The Journal of Choice Modelling ∎ (∎∎∎∎) ∎∎∎–∎∎∎

Contents lists available at ScienceDirect

The Journal of Choice Modelling

journal homepage: www.elsevier.com/locate/jocm

A simple method to estimate the roles of learning, inventories and category consideration in consumer choice ☆

Andrew T. Ching [a], Tülin Erdem [b], Michael P. Keane [c,*]

[a] Rotman School of Management, University of Toronto, Canada
[b] Stern School of Business, New York University, USA
[c] Nuffield College, University of Oxford, United Kingdom

ARTICLE INFO

Article history:
Received 10 June 2014
Received in revised form
13 October 2014
Accepted 15 November 2014

Keywords:
Learning
Inventory
Dynamics
Structural model
Expectations

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 et al. (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.

1. Introduction

The literature on dynamic discrete choice models of consumer demand has two main strands. In one, the main source of dynamics is consumer learning, while in the other it is inventories. Learning enables one to explain phenomena such as brand loyalty, strategic trial, slow diffusion of information, and heterogeneity in brand perceptions. Inventories allow one to explain phenomena such as reference prices, post-promotion dips and different elasticities for permanent vs. transitory price cuts. Both learning and inventory models are intrinsically dynamic. In the former, today's choice affects tomorrow's information set. In the latter, today's choice affects tomorrow's inventory. Thus, optimal behavior in each type of model requires that consumers solve a dynamic programming (DP) problem to find optimal choices.

Of course, we do not believe consumers can literally solve DP problems, so dynamic models are usually considered "as if" models (see Erdem and Keane, 1996). That is, we assume consumers understand that current choices affect future states, and that information and inventories have value. They may use heuristics (or rules of thumb) that take these values into account. The discrete choice DP literature assumes these heuristics are sufficiently sophisticated that, to a good approximation,

☆ Ching's work on this project was supported by the SSHRC. Keane's work was supported by Australian Research Council grants FF0561843 and FL110100247. We thank Masakazu Ishihara for helpful discussion, and seminar participants at University of Toronto for providing useful comments. We also thank two anonymous referees and Joffre Swait (the editor) for providing helpful feedback.

* Corresponding author.
E-mail addresses: Andrew.Ching@rotman.utoronto.ca (A.T. Ching), terdem@stern.nyu.edu (T. Erdem), michael.keane@nuffield.ox.ac.uk (M.P. Keane).

http://dx.doi.org/10.1016/j.jocm.2014.11.001
1755-5345/© 2014 Published by Elsevier Ltd.

consumers behave as if they solve a DP problem.

In this paper we consider a dynamic discrete choice model that relaxes the “as if” assumption. The model combines learning and inventories in a single framework. What is novel about our framework is that we do not assume consumers solve a DP problem. In order to do this, we extend a method developed by Geweke and Keane (2000) that allows one to relax many of the strong assumptions about choice behavior that underlie the conventional DP approach.

To model learning, we follow the structure in Erdem and Keane (1996). But we do not assume *a priori* that consumers value information optimally. Instead, we estimate the value consumers place on acquiring new information, based on their observed choice behavior. This enables us to test whether consumers engage in strategic trial – defined as the purchase of a brand that does not maximize current utility in order to gain information.

To model inventories, we follow the category consideration model of Ching et al. (2009). In an optimal inventory model, consumers examine prices in every period, and decide whether it is a good time to buy. Here, we assume instead that consumers have limited attention. In each period they follow a two stage process: In the first stage, they decide whether or not to consider a category. Consideration may be prompted by low inventory, advertising, displays, etc. In the second stage they make choices based on a conventional discrete choice model, except that the value of each brand is augmented by a future component that captures the perceived values of inventory and information. These valuations may or may not be optimal.

Thus, our model exhibits a form of antecedent volition: consumers exhibit “rational inattention,” as it would be too mentally burdensome to consider a product category in every period. They only consider a category when prompted by a cue like low inventory. But when they do consider the category, they may engage in sophisticated forward-looking behavior.

We emphasize however, that we do not impose forward-looking behavior. For certain parameter values, the model can imply that consumers are either forward-looking or myopic in the second stage of the choice process. The estimates also determine the frequency with which the category is considered, and the importance of different cues that may motivate consideration.

In summary, one goal of this paper is to develop a flexible model of the consumer choice process that can nest a wide range of different behaviors. These range from quite sophisticated forward-looking behavior that is a close approximation to what a DP solution would imply, all the way to complete myopia and frequent inattention to prices.

A second goal of the paper is to 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), Crawford and Shum (2006) and Ching (2010a, b); for inventory behavior see Erdem et al. (2003) and Hendel and Nevo (2006), and for category consideration see Ching et al. (2009).[1] Key open questions are whether learning or inventories provide a better explanation of consumer choice dynamics (see Erdem et al., 2008), and whether the category consideration mechanism is also an important aspect of choice behavior.

One approach to these questions 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. As we have already noted, both learning and inventories lead to models that are difficult to estimate, as they require solution of DP problems. The computational burden is greatly magnified 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. This is also made possible by extending the method of Geweke and Keane (2000), which enables one to estimate dynamic models without having to solve the DP problem. By using this method, we can, for the first time, estimate a model with both learning and inventories, and shed light on the role of each. At the same time, we relax many of the strong behavioral assumptions that underlie the conventional DP approach. We also test whether consumers engage in strategic trial as the forward-looking learning model implies.[2]

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, and 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 show 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.

[1] Seiler (2013) extends Ching et al. (2009) by introducing category consideration in an inventory stockpiling model.

[2] Note that 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. In order 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*.