

**DECOMPOSING PURCHASE ELASTICITY WITH A DYNAMIC
STRUCTURAL MODEL OF FLEXIBLE CONSUMPTION**

Tat Chan

Chakravarthi Narasimhan

Qin Zhang^{1,2}

November, 2006

¹ The authors are Assistant Professor of Marketing, Philip L. Siteman Professor of Marketing at Washington University in St. Louis, and Assistant Professor of Marketing at University of Texas at Dallas. Authors are listed in an alphabetical order. Corresponding author: Chakravarthi Narasimhan, narasimhan@wustl.edu, Tel: 314-935-6313.

² We thank Dick Wittink for his encouragement and wisdom. We also would like to thank Joel Huber and the two anonymous reviewers, as well as participants at the 2nd QME conference and Scott Neslin, for their excellent comments and suggestions.

ABSTRACT

It is well known that store level sales respond positively to short term price promotions. Under the assumption of constant consumption rate, previous literature has identified that the sources for this increase are brand switching, purchase acceleration, and stockpiling. However, recent research has shown that the consumption of households is also affected positively by price promotions. In this paper we offer a methodology to decompose the effects of price promotions into brand switching, stockpiling and change in consumption and explicitly allow for consumer heterogeneity in brand preferences and consumption needs. A dynamic structural model of a household that decides when, what, and how much to buy as well as how much to consume to maximize its expected utility over an infinite horizon is developed. By making certain simplifying assumptions we are able to reduce the dimensionality of the problem. We estimate the proposed model using scanner-panel datasets of household purchases in canned tuna and paper towels. The results from the model shed insights on the decomposition of price elasticity into its components. This could help managers to make inferences about which brands sales are most responsive to household stockpiling or consumption expansion as well as to understand how temporary price cuts affect future sales. Contrary to previous literature, we find that brand switching is not the dominant force for the increase in sales. We show that brand loyals respond to a price promotion mainly by stockpiling for future consumption while brand switchers do not stockpile at all. We also find that heavy users stockpile more, while light users mainly increase consumption when there is a price promotion.

Key words: flexible consumption, decomposition of price elasticity, consumer heterogeneity, dynamic structural model.

1. INTRODUCTION

It is well known that store level sales respond positively to short term price promotions. Since different households respond to promotions differently researchers have developed household-level models to understand the heterogeneity in purchase behavior. Following the seminal work of Guadagni and Little (1983) researchers (Gupta (1988), Chiang (1991), Chintagunta (1993), Bucklin, Gupta and Siddarth (1998) and Bell, Chiang and Padmanabhan (1999) etc.) have explored the effects of price promotions based on household scanner data and decomposed the short term effects into brand switching, purchase acceleration and purchase quantity with the latter generally arising out of stockpiling behavior by households. In all these papers the consumption rate of households was assumed to be invariant to changes in prices. However, recent research has shown that the consumption rate of households is also affected positively by price promotions (Ailawadi and Neslin (1998), Bell, Iyer and Padmanabhan (2002), Sun (2005)). The increase in consumption rate could be due to cross-category substitution as well as an increase due to the desire to consume more in the category. The latter effect may be caused by an income effect following a price promotion, or simply induced by having a larger inventory. The drivers of the increase in sales due to a temporary price promotion have to be identified to understand the impact on profits of manufacturers and retailers. If, for example, the increase in sales is primarily due to increase in consumption, both manufacturers and retailers could be better off from a temporary price promotion. If the effect is mostly brand switching, the manufacturer of the promoted brand is better off but the retailer may or may not be. If the increase in sales comes mostly from stockpiling the impact on profitability is ambiguous for both the manufacturer and retailers and further investigation is needed (Van Heerde, Leeflang,

and Wittink (2004). Furthermore, if households that have a high preference for a particular brand are those that stockpile the product for future use, then it represents a loss in profits to the manufacturer. Similarly, if households that already have a high consumption level stockpile rather than expand their consumption, this again could represent a loss in profits. Thus, modeling flexible consumption rate in household behavior is important since otherwise we may attribute, erroneously, an increase in sales to other factors leading to misleading implications for manufacturers and retailers.

The discussion above identifies two important research issues that need to be understood. First, a model of household behavior where consumption can be flexible based on external environment (prices, display, and feature ads) and internal resources (income, inventory holding costs) needs to be developed to quantify the sources of increase in sales of a product due to temporary price promotions. Second, the differences in household behavior (increasing consumption, brand switching, and stockpiling etc.) based on brand preferences and overall usage rate need to be incorporated to correctly identify which set of households respond in what way to temporary promotions. Such a model has potential implications for managers as to whether to offer any promotions, and if so what form it should be offered and to whom it should be potentially targeted. Our goal in this paper is to develop a household-level model that allows us to draw inferences on these important questions.

In this paper we develop a dynamic structural model of a household that maximizes discounted utility from its consumption in a category. Consumption in each period is endogenous and is based on current marketing mix, inventory levels, and preferences and the household's expectations about future prices. In deciding whether or

not to buy today, a household trades off inventory cost against the opportunity cost of buying the product in the future at a possibly higher price. Our model incorporates heterogeneity across households in inventory cost, price sensitivity, and underlying preferences. We propose a solution to overcome the dimensionality problem in our dynamic optimization model by imposing some simplifying assumptions. We apply our model to household level scanner-panel data in canned tuna and paper towel categories.

We run policy experiments using the estimates from our model. We are able to decompose the total effect of a price promotion into its components of increase in consumption, brand switching in current and future periods, and stockpiling. Our analysis reveals several interesting insights: (1) contrary to what is shown in most previous literature (e.g., Gupta (1988), Chiang (1991), Sun (2005) etc.)³, brand switching is not the dominant effect; (2) for larger share brands, the majority of promotion induced sales increase is attributed to stockpiling; (3) for smaller share brands, the consumption effect is greater. We further show that the household brand preference has a significant impact on its stockpiling and flexible consumption behavior: brand loyals mainly respond to a price promotion with stockpiling while brand switchers do not stockpile. We also find that heavy users stockpile more for future consumption while light users have larger consumption increase under the price promotion. These findings have important implications for pricing and promotion strategies for both manufacturers and retailers.

The rest of the paper is organized as follows. In the next section we describe related research and position our contribution with respect to the literature. Following this, we describe our dynamic structural model. Then we discuss the data and some details of

³ Our finding is consistent with that in Erdem, Imai, and Keane (2003), Van Heerde, Gupta, and Wittink (2003) and Van Heerde, Leeflang, and Wittink (2004).

our estimation model. Also in that section we discuss the estimation as well as the policy simulation results, and then provide some suggestions to managers based on our findings. Finally, we conclude the paper with some extensions and suggestions for future research. Technical details of the model estimation are provided in the appendixes.

2. LITERATURE REVIEW AND CONTRIBUTIONS

Gupta (1988) was the earliest to decompose the promotional responses into brand switching, purchase (incidence) acceleration and stockpiling effects using household-level data in the coffee category. He found that the dominant force was brand switching accounting for 84% of the change in response, while purchase acceleration accounted for 14% and stockpiling accounted for 2%. Similar results were reported by Chiang (1991) and in the cross-category study by Bell, Chiang and Padmanabhan (1999). Van Heerde, Gupta and Wittink (2003) propose a different decomposition measure based on unit sales. Using the same dataset as Gupta (1988) did but at the store level, they found that only 33% of unit sales increase is due to brand switching. These papers assume that consumption rate of households is a constant and does not change as a result of price reduction. Van Heerde, Leeflang and Wittink (2004) propose a regression model to decompose store level sales increase due to price promotions into cross-period (stockpiling), cross-brand (brand switching), and category expansion (consumption) effects. They find that each effect accounts for one third on average in four categories (tuna, tissue, shampoo and peanut butter). Such store-level models help us understand aggregate effects of price promotions, but lack the advantages of household-level models that can explicitly account for observed and unobserved household heterogeneity in inventory cost, price sensitivity and underlying preferences. Also, it may suffer from the

“over-parameterization” problem if one wants to estimate the cross-substitution patterns among many products. Our proposed household-level model enables us to explore the differences in promotional responses across households even when there are a large number of products in the choice set. Thus, we are able to add to this literature by documenting the link between promotional response, especially change in consumption, and household behavior.

There is a stream of literature that models flexible consumption under price uncertainty using dynamic structural models based on household-level scanner data. Erdem, Imai and Keane (2003) assume that households have an exogenous usage requirement in each period that is revealed to households after their purchases. The focus of their model is to study how inventory and future price expectations affect a household’s purchase decisions. Sun (2005) models consumption as an endogenous decision and explores how price promotions affect this. While we also model consumption as an endogenous decision our focus on price elasticity and its decomposition leads us to interesting insights about how heterogeneity in brand preferences and consumption needs affect promotional responses. Moreover, to overcome the problem of “curse of dimensionality” in the dynamic programming due to a large choice set and a large number of panel members, we adopt a hedonic approach in modeling households’ utility and invoke some few assumptions in the decision process of households to simplify the optimization problem.

3. MODEL

3.1 THE HOUSEHOLD’S PROBLEM

Assume that there are J products, and H households in the market. Let y_{ht} be a $J \times 1$ vector of household h 's consumption quantity, x_{ht} a $J \times 1$ vector of quantity purchased, both assumed to be continuous⁴, and $u_{ht}(\cdot)$ be the utility function of consumption, in period t .

At time t household h decides on whether to purchase, which product to purchase, how much to purchase, and how much to consume. Since products can be stocked and consumed in future periods, purchasing and consumption decisions in current period will affect the inventory that the household holds and change the implicit cost of future consumption. This creates a dynamic linkage among decisions across periods. Formally, the dynamic planning problem at time t for the household can be stated as follows:

$$\sup_{\{x_s, y_s\}} E_t \left\{ \sum_{s=t}^{\infty} \gamma^{s-t} [u_h(y_{hs}) - \lambda_h \cdot p_s' x_{hs} - c_h \cdot (\sum_{j=1}^J I_{hsj})] \mid \sigma_{ht} \right\} \quad (1)$$

$$s.t. \quad I_{hs} = x_{hs} + I_{hs-1} - y_{hs}$$

$$x_{hs}, y_{hs}, I_{hs} \geq 0$$

where p_s is a $J \times 1$ vector of prices, and I_{hs} a $J \times 1$ vector of the inventory level of products in period s . We use λ_h to denote the marginal utility of income and c_h the inventory cost of one standardized unit for household h , and γ is the discount factor that is common for all households. $E_t\{\cdot \mid \sigma_{ht}\}$ is the expectation operator conditional on the information set at t , σ_{ht} . The information set includes the inventory inherited from previous period, $I_{h,t-1}$, current and past marketing mix variables such as prices, features and displays, and household demographic variables. The endogenous decision variables

⁴ This is an approximation of the fact that consumers only make discrete unit purchases. Similar assumption has been made in literature such as Kim, Allenby and Rossi (2002) and Chan (2006).

in (1) include $\{x_{ht}, x_{h,t+1}, \dots; y_{ht}, y_{h,t+1}, \dots\}$, where each component is a $J \times 1$ vector, that are subject to the non-negativity constraints.

3.2 THE INDIRECT UTILITY FUNCTION

We use the hedonic approach to model a household's expected utility of consumption at some future period s evaluated at time t , $s \geq t$, as:

$$E_t u_h(y_{hs}) = (\Psi_{ht}' A' y_{hs} + \varphi_h)^{\alpha_h} \quad (2)$$

where, A is a $J \times (C + J)$ characteristic matrix. The first C columns represent observable product attributes such as brand names and flavors. The last $J \times J$ sub-matrix in A is an identity matrix each diagonal element of which indicates the existence of the corresponding unobserved product attribute. Ψ_{ht} is a vector of household-specific and time-varying coefficients consisting of a $C \times 1$ vector of household h 's preferences for the observed attributes, ψ_{ht} , and a $J \times 1$ vector of preferences for the unobserved attributes, ξ_{ht} . Hence, $\Psi_{ht} = (\psi_{ht} | \xi_{ht})$. We do not model state-dependence that might arise, for example, due to variety seeking behavior (Seetharaman 2004) to keep the model tractable and not let the dimensionality explode.

Given Ψ_{ht} and A , α_h and φ_h determine the curvature and intercept of the marginal utility – the marginal utility with respect to consumption of product j is $\alpha_h \cdot (\Psi_{ht}' A_j) \cdot (\Psi_{ht}' A' y_{hs} + \varphi_h)^{\alpha_h - 1}$ where A_j is the j -th row of the characteristic matrix, A . When the consumption is zero, i.e., $y_{hs} = 0$, it equals $\alpha_h \cdot \varphi_h^{\alpha_h - 1} (\Psi_{ht}' A_j)$. To guarantee diminishing marginal returns (i.e., concavity), α_h is restricted to be between 0 and 1. A household with a larger α is likely to consume more of the product category than those with a smaller α . To allow for corner solutions (i.e., zero purchases), φ_h is restricted to

be positive. Since it is difficult to separately identify both α_h and φ_h , we follow Kim, Allenby and Rossi (2002) and fix φ_h to be 1.

Let the household-specific coefficients in (1) and (2) be $\theta_h \equiv (\alpha_h, \lambda_h, c_h)$, and let Z_h be a vector of demographic variables for household h , we assume that $\theta_h = g(Z_h, \theta_0)$, where $g(\cdot)$ is a vector of functions, and θ_0 is a vector of parameters to be estimated. We will discuss the details in the empirical section.

3.3 PROPOSED SOLUTION TO THE DYNAMIC OPTIMIZATION PROBLEM

It is difficult to solve the dynamic optimization problem in (1) particularly when there is a large number of products in the choice set because of the “curse of dimensionality”. Previous empirical research typically relies on either product aggregation or some simplifying assumptions⁵. In our application to the tuna category we have 12 products with different combinations of product attributes. Product aggregation at brand level will mask some interesting cross-substitution patterns that exist at a more disaggregate level. We propose a solution to overcome this dimensionality problem. Assume $c_h > 0$ for all h , $p_{ij} > 0$ for all t and j . Let \bar{p}_j be the highest price that product j could possibly charge. As $\bar{p}_j - p_{ij}$ is finite for all t and j , we show in Appendix B that there exists a finite time period, T , such that, households do not expect to purchase at t and stockpile for the consumption in periods beyond $t+T$ no matter how much inventory they are holding. Thus, we can rewrite the problem in (1) to a finite horizon problem as below (for simplicity we omit the subscript h hereafter):

⁵ An example of the former is that Sun (2005) focuses on purchases of only two products (aggregated at brand level). An example of the latter is that Hendel and Nevo (2006) assume that the utility from a brand is derived entirely at the moment of purchase; hence, brand and quantity choices can be separated. Their assumption does not apply to our case since the utility from product attributes in our model is derived at the moment of consumption.

$$\begin{aligned}
& \sup_{\{x_s, y_s\}} E_t \left\{ \sum_{s=t}^{t+T} \gamma^{s-t} [u(y_s) - \lambda \cdot p'_s x_s - c(\sum_{j=1}^J I_{sj})] \mid \sigma_t \right\} \\
& s.t. I_s = x_s + I_{s-1} - y_s \\
& \quad x_s, y_s, I_s \geq 0
\end{aligned} \tag{3}$$

The optimal purchase and consumption levels in period t , $\{x_t^*, y_t^*\}$, in (3) are equivalent to the optimal solutions we would obtain from solving the infinite horizon problem in (1). This implies that empirical researchers could start from a reasonably large number for T and solve this finite horizon problem.

To solve the finite horizon dynamic optimization problem above one could use algorithms such as the backward induction method. But, with a large number of state variables, the problem is still too complicated to be solved. We therefore impose the following two assumptions to further simplify the problem:

(A.1) A household consumes each product in its inventory proportionately. That is, given its inventory after the purchase at time t , $I_{t-1} + x_t$, the household plans to consume a proportion δ_s of its inventory in a future period s , where $s \geq t$.

(A.2) In period t , after observing current prices, a household updates its expectation about future prices. Specifically, we assume that after observing the current price of product j in period t , p_{ij} , a household updates its expected prices for product j for future periods by the following:

$$p_{ij}^0 = p_{t-1,j}^0 + \omega \cdot (p_{ij} - p_{t-1,j}^0) \tag{4}$$

where $p_{t-1,j}^0$ is the expected price before the household observes p_{ij} , and p_{ij}^0 is the updated expected price after p_{ij} is observed⁶, and ω is a parameter to be estimated.

With these two assumptions, we can simplify and rewrite the dynamic problem in equation (3). In period t , under assumption (A.1) given previous inventory I_{t-1} and

⁶ We assume that a household's expected price for product j before it observes the price in the first period in our data, $p_{j,0}^0$, is the regular price charged by the household's most frequently visited store. The regular price is defined as the most frequently charged price in the store that is no less than its average price.

current purchases x_t , the household's consumption at t , y_t , can be written as $\delta_t \cdot (I_{t-1} + x_t)$, and the total inventory, $\sum_{j=1}^J I_{tj}$ after the consumption is $(1 - \delta_t) \cdot \sum_{j=1}^J (I_{t-1,j} + x_{t,j})$. For a future planning period s , since under assumptions (A.2) the expected future price is assumed to be constant over all planning periods, the expected purchase at s , x_s , will then only be used for consumption at s but not for stockpiling beyond s . Thus, the consumption at s , y_s , can be written as $\{\delta_s \cdot (I_{t-1} + x_t) + x_s\}$ and the inventory at s can be written as $(1 - \delta_t - \dots - \delta_s) \cdot (I_{t-1} + x_t)$. Hence, equation (3) now can be rewritten as follows:

$$\begin{aligned}
& \sup_{\{x_t, \dots, x_{t+T}, \delta_t, \dots, \delta_{t+T}\}} \{E_t u(\delta_t \cdot (I_{t-1} + x_t)) - \lambda p_t' x_t - c \cdot (1 - \delta_t) \cdot \sum_{j=1}^J (I_{t-1,j} + x_{t,j}) \\
& \quad \sum_{s=t+1}^{t+T} \gamma^{s-t} \cdot [E_t u(\delta_s \cdot (I_{t-1} + x_t) + x_s) - \lambda p_t^0 x_s - c \cdot (1 - \delta_t - \dots - \delta_s) \cdot \sum_{j=1}^J (I_{t-1,j} + x_{t,j})]\} \quad (5) \\
& \text{s.t. } x_t, \dots, x_{t+T}, \delta_t, \dots, \delta_{t+T} \geq 0, \sum_{s=t}^{t+T} \delta_s = 1
\end{aligned}$$

Comparing equation (5) with equation (3), first, we note that the space of consumption decisions is reduced from $J \times (T+1)$ to $I \times (T+1)$; second, the purchase in period s , x_s , only affects the expected consumption in that period and does not affect the inventory. Finally, the future expected price at s , p_s is the same for all planning periods and equal to p_t^0 , which is the expected price formed in period t .

We note that assumptions (A.1) and (A.2), while vastly simplifying the dynamic programming problem, impose some restrictions on the households' consumption and their price expectations. Below we examine the implications of these assumptions and compare them with those in literature.

Assumption (A.1) describes how different products in the inventory are depleted. This assumption will have impact on our inference about the quantity and identities of household inventory over time, and on the expected utility of consumption. In (A.1) we assume homogeneity in the consumption of inventory across products. Similar assumption is used in Erdem, Imai and Keane (2003) and Hendel and Nevo (2006). Alternatively, we could assume that a household consumes its most favorite product first (Sun (2005)), or that it consumes the inventory in the order they were bought, i.e., first-in-first-out (FIFO) rule. If the product category is highly perishable, FIFO would be the most reasonable assumption. However, this does not apply to our empirical applications in either canned tuna or paper towel category. The extent of impact of this assumption would depend on the existence of multiple products in household inventories. We further discuss this issue in section 4.3.

Assumption (A.2) is about a household's formation of price expectations. Since expected price $p_{j,t}^0$ is assumed to be constant over all future planning periods, this assumption implies that at time t it is the household's expectation that if it were to purchase at some time $s > t$ then it would be only for consumption in period s . To buy at time t and hold the product until s , a household trades off the inventory cost of holding until s to the savings that is realized in period t . In prior research (e.g., Erdem, Imai and Keane (2003) and Sun (2005)) consumer expectations are assumed to be fully rational. In this case one would first estimate the price generating process from data and assume that households' expectations conform to this process. In our model we do not assume this but allow households to update their price expectations every time they observe a new price. This vastly simplifies the dynamic optimization problem. However, if

households buy from promotion to promotion and stockpile on each price promotion just enough to last until the next promotion, our model will interpret this as an outcome of high inventory costs or households' lowering their expectations, i.e., a lower value for $p_{j,t}^0$.⁷ In this case, we would expect to see an unreasonably high estimate for the inventory holding cost from our estimation. In section 4.3 we conduct a simulation to examine the impact on our conclusions when households in fact expect frequent price promotions.

We provide more details on how to derive households' optimal decisions on whether to buy, which brand to buy, how much to buy and consume in Appendix C. We use the Simulated Method of Moments (SMM) in our estimation. The estimation procedure involves a nested algorithm for estimating θ : an "inner" algorithm that computes a simulated quantity purchased to solve the problem in (5) for a trial value of θ , and an "outer" algorithm that searches for the value of θ that minimizes a distance function between the simulated and observed quantity. We repeat the inner algorithm until the outer algorithm converges. Details of the method are discussed in Appendix C. Identification issues related to the model estimation are discussed in Appendix D.

4. EMPIRICAL ANALYSIS

4.1 DATA DESCRIPTION

We estimate the proposed model using the A. C. Nielsen scanner panel data on canned tuna from January 1985 to May 1987 in Sioux Falls, SD. The reason we choose

⁷ This occurs only when price promotions are frequent. If promotions are sparse, households will be more concerned about high inventory holding cost than about the potential saving from anticipating the next promotion. In that case our model will predict the same stockpiling behavior as the previous literature.

this category is that canned tuna is easily storable and potentially a good candidate for stockpiling and flexible consumption. The sample consists of 74,795 observations from 1,000 households drawn randomly from a panel of 3,250 households. The selected households made 13,394 purchases during the sample period of 123 weeks and bought exclusively the size 6.5 oz of canned tuna. We focus on purchases of 6.5 oz since 94.7% of the total quantity sold is of this package. There are three main product attributes: brand, water/oil based, and light/regular in fat content. The grouping of the total 33 SKUs by product attributes generates 12 product alternatives. The first 11 products are based on SKUs that share the same three attributes: brand name, water or oil, and light or regular, and the last product consists of SKUs that belong to other brands. Henceforth, we will use the term “product” to refer to one of these 12 alternatives. For each purchase occasion, we construct the price, feature and display of the product bought as the weighted average over the SKUs that belong to this product alternative. The weight used is the quantity of sold. For a product that a household does not purchase in a week, the price, feature and display are constructed as the numerical average over all the SKUs that belong to the product alternative in the household's most frequently visited store. Appendix A provides some summary statistics for these 12 products. The average number of units per purchase occasion was 2.15 units and the average inter-purchase time was 9.84 weeks.

The dataset also contains the demographic characteristics of the households such as family size, income, the employment status of the female head of the household, and type of residency, etc. We incorporate these variables in the estimation of our model. For the holiday effect on the purchase and consumption of tuna, we follow Chevalier, Kashyap

and Rossi (2003) and incorporate the impact of the religious holiday Lent, which we define as the period six weeks before Easter.

4.2 DETAILS OF THE MODEL

In our model, the parameter vector $\Psi_{ht} = (\psi_{ht} | \xi_{ht})$ represents the stochastic household preference for product attributes. To simplify computation, we assume that $\xi_{ht} \square normal(0, \sigma_\xi^2)$ and *i.i.d.* over households and time periods in our estimation. We use a random coefficient approach to model ψ_{ht} . Its first element, $\psi_{ht}^{[1]}$, represents the household consumption preference for tuna. We allow it to be affected by holiday Lent, i.e., $\psi_{ht}^{[1]} = \bar{\psi}^{[1]} + \phi \cdot HOLIDAY_t + \eta_{ht}^{[1]}$, where $HOLIDAY_t$ is an indicator of whether time t is in the weeks of Lent, $\eta_{ht}^{[1]}$ is a time-varying taste shock distributed as $normal(0, \sigma_1^2)$ that is *i.i.d.* over households and periods, and $\bar{\psi}^{[1]}$ and ϕ are parameters to be estimated. For simplicity we assume ϕ to be homogeneous across households⁸. Other elements of ψ_{ht} , $\psi_{ht}^{[k]}$, where $k \neq 1$, represent household preferences for brands, “Water” (vs. “Oil”) and “Light” (vs. “Regular”). To simplify the analysis we assume that these parameters are time-invariant: $\psi_{ht}^{[k]} = \bar{\psi}^{[k]} + \eta_{ht}^{[k]}$ for all t , where $\eta_{ht}^{[k]}$ is again distributed as $normal(0, \sigma_k^2)$, and $\bar{\psi}^{[k]}$ is a parameter to be estimated. To reduce the computational burden we assume that preferences for product attributes are independent. Also, given that the market share of products with the attribute of “regular” is extremely small (0.5%), it will be very difficult to identify, for example, the correlation in preferences between “water” and “light”.

⁸ We estimate an alternative model in which the heterogeneity in ϕ is allowed and assumed to be normally distributed. All parameter estimates are very close to the results we report in Table 1. Results are available from the authors upon request.

The parameter α in equation (2) is restricted to between 0 and 1. To allow for unobserved household heterogeneity, we assume that there exist K discrete segments with different values of α . But within each segment we allow α to vary with family size which seems to be the only demographic variable in our data that could affect the consumption rate. For household h that belongs to type k , the specification is as follows:

$$\alpha_{h,k} = \frac{\exp(\alpha_{0k} + \alpha_{1k} \cdot FMY_h)}{1 + \exp(\alpha_{0k} + \alpha_{1k} \cdot FMY_h)} \quad (6)$$

where FMY_h represents the family size of household h , and $\alpha_{k,0}$ and $\alpha_{k,1}$ are parameters to be estimated. Unlike the heterogeneity specification for ψ_{ht} , we assume that the heterogeneity for α_h has a latent structure. If, for example, we assume $\alpha_{h,0}$ and $\alpha_{h,1}$ to be normally distributed, then after the transformation in (6) α_h would be a non-standard distribution which we find difficult to justify. Latent structure specification seems to be a more robust assumption in this case.

Households are likely to be heterogeneous in the price sensitivity parameter λ . We follow the extant empirical literature and assume that price sensitivity is a function of the household income level. Moreover, a household with working female head may be less price-sensitive than those without. Hence, we use the following specification for λ

$$\lambda_h = \frac{1}{\left(\frac{\ln(INCOME_h)}{\ln(minINCOME)}\right)^{\lambda_1} \cdot EMPLOY_h^{\lambda_2}} \quad (7)$$

where $INCOME_h$ is the income of household h , $minINCOME$ is the minimum household income in data, and $EMPLOY_h$ is the employment status of the female household head. This specification restricts the parameter λ_h to positive. A positive value for λ_1 or λ_2 implies that households with a higher income or a working female head are less price

sensitive. Since the product category preference ($\bar{\psi}^{[1]}$ above) is estimated in the model, we normalize the price sensitivity of a household with the minimum income level and non-working female head to one, as implied in (7). Also, because of this normalization the unobserved heterogeneity in λ_h cannot be identified.

Inventory cost may be affected by a household's income level (more affluent households may have larger houses hence lower inventory cost) and the residence type (living in single family houses may imply lower inventory cost compared to living in apartments). To allow for the heterogeneity in inventory costs, c_h , among households, we assume that

$$c_h = c_0 + c_1 \cdot [\ln(INCOME_h) - \ln(minINCOME)] + c_2 \cdot RESIDENCE_h \quad (8)$$

where $RESIDENCE_h$ is the residence type of the household h . It is one if household h lives in a single family house, and zero otherwise. We do not expect other demographic variables to have significant impacts on the inventory costs. Because inventory holdings are not directly observed from data, but have to be inferred from household purchases, it is difficult to estimate the unobserved heterogeneity in c_h and the parameter ω in equation (4).

We assume that in-store marketing activities such as displays and features will change the marginal utility of consumption by a constant amount, β_d and β_f , respectively. We further assume that household's expected utility at some future period s is affected by the current displays and features. This seems to be a reasonable assumption as what we model is the current purchasing decisions at time t . Hence, the indirect utility function in equation (2) can be rewritten to include effects of displays and features as

$$E_t u_h(y_{hs}) = (\Psi_{ht}' A' y_{hs} + \varphi_h)^{\alpha_h} + \beta_d \cdot display_t' y_{hs} + \beta_f \cdot feature_t' y_{hs} \quad (9)$$

where β_d and β_f are parameters to be estimated.

We do not observe the inventory that household h has when it makes its first purchase in our observational period, $I_{h,0}$. This is inferred from the estimation. Intuitively, if household h makes a purchase when the price in that period is high, it could indicate that its inventory is low. Therefore, we assume that

$$I_{h,0} = \exp(v_h), \text{ and } v_h \square normal(\rho \cdot (p_{h1,j} - p_j^a), 1) \quad (10)$$

where $p_{h1,j}$ is the observed price of the chosen product j when household h made the first purchase in data, p_j^a is the average price of product j over the whole sample period, and ρ is a parameter to be estimated. If ρ is negative, a higher first purchase price will indicate a lower initial inventory.

Finally, we fix the number of finite periods to $T = 12$. Given that our estimated inventory cost per week (we will discuss in the result section) is 2 cents, and the differences between the regular price and the minimum price for all products in our data is smaller than 15 cents, our results should not change even if we increase the value of T . Note that for any household inventory on hand may last for more than 12 weeks if the realized consumption is less than what is expected. This implies that, when making the purchase decisions, households do not *plan* to stock up for more than 12 weeks.

4.3 ESTIMATION RESULTS

We report the results of the estimation in Table 1. Starkist and CKN are the two most preferred brands. Attributes “water” and “light” also on average increase consumption utility. However, there is a large heterogeneity in household preferences.

For example, the mean preference for the category is -0.12, with a standard deviation of 0.48. The negative mean preference is consistent with our data in which most households only buy tuna occasionally. Similar extent of preference heterogeneity also exists for preferences for other product attributes. Household responses to marketing activities like feature advertising and store display are significantly positive. Income does not have an impact on price sensitivity but if the female head of the household is employed then the household is less price sensitive which is consistent with past research. The estimated average inventory cost of a can of 6.5 oz tuna is 2 cents per week, which does not appear to be exceedingly large. The unit inventory cost of a household is not affected by its income level or residence type. One explanation for this could be that because canned tuna is relatively easy to store households do not run into space constraints caused by income or residency types.

The consumption level in our model is determined by $\alpha_{h,k} = \frac{\exp(\alpha_{0k} + \alpha_{1k} \cdot FMY_h)}{1 + \exp(\alpha_{0k} + \alpha_{1k} \cdot FMY_h)}$. We assume two latent segments and estimate their α_0 and α_1 respectively. The first segment, accounting for 79.9% of all the households, has a smaller value of α_0 . This is the segment of light users of tuna. It seems counter intuitive that α_1 is negative for both segments implying that smaller families have a larger consumption rate for this category. An obvious explanation is that single households use tuna as a substitute for cooked meals and are likely to purchase it more. The negative ρ implies that if a household buys at a higher price during its first observed purchase the household might have a lower starting inventory level. This is consistent with our intuition. The estimate of ω is 0.20 and significant. Its positive sign confirms our

expectation. Since $0 < \omega < 1$, it implies that households will adjust their expectations downward when they find the current prices lower than their prior expectations; however, the updated expected prices will be higher than the current prices. As expected the religious holiday of Lent has positive impact on the consumption of tuna. During Lent, households consume more tuna to substitute for meat.

[Insert Table 1]

An advantage in using the hedonic approach to model households' preferences to products is to succinctly reveal the substitutability across a large number of products. Without using the hedonic approach, to recover the cross-substitution patterns we would have to estimate the whole variance-covariance matrix corresponding to the random coefficient preferences for 12 products. Here we only need to model the distribution of household preferences for the major product attributes, which vastly reduces the dimensionality. We simulate cross-elasticities among the 12 products based on our estimates and find some interesting results. For example, for the product "CKN, Water Light", "CKN, Oil, Light" and "Starkist, Water, Light" are its close substitutes, while "Starkist, Oil, Light" is not (cross-elasticities are 0.16, 0.13 and 0.07, respectively). And for the product "Starkist, Oil, Light", "Starkist, Water, Light" and "CKN, Oil, Light" are equally substitutable (the cross-elasticities are 0.07, 0.06 respectively). As we can see "CKN, Water Light" and "CKN, Oil, Light" are more substitutable (0.16) than "Starkist, Oil, Light" and "Starkist, Water, Light" (0.07). This suggests that there is an interaction effect between brand name and other attribute on substitutability. These results will help managers to make strategic decisions when they consider price promotions at SKU level instead of brand level.

Using the estimation results, we calculate the expected total purchases of all 1,000 households in each period. We compute the R^2 statistics by comparing the expected and observed total household purchases in each week. The overall R^2 is 0.86. For individual products, for example, the R^2 for the two products that have the largest market shares, “StarKist, Water, Light” and “CKN, Water, Light”, are 0.81 and 0.76 respectively. All these suggest a high level of model fit with the data.

Note that assumption (A.1) will only have an impact on the estimation results if households in our data stockpile multiple products in their inventory. Our data shows that only about 30% of the purchases involve an inter-purchase time that is less than 6 weeks and a brand choice that is different from the one on the previous purchase. Unless households stockpile for more than 5 weeks our results will not be sensitive to this assumption, and that indeed appears to be the case from our estimation.

Next, we conduct a simulation to examine whether and how much the estimates from our model will be biased when assumption (A.2) is violated. We first simulate the consumption and purchases of 100 hypothetical households in 100 weeks using traditional dynamic optimization techniques, under the assumption that all those households are forward looking and correctly anticipate the discounting process. In order to make the problem tractable, we assume that there is only one brand in the market and households are only heterogeneous in their brand preferences. We further assume that in each week the brand charges a regular price of \$0.8 with 80% probability, and a promotional price of \$0.6 with 20% probability. The 100 households correctly forecast that there is a 20% chance of a promotion in next week, 36% chance of at least one promotion in next two weeks, and 49% change of at least one promotion in next three

weeks, and so on. We make these assumptions so we can evaluate the bias of the estimates from our model, if any, when the actual household price expectation behavior is not consistent with our assumption. Finally we estimate our model based on the simulated data. Table 2 shows the “true” parameters that we used to simulate the data and the estimation results from our proposed model. The results show that our model can largely recover the true parameters except the inventory cost.⁹

We are more concerned about how the decomposition of promotional effects (more details in section 4.4) is affected when assumption (A.2) is violated. To address this issue, we compare the purchases and consumption from our simulation with what our model would predict based on the estimation results. The correlations between simulated and predicted purchases and consumption are 0.99 and 0.98 respectively, suggesting that our model recovers quite accurately these quantities. Thus, even if assumption (A.2) is violated, our results on the decomposition of promotional effects remain valid.¹⁰

[Insert Table 2]

4.4 DECOMPOSING THE PROMOTIONAL EFFECTS

In order to answer the questions that we raised earlier regarding the effects of price promotions on increase in consumption, brand switching and stockpiling¹¹, we conduct

⁹ The biased estimate for inventory cost reflects the fact that our model ignores the full and correct expectations of periodic discounts by the households. However, note that the estimated inventory cost from our data of tuna purchases is 2 cents, which does not seem unreasonably high, suggesting that assumption (A.2) may not be much violated in our data.

¹⁰ To check the sensitivity of household decisions to various specifications of price expectations, in empirical analysis we also estimate two other specifications: (1) $p_{j,t}^0 = p_j^r + \omega \cdot (p_{j,t} - p_j^r)$; (2) $p_{j,t}^0 = p_j^a + \omega \cdot (\frac{p_{j,t-1} + p_{j,t}}{2} - p_j^a)$, where p_j^r is the regular price of product j , p_j^a is the average price of product j . Estimation results are similar to those of the specification in equation (4), but these two models perform worse in data fit. Results are available from the authors upon request.

¹¹ We cannot identify the store switching effects in our model. Because households visit stores for basket shopping, to address this issue we will have to study the purchasing patterns across multiple categories, which is a very difficult task using our household-level data. Moreover, we suspect that the store switching effect is not significant in this category, as purchases in tuna usually account for a relatively small portion of total basket expenditure and is hardly a major

simulations to run policy experiments using the estimates from our model. In the simulations, we use the 1,000 households used in our estimation and for each household we draw 10 random samples from the estimated distribution of the stochastic components in the utility function. We first simulate the purchases and consumption quantity of these households when they face the regular prices for all products for $T = 12$ successive weeks. Then, we simulate the purchases and consumption quantity for the same households over the 12 weeks when each product has a price discount of 10 cents¹² in the first week. Comparison of the sales before and after the price cut allows us to decompose the effects of price promotions into those of consumption increase, brand switching and stockpiling. We define the total effect of a price cut as the sales increases in week one. That is:

$$TE_j = \left[\sum_{h=1}^{10000} (x_{hj1}(p_j^1) - x_{hj1}(p^0)) \right] \quad (11)$$

where p^0 is the price vector when all products charge their regular prices, and p_j^1 is the price vector when product j has a price cut while other products remain at their regular prices.

The consumption effect is defined as the difference in the total consumption in all 12 weeks before and after the price cut. Households could change their consumption rate as prices change due to the following effects: first, there is a category substitution effect as households switch from consuming other categories such as meat to tuna; second, having more inventory may induce households to consume more tuna than they usually do; third, there is an income effect due to price promotions.

factor in store choice decisions. If households switch from a store that is not in our data to a store in our data due to a price promotion, the increase in purchase will be counted as an increase in consumption.

¹² We also conduct the simulations under a price discount of 20 cents, and the results are robust.

Since prices only change in the first week in our simulation and we assume that the household planning horizon is 12 weeks, i.e., households will not purchase in week one and hold them as inventory for future consumption beyond the 12th week, the consumption effect is exactly equal to the difference of total purchases in 12 weeks before and after the price cut. That is:

$$\begin{aligned}
CE_j &= [\sum_{t=1}^{12} \sum_{h=1}^{10000} (y_{hjt}(p_j^1) - y_{hjt}(p^0))] \\
&= [\sum_{t=1}^{12} \sum_{h=1}^{10000} (x_{hjt}(p_j^1) - x_{hjt}(p^0))]
\end{aligned} \tag{12}$$

where $y_{hjt}(\cdot)$ and $x_{hjt}(\cdot)$ are household h 's consumption level and purchase quantity for product j in period t .

The brand switching effect in all 12 weeks is measured by the decrease in total sales of other brands (whose prices remain as their regular prices) as the focal brand cuts its price in week one. We compute two components: brand-switching that occurs in the first week (BS_{j1}) and in later weeks (BS_{jL}). That is:

$$\begin{aligned}
BS_{j1} &= [\sum_{h \in H_{-j1}} (x_{h1}(p^0) - \sum_{h \in H_{-j1}} x_{h1}(p_j^1))] \\
BS_{jL} &= [\sum_{t=2}^{12} \sum_{h \in H_{-jt}} x_{ht}(p^0) - \sum_{t=2}^{12} \sum_{h \in H_{-jt}} x_{ht}(p_j^1)]
\end{aligned} \tag{13}$$

where H_{-jt} are those households who purchase brands other than j in week t . Brand-switching in later weeks may exist when some households switch to buying j under price promotion in week one and stockpile for later consumption; hence, they will buy less of other brands in later weeks. Notice that since these households are paying the same prices for the other brands under p^0 or p_j^1 , the consumption and stockpiling effect of the other brands are zero.

The difference between TE_j and the sum $CE_j + BS_{j1} + BS_{jL}$ is the residual component of sales increase of product j that is neither from increase in consumption nor from the substitution of other products in current and future weeks. This can only be due to the stockpiling (cross-period substitution in Van Heerde, Leeflang and Wittink 2004) effect, which comes from the behavior of those households that would purchase product j under the regular price p^0 in later weeks but are taking advantage of the price cut in week one. Note that those households that would not buy product j under p^0 but switch to product j under prices p_j^1 may also stockpile but this effect is accounted as brand switching (in later weeks) under our definition. Therefore, the stockpiling effect is:

$$SP_j = TE_j - CE_j - BS_{j1} - BS_{jL} \quad (14)$$

4.4.1. DECOMPOSING THE AGGREGATE PROMOTIONAL EFFECTS

Dividing the total effects of the price promotion in equation (11) by the total unit sales of product j under the original (regular) price and further by the percentage of price change, we have the total price elasticity of quantity sales. We decompose the elasticity into consumption, brand switching, and stockpiling effects by dividing equation (12) - (14) also by the total unit sales of product j under the original price and then by the percentage of price change. The decomposition is given in Table 3.

[Insert Table 3 here]

First, the average consumption effect is 29%, the brand switching effect 28% (the brand switching in the first week accounts for 22%), and the stockpiling effect 43%. Van Heerde, Leeflang and Wittink (2004) report a similar decomposition result obtained from

their store level model for the same category (31%, 31% and 38%, respectively)¹³. Compared with Sun (2005) (33%, 42% and 25%, respectively), our consumption effect is similar but the brand switching effect is lower and stockpiling effect is higher.

Second, we notice that the effects are different between larger share brands (i.e., StarKist and CKN) and smaller share brands (i.e., 3 Diamond and CTL):

1. Smaller share brands have higher total price elasticity than larger share brands. This is consistent with prior findings (Chintagunta 1993, Kopalle, Mela and Marsh (1999)).
2. The stockpiling effect for the two larger share brands (53% for Starkist and 51% for CKN) is larger than for the smaller share brands (32% for 3 Diamond and 35% for CTL). This is consistent with what is found in Mace and Neslin (2004).
3. Brand switching effect is relatively small for larger share brands but substantially higher for smaller share brands.
4. Consumption effect is substantial for all brands. But this effect is higher for smaller share brands than for larger share brands.

The comparisons imply that, from a manufacturer's perspective, the strategy of temporarily cutting prices to steal sales from other brands might not be very effective for large share brands. Unlike in the one period game, a larger brand's profits could be hurt in the long run in our case since a large portion of its sales increase comes at the expense of future sales. This is because these brands have more brand loyalists who may tend to stockpile more during promotions (we will provide supporting evidence in section 4.4.2).

¹³ They further decompose the consumption effect ("category expansion effect") into "cross-store effect" and "market expansion effect". Since our paper does not model the store switching behavior, we are not able to replicate their exercise.

Similar conclusion has been reached by Kopalle, Mela and Marsh (1999). They show that larger share brands could increase profits by reducing the frequency of price promotions.

4.4.2. HOUSEHOLD HETEROGENEITY IN RESPONSE TO PROMOTIONS

We first investigate the impact of brand preferences on households' responses to price promotions. For each brand the households are grouped into two segments based on their purchases at regular prices. Households in the first segment ("brand loyals") purchase the focal brand at its regular price. The households in the second segment ("brand switchers") consist of two types, one that purchases only on promotion and the second that purchases other brands also at regular price but will buy the focal brand on promotion. To save space we only report the weighted average elasticity of all 12 products, weighted by their market share at regular prices. The decomposition results of these two segments are shown in Table 4.

[Insert Table 4]

Comparing the elasticity decomposition between the two segments generates the following insights:

1. For brand loyals, the majority of the increase in purchases from price promotions can be attributed to stockpiling while brand switchers do not stockpile.
2. Brand loyals increase their consumption more than brand switchers do.
3. Brand loyals are more price elastic than brand switchers due to flexible consumption and stockpiling.

These findings suggest that it may be less beneficial for a firm with large number of loyals to offer promotion since it will lose profits from loyals not only in the current

period (charging a lower price to loyals who would have bought the product anyway) but also in future periods (heavy stockpiling by loyals). On the other hand, in a very competitive market (with a high proportion of brand switchers), temporary price cuts will generate the benefit of stealing sales from other brands as well as an increase in consumption at no expense of future sales. Therefore, knowing the market composition in terms of households' brand preferences will help managers to decide whether it is appropriate to use price promotion strategies.

Next, we investigate the impact of household heterogeneity in category consumption preferences on their promotional responses. We divide all households into two groups, "heavy users" and "light users". A heavy user is a household whose total quantity purchased in all 12 weeks is above the average. Otherwise it is a light user. Table 5 shows the decomposition results.

[Insert Table 5]

Comparisons of elasticity decomposition between the two groups show that:

1. Light users show a dominant consumption effect, while heavy users show a larger stockpiling effect.
2. Light users are more responsive to price promotions.

The smaller consumption effect for heavy users is consistent with the satiation effect. Considering the fact that the canned tuna category has relatively flexible consumption, when induced by lower prices, light users have more room to increase consumption compared to heavy users (perhaps from switching consumption from other categories). The larger stockpiling effect for heavy users suggests that heavy users are more strategic in terms of planning for future consumption. This can be attributed to the

different importance of the tuna category to the two groups of households and is consistent with the findings in Zhang, Seetharaman and Narasimhan (2003).

The two important findings that we have identified for canned tuna category are: when facing a price promotion (1) brand loyals mainly respond with stockpiling, while brand switchers switch brands and increase consumption; (2) light users mainly increase consumption while heavy users stockpile more. This can help managers to design better pricing and promotion strategies. Though a rigorous analysis of the optimal pricing strategy that a firm should use under these settings is out of the scope of our paper, we will provide some general guidance here. In general a firm can increase its long term profits by discouraging brand loyals and heavy users from responding to price promotions while encouraging brand switchers and light users to do so. For example, a firm can give price discounts on small package sizes to induce brand switchers and light users to switch. To more efficiently price discriminate and discourage brand loyals and heavy users from taking advantage of these price discounts the per unit price of large package sizes should not be higher than the discounted per unit price of smaller package size. This shows that the presence of loyals and heavy users constrains a firm's strategies in competing for light users and brand switchers.

4.5 REPLICATING THE ANALYSES IN PAPER TOWELS

We have shown that in a category like canned tuna where households are likely to be flexible in changing their consumption rates, price promotions have substantial consumption as well as stockpiling effects. For non-food categories where the consumption rate is relatively fixed we expect that the ability of households to increase consumption is limited so price promotions will mainly induce them to stockpile for

future consumption. To demonstrate the validity of our model, we estimate our model using household purchase data from a non-food category – paper towels – and then conduct the elasticity decomposition. We employ IRI’s scanner panel data on household purchases of single-roll paper towels in a large U.S. city from June 1991 to June 1993. We focus on 99 households that only purchased single-roll paper towels¹⁴ and made at least 5 purchases during the 104 weeks. The data consists of 8,714 observations which include 1,795 purchases. We group the SKUs into four product alternatives, which include top-selling brands Scott, Viva and Scott (accounted for 82% of total purchases), and other brands. The average number of rolls per purchase was 1.69 and the average inter-purchase time was 8.72 weeks. The rules for data construction for paper towels are the same as those used for canned tuna. The household demographic characteristics used in the estimation are family size, income, the employment status of the female head of the household.

The estimation results are reported in Table 6. We then use these estimates to run simulations to obtain the changes on quantity purchased and consumption for all 99 households during the planning horizon of 12 weeks before and after a price cut of 10 cents in the first week. We calculate the unit sales decomposition according equation (12) - (14) and then compute the corresponding elasticity components. The elasticity decomposition for all households for the three brands is reported in Table 7.

[Insert Table 6 and Table 7]

Consistent with what we find in canned tuna, brand switching effect is lower than

¹⁴ This may imply that the households that we use in the data are relatively light users, as households with larger consumption needs tend to buy paper towel in packs. We do not use those households since choices on multi-roll packages will pose a server problem of discrete unit purchases that our model cannot handle. Our purpose for this exercise is only to see how well our model can recover the decomposition of promotional effects in the case when the consumption rate is relatively fixed.

what most previous literature has identified -- an average of 14% (with the brand switching occurring in the first week being 10%). As expected, compared with elasticity decomposition in canned tuna, the consumption effect is much lower (an average of 18% vs. an average of 29% in canned tuna). We notice that the consumption effect is not zero although the consumption for paper towels is considered to be relatively fixed. This could be primarily due to the category substitution effect (e.g., households may switch from cloth rags to paper towels). Alternatively, this may also be due to the income effect (because of the price promotion) or the fact that households may be induced to use more paper towels with more inventory in hand. Finally, the stockpiling effect (an average of 69%) is much higher than the tuna category and it is the majority impact of price promotions in this category. This is consistent with our intuition since paper towel is a non-food category that is more storable than food categories.

5. CONCLUSIONS AND FUTURE RESEARCH

We propose a dynamic structural model to understand the impact of temporary price promotions on households' behavior and to identify the relative influences of consumption increase, brand switching and stockpiling on the total impact. Using a hedonic approach to model a household's utility function and imposing some simplifying behavioral assumptions, we reduce the dimensionality of the problem and apply our model to a market in the presence of large number of households and product alternatives. Our methodology helps us to understand the substitution patterns among different product attributes. Our decomposition of price elasticity in the canned tuna and paper towel categories shows a much smaller brand switching effect than what has been documented earlier. We also find that for larger share brands the majority of promotion

induced sales increase is attributable to stockpiling for future consumption. We then investigate household heterogeneity in terms of brand preferences, and find that with price promotions brand loyalists mainly stockpile for future consumption, while brand switchers almost do not stockpile. We also find that in response to a price promotion, light users increase consumption more while heavy users stockpile more for future consumption. These findings can help managers to better design their pricing and promotion strategies.

In this paper, we did not consider state dependence, i.e., households might seek variety from one purchase to another or be inertia. It might be of interest to explore the impact of price promotions when both variety-seeking (or inertia) and flexible consumption behavior are present. The consumption increase we capture in this paper is more in line with total market expansion which is different from spillover effects (e.g., price promotions of a product benefit the non-promoted products (Van Heerde, Gupta and Wittink 2003)). Aggregate increase in consumption at the store level could also be due to store switchers that we do not consider in our model. In future research, it would be of interest to explicitly model store switching behavior, and further study how much of the category expansion in one store is due to store switching vs. households' increasing their consumption.

Appendix A. Some Summary Statistics of Tuna Products

Product alternative	Market Share (%)	Avg. Price (\$/Unit)	Avg. Feature	Avg. Display
StarKist, Water, Light	36.32	0.78	0.14	0.04
CKN, Water, Light	32.56	0.81	0.13	0.05
StarKist, Oil, Light	13.00	0.70	0.16	0.04
CKN, Oil, Light	10.05	0.82	0.12	0.05
CTL, Water, Light	3.44	0.65	0.02	0.04
3 Diamond, Water, Light	1.66	0.62	0.01	0.08
CTL, Oil, Light	1.24	0.65	0.03	0.04
3 Diamond, Oil, Light	0.71	0.62	0.01	0.08
Other Brands	0.54	1.14	0.00	0.00
CTL, Oil, Regular	0.27	0.64	0.02	0.11
StarKist, Water, Regular	0.19	1.54	0.00	0.00
CTL, Water, Regular	0.01	1.22	0.00	0.00

Appendix B. Reduce a Infinite-Horizon Problem to a Finite-Horizon Problem

Our objective is to reduce the infinite-horizon problem in (1) to a finite-horizon problem. First, we can replace the purchase and the inventory in period t , x_t and I_{t-1} , with the consumption at current and all future periods, i.e., $x_t = \sum_{s=t}^{\infty} y_{t,s}$ and $I_t = \sum_{s=t}^{\infty} y_{t-1,s}$. The subscript “ t ” in $y_{t,s}$ denotes the time of purchase, and “ s ” the time of consumption, and “ $t-1$ ” in $y_{t-1,s}$ indicates that it is quantity inherited from inventory at period $t-1$. Then we can rewrite the infinite horizon planning problem into

$$\begin{aligned} & \sup_{\{y\}} E_t u(y_{t-1,t} + y_{t,t}) - \lambda \cdot p_t' \left(\sum_{s=t}^{\infty} y_{t,s} \right) - c \left\{ \sum_{j=1}^J \left[\sum_{s=t+1}^{\infty} (y_{t-1,s,j} + y_{t,s,j}) \right] \right\} \\ & + \sum_{s=t+1}^{\infty} \gamma^{s-t} E_t \left\{ \begin{aligned} & u(y_{t-1,s} + y_{t,s} + \dots + y_{s,s}) - \lambda \cdot p_s^{0t} \left(\sum_{u=s}^{\infty} y_{s,u} \right) \\ & - c \left[\sum_{j=1}^J \sum_{u=s+1}^{\infty} (y_{t-1,u,j} + y_{t,u,j} + \dots + y_{s,u,j}) \right] \mid \sigma_t \end{aligned} \right\} \\ \text{s.t. } & I_{t-1} = \sum_{s=t}^{\infty} y_{t-1,s} \\ & y_{s,u,k} \geq 0, u \geq s, k = 1, \dots, J \end{aligned} \tag{B.1}$$

Since $\bar{p}_j - p_{ij} < b < \infty$, for all t and j , that is the difference between the highest possible price in data and the observed price of product j is finite, there exists a finite time period, T_1 , such that households will not purchase at time t for the consumption of time u , for all $u > T_1$, under the condition that the cost from purchasing the product today and stockpiling for consumption at period T_1 is greater than the cost of having to buying the product at period T_1 and consume it at the same period (even at the highest price), i.e.,

$$\begin{aligned} & \lambda \cdot p_{ij} + c \left(\sum_{u=t}^{T_1} \gamma^{u-t} \right) > \lambda \cdot (\gamma^{T_1-t} \bar{p}_j) \\ & \Rightarrow T_1 - t < \frac{\log \left(\left(\frac{c}{\lambda} + p_{ij} \right) / \bar{p}_j \right)}{\log(\gamma)} \end{aligned} \tag{B.2}$$

In this case, $y_{t,u,j} = 0$, for all $u > T_1$. That is, no household will purchase product j at t and hold it as inventory for the consumption after period T_1 .

Now let us consider the endogenous variables $\{y_{t-1,b}, y_{t-1,t+1}, \dots\}$, that is, the inventory inherited from period $t-1$ which will be used for future consumption. By assumption, marginal utility (i.e., $\frac{\partial E_t u_h(y_{hs})}{\partial y_j} \Big|_{y_{hs}=0} = \alpha_h \cdot \Psi_{ht}' A_j'$) is finite at zero consumption level. Then there exists a T_2 such that

$$\begin{aligned} \alpha \cdot \psi' A_j \cdot \gamma^{T_2-t} &< c \cdot \sum_{u=t}^{T_2} \gamma^{u-t} \\ \Rightarrow T_2 - t &< \frac{\log(c) - \log(\alpha \cdot \psi' A_j \cdot (1-\gamma) + c)}{\log \gamma} \end{aligned} \quad (\text{B.3})$$

This implies that a household will not stockpile at t for the consumption after period T_2 due to the existence of its inventory cost and the discount rate.

Let $T = \max\{T_1, T_2\}$. We can write a finite horizon problem equivalent to the infinite horizon problem in (B.1)

$$\begin{aligned} \sup_{\{\}} E_t u(y_{t-1,t} + y_{t,t}) - \lambda \cdot p_t' \left(\sum_{s=t}^T y_{t,s} \right) - c \left\{ \sum_{j=1}^J \left[\sum_{s=t+1}^T (y_{t-1,s,j} + y_{t,s,j}) \right] \right\} \\ + \sum_{s=t+1}^T \gamma^{s-t} E_t \left\{ u(y_{t-1,s} + y_{t,s} + \dots + y_{s,s}) - \lambda \cdot p_s^{o'} \left(\sum_{u=s}^T y_{s,u} \right) - c \left[\sum_{j=1}^J \sum_{u=s+1}^T (y_{t-1,u,j} + y_{t,u,j} + \dots + y_{s,u,j}) \right] \mid \sigma_t \right\} \\ \text{s.t. } I_{t-1} = \sum_{s=t}^T y_{t-1,s} \\ y_{s,u,k} \geq 0, u \geq s, k = 1, \dots, J \end{aligned} \quad (\text{B.4})$$

Solutions of $y_{t-1,s}$ and $y_{t,s}$, for all $s \geq t$, in (B.4) are equivalent to solutions in (B.1). Therefore, our infinite horizon problem has been reduced to a finite horizon planning problem.

Appendix C. Model Estimation (Simulated Method of Moments)

First, we examine how a household decides whether to buy, which product to buy, how much to consume and how much to buy. To simplify the analysis, suppose that the household does not hold inventory in week t . The utility function in (2) implies that, after observing p_t , a household will choose at most one product, j^* , such that, for all $k = 1, \dots, J$,

$$\begin{aligned} \frac{MU_{j^*}(0)}{\lambda p_{t,j^*}} &= \frac{\alpha \cdot \psi' A_{j^*}}{\lambda p_{t,j^*}} \geq \frac{\alpha \cdot \psi' A_k}{\lambda p_{t,k}} = \frac{MU_k(0)}{\lambda p_{t,k}} \\ \Rightarrow \frac{\psi' A_{j^*}}{p_{t,j^*}} &\geq \frac{\psi' A_k}{p_{t,k}} \end{aligned} \quad (C.1)$$

where $MU_{j^*}(0)$ is the marginal utility level with respect to j^* at $y_k = 0, \forall k = 1, \dots, J$.

Corner solution exists when $\max_{\{k\}} \left\{ \frac{\alpha \cdot \psi' A_k}{\lambda p_{t,k}} \right\} < 1$, for all k . In this case we have $x_t = 0$. This

occurs when the household finds current prices too high to purchase the product category at week t

With the expected price p_t^0 and the utility function as specified in (2), the household expects to choose at most one product j^0 in any future period s , such that, for all k ,

$$\begin{aligned} \frac{MU_{j^0}(0)}{\lambda p_{t,j}^0} &= \frac{\alpha \cdot \psi' A_{j^0}}{\lambda p_{t,j}^0} \geq \frac{\alpha \cdot \psi' A_k}{\lambda p_{t,k}^0} = \frac{MU_k(0)}{\lambda p_{t,k}^0} \\ \Rightarrow \frac{\psi' A_{j^0}}{p_{t,j}^0} &\geq \frac{\psi' A_k}{p_{t,k}^0} \end{aligned} \quad (C.2)$$

Again, when $\max_{\{k\}} \left\{ \frac{\alpha \cdot \psi' A_k}{\lambda p_{t,k}^0} \right\} < 1$, for all k , the expected purchase in period s will be zero.

This occurs when the household does not normally purchase or consume the category, and will only purchase when there is a big price promotion.

As the above discussion implies, the household will purchase in week t and hold it for consumption in week s , where $s > t$, only if the following two conditions are satisfied:

$$(i) \quad \gamma^{s-t} \cdot MU_{j^*}(0) - \lambda p_{t,j^*} - c \cdot \sum_{u=0}^{s-t} \gamma^u \geq 0$$

and

$$(ii) \quad \gamma^{s-t} \cdot MU_{j^*}(0) - \lambda p_{t,j^*} - c \cdot \sum_{u=0}^{s-t} \gamma^u \geq \gamma^{s-t} \cdot [MU_{j^0}(0) - \lambda p_{t,j^0}^0]$$

Condition (i) ensures that discounted consumption utility in week s net of purchasing costs in week t and discounted inventory cost is non-negative. Condition (ii) ensures that it is worthwhile to buy now and hold inventory until week s .

Suppose the above two conditions are satisfied, *i.e.*, the household purchases in week t and holds it for consumption in week s . The optimal purchase quantity then satisfies the third condition that is derived from the first-order condition:

$$(iii) \quad \gamma^{s-t} \cdot MU(y_{t,s,j^*}) - \lambda p_{t,j^*} - c \cdot \sum_{u=0}^{s-t} \gamma^u = 0$$

where y_{t,s,j^*} is the optimal quantity purchased in week t for the consumption in week s . The optimal level of x_t is equal to the sum of y_{t,s,j^*} , $s=t, \dots, T$, in the j^* -th row and zero

elsewhere. Thus, the optimal proportion of consumption $\delta_s^* = \frac{y_{t,s,j^*}}{x_{t,j^*}}$.

When the household holds positive inventory in week t , solution y_{t,s,j^*} for all s cannot be solved separately as in conditions (i) to (iii), since the household has to additionally consider the benefit and cost of consuming the inventory *vs.* buying in current week. However, basic principles of the solution concept discussed above still apply. In our algorithm, we directly solve the non-linear constrained optimization problem in (5), given parameters θ and inventory level I_{t-1} .

Based on the above discussion, an appropriate technique is to use the Method of Moments by matching the expected quantity purchased obtained from the maximization problem with the observed purchases. The estimation procedure involves a nested algorithm for estimating parameters θ : an “inner” algorithm computes a simulated quantity purchased which solves the problem in (5) for each trial value of θ , and an “outer” algorithm searches for the value of θ that minimize a criterion function. The inner algorithm is repeated every time when θ is updated by the outer algorithm.

Procedures of the inner algorithm are the following: Given the conditional distribution function of random variables ε ¹⁵, $F(\cdot | X_t)$, where X_t represents all explanatory variables including marketing variables such as prices, product features and displays, as well as demographic variables like household size, income level, residence type, female employment status, we can solve the expected values of $\{x_t, \dots, x_{t+T}; \delta_t, \dots, \delta_{t+T}\}$ from (5). That is,

$$\{x_t^*, \dots, x_{t+T}^*, \delta_t^*, \dots, \delta_{t+T}^*\} = \int \arg \max_{\{x_t, \dots, x_{t+T}, \delta_t, \dots, \delta_{t+T}\}} \{\text{problem in (5)} F(d\varepsilon | X_t)\} \quad (\text{C.3})$$

However, x_t^* in (C.3) is non-tractable because of the non-negativity constraints. Instead, we use simulation method: for each household h and period t , we draw $\tilde{\varepsilon}_{ht,1}, \dots, \tilde{\varepsilon}_{ht,ns}$ from the distribution function $F(\cdot | X_{ht})$, where ns is the number of simulation draws. Conditional on each simulation draw $\tilde{\varepsilon}_{ht,s}$, we solve a non-linear constrained optimization problem for the optimal quantity purchased levels at time t , $\tilde{x}(X_{ht}; \theta; \tilde{\varepsilon}_{ht,s})$, using a derivative search procedure called the Sequential Quadratic Programming. In this method variables are updated in a series of iterations beginning with a starting value that satisfies the constraints in (5). If \tilde{x}_n is the current value at iteration n , then its successor is $\tilde{x}_{n+1} = \tilde{x}_n + \rho\delta$, where δ is a direction vector, and ρ a scalar step length. The procedure is repeated for every simulation draw and this generates a simulated dataset $\tilde{x}(X_{ht}; \theta) = (1/ns) \cdot \sum_{s=1}^{ns} \tilde{x}(X_{ht}; \theta; \tilde{\varepsilon}_{ht,s})$. When the utility function is concave and its Jacobian and Hessian matrices can be written down in analytical forms, convergence of \tilde{x} is fast.

The outer algorithm searches for the estimator θ . We make a major identification assumption that there are no unobserved characteristics in the model. Therefore, there is no endogeneity issue for marketing variables such as prices. Although this assumption can be challenged, it solves the data problem as good instruments for weekly prices are not available. Under the identification assumption, this yields a moment condition:

¹⁵ Here ε includes (i) all stochastic components in the utility function such as individual product preferences ξ_h and η_h , (ii) a standard normal variable that relates to the initial inventory level $I_{h,0}$, and (iii) discrete distributed stochastic variables relates to the types of α_h a household has in the utility function.

$$E[x_{ht} - \tilde{x}(X_{ht}; \theta_0)] = 0 \quad (\text{C.4})$$

where θ_0 is the true parameters. The estimator θ_n is obtained by the following non-linear least-square estimator

$$\theta_n = \arg \min_{\theta \in \Theta} Q_n(\theta) = \arg \min_{\theta \in \Theta} \frac{1}{H \times T_h} \sum_{h=1}^H \sum_{t=1}^{T_h} [x_{ht} - \tilde{x}(X_{ht}; \theta)]^2 \quad (\text{C.5})$$

We use the Nelder-Meade (1965) nonderivative simplex method to search for θ_n . Estimators based on this moment condition are called the Simulated Method of Moment (SMM) estimators (Pakes (1986), Pakes and Pollard (1989), McFadden (1989)). One major advantage of using the SMM is that ns can be finite (even when $ns = 1$) and we still obtain consistent estimators. This helps to reduce the computational burden in model estimation. Our methodology here is very similar to Chan (2006). However, his model is to estimate multiple-product, multiple-unit purchase decisions, while our model is to estimate the purchase and consumption decisions over multiple periods.

Appendix D. Identification of Different Types of Consumption Behavior

Depending on the exogenous (1) preferences for the attributes of different products, (2) inventory cost, and (3) the slope coefficient α , a price promotion will have different effect on different households. A major identification issue in our estimation is that we, as researchers, do not observe consumption and inventory of households in the data. We only observe whether a household makes a purchase and if so which product it buys and the quantity it purchases in each period. Still, we can identify the parameters of a household's utility function from the observed variations in its purchasing pattern over time: Brand switching patterns of households over time help to identify the differences in product preferences of households, and variations in purchase quantity and time-intervals between purchases help to identify the inventory cost and consumption rate changes. For example, suppose that there are two households, A and B , who buy one unit of the product in each period at the same price. Suppose the price is cut by 10 percent in the current period, and both A and B increase their purchase from 1 to 2 units. If household A comes back to purchase 1 unit again in the next period, but household B does not make a purchase until period 3, we can infer that A increases its consumption and does not stockpile in the current period, while household B does the opposite. In this case household A has a flexible consumption rate but a higher inventory cost than household B . Suppose there is another household, C , which buys 4 units of the product during promotion, and only comes back to market in period 3. Then we infer that household C may have a more flexible consumption rate than household B but a lower inventory cost than household A . Hence, the parameters are identifiable if there are enough variations on purchases quantity and inter-purchase time-intervals in responses to prices. Given these parameter values, effects of stockpiling, brand-switching and consumption increase due to temporary price promotions can then be identified. For example, when there is a price promotion household A will show a larger consumption effect but a smaller stockpiling effect than household B does, while household C will show a larger consumption as well as stockpiling effect.

Another identification issue arises from the fact that we do not observe households' price expectations, and the parameter ω in equation (4) and (5) has to be estimated. Suppose in week t household h does not have any inventory. Assume that household h

observes a price higher than its expected price $p_{h,t}^0$ and decides to buy one unit. If in week $t+1$ the price is the same as last week (i.e., $p_{t+1} = p_t$) and h decides to buy more. In this case one will infer that h 's expected price $p_{h,t+1}^0$ must be adjusted upward, and ω is positive (this is because h would not purchase for stockpiling purpose in week $t+1$ if it was the case that $p_{h,t+1}^0 = p_{h,t}^0 < p_{t+1} = p_t$). On the other hand, if h is buying less than one unit in week $t+1$, one may infer that it is expecting discounts in the near future; hence, with the same argument, ω should be negative.

In order to demonstrate that the model parameters can be identified, we conduct a simulation by fixing the parameters and stochastic components of the objective function in equation (5) to generate the household purchasing and consumption data. Then we estimate the model by using another set of randomly drawn stochastic numbers. The results are available from the authors upon request. Overall the estimated parameters are very close to the “true” ones, rendering support for the above arguments.

Table 1. Estimation Results in Tuna

Parameters	Est.	Std. Error
Tuna	-0.12	0.01
StarKist	3.01	0.01
CKN	2.15	0.01
3 Diamond	-2.87	0.03
CTL	-2.31	0.03
Water	2.79	0.01
Light	0.35	0.01
Feature	1.97	0.01
Display	2.65	0.01
Price*Income	0.001	0.01
Price*Employ	0.02	0.01
C ₀ (Inventory)	0.02	0.001
C ₁ (Inventory*Income)	0.002	0.001
C ₂ (Inventory*ResidenceType)	-0.0001	0.004
α_{01}	-1.97	0.02
α_{02}	-0.0003	0.01
α_{11}	-1.01	0.02
α_{12}	-0.16	0.01
Probability of Segment 1	0.80	0.08
σ_1 (Tuna)	0.48	0.01
σ_2 (StarKist)	2.94	0.01
σ_3 (CKN)	3.68	0.01
σ_4 (3 Diamond)	0.47	0.03
σ_5 (CTL)	0.09	0.03
σ_6 (Water)	3.09	0.01
σ_7 (Light)	0.48	0.01
σ_8 (Feature)	0.27	0.01
σ_9 (Display)	0.23	0.01
σ (Product-Specific Household Preferences)	2.73	0.01
ρ	-2.19	0.10
ω	0.20	0.01
Holiday Effect	0.04	0.01
Function Value	477.72	

Table 2. Comparison of True and Estimated Parameters from a Calibration

	True Parameter	Estimated Parameter	Std. Error
Brand Preference	3.00	2.79	0.12
Inventory Cost	0.03	0.12	0.01
Power Term (α)	0.50	0.52	0.003
Brand Preference Heterogeneity (σ)	1.00	0.94	0.04
Updated Price Expectations (ω)	N/A	0.12	0.56

Table 3. Decomposition of Price Elasticity for All Households in Tuna

	Total Price Elasticity in 1st Week	Consumption Elasticity	Brand Switching Elasticity in 1st Week	Brand Switching Elasticity in later Weeks	Stockpiling Elasticity
Starkist	-1.17 (100%)	-0.31 (26%)	-0.18 (16%)	-0.06 (5%)	-0.62 (53%)
CKN	-1.32 (100%)	-0.30 (23%)	-0.25 (19%)	-0.10 (7%)	-0.67 (51%)
3Diamond	-1.66 (100%)	-0.60 (36%)	-0.40 (24%)	-0.13 (8%)	-0.53 (32%)
CTL	-1.53 (100%)	-0.49 (32%)	-0.43 (28%)	-0.07 (5%)	-0.53 (35%)

Note: The elasticities reported here are weighted average elasticities at brand level based on the sales at regular prices. The percentages relative to the total elasticities are in parentheses.

Table 4. Decomposition of Price Elasticity for Brand Loyals and Switchers in Tuna

	Total Price Elasticity in 1st Week	Consumption Elasticity	Brand Switching Elasticity in 1st Week	Brand Switching Elasticity in later Weeks	Stockpiling Elasticity
Brand Loyals	-0.90 (100%)	-0.27 (29%)	- -	- -	-0.63 (71%)
Brand Switchers	-0.04 (100%)	-0.01 (17%)	-0.02 (63%)	-0.01 (21%)	0.00 (0%)

Note: The elasticities reported here are weighted average elasticities based on the sales at regular prices. The percentages relative to the total elasticities are in parentheses.

Table 5. Decomposition of Price Elasticity for Heavy and Light Users in Tuna

	Total Price Elasticity in 1st Week	Consumption Elasticity	Brand Switching Elasticity in 1st Week	Brand Switching Elasticity in later Weeks	Stockpiling Elasticity
Heavy Users	-1.25 (100%)	-0.32 (26%)	-0.23 (18%)	-0.07 (6%)	-0.63 (50%)
Light Users	-2.22 (100%)	-1.28 (57%)	-0.40 (18%)	-0.11 (5%)	-0.43 (20%)

Note: The elasticities reported here are weighted average elasticities based on the sales at regular prices. The percentages relative to the total elasticities are in parentheses.

Table 6. Estimation Results in Paper Towels

Parameters	Est.	Sta. Err.
Paper Towels	4.85	0.01
Bounty	8.46	0.01
Viva	8.07	0.01
Scott	5.08	0.01
Feature	0.85	0.03
Display	3.66	0.02
Price*Income	0.22	0.004
Price*Employ	-0.60	0.01
C_0 (Inventory)	0.02	0.001
C_1 (Inventory*Income)	0.02	0.002
α_{01}	-8.18	0.01
α_{02}	-1.04	0.01
α_{11}	0.85	0.01
α_{12}	0.07	0.005
Probability of Segment 1	0.55	0.002
σ_1 (Paper Towels)	6.25	0.01
σ_2 (Bounty)	8.58	0.01
σ_3 (Viva)	12.53	0.01
σ_4 (Scott)	5.76	0.01
σ_5 (Feature)	0.45	0.02
σ_6 (Display)	1.21	0.01
ρ	1.13	0.62
ω	0.14	0.005
Function Value	29.25	

Table 7. Decomposition of Price Elasticity for All Households in Paper Towels

	Total Price Elasticity in 1st Week	Consumption Elasticity	Brand Switching Elasticity in 1st Week	Brand Switching Elasticity in later Weeks	Stockpiling Elasticity
Bounty	-1.14 (100%)	-0.20 (18%)	-0.13 (12%)	-0.05 (4%)	-0.76 (66%)
Viva	-1.04 (100%)	-0.19 (19%)	-0.07 (7%)	-0.02 (2%)	-0.75 (72%)
Scott	-1.57 (100%)	-0.25 (16%)	-0.17 (11%)	-0.09 (5%)	-1.06 (68%)

Note: The elasticities reported here are based on the sales at regular prices. The percentages relative to the total elasticities are in parentheses.

References

- Ailawadi, K. L. and S. A. Neslin (1998), "The Effect of Promotion on Consumption: Buying More and Consuming It Faster," *Journal of Marketing Research*, 35 390-398.
- Bell, D. R., J. Chiang. and V. Padmanabhan (1999), "The Decomposition of Promotional Response: An Empirical Generalization," *Marketing Science*, 18 504-526.
- Bell, D. R., G. Iyer and V. Padmanabhan (2002), "Price Competition under Stockpiling and Flexible Consumption," *Journal of Marketing Research*, 39 292-303.
- Bucklin, R. E., S. Gupta, and S. Siddarth (1998), "Determining Segmentation in Sales Response Across Consumer Purchase Behaviors," *Journal of Marketing Research*, 35 189-197.
- Chan, T. (2006), "Estimating A Continuous Hedonic Choice Model with An Application to Demand for Soft Drinks," *RAND Journal of Economics*, forthcoming.
- Chevalier, J.A., Kashyap, A., and Rossi, P. (2003), "Why Don't Prices Rise during Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review*, 93 15-37.
- Chiang, J. (1991), "A Simultaneous Approach to the Whether, What and How Much to Buy Questions," *Marketing Science*, 10 297-315.
- Chintagunta, P. K. (1993), "Investigating Purchase Incidence, Brand Choice and Purchase Quantity Decisions of Households," *Marketing Science*, 12 184-208.
- Erdem, T., S. Imai and M. P. Keane (2003), "Consumer Price and Promotion Expectations: Capturing Consumer Brand and Quantity Choice Dynamics under Price Uncertainty," *Quantitative Marketing and Economics*, 1 5-64.
- Guadagni, P. M. And J. D. Little (1983), "A Logit Model of Brand Choice Calibrated on Scanner Data," *Marketing Science*, 2 203-238.
- Gupta, S. (1988), "Impact of Sales Promotions on When, What, and How Much to Buy," *Journal of Marketing Research*, 25 342-355.
- Franses, P. H. (2005), "On the Use of Econometric Models for Policy Simulation in Marketing," *Journal of Marketing Research*, 42 4-14.
- Hendel, I. and A. Nevo (2006), "Measuring the Implications of Sales and Consumer Stockpiling Behavior," *Econometrica*, forthcoming.
- Kim, J., G. M. Allenby and P. E. Rossi (2002), "Modeling Consumer Demand for Variety," *Marketing Science*, 21 229-250.

Kopalle, P. K., C. F. Mela and L. Marsh (1999), "The Dynamic Effect of Discounting on Sales: Empirical Analysis and Normative Pricing Implications," *Marketing Science*, 18 317-332.

Mace S. and S. A. Neslin (2004), "The Determinants of Pre- and Postpromotion Dips in Sales of Frequently Purchased Goods," *Journal of Marketing Research*, 41 339-350.

McFadden, D. (1989), "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," *Econometrica*, 57, 995-1026.

Nelder, J. A. and R. Mead (1965), "A simplex method for function minimization," *Computer Journal*, 7 308-313.

Pakes, A. (1986), "Patents as Options: Some Estimates of the Value of Holding European Patent Shocks," *Econometrica*, 54, 755-785.

Pakes, A. and Pollard, D. (1989), "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 57, 1027-1057.

P.B. Seetharaman (2004), "Modeling Multiple Sources of State Dependence in Random Utility Models: A Distributed Lag Approach." *Marketing Science* 23 (2), 263-271.

Sun, B. (2005), "Promotion Effect on Endogenous Consumption," *Marketing Science*, 24 430-443.

Van Heerde, H. J., S. Gupta and D. R. Wittink (2003), "Is 75% of the Sales Promotion Bump Due to Brand Switching? No, Only 33% is," *Journal of Marketing Research*, 40 481-491.

Van Heerde, H. J., P. S. H. Leeflang and D. R. Wittink (2004), "Decomposing the Sales Promotion Bump with Store Data," *Marketing Science*, 23 317-334.

Zhang, Q., P.B. Seetharaman and C. Narasimhan (2003), "Incorporating Future Price Expectations in a Model of Whether, What and How Much to Buy Decisions," working paper, University of Texas at Dallas.