On the Timing and Depth of a Manufacturer's Sales Promotion Decisions with Forward-looking Consumers

Yan Liu

Texas A&M University-MKTG
220G Wehner Bldg.
4112 TAMU
College Station, TX 77843-4112
Phone: (979) 845-2547
Email: yliu@mays.tamu.edu

Subramanian Balachander

Krannert School of Management
Purdue University, West Lafayette, IN, 47907
Phone: (765) 494-4466
E-mail: sbalacha@purdue.edu

Sumon Datta

Krannert School of Management
Purdue University, West Lafayette, IN, 47907
Phone: (765) 496-7747
E-mail: sdatta@purdue.edu

September 2011

Abstract

This paper investigates a manufacturer's optimal timing and depth of price promotions over a planning horizon in a frequently purchased packaged goods context. Our empirical analysis comprises of two steps. In the first step, we obtain heterogeneous demand side parameters with a dynamic structural model. In this model, consumers decide whether to buy, which brand to buy and how much to buy conditional on their rational expectations of future promotions. In the second step, we specify a dynamic game between consumers and a focal manufacturer and solve for the optimal promotion policy, taking the structural demand-side parameters from the first step as given. We obtain the optimal promotion policy as the Markov-perfect equilibrium outcome of the dynamic game. In our empirical application, we develop the optimal promotion schedule for the StarKist brand in the canned tuna category using household-level panel data. We find that it is optimal for the manufacturer to promote when the mean inventory for brand switchers is sufficiently low and that the optimal discount depth decreases in the mean inventory for brand switchers. We also find that StarKist could increase profit by offering more frequent but shallower price promotions. Interestingly, we find that the manufacturer's profit increases as consumers become more forward-looking (discount the future less).

Key words: Optimal Promotion Schedule, Forward-Looking Consumers, Promotion Expectations, Dynamic Game, Structural Model

1. Introduction

In the frequently purchased consumer packaged goods (CPG) industry, a manufacturer's sales promotion planning process typically involves the development of a promotion calendar that details the timing and description of promotion events over a planning horizon such as a year (Blattberg and Neslin 1990). Such a promotion calendar is typically an important part of a brand's marketing plan. In this paper, we study the optimal timing and depth of in-store price promotions to be offered by a manufacturer over a planning horizon. Many recent studies have shown that consumers' forward-looking behavior influences their response to promotions, and affects the incremental sales from promotions (Sun, Neslin and Srinivasan 2003, Erdem, Imai, and Keane 2003). However, the implications of consumer's forward-looking behavior for the optimal timing and depth of price promotions has not been explored at length in the literature. In this paper, we demonstrate how a manufacturer in a consumer packaged goods (CPG) category can design an in-store price promotion calendar that best accommodates consumer forward-looking behavior. To the best of our knowledge, this is the first paper that integrates consumer's dynamic response to promotions into a solution process for a manufacturer's optimal price promotions.

We highlight two critical challenges in deciding the optimal price promotion calendar for a manufacturer in a CPG category. First, an effective price promotion schedule should discriminate between consumers who are willing to buy the manufacturer's brand at the regular price and/or are less likely to stockpile during a promotion event, and price sensitive consumers who are more likely to switch between brands and/or stockpile during a promotion event. Unlike coupons which can be targeted towards specific consumer segments, in-store price promotions are available to all shoppers at the time of the offering. Therefore, in-store price promotions have to achieve price discrimination by the choice of their timing and depth of discounts. In frequently purchased products, consumers' purchase decisions not only depend on their preferences and their price and promotion sensitivities but also on their inventory level; consumers are more likely to respond to promotions and make purchases when they have low inventory. Therefore, a manufacturer could design the timing and depth of price promotions for maximum discriminative effect by exploiting differences in anticipated inventory levels over time across consumer segments.

For the purposes of illustration, consider the example shown in Figure 1A. This figure shows the mean inventory over time of a frequently purchased product for two hypothetical segments of households. The first segment, which we refer to as *loyals* for simplicity, has a strong preference for brand *A* and is also less interested in stockpiling the product (has high inventory holding costs). Consequently, households in this segment are not sensitive to promotions and buy just enough of their preferred brand, *A*, to meet their weekly consumption regardless of its promotion status. Thus, in Figure 1A the mean inventory for *loyals* is relatively stable over time. In contrast, the second segment, named *switchers*, is more willing to switch brands to avail of a lower price and is also more willing to stockpile the product. So, the mean inventory for *switchers* increases after a promotion by brand *A* (weeks 1 and 5) as they stockpile during promotion events.

Now, consider brand *A*'s decision after the promotion in week 1 about the timing and depth of its next price promotion. Brand *A* can potentially gain additional sales and profit from the *switchers* by offering a deep discount in a subsequent week but the tradeoff is that a price cut would reduce profit from *loyals* who would have anyhow bought at the regular price. In weeks when there is no promotion, the inventory of *switchers* depletes due to consumption and they become progressively more responsive to a price promotion so that over time, the manufacturer's tradeoff turns in favor of offering a promotion (assuming competing brands have not offered a promotion by then). In week 5, say, when the *switchers* become sufficiently responsive to a promotion because of low inventory brand *A* may find that it is profitable to offer a temporary discount. More generally, it is desirable to have the timing of promotions coincide with weeks when *loyals* have sufficient inventory and are less likely to make a purchase whereas the *switchers* have low inventory and are ready to purchase and stockpile. This example shows that to design a promotion calendar that can discriminate between different kinds of consumers, the manufacturer needs to forecast consumers' inventory over time. We refer to this as the *price discrimination problem*.

The second challenge for a manufacturer relates to consumers' forward-looking behavior. Consumers have expectations about future promotions in the category, and these expectations influence their purchase decisions (Erdem and Keane 1996, Gönül and Srinivasan 1996, Erdem, Imai, and Keane 2003, Liu and Balachander 2011). For example, if a manufacturer offers promotions less frequently, then consumers expect to wait longer between the manufacturer's successive promotions, and may therefore stockpile more during a promotion event by the manufacturer. Further, depending on their inventory holding costs, forward-looking consumers may accelerate purchases to coincide with promotion events (Erdem et al. 2003).

In our illustrative example, suppose brand *A* decides to offer price promotions less frequently at approximately every six weeks instead of every four weeks. In such a case, brand *A*'s second promotion would be expected in week 7 instead of week 5. With the new infrequent promotion schedule, *switchers* may stockpile more during promotion weeks as they expect to wait for a couple of weeks more to get another deal price from brand *A*. Therefore, in week 1 their inventory may shift upwards as shown in Figure 1B. A reduced-form demand side model, which measures consumers' response to promotion entirely based on relationships observed in historical data (Lucas 1976), would not predict such a change in consumers' stockpiling behavior. Thus, when a manufacturer decides on the optimal promotion schedule, it has to take into account that consumers are forward-looking, and that consumers' purchase decisions and the optimal promotion policy are endogenous. We refer to this as the *forward-looking consumer problem*.

To deal with these issues in our empirical development of an optimal promotion schedule for a focal manufacturer of packaged goods, we estimate first a dynamic structural demand-side model incorporating latent classes of households who are strategic and forward-looking. In this demand model, households in each period trade off the benefits of stockpiling at the current period price against the opportunity cost of delaying purchase for a future promotional price. Thus, households decide whether to buy (purchase incidence), which brand to buy (brand choice) and how much to buy (discrete quantity choice) conditional on their inventory levels and their rational expectations of future prices and promotions in the product category. Following this first step of estimating a structural demand model, we formulate, in a second step, a dynamic game between households and the focal manufacturer using the estimated demand model to develop the manufacturer's optimal promotion schedule. In this game, consumers' expectation about the promotion policy of the focal manufacturer is endogenous to the manufacturer's chosen promotion policy. Concurrently, the manufacturer maximizes the discounted sum of future profits over an infinite time horizon by taking into account consumers' responsiveness to promotions based on their projected inventory levels, their forward-looking behavior and competitors' promotion policy. The Markov Perfect Nash Equilibrium (MPNE) solution of the game gives us the manufacturer's promotion policy that is optimal given consumers' purchase decisions, which are in turn optimal given the manufacturer's promotion policy. In obtaining the equilibrium solution, we factor in the role of the retailer by assuming a pass-through rate for manufacturer's promotions and by imposing a minimum retailer profit constraint (cf. Silvia-Risso et al. 1999). We discuss our two-step approach in more detail in the next section.[1]

In contrast to an alternative approach of simultaneously estimating demand and supply side parameters, the two-step approach has the following advantages. First, we can estimate the demand parameters without making strong assumptions about manufacturer's behavior thus avoiding potential misspecification errors. Second, because we do not impose manufacturer optimal behavior in the estimation, we are able to recommend an optimal promotion schedule for the manufacturer. Third, the two-step approach significantly reduces the computational cost of estimation since the demand side model parameters are estimated in the first step and not as part of the equilibrium solution of the game.

Combining the focal manufacturer's equilibrium promotion policy solution obtained in our second step analysis with the estimated starting distribution of inventory across households and the competitor's anticipated promotion schedule, we determine the optimal promotion timing and discrete promotion depth decisions for the focal manufacturer. Thus, our approach, similar to Silvia-Risso et al. (1999), requires information about the competitor's anticipated promotion schedule to come up with a promotion plan for the focal manufacturer.[2] In our empirical application, we do a sensitivity analysis and find that the proposed optimal promotion schedule of the manufacturer is fairly robust to small amounts of error in anticipating the competitor's promotion schedule. We estimate the starting inventory distribution across households using past purchase data available for a scanner panel of households. Thus, our approach requires access to such data.

---

[1] Similar two-step approaches have been used to study optimal pricing, entry and exit in the commercial aircraft industry (Benkard, 2004), optimal advertising in the Frozen Entrée category (Dube et al., 2005), and optimal intertemporal pricing for video-games (Nair, 2007).

[2] In our discussions with brand managers, we found that they believed they generally had a good idea of competitors' promotion schedules when they designed their brand's promotion calendar. The brand managers made these judgments based on past promotion practices and market intelligence.

In our empirical application, we solve for the equilibrium promotion policy of a leading manufacturer in the canned tuna category. The substantive insights from our model are that this manufacturer can increase profits by offering more frequent and shallower promotions. Using counterfactual experiments, we solve for the equilibrium promotion policy with different consumer discount rates and find, counter intuitively, that the manufacturer can be better off if the consumers are more forward-looking (discount the future less). This result arises because as consumers become more forward-looking, they stockpile more during promotions, leading to incremental sales and profit. Theoretical analysis shows a similar result for retailer stockpiling as being beneficial to the manufacturer (Lal et al. 1996; Desai et al. 2010). Interestingly, we obtain a similar result for consumer stockpiling. Our results also suggest why manufacturer price promotions intended to "load the consumer" can be profitable (Blattberg and Neslin 1990).

Table 1 positions this paper's contribution against the related literature on promotion policies of manufacturers or retailers. Neslin et al. (1995) use monthly aggregate data to study how the responses of retailers and consumers influence a manufacturer's optimal advertising and promotion plan by using a linear consumer and retailer side model and a dynamic supply side optimization model.[3] They find that it is in the manufacturers' interest to promote more steeply as consumer response to retailer promotions and retailer pass-through rates increase. Silva-Risso et al. (1999) use disaggregate data and a static consumer demand model to develop a manufacturer's weekly promotion planning model over a time horizon of one year. In their empirical application, they find that the manufacturer could substantially improve the profitability of its price promotion with lower discounts and a higher promotion frequency. In other related work, Kopalle et al. (1999) consider the potentially negative long-term effects of

[3] In related work, Sriram and Kalwani (2007) study the manufacturer's optimal advertising and promotion policies for determining spending levels for these variables.

price discounting in developing normative implications for price promotion schedules. With the SCAN*PRO model, they find that promotions have positive contemporaneous effects on sales and negative future effects on baseline sales. Tellis and Zufryden (1995) investigate a retailer's category promotion timing and depth taking the manufacturer's action as given. They show that the retailer's price discount is higher for a market with *switchers* only than for a market with *loyals* only. More recently, Shankar and Krishnamurthi (2009) study a retailer's decisions on both regular price changes and promotions with a retailer decision support model, RETPRICE. They investigate two different pricing policies, *Every Day Low Pricing* (*EDLP)* and *High-Low Pricing (HILO)*.

Our study differs from the above papers in the following ways. First, we consider the dynamic evolution of inventory of different consumer types so that a manufacturer can forecast consumers' response to alternative promotion schedules, and thus more effectively address the price discrimination problem. Second, we consider consumers' forward-looking behavior and its influence on manufacturer's optimal policy decisions. While the above papers model consumers' response to promotions with a reduced-form model, we adopt a dynamic structural demand model where consumers maximize the sum of discounted utilities over time.[4] Third, and most importantly, we allow the manufacturer's promotion decisions and consumers' purchase decisions to be endogenous to each other through a dynamic game between the manufacturer and consumers. For a durable product (video game), Nair (2007) solves for the equilibrium dynamic pricing strategy of a monopolist in a game with consumers. However, in contrast to durable goods, consumer's repurchasing and stockpiling behavior are critically important in solving the price discrimination problem of a manufacturer of packaged goods. Thus, interestingly, while we find that households'

[4] Ailawadi, Kopalle and Neslin (2005) also use a dynamic structural model but focus on manufacturer-retailer interactions while our structural model focuses on more accurately capturing consumer-manufacturer interaction.

forward-looking behavior is advantageous to the manufacturer, Nair (2007) finds that the manufacturer is worse off when consumers are more forward-looking. His result is consistent with the classic Coase conjecture for durable goods (Coase 1960). In contrast, we consider a frequently purchased product category, in which forward-looking behavior can enable better price discrimination between price-insensitive brand loyals and brand switchers, who may be more willing to stockpile the product during a promotion if they are forward-looking.

The rest of the paper comprises of the following steps. First, we lay out a framework for promotion decisions by a manufacturer and specify the model. Next, we present the two-step empirical estimation approach consisting of demand estimation and the numerical solutions algorithm of the dynamic game. Then we describe the data used for our empirical application. Subsequently, we discuss the results of our demand estimation, followed by a discussion of the equilibrium promotion policy derived from solving the game between consumers and the manufacturer. Finally, we run counterfactual experiments and conclude with a summary of our findings.

2 A Framework of Promotion Decision

In this section we lay out the two-step approach to address the promotion scheduling problem of a manufacturer. We first describe forward-looking consumers' purchase decisions with a dynamic demand model. We then specify a dynamic manufacturer-consumer game taking the estimates from the demand side model as given.

2.1 Structural Dynamic Demand Model Specification

We begin by specifying consumers' utility followed by a specification of the promotion and price process in the focal category. Then, we present consumers' dynamic programming problem of maximizing the discounted sum of current and future utilities.

2.1.1 Consumer Utility Function

Similar to Sun et al. (2003) and Erdem et al. (2003), for each week t, we define household i's utility as originating from the consumption of the focal category and outside goods subject to a budget constraint. We can then derive the indirect utility of consumer i purchasing quantity Q_{it} of brand $j \in \{1,2,...,J\}$ as follows (see details in Appendix A):

$$\begin{aligned} U_{it}(j,Q_{it}) = \sum_{j=1}^{J} \psi_{ij} C_{ijt} + \alpha_{1i}\left(price_{jt}, d_{ijt}, Q_{it}\right) + \alpha_{2i}\left(prom_{jt}, Dscnt_{jt}, d_{ijt}, Q_{it}\right) + \alpha_{3i}\overline{Inv_{it}} \\ + \alpha_{4i}\overline{Inv_{it}}^{2} + \alpha_{5i} y_{it} \end{aligned} \tag{1}$$

In the above equation, C_{ijt} is the consumption quantity of brand j in the focal category and ψ_{ij} denotes the consumption benefit from one unit of brand j. $price_{jt}$ is the regular retail price for brand j in week t, $prom_{jt}$ is an indicator which equals one if there is a price promotion for brand j in that week, and $Dscnt_{jt}$ is the value of the discount seen by consumers. d_{ijt} is a choice indicator that equals one if the household chooses brand j and is zero otherwise, and Q_{it} is the discrete quantity purchased by the household.[5] The terms containing $\overline{Inv_{it}}$ capture linear and quadratic components of inventory holding costs (cf. Erdem et al. 2003), where $\overline{Inv_{it}}$ is the simple average of the category inventory levels at the beginning and end of week t. y_{it} is household i's income in week t. We set y_{it} equal to zero without loss of generality since it is a constant across all choices and does not affect the household's purchase decision (cf. Erdem Imai and Keane 2003). $\alpha_{1i}, \alpha_{2i}, \alpha_{3i}$ and α_{4i} are the parameters corresponding to price, promotion and linear and quadratic inventory holding cost respectively. The utility when the consumer does not purchase any of the brands ($k = 0$, $Q_{it} = 0$ and $d_{i0t} = 1$) is:

$$U_{it}(0,0) = \left(\sum_{j=1}^{J} \psi_{sj} C_{ijt}\right) + \alpha_{3s}\overline{Inv_{it}} + \alpha_{4s}\overline{Inv_{it}}^{2} \tag{2}$$

[5] In our data, there were no instances where a household purchased multiple brands in a week.

If a household's inventory of brand j at the beginning of week t is Inv_{ijt} then the inventory at the end of week t (beginning of week $t+1$) is:

$$Inv_{ijt+1} = Inv_{ijt} + d_{ijt}Q_{it} - C_{ijt} \tag{3}$$

By using the following formulation for C_{it} (cf. Ailawadi and Neslin 1998), we allow the total weekly consumption in the category to be flexible so that it may increase with the household's inventory level:

$$C_{it} = \left(\sum_{j=1}^{J}\left(Inv_{ijt} + d_{ijt}Q_{it}\right)\right) * \left[\frac{\bar{C}_i}{\bar{C}_i + \sum_{j=1}^{J}\left(Inv_{ijt} + d_{ijt}Q_{it}\right)^{\varphi_s}}\right] \tag{4}$$

where, $\bar{C}_i$ is the average category consumption rate that can be inferred from the long-term purchase data of household i (Sun et al. 2003), and φ_s is the segment-level parameter to be estimated. Next, we determine a household's weekly consumption quantity for each brand j, C_{ijt}, by assuming that a household consumes brands from their inventory in the order of their benefits, ψ_{ij}, until their consumption reaches C_{it}. We model consumers' unobserved heterogeneity by allowing for latent segments. Thus, in the remainder of the paper, we replace subscript i in the model parameters with s (s = 1, 2,…, S) and also include subscript s in U_{it} to denote that household i belongs to a latent segment s of consumers.

2.1.2 Promotion and Price Process and Consumers Expectations

The observed sales are influenced by forward-looking consumers' expectations about each manufacturer's future price promotions. Consistent with the literature on forward looking behavior of consumers (Erdem et al. 2003; Nair 2007), we estimate a price promotion process from observed data and assume that consumers have rational expectations of this process when making purchase decisions. Note that a price promotion process constitutes of two components – promotion incidence and promotion depth (actual price).

Extant research on consumers' forward-looking behavior has modeled consumers' promotion incidence expectations using a First-order Markov (FOM) process where consumers' expectations depend on manufacturers' promotion activity only in the previous one period (Gonul and Srinivasan, 1996; Sun et al., 2003; Erdem et al., 2003; Hendel and Nevo 2006). In contrast, Liu and Balachander (2011) use a Proportional Hazards Model (PHM) that allows consumers' expectations to depend on promotion activity in the past several periods. They find that a PHM fits the promotion incidence data better, and that a structural model employing a PHM specification for consumers' promotion expectations explains sales data better than the FOM specifications of Sun et al., (2003) and Erdem et al., (2003).

We follow Liu and Balachander (2011) and assume that probability of brand *j* on promotion in week *t* is a function of the time elapsed since brand *j*'s previous promotion, T_{jt}. Thus we have:

$$T_{jt} = \begin{cases} 1 & if \quad prom_{jt-1} = 1 \\ T_{jt-1} + 1 & if \quad prom_{jt-1} = 0 \end{cases} \tag{5}$$

Further, we model the promotion incidence for brand *j* in week *t* using a PHM with a Weibull hazard specification and a covariate function. The covariate is a vector, T_{-jt}, of the time since last promotion of each of brand *j*'s competitors. This covariate vector allows us to capture the effects of competitive promotion activities. Thus, the PHM hazard function is

$$h(T_t) = h(T_{jt}) e^{T_{-jt}\beta_j} = \rho_j \gamma_j T_{jt}^{\gamma_j - 1} e^{T_{-jt}\beta_j} \tag{6}$$

The promotion probability of brand *j* in week *t* is given by:

$$Prob(prom_{jt} = 1 \mid T_t) = h(T_{jt}) * e^{T_{-jt}\beta_j} e^{-\int_0^{T_{jt}} h(u) e^{T_{-jt}\beta_j} du} \tag{7}$$

The brand-specific parameters β_j, ρ_j and γ_j are estimated using maximum-likelihood methods from the observed promotion incidence in the data.

Conditional on a promotion for brand j in week t, we assume that consumers' expectation of the promotion depth, $Dscnt_{jt}$, follows a normal distribution. The mean and variance of this distribution correspond to the brand's promotion depths observed in the data. Likewise, consumers' expectation of the brand's regular price, $price_{jt}$, is also assumed to follow a normal distribution, whose parameters are estimated from the data. Similar assumptions for promotion depth and regular price can be found in Sun et al. (2003).

2.1.3 Consumers' Dynamic Programming Problem

Forward-looking households make brand and quantity choices that maximize their discounted sum of utilities from the current and future periods. A household's choice at time t depends on the state of two sets of variables - a household-specific vector of inventory levels of the various brands at the beginning of each week, $Inv_{it} = \left[Inv_{i1t},\ Inv_{i2t}, \ldots,\ Inv_{iJt}\right]$, and a vector of times since the previous promotion for all brands, $T_t = \left[T_{1t},\ T_{2t}, \ldots,\ T_{Jt}\right]$, which determines expectations about future price promotions based on the PHM. Hence, a household's dynamic programming problem for purchase decisions can be described by a value function with these two state variables. When the household belongs to consumer segment s, its value function, VH_{ist}, must satisfy the following *Bellman equation*:

$$VH_{ist}\left(Inv_{it}, T_t\right) = \max_{k\in\{0,1,\ldots,J\},q} \begin{pmatrix} U_{ist}\left(k,q\right) + \varepsilon_{it}\left(k,q\right) \\ +\ \delta_H E_{prom_{t+1}(T_{t+1})}\left[VH_{ist+1}\left(Inv_{it+1}, T_{t+1}\right)\right] \end{pmatrix} \qquad (8)$$

where $\varepsilon_{it}\left(k,q\right)$ is the choice-specific unobserved utility of household i in week t and δ_H is the common discount factor for all households. Note that the choice options, k, include choosing a specific brand as well as the no purchase option (k = 0). Note that the expectation operator in equation (8) is calculated over the likelihood of future promotion incidence, depth of price promotions, and the regular prices as described in section 2.1.2.

We now introduce the notation $VH_{ist}(k,q|Inv_{it},T_t)$ for the observable component of the household's choice-specific value function when it chooses option *k* and quantity *q*.

$$VH_{ist}\left(k,q\middle|Inv_{it},T_t\right)=U_{ist}\left(k,q\right)+\delta_H E_{prom_{t+1}(T_{t+1})}\left[VH_{ist+1}\left(Inv_{it+1},T_{t+1}\right)\right] \qquad (9)$$

Thus, equation (8) can be rewritten as:

$$VH_{ist}\left(Inv_{it},T_t\right)=\max_{k\in\{0,1,\dots,J\},q}\left(VH_{ist}\left(k,q\middle|Inv_{it},T_t\right)+\varepsilon_{it}\left(k,q\right)\right) \qquad (10)$$

Consistent with previous marketing literature, we assume that $\varepsilon_{it}\left(k,q\right)$ follows an i.i.d. extreme value distribution. Then conditional on belonging to segment *s*, the probability that household *i* will choose option *l* and quantity *q* in week *t* is given by the standard multinomial logit formula:

$$\text{Prob}\left(d_{ilt}=1,Q_{it}=q\mid Inv_{it},T_t;\theta_s\right)=\frac{\exp\left(VH_{ist}(l,q|Inv_{it},T_t)\right)}{\sum_{k\in\{0,1,\dots,J\}}\sum_{Q}\exp\left(VH_{ist}(k,Q|Inv_{it},T_t)\right)} \qquad (11)$$

where, θ_s is the set of all model parameters corresponding to consumer segment *s*.

2.2 Dynamic Game between Focal Manufacturer and Households

We first specify the focal manufacturer's weekly profit and then describe the solution process for the dynamic game between the manufacturer and consumers.

2.2.1 Manufacturer's Weekly Profit and State Variables

Using households' choice probabilities from Equation (11), the expected sales of brand *l* in week *t* from a single household *i* in segment *s* that has a vector $\boldsymbol{\nu}=(\nu_1,\ \nu_2,\ \dots,\ \nu_J)$ of inventory levels of the various brands is:

$$Sales_{ilt}\left(Inv_{it}=\nu,T_t;\theta_s\right)=\sum_{Q}Q*\text{Prob}\left(d_{ilt}=1,Q_{it}=Q\mid Inv_{it}=\nu,T_t;\theta_s\right) \qquad (12)$$

Note that the manufacturer's expected sales in equation (12) depends on two groups of state variables - the time since previous promotion (which drives promotion incidence expectations), and household inventories, Inv_{it}. We assume that the manufacturer uses historical purchase data of a panel of households to estimate household inventories (using equation (3)) at the beginning of the promotion planning period. The use of estimated household inventory is common in many marketing models (e.g. Silvia-Risso et al. 1999; Erdem et al. 2003). Given the starting inventory, the manufacturer can estimate the probability distribution of future household inventories conditional on their promotion policy using equations (3) and (11). For computational purposes, we discretize the state space of households' inventories. With I discrete inventory levels for each of the J brands, the number of possible combinations of inventory levels is I^J. With S consumer segments, the manufacturer's state variables would be the number of households in each of $S*(I)^J$ categories that each represents households that belong to a particular segment and have a particular combination of inventory levels. Thus the dimensions of the state space increases exponentially with I making it difficult to increase I for greater precision.[6]

To reduce the state space dimensions to a manageable number, we assume that in each consumer segment s, the discrete inventory level of each brand j in week t follows a Poisson distribution with parameter λ_{sjt}. We denote the set of $S*J$ parameter values (J parameters in each of the S consumer segments) by the vector $\Lambda_t = \left[\lambda_{11t},...,\lambda_{1Jt},...,\lambda_{S1t},...,\lambda_{SJt}\right]$. The elements of Λ_t evolve over time with households' inventories and thus may be used as the manufacturer's state variables. This approach not only reduces the state space dimensions but it also allows us to improve accuracy by using a large number of discrete inventory levels.

[6] For example, with two consumer segments (S = 2), two brands (J = 2), and just three inventory levels (I = 3), high, medium and low, the number of segment-inventory household types is 2*3*3=18.

If we denote the set of $(I)^J$ possible combinations of inventory levels by $\{\nu\}$ then the number of households, $n_{\nu st}$, who belong to segment s and have a specific combination of the discrete inventories, $Inv_t = \nu$ in week t is given by:

$$n_{\nu st}\left(\Lambda_t\right) = n_s * \left(\prod_{j=1}^{J} f\left(Inv_{jt} = \nu_j \mid \lambda_{sjt}\right)\right); \quad \nu = (\nu_1, \ldots \nu_j, \ldots \nu_J) \in \{\nu\} \qquad (13)$$

where n_s is the number of households in consumer segment s and $f(.)$ is the pdf of the Poisson distribution. Now, the total expected sales of brand l in week t is:

$$Sales_{lt}\left(\Lambda_t, T_t\right) = \sum_{s} \sum_{\nu \in \{\nu\}} Sales_{ilt}\left(Inv_{it} = \nu, T_t; \theta_s\right) * n_{\nu st}\left(\Lambda_t\right) \qquad (14)$$

Next, we factor in the role of the retailer by making some standard assumptions about retailer behavior that are similar to those in Silvia-Risso et al. (1999). Specifically, we assume a "pay-for-performance" environment where the retailer has no incentive to indulge in forward buying. We also assume that the retailer passes through a certain proportion ($pthru\%$) of manufacturers' price promotion to consumers. Similar to Silva-Risso et al. (1999), we assume that a manufacturer of brand j may only offer a discount that is a multiple of 5 cents, $Mdscnt_{jt} \in \{0, 5, 10, \ldots, 30\}$. Hence, when consumers buy from the retailer, they see a discount of $Dscnt_{jt} = Mdscnt_{jt} pthru\%$. The manufacturer discount that is passed through to consumers comes off a regular retail price set by the retailer, $price_{jt} = Mprice_{jt}\left(1 + mkup\%\right)$, where $mkup\%$ is the retailer markup over manufacturers' regular wholesale price, $Mprice_{jt}$. We assume that the retailer's pass-through rate and markup are fixed and that they are the same for all brands. We also assume that a manufacturer has a constant marginal cost, $Mcost_j$ which is known.

We assume that the manufacturer's regular wholesale prices are predetermined and focus on optimizing the timing and depths of the promotional discounts to be offered by the manufacturer. Given our assumptions, when brand l's manufacturer gives a discount of $Mdscnt_{lt}$ in week t, its profit, π_{lt}, and the retailer's category profit, π_{Rt}, are:

$$\pi_{lt}\left(Mdscnt_{lt}\,|\Lambda_t, T_t\right) = Sales_{lt}\left(\Lambda_t, T_t\right) * \left(Mprice_{lt} - Mdscnt_{lt} - Mcost_l\right) \quad (15a)$$

$$\pi_{Rt}\left(Mdscnt_{lt}\,|\Lambda_t, T_t\right) = \sum_{j=1}^{J} Sales_{jt}\left(\Lambda_t, T_t\right) * \left(price_{jt} - Dscnt_{jt} - \left(Mprice_{jt} - Mdscnt_{jt}\right)\right) \quad (15b)$$

2.2.2 Dynamic Game between Manufacturer and Consumers

In order to obtain the focal manufacturer's optimal price promotion policy, we solve for the stationary Markov-perfect equilibrium, which is the fixed point of the dynamic game between the focal manufacturer and the set of households. The equilibrium has the following requirements: (1) For every realization, ($\Lambda_t = \Lambda$, $T_t = T$), of the state variables, the focal manufacturer's price promotion policy maximizes its discounted sum of future profits, subject to any constraints, given households' purchase strategy, and every household maximizes its discounted sum of future utilities given the focal manufacturer's promotion policy. (2) The expectations about future states that the focal manufacturer and households use to maximize their respective payoffs is consistent with the strategies of the focal manufacturer and the households. (Note that the strategy remains the same at (Λ, T) regardless of the time t because we solve for a stationary equilibrium.) An implication of requirement (2) above is that consumers' expectations of future promotions by the manufacturer are formed endogenously with the promotion policy chosen by the manufacturer. We had referred to this earlier as the *forward-looking consumer problem*. Further, consistent with requirement (1), Λ_t, representing the distribution of household inventory levels influences the manufacturer's strategy. We had referred to this earlier as the *price discrimination problem*.

We now specify the objective functions for the focal manufacturer and the households in the dynamic game. A promotion policy designed by the manufacturer may not be implementable if it does not gain retailer acceptance. In particular, as Equation 15b shows, the retailer is concerned about its profit from all brands and not just the profit from the focal manufacturer's brand. In order to incorporate retailer acceptance issues in the promotion policy, we require that the retailer's long-term discounted category profit under any new promotion policy chosen by the manufacturer be at least as much as the category profit under the existing promotion schedule, $\bar{\pi}_R$. Other researchers (e.g. Montgomery 1998, Silvia-Risso et al. 1999) have used similar constraints to incorporate retailer acceptance. Denote the equilibrium promotion policy of the focal manufacturer l as $Mdscnt_l^*|\Lambda,T$. Again, note that the policy does not have time subscripts because it is a stationary policy. Then, the equilibrium policy, $Mdscnt_l^*|\Lambda,T$, and the discounted value of profits or the value function, VM_{lt}, of brand l's manufacturer must satisfy the following constrained *Bellman* equation:

$$VM_{lt}\left(\Lambda_t, T_t\right) = \max_{Mdscnt_{lt}} \left(\pi_{lt}\left(Mdscnt_{lt} \middle| \Lambda_t, T_t\right) + \delta_M E_{prom_{t+1}(T_{t+1})}\left[VM_{lt+1}\left(\Lambda_{t+1}, T_{t+1}\right)\right]\right)$$

$$\textit{subject to:} \qquad \sum_{\tau=t}^{\infty} \delta_R^{\tau-1} \pi_{R\tau}\left(Mdscnt_{l\tau}\right) \geq \bar{\pi}_R \tag{16}$$

where δ_M and δ_R are the discount factors of the manufacturer and the retailer respectively.

In line with equation (8), the objective or value function in the dynamic game for household i belonging to segment s is given by the following *Bellman* equation:

$$VH_{ist}\left(Inv_{it}, \Lambda_t, T_t\right) = \max_{k \in \{0,1,\ldots,J\}, q} \begin{pmatrix} U_{ist}\left(k,q\right) + \varepsilon_{it}\left(k,q\right) \\ +\ \delta_H E_{prom_{t+1}(T_{t+1})}\left[VH_{ist+1}\left(Inv_{it+1}, \Lambda_{t+1}, T_{t+1}\right)\right] \end{pmatrix} \tag{17}$$

Unlike in Equation (8), when solving the game, we use the demand parameters from our first-step estimation to calculate $U_{ist}(k, q)$ in Equation (17). Note that the households' state variables in Equation (17) also includes the vector of parameters, Λ_t, that represent the distribution of household inventory levels. The rationale is that because Λ_t drives the focal manufacturer's strategy, expectation of the evolution of Λ_t allows consumers to anticipate future promotions by the focal manufacturer. On the other hand, for competing brands other than that of the focal manufacturer, household expectations of future promotions is driven by the vector of times since last promotion, T_t. The evolution of Inv_{it} is given by equation (3) while the evolution of the elements of T_t is given by equation (5). We determine the evolution of Λ_t empirically as part of the solution algorithm (see Appendix B and Technical Appendix C). Equation (17) leads to an equation similar to Equation (11) giving the household's probability of choice for each alternative at every state:

$$\text{Prob}\left(d_{ilt}=1, Q_{it}=q \mid Inv_{it}, \Lambda_t, T_t; \theta_s\right) = \frac{\exp\left(VH_{ist}(l,q \mid Inv_{it}, \Lambda_t, T_t)\right)}{\sum_{k\in\{0,1,\ldots,J\}} \sum_{Q} \exp\left(VH_{ist}(k,Q \mid Inv_{it}, \Lambda_t, T_t)\right)} \quad (18)$$

where $VH_{ist}(k,q|Inv_{it},\Lambda_t,T_t)$ is given by an equation analogous to Equation (9) as follows:

$$VH_{ist}\left(k,q \mid Inv_{it}, \Lambda_t, T_t\right) = U_{ist}\left(k,q\right) + \delta_H E_{prom_{t+1}(T_{t+1})}\left[VH_{ist+1}\left(Inv_{it+1}, \Lambda_{t+1}, T_{t+1}\right)\right] \quad (19)$$

3 Estimation and Numerical Solution Algorithm

3.1 First Step: Structural Dynamic Demand Model Estimation

Following past literature (e.g. Erdem et al 2003, Nair 2007), we set the households' common discount factor, δ_H, to 0.98. The likelihood of household i's brand and quantity choice decisions in the observed data between week t=1 and week $t=T$ is:

$$L_i = \sum_s \left(\omega_s \prod_t \prod_j \prod_q Prob(d_{ijt}=1, Q_{it}=q \mid Inv_{it}, T_t, \theta_s)^{d_{ijt} Ind(Q_{it}=q)} \right) \quad (20)$$

where, ω_s denotes the population proportion of segment s and $Ind(Q_{it}=q)$ is an indicator of household's observed quantity choice.

We solve the consumer's dynamic programming problem by using value iteration (Rust 1996) in conjunction with the Multidimensional Simplicial Interpolation method (Judd 1998; Keane and Wolpin 1994) to enumerate households' value function over the state space.

We calculate the value functions at a large number of Chebychev quadrature grid points in the state space and approximate the value at the other state space points.

3.2 Numerical Solution for Manufacturer’s Optimal Promotion Policy

Using the estimates of households’ demand parameters obtained in the first step, we solve for the fixed-point equilibrium of the game between the manufacturer and consumers numerically through policy iterations (Rust 1996, Nair 2007). We assign each household to a consumer segment s in a Bayesian fashion (Mazumdar and Papatla 2000). Again, we calculate the value functions at a large number of Chebychev quadrature grid points in the state space and approximate the households’ value function VH (Equation 17) and the manufacturer’s value function VM (equation 16) at the other state space points. We set the manufacturers and the retailer’s discount factors equal to the households’ discount factor $\left(\delta_M = \delta_R = 0.98\right)$. A unique Markov-Perfect equilibrium solution in pure strategies of the dynamic game need not exist. So we check the uniqueness of our solution by repeating the numerical solution algorithm with different starting values. The detailed steps of the algorithm can be found in the Appendix.

4 Data

We implement our model in the canned tuna category. We use store scanner data and household-level scanner panel purchase data from A. C. Nielsen for a period of 100 weeks from 1986 to 1987 in Sioux Falls, SD. The store scanner data provides information about prices and promotions and shows that the canned tuna category was heavily promoted during the data period. We restrict our attention to only one retail chain which implemented the same price promotion schedule across all its stores. Two leading brands accounted for over 92% of the market share: *StarKist 6.5 Oz.* and *Chicken of the Sea (COS) 6.5 Oz.* Table 2 presents some descriptive statistics for the two brands.

We estimate the demand side model with the first 50 weeks of data. In our estimation sample, we include households that made at least two purchases during those 50 weeks. The final sample consists of 11035 observations of 836 households. The maximum quantity of tuna purchased by any household in a given week is seven cans. Households bought one can on 27% of the purchase occasions and two cans on 51% of the purchase occasion. We set each household's average consumption rate $\bar{C}_i$ in equation (4) to equal the household's observed average weekly consumption during the first 50 weeks.[7] After estimating the demand side model parameters, we solve for the equilibrium promotion policy of *StarKist*. Based on COS's promotion schedule for the second 50 weeks, we develop *StarKist*'s optimal promotion schedule for that period using the equilibrium promotion policy. We chose *StarKist* as the focal manufacturer for illustrative purposes; a similar analysis can be carried out from the point of view of *COS*.

5 Results

5.1 Promotion Process and Consumer Demand

We first used a PHM that had competitors' promotion activities as covariates(equation 6). However, the promotion incidence data in our empirical application showed that the parameters corresponding to competitors are not significant. Therefore, we drop the covariates function when modeling consumer's promotion expectation. Table 3 reports the maximum-likelihood estimates (MLE) of the Weibull hazard model without covariates for the promotion process for the two brands in the tuna category. Figure 2 plots the expected promotion probability of the *StarKist* brand over the estimation horizon (week 1 to week 50) according to the fitted Weibull model. The dotted line represents the promotion indicator

[7] We simulate the starting inventory of households at the beginning of the first 50 weeks by using a procedure similar to Erdem et al. 2003. Specifically, we assume that households start with zero inventory t_0 weeks prior to the start of the first 50 weeks and simulate the households' purchases and consumption for the t_0 weeks to arrive at an expected starting inventory at the beginning of the first 50 weeks. We set t_0=100.

which equals one (spikes of the dotted line) when a promotion is actually observed for *StarKist*. The figure shows that in the week right after a promotion, there is a very low promotion probability for another promotion by the brand. This probability increases gradually during inter-promotion periods, increasing with the number of weeks since the previous promotion. Similar outcomes are obtained for the COS brand.

We estimate our demand side model by assuming that consumers' promotion incidence expectations follow the fitted Weibull models. Recall that consumers' expectation of the promotion depth and the regular price for each brand are assumed to be based on independent normal distributions whose means and variances correspond to the promotion depths and prices observed in the data. Table 4 presents the estimation results of the dynamic household choice model. A comparison between models with different number of latent consumer segments showed that the model with two consumer segments has the highest BIC. Thus, we only report the estimation results with two consumer segments.

The results show that about 21% of the households belong to segment one and have a strong preference for the *StarKist* brand (0.32). The negative sign for the *COS* brand in this segment (-0.63) indicates that c*eteris paribus*, these households would rather not consume any tuna than consume the *COS* brand. The households in segment two only have a marginal preference for the *StarKist* brand over the *COS* brand (0.65 vs. 0.39) but unlike segment one these households prefer consuming the *COS* brand over not consuming any tuna. The regular price and promotion coefficients for both segments are significant and have the expected sign. But the promotion sensitivity of segment two is far greater than that of segment one. Based on these findings, we refer to consumers in segment one as (*StarKist*) *loyals* as they have high brand preference and low promotion sensitivity, and consumers in segment two as *switchers* as they have weak brand preference and high promotion sensitivity.

We now turn our attention to households' inventory carrying cost. The estimates of the linear and quadratic inventory cost terms imply that the cost of carrying a stock of one can of tuna is about 1.0 cent and 3.1 cent per week for *loyals* and *switchers*, respectively. But at five cans, the holding costs become 17.5 cents and 15.0 cents per week, and at seven cans, it becomes 31.5 cents and 21.0 cents per week. Thus, in general, households with high inventory are less likely to make purchases or respond to promotions. However, we find that even at high inventory levels, the *switchers* have a relatively lower holding cost. Hence, they may be more likely to stockpile during promotion events even when they have high inventory. Finally, the flexible consumption rate parameters are close to one for both segments, which indicates low consumption flexibility (cf. Ailawadi and Neslin 1998). This result is to be expected as canned tuna has a long shelf life and consumers are less likely to increase their consumption due to inventory pressure.

5.2 Promotion Planning Implications

Given the first-step demand estimates, we numerically solve for the equilibrium promotion policy for *StarKist*. In solving the game, we assume the following: (1) The retailer's pass-through rate is 80% and the retailer markup is 20% (cf. Dhar and Hoch 1996, Jedidi, Mela and Gupta 1999); (2) The manufacturer's marginal cost is \$0.20 per unit (cf. Jedidi, Mela and Gupta 1999 for a similar assumption). We first present the equilibrium promotion policy for *StarKist*.[8] Next, we simulate *StarKist's* optimal promotion schedule over 50 weeks with the initial inventories of each household to be their inventories of the last period of the estimation sample and assuming the competitor *COS*'s promotion schedule to be its schedule observed in our data. As discussed in the introduction, this assumption implies that the focal manufacturer, *StarKist*, correctly anticipates *COS*'s promotion schedule. We

[8] We ran our solution algorithm with different starting values for the equilibrium promotion policy and households' value function. We obtained the same solution in all cases suggesting that our equilibrium solution may be unique.

subsequently analyze the sensitivity of *StarKist*'s equilibrium promotion schedule to errors in anticipating *COS*'s promotion schedule.

5.2.1. Optimal Promotion Policy

Figure 3 plots *StarKist's* equilibrium promotion policy as a function of households' mean inventory of *StarKist* when the mean inventory of *COS* equals 0.5 for both consumer segments and when $T_{COS,t+1} > 1$.[9] The figure indicates that it is optimal for *StarKist* to offer a 15-cent discount (*Mdscnt* = 15) when *switcher*s have low inventory and *loyals* have high inventory. This is the best situation for *price discrimination* between the two consumer segments as then the *switchers* are very responsive to a promotion while the *loyals* are unlikely to be in the market for a purchase. In general, the figure shows that *StarKist*'s equilibrium promotion depth is increasing with *loyals*' mean inventory and decreasing with *switchers*' mean inventory.[10]

We also find that irrespective of households' current inventory levels, Λ_t, *StarKist*'s equilibrium policy is to not offer a promotion in the same week when *COS* is expected to be on promotion ($T_{COS,t+1} = 1$). To understand the intuition behind this result, we solved for the equilibrium promotion policy without the retailer's incentive compatibility constraint in Equation (17). In this unconstrained case, we found that it is optimal for *StarKist* to offer a promotion for some states of Λ_t even when *COS* is on promotion in the same week. This result suggests that the constraint of ensuring the retailer at least his status quo category profit forces StarKist not to entertain a promotion policy which results in both brands being

[9] In the equilibrium solution of the game between *StarKist* and consumers, we find that when *COS* is not on promotion in a period ($T_{COS,t+1} > 1$), the equilibrium promotion policy stays constant across different values of $T_{COS,t+1}$. In other words, consumers' expectation of *COS*'s future promotions has little impact on *StarKist's* promotion decisions. This may be due to the relatively small market share of *COS* as all households have a preference for *StarKist* over *COS*.

[10] The displayed policy is representative of those obtained for other *COS* inventory levels.

promoted by the retailer in the same week. This finding highlights the importance of accounting for the retailer's interests when designing the optimal promotion schedule.

5.2.2. Optimal Promotion Schedule

After solving for the equilibrium promotion policy with the first 50 weeks of data, we simulate *StarKist's* optimal promotion schedule for the following 50 weeks by assuming the same promotion schedule for *COS* as is observed in the data.

Figure 4 juxtaposes the simulated optimal promotion schedule and the observed promotion schedule for *StarKist*, and rows 1 and 2 in Table 5 compare the summary statistics of the two schedules. The optimal promotion schedule recommends more frequent promotions and shallower discounts.[11] The average optimal inter-promotion interval is 29% shorter than the average inter-promotion interval observed in the data for *StarKist* (7.1 weeks vs. 10.0 weeks). Also, the average optimal promotion depth is 51% lower than the average observed promotion depth (10.3 cents vs. 21.2 cents). *StarKist's* sales and profit increase by about 8% under the optimal schedule, suggesting that the observed promotion schedule may be sub-optimal.[12]

Figure 5 shows *StarKist's* optimal promotion schedule, *COS*'s observed promotion schedule and the simulated mean inventory of the two consumer segments. Note that the mean inventory for *loyals* is relatively stable over time whereas the mean inventory for *switchers* greatly increases whenever a brand is on promotion and it gradually decreases in the inter-promotion periods. This is because *switchers* strategically adjust their purchase incidence to match with promotion events and then stockpile if there is a deep discount. Therefore, under the optimal schedule, *StarKist* has a price promotion when the mean

---

[11] Interestingly, Silvia-Risso et al. (1999) reach a similar conclusion in their empirical study.

[12] These outcomes are conditional on the assumed pass-through rate, markup and manufacturer's marginal cost. A sensitivity analysis with other reasonable values of pass-through rate showed that the manufacturer finds it more effective to promote when the pass-through rate is high, resulting in more frequent promotions. This finding is consistent with Silva-Risso (1999). Sensitivity analyses involving different values of markup and marginal cost gave very similar substantive results.

inventory for *switchers* is sufficiently low so that the profit from incremental sales to *switchers* is greater than the loss from the lower net margin to *loyals*. Also, consistent with the equilibrium promotion policy, the depth of the discount is lower for higher levels of *switchers'* inventory.

6 Counterfactual Simulations

6.1 Impact of Consumers' Forward-Looking Behavior

To understand the impact of consumers' forward-looking behavior on the manufacturer's optimal promotion policy and associated profits, we solve for *StarKist's* equilibrium policy in the second step when households are more or less forward-looking, i.e. $\delta_H = \{0, 0.50, 0.98\}$. Lower values of the discount factor denote that households are less forward-looking (more myopic). Thus, $\delta_H = 0$ corresponds to the case when households are fully myopic and only interested in the utility in the current period. Figure 6 plots the simulated promotion schedules with different δ_H, and Table 5 compares these schedules. The results indicate that *StarKist* should promote less often and offer shallower discounts when consumers are more myopic. The rationale is that myopic consumers stockpile less during price promotions resulting in lower incremental sales from switchers. Consequently, price promotions become less attractive leading to a reduction in their frequency along with a preference for shallower discounts whenever price promotions are profitable.

Figure 7 plots *StarKist's* value function (*VM*) for the different values of δ_H. The value function increases with households' discount rate. In other words, *StarKist* is better off when consumers are more forward-looking. At first glance this may seem counterintuitive as one might expect that a manufacturer would prefer to have myopic and naïve consumers from whom it may potentially extract more surplus. The reason behind this interesting finding is as follows. When consumers are forward-looking, sales decrease during regular-price weeks as consumers (most of whom are switchers) adjust their purchase incidence to try and coincide with promotion events. The consequent lower sales in regular-price weeks has a negative effect on *StarKist's* discounted profit. However, forward-looking *switchers* also stockpile and purchase large quantities during promotions, which leads to incremental sales due to brand switching. This has a positive effect on *StarKist's* discounted profit. In our application, the net effect of consumers becoming more forward-looking is positive when *StarKist* implements the appropriate optimal promotion schedule. Therefore, *StarKist* actually benefits from consumers' forward-looking behavior. As noted earlier, this result is consistent with manufacturers offering sales promotion with the objective of loading up the consumer's pantry and thus taking the (forward-looking) consumer out of the market (Blattberg and Neslin 1990). Further, our results represent a contrast from the opposite result obtained by Nair (2007) in a durable product category.

We now investigate *StarKist's* profit when it wrongly assumes that consumers are not forward-looking. That is, we compare the outcomes when *StarKist* implements promotion schedules that are developed by wrongly assuming that households are fully myopic ($\delta_H = 0$) or partially myopic ($\delta_H = 0.5$) when, in fact, they are forward-looking ($\delta_H = 0.98$). Table 6 shows the comparative results. The sales of *StarKist* are 15.9% lower and the present discounted value (PDV) of profit is 9.4% lower when *StarKist* assumes that δ_H is 0.5. The situation is worsened when *StarKist* assumes that δ_H is zero. Then sales are 28.6% lower and PDV of profit is 20.9% lower. The sales and profit are lower because of the lost opportunity to facilitate stockpiling by switchers with more frequent price promotions or deeper discounts and because such loading up of switchers leads to more effective price discrimination. This result highlights the importance of incorporating a structural model of consumers' forward-looking purchasing behavior when developing the manufacturer's promotion schedule.

6.2 Errors in Anticipating Competing Brand's Promotion Schedule

In our application, we assumed that *StarKist* correctly anticipates *COS*'s future promotion schedule. As indicated earlier, our discussions with brand managers reveal that they have a reasonable idea of when to expect a competitor's promotion based on past history or from market intelligence. However, in general, the brand manager's expectation about a competitor's future promotion schedule may be subject to error. Therefore, we test the robustness of *StarKist's* promotion policy when there are relatively small errors in anticipating *COS*'s promotion schedule. Specifically, we consider scenarios where the entire 50-week promotion schedule of *COS* may get shifted either forward by one week or backward by one week compared to the anticipated *COS* promotion schedule. Let us denote the anticipated *COS* promotion schedule used for planning purposes by *StarKist* to be C0 (referred to as the "Plan" Schedule in Table 7). The promotion schedule of *COS* that is shifted one week forward (backward) from the Plan schedule will be referred to as C0+1 (C0-1). Table 7 compares *StarKist's* profits under the different cases. The first row corresponds to the case where *StarKist* develops its promotion schedule by believing that *COS* will implement C0 but then *COS* actually goes on to implement C0+1 or C0-1. The second and third rows correspond to the cases where *StarKist* believes that *COS* will implement C0+1 and C0-1 respectively, but then *COS* actually goes on to implement C0. As Table 7 shows, there are only small changes in *StarKist's* profit across these different cases (+0.03% to +1.18%). This suggests that the optimal promotion policy may be fairly robust to small amounts of uncertainty about the competitor's promotion schedule.

7 Conclusion and Limitations

A significant element of a brand's marketing plan in the CPG industry is the promotion calendar for the planning period. An important concern in developing the promotion calendar is that consumer expectations of price promotions can induce stockpiling and deal-to-deal buying thereby reducing the profitability of price promotions (Blattberg and Neslin 1990). In this paper, we demonstrate how a CPG manufacturer can incorporate consumers' forward-looking behavior and dynamic response to promotions into a solution process for developing the optimal timing and depth of in-store price promotions over a planning horizon. Using inputs from household-level panel data, we obtain the optimal promotion schedule of a focal manufacturer as the equilibrium solution of a dynamic game between households and the manufacturer. We formulate the dynamic game using demand parameters obtained in an estimation step incorporating a dynamic structural model of consumer choice. Our solutions process can enable CPG manufacturers to enhance the productivity of price promotions thereby increasing the efficiency and effectiveness of marketing activities.

Our solution framework is applied to the promotion depth and timing decision problem for a leading brand, *StarKist*, in the canned tuna category. We find that the optimal promotion schedule for the assumed cost and pass-through parameters suggests that *StarKist* should offer more frequent but shallower promotions over the 50-week planning horizon. We also find that *StarKist* could have substantially increased its profit if, instead of the actual schedule, it had adopted the optimal schedule. Interestingly, our simulated results show that consumers' forward-looking behavior has a positive impact on *StarKist*'s discounted profit. In other words, if forward-looking behavior leads some consumers to stockpile during deals, the result can mean an increase in profit for the manufacturer. In other results, we find that *StarKist* should promote less frequently and less steeply as consumers become more myopic. Not surprisingly, if *StarKist* incorrectly assumes forward-looking consumers to be myopic, its profit is hurt because of a suboptimal promotion schedule. Finally, our sensitivity analysis shows that even though the recommended promotion schedule for *StarKist* is based on an anticipated promotion schedule for the competing *COS* brand, uncertainty about *COS*'s promotion schedule has little impact on the profitability of the optimal promotion schedule.

Our study has the following limitations. First, we solve for the optimal promotion policy for a single manufacturer and do not simultaneously consider the strategic choice by the competing brand of its promotion policy. An advancement would be to model the strategic choices of multiple manufacturers simultaneously as the solution of a game. This interesting but difficult empirical problem is left for future research. Second, although we incorporate the role of the retailer in our setup, we do not structurally model the retailer's decisions. Specifically, we assume that the retailer passes through a certain proportion of the trade discount to the consumers and this assumed pass-through rate can be based on the manufacturer's experience with the retailer. A possible future extension of our study is to explicitly model the retailer's decision thereby allowing the retailer to potentially coordinate the in-store promotions of competing brands. Finally, another interesting extension would be to incorporate the decision of other elements of promotion such as displays, feature advertisements and coupons into the promotion planning decision model.

Table 1: Positioning this Paper in the Literature

	Structurally model forward-looking consumer's purchase decision	Solve dynamic game	Marketer(s)	Optimal policy	Industry
Erdem et al. (2003)	Yes				CPG
Sun et al. (2003)	Yes				CPG
Neslin et al. (1995)	No	No	Manufacturer	Advertising expense and promotion discount	CPG
Tellis et al. (1995)	No	No	Retailer	Timing and depth of promotion	CPG
Silva-Risso et al. (1999)	No	No	Manufacturer	Promotion calendar	CPG
Kopalle et al. (1999)	No	No	Manufacturer and Retailer	Prices	CPG
Ailiwadi et al. (2005)	No	Yes	Manufacturer and Retailer	Regular price and promotion	CPG
Shankar et al.(2009)	No	No	Retailer	Regular price and promotion	CPG
Nair (2007)	Yes	Yes	Manufacturer	Diminishing prices	Durable goods
This paper	Yes	Yes	Manufacturer	Timing and depth of promotion	CPG

Table 2: Descriptive Statistics of Tuna Data

Brand Name	Market Share	Mean Price per Can (6.5 Oz) $.	Promotion Frequency	Average inter promotion time (weeks)	Mean Promotion Depth($)
StarKist	67.79%	0.755	0.129	8	0.185
COS	24.35%	0.753	0.169	6	0.199

Table 3: Estimation Results of Promotion Incidence Process

Parameters	***StarKist***	***COS***
Baseline- γ	0.13(0.001)***	0.19(0.01)***
Baseline-ρ	1.61(0.82)**	3.74(0.81)***
-LL	44.46	

Notes: Standard errors in brackets
****: significant at p = 0.01
****: significant at p = 0.05

Table 4: Estimates of Structural Household Demand Model

Parameters	**Segment one (*loyals*)**	**Segment two (*switchers*)**
StarKist	0.323(0.083)***	0.654(0.139)***
COS	-0.632(0.083)***	0.394(0.138)***
Regular Price ($)	-1.860(0.124)***	-2.547(0.204)***
Promotion Discount ($)	1.149(0.107)***	5.504(0.162)***
Inventory cost_linear ($/Can)	-0.010(0.004)***	-0.031(0.013)***
Inventory cost_quadratic ($/Can^2)	-0.005(0.002)***	-0.005(0.007)
Flexible consumption(ϕ)	0.822(0.044)***	0.956(0.044)***
Segment probability	0.291	0.709(0.017)***
-LL	395.638	

Notes: Standard errors in brackets
****: significant at p = 0.01
**: significant at p = 0.05
*: significant at p = 0.10

Table 5: Comparative Statistics of Equilibrium Promotion Schedules with Different Household Discount Rates

		Promotion frequency	**Mean promotion depth (retail discount) (cents)**	**Average inter-promotion time (weeks)**
Observed Schedule		0.10	21.2	10.0
Equilibrium Schedule with	$\delta_H = 0.98$	0.14	10.3	7.1
	$\delta_H = 0.50$	0.12	10.0	8.3
	$\delta_H = 0.00$	0.10	9.6	10.0

Table 6: Profit Implications of *StarKist* Misperceiving Household Discount Rate (True Household Discount Rate = 0.98)

	Observed schedule	Equilibrium Schedule Assuming		
		δ_H = 0.98	δ_H = 0.50	δ_H = 0.00
Sales of *StarKist* (Cans)	4532	4889	4114	3491
PDV of profit ($)	1205.25	1302.89	1179.61	1030.52

Table 7: *StarKist*'s Profits under Uncertainty about *COS*'s Promotion Schedule

	Actual Schedule Used by *COS*		
	COS "Plan" Schedule	*COS* Schedule Shifts Forward 1 week from "Plan"	COS Schedule Shifts Backward 1 week from "Plan"
StarKist Profit ($) Assuming *COS* Uses "Plan" Schedule	1302.89	1318.31 (1.18%)	1315.72 (0.98%)
StarKist Profit ($) Assuming *COS* Schedule Shifts Forward 1 Week from "Plan"	1304.60 (0.13%)		
StarKist Profit ($) Assuming *COS* Schedule Shifts Forward 1 Week from "Plan"	1303.32 (0.03%)		

Notes: Percentage change from profit with actual COS schedule in brackets

Figure 1A: Illustrative Example: Dynamic Evolution of Household Inventory

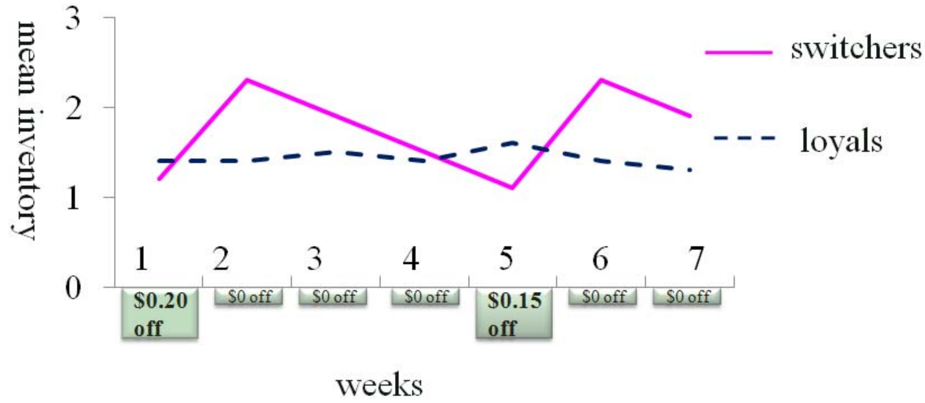

Figure 1B: Illustrative Example: Dynamic Evolution Of Household Inventory After Policy Change

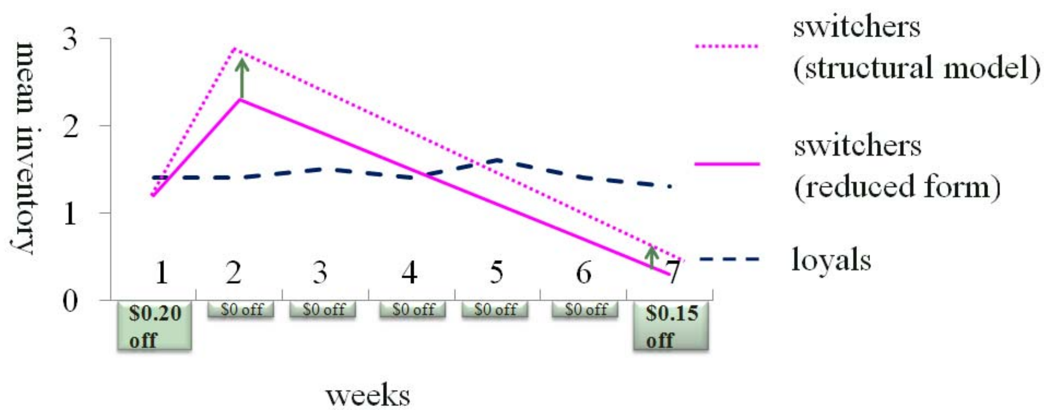

Figure 2: Promotion Incidence Process – PHM Model

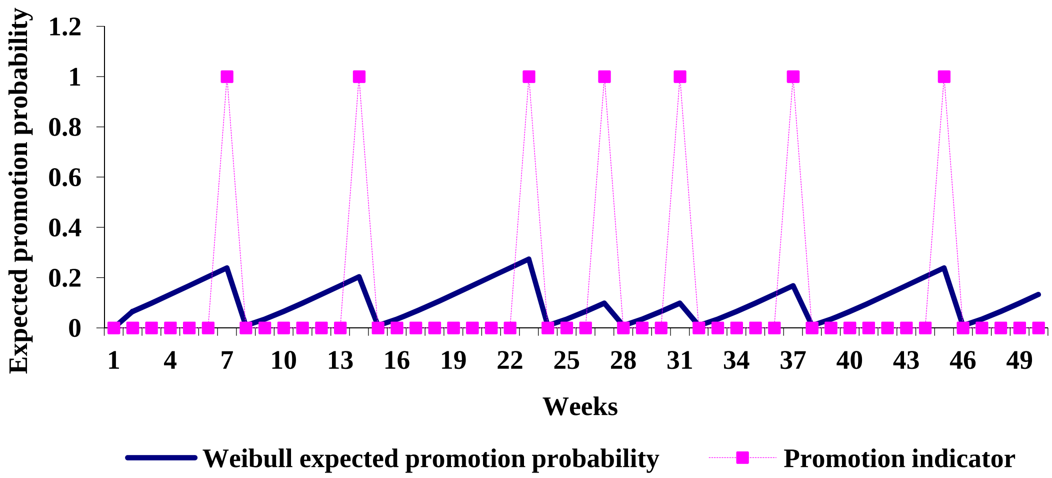

Figure 3: Equilibrium Promotion Policy With Forward-Looking Households (δ_H =0.98) (Mean Inventory Of *COS* Is 0.5 In Segment One And Segment Two)

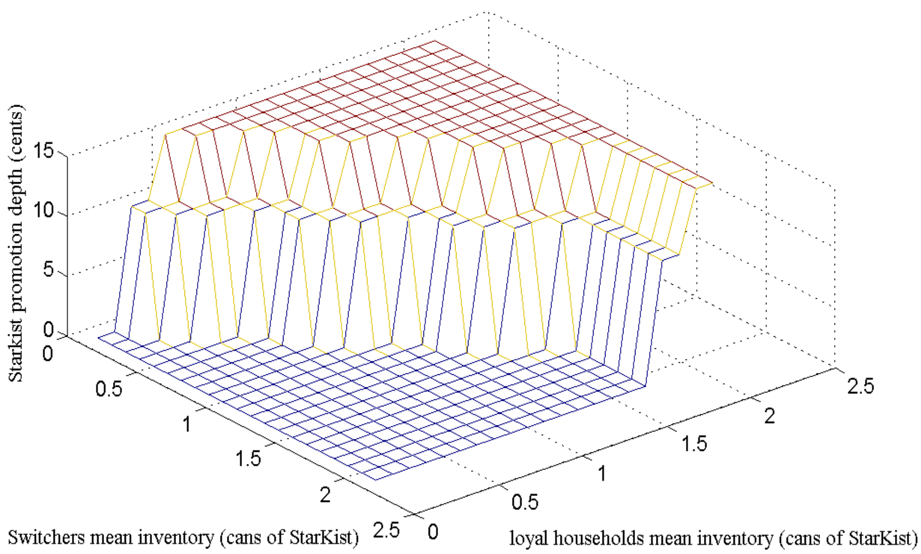

Figure 4: Equilibrium Promotion Schedule vs. Observed Schedule

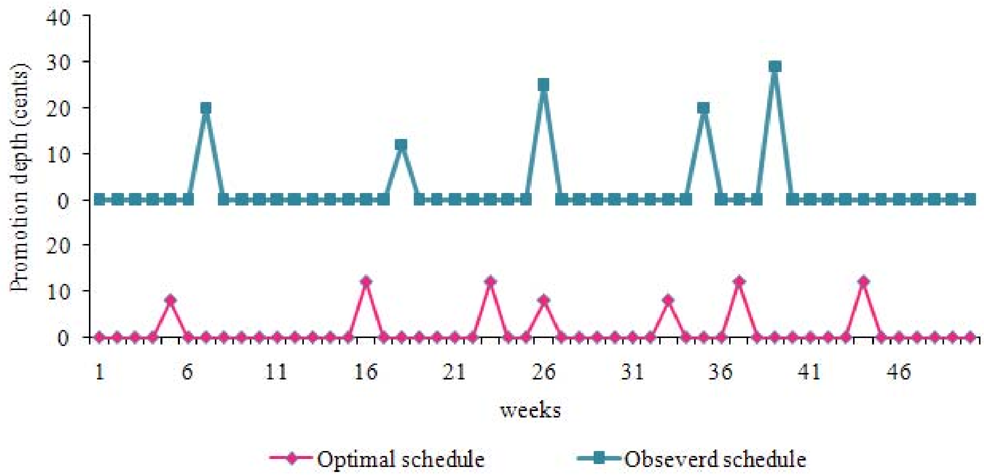

Figure 5: Mean Inventory With Simulated Equilibrium Promotion Schedule

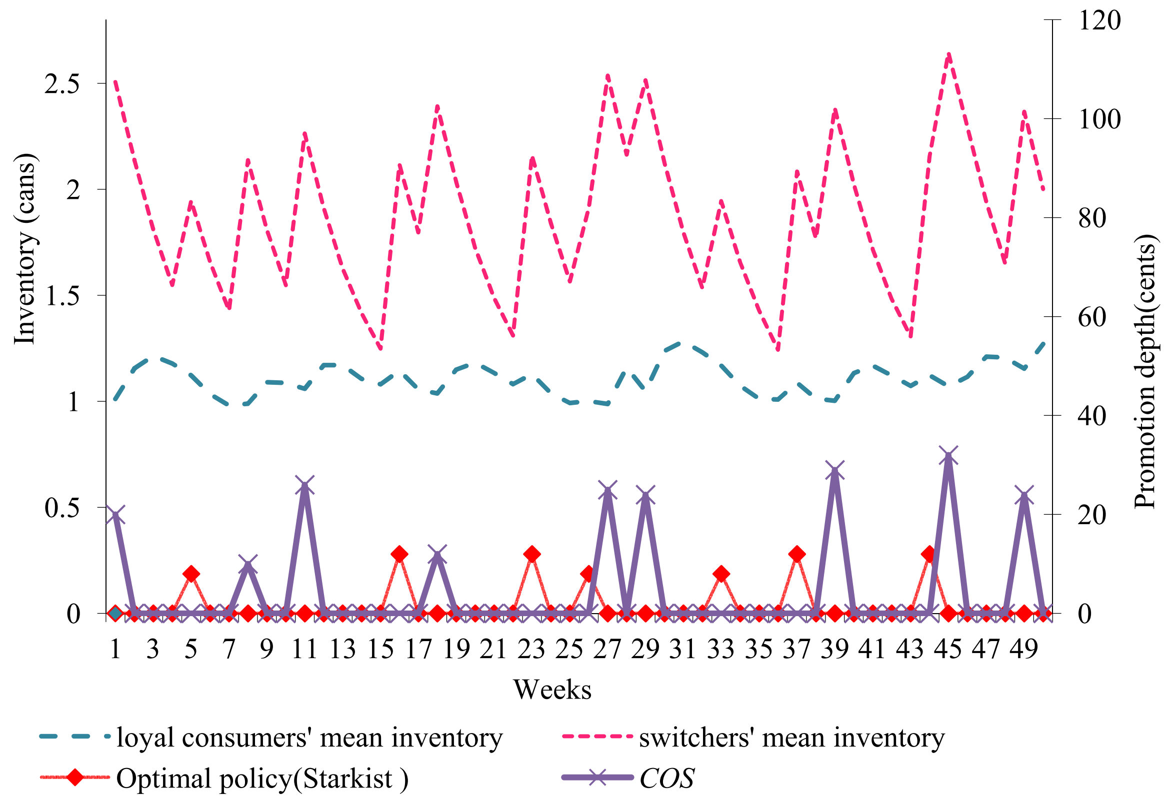

Figure 6: Equilibrium Promotion Schedule With Different δ_H

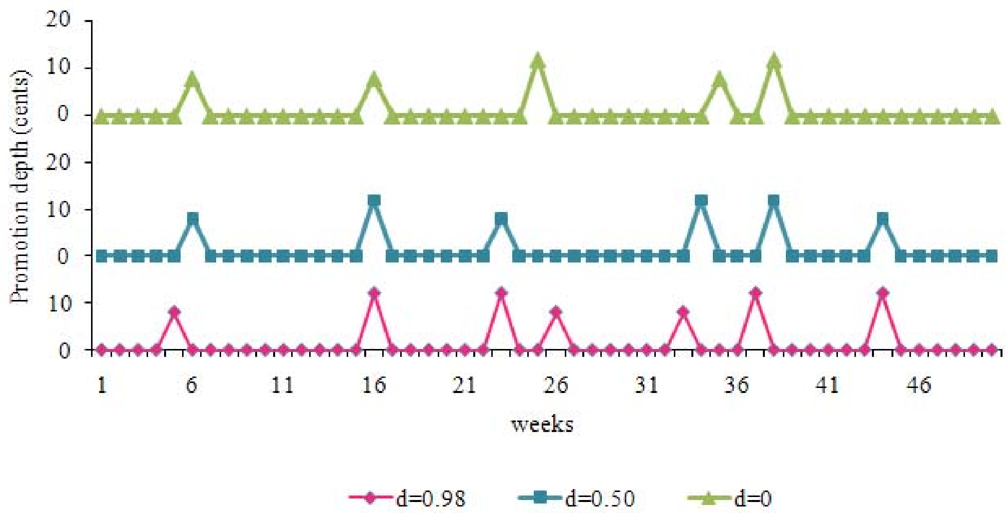

Figure 7: *StarKist*'s Value Function With Different δ_H Levels
(Mean Inventory Of *COS* is 0.5 In Segment One And Segment Two)

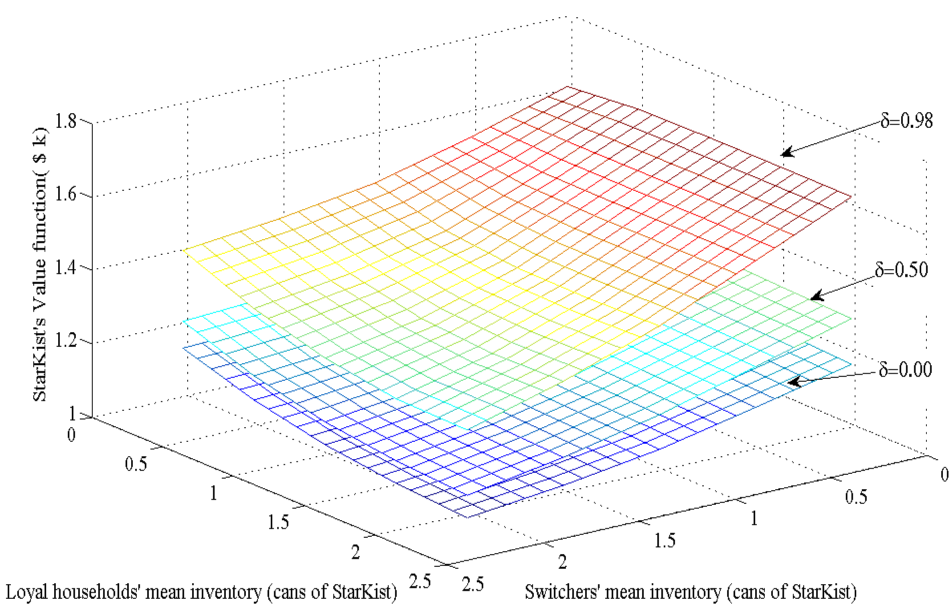

References

Ailawadi, K. L., Kopalle, P. K. and Neslin,S.A. 2005, "Predicting Response to a Major Policy Change: Combining Game-Theoretic and Empirical Analysis," *Marketing Science*, **24** (1), 12-24.

Ailawadi, K. L., and Neslin,S.A.. 1998 “The Effect of Promotion on Consumption: Buying More and Consuming It Faster.” *Journal of Marketing Research*, Vol. 35, No. 3. 390-398.

Benkard, C. L. (2004): “A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft,” *Review of Economic Studies*, 71, 581–611.

Blattberg, R. C. Neslin, S. A.. 1990. “Sales Promotion: Concept, Methods, and Strategies.” Prentice-Hall, Englewood Cliffs, NJ.

Coase, R. 1960. “The Problem of Social Cost.” *Journal of Law and Economics*, 1, 1-44.

Desai, P. S., O. Koenigsberg and D. Purohit (2010). “Forward Buying by Retailers,” *Journal of Marketing Research*, 67 (February), 90-102.

Dhar, S. K. and Hoch, S. J.. 1996. “Price discrimination using in-store merchandising.” *Journal of Marketing* **60** (1) 17–31.

Dubé, J. P., Hitsch, G., & Manchanda, P. (2005). An empirical model of advertising dynamics. *Quantitative Marketing and Economics*, 3(2), 107–144.

Erdem, T., and Keane, M. P.. 1996. ‘‘Decision-Making Under Uncertainty: Capturing Choice Dynamics in Turbulent Consumer Goods Markets,’’ *Marketing Science*, **15**(1), 1–21.

Erdem, T., Imai, S., Keane, M.P.. 2003. “Consumer price and promotion expectations: Capturing consumer brand and quantity choice dynamics under price uncertainty.” *Quantitative Economic and Marketing* **1** 5–64.

Gönül, F. and Srinivasan, K. 1996. “Impact of consumer expectations of coupons on purchase behavior.” *Marketing Science*. **15** 262–279.

Hendel, Igal and Aviv Nevo. 2006. “Measuring the implications of sales and consumer stockpiling behavior.” *Econometrica*, Vol. 74, No. 6, pp. 1637-1673.

Jedidi, Kamel, Carl F. Mela, Sunil Gupta. 1999. “Managing advertising and promotion for long-run profitability.” *Marketing Science* **18** (1) 1–22.

Judd, Kenneth L. 1998. *Numerical Methods in Economics*, The MIT Press.

Keane, Michael P., Kenneth I. Wolpin. 1994. “Solution and estimation of dynamic programming models by simulation.” *Review of Economic Statistics*. **76** 684–672.

Kopalle, P., Mela, C.F. and Marsh, L.. 1999. "The Dynamic Effect of Discounting on Sales: Empirical Analysis and Normative Implications", *Marketing Science* 18(3) 317-332.

Lal, R., J. D. C. Little and J. M. Villas-Boas. 1996. "A Theory of Forward Buying, Merchandising, and Trade Deals," *Marketing Science*, 15 (1) 21-37.

Liu, Y. and Balachander, S. 20011. "Dynamic Brand and Quantity Choice with a Hazard Model of Promotion Expectation." Purdue University, working paper.

Lucas, Robert 1976. "Econometric Policy Evaluation: A Critique.". *Carnegie-Rochester Conference Series on Public Policy* **1**: 19–46.

Mazumdar, Tridib, Purushottam Papatla. 2000. "An investigation of reference price segments." *J.Marketing Res.* **37** 246-258.

Nair, H.. 2007 "Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games" Quantitative Marketing and Economics 5(239-292).

Neslin, S.A., Powell, S. and Schneider Stone, L., 1995. "The Effects of Retailer and Consumer Response on Optimal Manufacturer Advertising and Trade Promotion Strategies." *Management Science* **41** 5, pp. 749–766.

Rust, J.. 1996: "Numerical Dynamic Programming in Economics," in Handbook of Computational Economics, ed. by H. Amman, D. Kendrick, and J. Rust. Amsterdam: Elsevier-North Holland, Ch. 14, 620-729.

Silva-Risso, J., Bucklin, R. and Morrison, D., 1999. "A Decision Support System for Planning Manufacturers' Sales Promotion Calendars." Marketing Science 18 3, pp. 274–300.

Shankar, Venkatesh and Lakshman Krishnamurthi (2009), "RETPRICE: A Retailer Pricing and Promotion Decision Support Model," Working Paper, Texas A&M University, College Station, TX.

Sriram, S., and Kalwani, M. (2007), "Optimal Advertising and Promotion Budgets in Dynamic Markets with Brand Equity as a Mediating Variable," *Management Science*, Vol. 53 (1),46-60.

Sun, B, Neslin,S.A. and Srinivasan,K. 2003 "Measuring the impact of promotions on brand switching when consumers are forward looking." Journal of Marketing Research. XL 389-405.

Tellis, G. J. and Zufryden, F.S.. 1995. "Tackling the retailer decision maze: Which brands to discount, how much, when and why?" *Marketing Science* **14**(3) 271-299.

APPENDIX A

Consumer Utility Function

Similar to Sun et al. (2003) and Erdem et al. (2003), for each week t, we define household i's observable component of utility, U_{it}, from the consumption of the focal category and outside goods as follows:

$$U_{it} = \left(\sum_{j=1}^{J} \psi_{ij} C_{ijt} \right) + \zeta_i Z_{it} \tag{A1}$$

where, C_{ijt} is the consumption quantity of brand j (j = 1, 2,…,J) in the focal category, and Z_{it} is the consumption quantity of outside goods. The parameters ψ_{ij} and ζ_i denote the consumption benefits from one unit of brand j and one unit of the outside good, respectively.

Each week, the household makes a brand-quantity purchase decision in the focal category which we denote by $d_{ijt}Q_{it}$, where d_{ijt} is a choice indicator that equals one if the household chooses brand j and is zero otherwise, and Q_{it} is the discrete quantity purchased by the household.[13] A household may also choose not to purchase in the focal category in which case, $d_{ijt} = 0$ for $j \in \{1, 2, ..., J\}$. If we denote the household income in week t by y_{it}, then the following budget constraint must be satisfied:

$$y_{it} = \left(\sum_{j=1}^{J} \left(price_{jt} - prom_{jt} Dscnt_{jt} \right) d_{ijt} Q_{it} \right) + \left(\iota_{1i} \overline{Inv_{it}} + \iota_{2i} \overline{Inv_{it}}^{2} \right) + Z_{it} \tag{A2}$$

The first component on the right-hand side of Equation (A2) corresponds to the purchase in the focal category where $price_{jt}$ is the regular retail price for brand j in week t, $prom_{jt}$ is an indicator which equals one if there is a price promotion for brand j in that week, and $Dscnt_{jt}$ is the value of the discount seen by consumers. The second component in equation (A2) is the

[13] In our data, there were no instances where a household purchased multiple brands in a week.

inventory holding cost for the focal category, which is quadratic in the average inventory level (cf. Erdem et al. 2003). $\overline{Inv_{it}}$ is the simple average of the category inventory levels at the beginning and end of week *t*. Further, if a household's inventory of brand *j* at the beginning of week *t* is Inv_{ijt} then the inventory at the end of week *t* (beginning of week *t+1*) is:

$$Inv_{ijt+1} = Inv_{ijt} + d_{ijt}Q_{it} - C_{ijt} \tag{A3}$$

and the household's average category inventory level in week *t*, $\overline{Inv_{it}}$, is:

$$\overline{Inv_{it}} = \sum_j Inv_{ijt} + 0.5(\sum_j d_{ijt}Q_{it} - \sum_j C_{ijt}) \tag{A4}$$

The last term in Equation (A2) is the expenditure on outside goods with the price per unit of outside good normalized to one. Replacing this term by combining Equations (A1) and (A2) gives us the following expression for U_{it}:

$$U_{it} = \left(\sum_{j=1}^{J} \psi_{ij} C_{ijt}\right) - \zeta_i \left(\sum_{j=1}^{J} \left(price_{jt} - prom_{jt} Dscnt_{jt}\right) d_{ijt} Q_{it}\right) - \zeta_i \left(\iota_{1i} \overline{Inv_{it}} + \iota_{2i} \overline{Inv_{it}}^2\right) + \zeta_i y_{it} \tag{A5}$$

We now make a few adjustments to Equation (A5). We model consumers' unobserved heterogeneity by allowing for latent segments. Thus, we replace subscript *i* in the model parameters with *s* (*s* = 1, 2,…, *S*) and also include subscript *s* in U_{it} to denote that household *i* belongs to a latent segment *s* of consumers. Second, we allow price and promotion to have different effects on consumers' utility. Finally, we set y_{it} equal to zero without loss of generality since it is a constant across all choices and does not affect the household's purchase decision (cf. Erdem Imai and Keane 2003). Hence, Equation (A5) is revised to:

$$U_{ist} = \left(\sum_{j=1}^{J} \left(\psi_{sj} C_{ijt} + \alpha_{1s}\left(price_{jt} d_{ijt} Q_{it}\right) + \alpha_{2s}\left(prom_{jt} Dscnt_{jt} d_{ijt} Q_{it}\right)\right)\right) + \alpha_{3s} \overline{Inv_{it}} + \alpha_{4s} \overline{Inv_{it}}^2 \tag{A6}$$

The choice-specific utility for brand $k \in \{1, 2, ..., J\}$ and quantity $Q_{it} = q$ is:

$$U_{ist}(k,q) = \left(\sum_{j=1}^{J} \psi_{sj} C_{ijt}\right) + \alpha_{1s}\left(price_{kt} q\right) + \alpha_{2s}\left(prom_{kt} Dscnt_{kt} q\right) + \alpha_{3s} \overline{Inv_{it}} + \alpha_{4s} \overline{Inv_{it}}^2 \tag{A7}$$

The utility of the no-purchase option ($k = 0$, $Q_{it} = 0$ and $d_{i0t} = 1$) is:

$$U_{ist}(0,0)=\left(\sum_{j=1}^{J}\psi_{sj}C_{ijt}\right)+\alpha_{3s}\overline{Inv_{it}}+\alpha_{4s}\overline{Inv_{it}}^{2} \tag{A8}$$

APPENDIX B

Numerical Algorithm of the Dyamic Game

The numerical solution algorithm for the Markov-Perfect equilibrium in pure strategies of the dynamic game presented in Section 2.2.2 is as follows:

Step 1: Start with a guess $Mdscnt^{0}(\Lambda,T)$ of $Mdscnt^{*}(\Lambda,T)$, and a guess, $VH_{s}^{0}(k,q\,|\,Inv,\Lambda,T)$, of the corresponding household's value function in equation (19) for each combination of values of the state variables and for each segment. For each combination of values of the state variables and for each segment, compute households' purchase probability (equation 18), update consumers' brand specific inventory for all segments using equation (3), and fit them into Poisson distribution to yield Λ^{0}, the initial estimate of the transition state from Λ.

Step 2: Let $Mdscnt^{m}(\Lambda,T)$ be the current (m^{th}) iterate of $Mdscnt^{*}(\Lambda,T)$, and $VH_{s}^{n}(Inv,\Lambda,T)$ represent the current (n^{th} iterate) of the household value function (equation 17). Using $VH_{s}^{n}(Inv,\Lambda,T)$, compute consumers' purchase probability (using equations 18 and 19), and update consumers' brand specific inventory for all segments using equation (3), and fit them into Poisson distribution to yield Λ^{n}, the transition state from Λ (see online technical appendix for details)·

Step 3: Solve for $VH_{s}^{n+1}(Inv,\Lambda,T)$ using Equation (17).

Step 4: Iterate on Steps 2 and 3 till $|VH_{s}^{n+1}-VH_{s}^{n}|<\varepsilon$ for all s and for all grid points and set $VH_{s}^{*}=VH_{s}^{n+1}$ and $\Lambda^{*}=\Lambda^{n}$.

Step 5: Solve for the value functions, $VM_l(\Lambda,T)$ and $VR(\Lambda,T)$, for the focal manufacturer l and the retailer respectively that satisfy their respective Bellman equations:

$$VM_l(\Lambda,T) = \pi_l\left(Mdscnt^m \middle| \Lambda,T\right) + \delta_M E_{prom(T')}\left[VM_l\left(\Lambda^*,T'\right)\right] \quad \forall T'$$

$$VR\left(\Lambda,T\right) = \pi_R\left(Mdscnt^m \mid \Lambda,T\right) + \delta_M E_{prom(T')}\left[VR\left(\Lambda^*,T'\right)\right] \quad \forall T'$$

π_l and π_R are given by equations (15a) and (15b) respectively. In computing *Sales*(.) for these equations, equations (12), (14), (17), (18) and (19) are used in conjunction with VH_s^* from Step 4.

Step 6: Compute improved policy, $Mdscnt^{m+1}(\Lambda,T) = \arg\max_{Mdscnt \geq 0}[VM_l(\Lambda,T)]$, subject to $VR(\Lambda,T) \geq \bar{\pi}_R$.

Step 7: If $Mdscnt^{m+1}(\Lambda,T) = Mdscnt^m(\Lambda,T)$, stop, and set $Mdscnt^*(\Lambda,T) = Mdscnt^{m+1}(\Lambda,T)$; else, go back to Step 2 with $Mdscnt^{m+1}(\Lambda,T)$, VH_s^* and Λ^* as initial guesses.