase decision stage. As Ching, Erdem and Keane (2009) note, this is precisely how the PC model is able to explain the empirical generalization that brand choice is much more sensitive to price than it is to purchase timing. Second, as display now matters through the consideration stage (Table 5), it is not surprising that it becomes less important in the utility function.

Finally, as we see in Table 5, bottom row, π_{pg} is now negative and significant. This, together with positive γ_{pg} in the consideration stage, appears consistent with a view that stock-out concerns primarily affect the consideration decision. Once this factor is controlled for, the parameter π_{pg} primarily captures the influence of storage/carrying costs, whereby consumers prefer to have less inventory.

Summary of Results. Comparing the full model (learning, strategic trial, consideration) with the most basic model (myopic learning), the BIC improvement is 26 points. An interesting question is whether most of this improvement could have been achieved by adding the consideration stage alone. Is strategic trial (i.e., forward-looking behavior) really important?

To answer this question, we also estimated a myopic learning model with a consideration stage. The fit is reported in Table 6 column 3 (“myopic with category consideration”). The log-likelihood of this myopic model with consideration is 43 points worse than the full model, the AIC is 69 points worse, and the BIC is 18 points worse. Moreover, the majority of the improvement of the full model over the basic over the basic model is achieved by including strategic trial. Of the 26 point improvement, 8 points is achieved by adding consideration (compare columns #1 and #3) while an additional 18 points is achieved by adding strategic trial (compare columns #3 and #4). Thus, we now have clear evidence that strategic trial is important.

Recall that when we compared models #1 and #2 (i.e., the myopic and forward-looking learning models), the AIC and BIC generated contradictory results about the importance of strategic trial. But, if we start from the myopic model that includes the consideration stage, and

add forward-looking behavior (strategic trial), we get an unambiguous improvement in model fit (compare columns 3 and 4 of Table 6). Clearly then, there is an interaction such that the evidence for strategic trial is much stronger if we also account for consideration.

This result is very interesting and also intuitive. Obviously a consumer cannot operate in a forward-looking mode in periods when he/she does not even consider the category. Put simply, how can a consumer contemplate strategic trial in periods when he/she does not even look at the category? However, once the purchase mindset is on, consumers devote more mental resources to purchase decisions, and strategic trial does come into play.

Our model #2 (learning and no consideration stage) imposes that, if consumers engage in forward-looking behavior, they must do so in *every* period. But our more general model #4, which includes the consideration stage, permits consumers to only contemplate strategic trial in those periods when they actually consider the category.

Ching, Erdem and Keane (2013) note that forward-looking learning models have often been found to fit only slightly better than myopic learning models. We find that here as well when we compare models #1 and #2. But here we find that allowing for forward-looking behavior leads to a substantial improvement in fit over a myopic model, as long as both models include a consideration stage (compare models #3 and #4). This suggests that the reason prior work has not found very strong evidence of forward-looking behavior is that it is too strong to assume that consumers are forward-looking in every period. Instead, consumers appear to act in a forward-looking manner only in those periods when the purchase mindset is on.

In summary, our results suggest that learning, strategic trial, inventories, and category consideration are all important factors driving choice dynamics in the diaper category.

9. Conclusion

In this paper, we have (1) developed a new and simpler approach to estimating dynamic models, and (2) used this new approach to estimate a model that includes learning, inventory and category consideration effects simultaneously. In an application to the diaper category, we found that learning, strategic trial, inventories and category consideration are all important factors driving choice dynamics. Our results favor a two stage decision model where, in stage 1, consumers decide whether to consider a category, and inventories play an important role in this decision. In stage two, consumer behavior is well described by a forward-looking learning model

with strategic trial. This two-stage model fits better than a standard learning model, presumably because it relaxes the strong assumption that consumers engage in forward-looking behavior in every period. Instead, it assumes that consumers only engage in this higher level of mental processing in a subset of periods when they consider the category.

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

Typically, we cannot fully understand the empirical implications of incorporating different sources of dynamics in a model until we solve the model and estimate its structural parameters. But the computational burden of solving for the optimal solution via dynamic programming has been the main hurdle for such a research agenda. Our approach should assist in this process by allowing us to estimate more complex models without the need to fully solve the DP problem. This allows us to at least test whether various sources of dynamics are important (e.g., strategic trial, inventories) without the need for fully structural estimation.

Finally, we should also highlight the drawbacks of our proposed approach. Because we do not impose all the structure implied by a dynamic model, our “semi-structural” model has more parameters than a fully structural model. This is because, in addition to the structural parameters of the utility function, we must also estimate the reduced form parameters that describe how the expected value functions depend on the state variables.

Our approach is also more data intensive than a fully structural approach, as it requires exclusion restrictions (i.e., variables that shift the future component of the value functions but do not affect current payoffs). Still, given the identification results of Thesmar and Magnac (2002) and Rust (1994), it seems clear that one would like to have such exclusions when estimating a dynamic model, even if the model is formally identified from functional form restrictions.

Finally, our semi-structural approach cannot be used for counterfactual experiments that would change the functional form of the expected future value functions. This requires full solution of the DP problem. However, as we recover the expected future value functions as a flexible function of the state variables, we believe our approach should provide useful guidance about how consumers form expectations.

References

- Akerberg, D. (2003), "Advertising, learning, and consumer choice in experience good markets: A structural empirical examination," *International Economic Review*, 44: 1007-1040.
- Chan, T., C. Narasimhan, and Y. Xie (2013), "Treatment Effectiveness and Side Effects: A Model of Physician Learning," *Management Science*, vol.59(6), 1309-1325.
- 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. and M. Ishihara (2010), "The effects of detailing on prescribing decisions under quality uncertainty," *Quantitative Marketing and Economics*, 8(2), 123-165.
- Ching, A., T. Erdem and M. Keane (2009) "The Price Consideration Model of Brand Choice," *Journal of Applied Econometrics*, 24(3), 393-420.
- Ching, A., T. Erdem and M. Keane (2013). "Learning Models: An Assessment of Progress, Challenges and New Developments," *Marketing Science*, 32:6 (Nov/Dec), 913-938.
- Erdem, T., M. Keane (1996), "Decision-making under uncertainty: capturing dynamic brand choice processes in turbulent consumer goods markets," *Marketing Science* 15 1-20.
- Erdem, T., S. Imai, M. Keane (2003), "Brand and quantity choice dynamics under price uncertainty," *Quantitative Marketing and Economics* 1 5-64.
- Erdem, T., M. Keane and B. Sun (2008), "A Dynamic Model of Brand Choice when Price and Advertising Signal Product Quality," *Marketing Science*, 27:6, 1111-25.
- Geweke, J. and M. Keane (2000), "Bayesian Inference for Dynamic Discrete Choice Models without the Need for Dynamic Programming," in Mariano, Schuermann, and Weeks, editors, *Simulation Based Inference and Econometrics: Methods and Applications*, 100-131. Cambridge University Press.
- Geweke, J. and M. Keane (2001), "Computationally Intensive Methods for Integration in Econometrics," in *Handbook of Econometrics: Vol. 5*, Heckman and Leamer (eds.), Elsevier Science B.V., 3463-3568.
- Hendel, I. and A. Nevo (2006), "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 74(6), 1637-1673.

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

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

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

Magnac, Thierry and David Thesmar (2002), "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70:2 (March), 801-816.

Nerlove, M., K.J. Arrow (1962), "Optimal Advertising Policy under Dynamic Conditions," *Economica*, 29:114, 129-142.

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.

Seiler, Stephan (2013), "The Impacts of Search Costs on Consumer Behavior: A Dynamic Approach," *Quantitative Marketing and Economics*, 11:2, 155-203.

Table 1. Summary statistics of household characteristics

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

Table 2. Summary statistics of product characteristics

	No purchase	Huggies	Pampers	LUVS	Store Brand	Others	Multiple Brands
Share (%)	74.60%	7.80%	8.74%	2.77%	5.32%	0.15%	0.63%
mean(p_{it})	n.a.	26.44	27.48	22.03	19.58	n.a.	n.a.
mean($display_{it}$)	n.a.	0.001	0.002	0.001	0.003	n.a.	n.a.
mean(ad_{it}) (X1000)	n.a.	6.30	9.78	3.29	0.255	n.a.	n.a.
mean(interpurch spell)*	n.a.	4.04	3.84	2.95	4.28	8.38	4.03
mean(normalized spell)**	n.a.	0.060	0.056	0.057	0.074	0.189	0.033

* When calculating the "brand-specific" spell, we count a spell belongs to the brand that someone bought at the beginning of the spell.

** Normalized spells are calculated by dividing each spell with the number of diapers bought at the beginning of the spell.

Table 3. Estimates of learning related parameters

	Without category consideration				With category consideration			
	Myopic		Forward-looking		Myopic		Forward-looking	
parameter	estimate	s.e.	estimate	s.e.	estimate	s.e.	estimate	s.e.
Q_{10}	1.961*	0.720	3.874*	0.165	4.516*	1.091	12.28*	0.643
Q_{20}	1.932*	0.722	3.605*	0.22	4.476*	1.091	12.29*	0.644
Q_{30}	1.594*	0.706	3.716*	0.146	4.130*	1.089	11.66*	0.63
Q_{40}	1.802*	0.709	3.797*	0.191	4.359*	1.091	12.58*	0.645
Q_1	2.238*	0.703	6.333*	0.359	5.340*	1.175	16.78*	0.767
Q_2	1.867*	0.633	6.407*	0.514	5.428*	1.285	16.91*	0.754
Q_3	3.170*	0.886	8.251*	0.491	7.083*	1.252	16.00*	0.714
Q_4	1.627*	0.616	6.051*	0.438	4.930*	1.221	17.53*	0.942
λ_1	0.003	0.007	0.001	0.003	0.003	0.006	0.001	0.002
λ_2	-0.005	0.007	-0.002	0.004	-0.004	0.006	-0.002	0.002
λ_3	-0.004	0.011	-0.002	0.006	-0.004	0.011	-0.005	0.004
σ_0	1.626*	0.158	1.029*	0.034	1.702*	0.123	0.888*	0.018
σ_δ	1.941*	0.254	3.028*	0.07	2.914*	0.338	5.305*	0.118
σ_A	784.6*	1.259	782.4*	1.292	844.2*	1.000	846.2*	1.032

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

Table 4: Estimates of Expected Utility Function Parameters

	Without category consideration				With category consideration			
	Myopic		Forward-looking		Myopic		Forward-looking	
parameter	estimate	s.e.	estimate	s.e.	estimate	s.e.	estimate	s.e.
Utility parameters for control variables								
ϕ_p (price)	-0.010*	0.005	-0.054*	0.006	-0.003	0.005	-0.083*	0.007
ϕ_f (feature)	0.326	0.412	0.260	0.414	0.325	0.420	0.240	0.469
ϕ_d (display)	2.597*	0.984	2.764*	0.990	2.061*	0.998	2.04*	0.993
ϕ_{ad} (ad exp)	-0.0008	0.002	-0.002	0.002	-0.001	0.002	-0.004	0.003
Utility of other small brands (just brand intercepts)								
α_5	-6.419*	0.372	-6.386*	0.391	-6.003*	0.374	-5.532*	0.376
α_6	-5.001*	0.169	-4.969*	0.190	-4.596*	0.199	-4.123*	0.216

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

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

	Without category consideration				With category consideration			
	Myopic		Forward-looking		Myopic		Forward-looking	
parameter	estimate	s.e.	estimate	s.e.	estimate	s.e.	estimate	s.e.
Category Consideration Stage								
γ_0 (intercept)					0.391	0.515	-0.390	0.27
γ_{et} (time since 1st buy)					-0.032*	0.005	-0.018*	0.003
γ_{age} (age)					0.953**	0.541	1.421*	0.393
γ_{inc} (income)					0.263	0.331	0.149	0.196
γ_{edu} (I_(fedu=low))					0.203	0.481	0.214	0.237
γ_f (sum_feature_j)					-0.306	0.686	-0.357	0.358
γ_d (sum_display_j)					169.9*	1.001	153.2*	2.896
γ_{pg} (pgap)					0.908*	0.268	0.482*	0.112
Expected Future Components								
π_{01} ($\Delta\sigma_{i1}$)			-70.35*	10.46			-73.27*	1.325
π_{02} ($\Delta\sigma_{i2}$)			-29.97*	4.288			-39.56*	1.025
π_{03} ($\Delta\sigma_{i3}$)			-83.87*	13.08			-103.86*	1.308
π_{03} ($\Delta\sigma_{i4}$)			-42.94*	10.27			-37.13*	1.017
π_{11} ($Q_{i1}*\Delta\sigma_{i1}$)			20.98*	3.104			37.30*	2.985
π_{12} ($Q_{i2}*\Delta\sigma_{i2}$)			13.97*	2.019			34.27*	2.848
π_{13} ($Q_{i3}*\Delta\sigma_{i3}$)			22.63*	3.505			64.47*	4.658
π_{14} ($Q_{i4}*\Delta\sigma_{i4}$)			12.89*	3.026			21.53*	1.854
π_{pg} (pgap)	0.011	0.0078	0.010	0.009	-0.028*	0.01	-0.051*	0.011

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

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

	Without category consideration		With category consideration	
	Myopic	Forward-looking	Myopic	Forward-looking
log-likelihood	-3600.95	-3584.15	-3563.24	-3520.55
(-2*log-likelihood)	7201.90	7168.30	7126.48	7041.10
AIC	7243.90	7226.30	7184.48	7115.10
BIC	7378.60	7412.80	7370.98	7353.05
#parameters	21	29	29	37
#obs	4588	4588	4588	4588