MICHAEL LEWIS*

Despite the proliferation of loyalty programs in a wide range of categories, there is little empirical research that focuses on the measurement of such programs. The key to measuring the influence of loyalty programs is that they operate as dynamic incentive schemes by providing benefits based on cumulative purchasing over time. As such, loyalty programs encourage consumers to shift from myopic or single-period decision making to dynamic or multiple-period decision making. In this study, the author models customers' response to a loyalty program under the assumption that purchases represent the sequential choices of customers who are solving a dynamic optimization problem. The author estimates the theoretical model using a discrete-choice dynamic programming formulation. The author evaluates a specific loyalty program with data from an online merchant that specializes in grocery and drugstore items. Through simulation and policy experiments, it is possible to evaluate and compare the long-term effects of the loyalty program and other marketing instruments (e.g., e-mail coupons, fulfillment rates, shipping fees) on customer retention. Empirical results and policy experiments suggest that the loyalty program under study is successful in increasing annual purchasing for a substantial proportion of customers.

The Influence of Loyalty Programs and Short-Term Promotions on Customer Retention

Loyalty programs have long been an important element of customer relationship management for firms in travel-related industries such as airlines, hotels, and rental cars. Information technology that enables firms to practice individual-level marketing has facilitated the spread of loyalty programs into such diverse industries as gaming, financial services, and retailing (Deighton 2000). Accordingly, academic researchers have begun to study loyalty programs. Behaviorally oriented researchers, such as Soman (1998) and Kivetz and Simonson (2002), have studied the effect of delayed incentives on consumer decisions. Zhang, Krishna, and Dhar (1999), Kim, Shi, and Srinivasan (2001), and Kopalle and Neslin (2003) have proposed analytical models to study the impact of loyalty programs in categories with different structures. This study contributes to the literature that is focused on empirically measuring response to loyalty programs (Drèze and Hoch 1998; Sharp and Sharp 1997).

Loyalty programs that base rewards on cumulative purchasing are an explicit attempt to enhance retention. Such programs encourage repeat buying and thereby improve retention rates by providing incentives for customers to purchase more frequently and in larger volumes. However, dynamically oriented promotions, such as loyalty programs, represent just one possible technique for increasing customer retention. Repeat buying may also be encouraged through various means such as short-term discounts on merchandise or reduced shipping charges. Therefore, it is important to develop models that can simultaneously estimate the influence of dynamic and current factors on long-term customer behavior. In this article, I report on a methodology for assessing the relative impact of loyalty-based promotions, short-term promotions, and individual-level factors on customer purchasing over time.

There is only limited and contradictory published empirical work on the value of loyalty programs. A relevant study by Sharp and Sharp (1997) analyzes individual-level data by using a one-period switching model to measure the ability of a loyalty program to alter normal repeat-purchase rates. Unfortunately, the study's results are inconclusive. In contrast, Drèze and Hoch (1998) report on a category-

specific loyalty program that results in increases for both the specific category and total store traffic. The contrasting findings are consistent with the lack of consensus on the ability of loyalty programs to increase customer retention (see Dowling and Uncles 1997). It should be emphasized that studies that question the value of loyalty programs (e.g., Dowling and Uncles 1997; Sharp and Sharp 1997) are largely based on research that uses single-period switching models. Additional research with models that more fully replicate the dynamics of consumer response is needed to judge the effectiveness of dynamically oriented loyalty programs.

For a frequency program to be effective in increasing loyalty, it must have a structure that motivates customers to view purchases as a sequence of related decisions rather than as independent transactions. That is, the structure must give customers an incentive to adopt a dynamic perspective. O'Brien and Jones (1995) suggest that the major factors that customers consider when evaluating programs are the relative value of awards and the likelihood of achieving a reward. Furthermore, the likelihood of achieving a reward is a function of cumulative buying thresholds and time constraints. These design elements (e.g., thresholds, rewards, time constraints) combine with individual-level requirements and preferences to determine the customer's expected benefits of participating in a loyalty program.

A special characteristic of loyalty programs is that their attractiveness may change dynamically with a customer's decisions. As purchases are made, both the customer's investment in the program and the customer's likelihood of earning a reward increase. Conversely, when a customer decides not to purchase in a given period, the likelihood of earning a reward decreases, because the customer moves no closer to the reward threshold, and the time left to earn rewards shrinks. The assessment of a program's attractiveness is further complicated because customers usually have imperfect knowledge of their future requirements and of the marketing policies of the firm. These dynamic factors are a challenge in the modeling of customer response to loyalty programs.

This study empirically estimates the impact of a reward program and other elements of the marketing mix on customer buying behavior over time by developing a model that replicates dynamic consumer response to a loyalty program. In contrast to previous models, the current model considers the impact of previous purchasing activity and customer expectations. The underlying behavioral assumption is that a reward program can motivate customers to base their purchasing decisions both on the current environment and on a long-term goal of achieving a frequent buyer reward. In other words, an effective reward program can encourage customers to make decisions that maximize expected utility over an extended time horizon rather than at each purchase occasion. This assumption is consistent with previous findings in the literature that expectations of the future can affect consumers' current-period decisions (e.g., Boulding et al. 1993; Lemon, White, and Winer 2002).

The empirical section of this article uses individual-level customer data from an Internet grocer to develop a dynamic model of customer retention. The model identifies the key factors that influence customers to make repeat purchases over time. A model based on an appropriate dynamic behavioral specification can assist an overall customer management strategy in two respects. First, it is a means for assessing the influence of dynamic promotions such as loyalty programs that operate over an extended period. Second, by controlling for the dynamic elements of a customer's decision problem, the model removes a possible source of bias and may provide better estimates of the effects of short-term marketing tactics.

The analysis employs a research methodology known as discrete-choice dynamic programming (Eckstein and Wolpin 1989). The approach is based on the idea that a customer's observed sequence of decisions may be interpreted as the solution to a dynamic optimization problem. These methods are established in economics (Eckstein and Wolpin 1989; Rust 1994) and are becoming more common in marketing (Erdem, Imai, and Keane 2003; Erdem and Keane 1996; Gönül and Shi 1998; Gönül and Srinivasan 1996; Sun, Neslin, and Srinivasan 2004). Dynamic programming methods are ideal for analyzing individual choices that are based on both current and future expected benefits. A loyalty program that bases awards on the level of purchasing over a specified period is a prime example of such a decision problem. A further benefit of dynamic programming methods is that the estimated coefficients can be used to conduct simulations that replicate the consumer's dynamic decision process.

The primary contribution of this research is a framework for modeling the influence of a reward program and other marketing instruments on customer retention. Firms have multiple options for their promotional budgets, so models that can quantify the long-term effects of loyalty programs and other options (e.g., pricing, coupons, shipping fees) can help the firm justify its choices. Although most database marketing applications focus on tasks such as customer scoring that are designed to maximize the profitability of single-period mailing efforts (Bult and Wansbeek 1995), the current research focuses on customers' response to a range of marketing instruments over an extended period. The model provides the means to support multicampaign direct marketing in environments in which customers have a dynamic orientation.

In terms of substantive findings, the results suggest that the loyalty program under examination is successful in changing customer behavior and in motivating customers to increase purchasing. In addition to a statistically significant estimate for the loyalty reward parameters, formulations that assume that customers are dynamically oriented fit better than do models that do not include a dynamic structure. The simulation experiments also provide a means for estimating the magnitude of response to the program.

The remainder of the article is organized as follows: In the next section, I develop a dynamic programming model of customer response to a loyalty program. Next, I discuss model estimation issues and present empirical results. I then report the results of policy experiments that are designed to assess the impact of the loyalty program and other promotions on customer retention. I subsequently address limitations and future research opportunities and conclude with a discussion of key findings and managerial applicability.

MODELING THE DYNAMICALLY ORIENTED CUSTOMER

The model is intended to replicate the decision-making process of a customer when dynamic incentives exist. As

such, the model needs to include the factors that affect the attractiveness of a customer's options at a given time and a structure that captures dynamic considerations. I begin with a customer who makes repeated choices from among J mutually exclusive alternatives in each of T discrete periods, where the choice set is defined in terms of purchase quantity (I dropped customer-specific indexes for clarity). The choice set also includes a no-purchase option. This model allows a customer to choose the extent to which he or she participates in a program over a sequence of occasions.

The quantity decision is likely to be based on marketing factors such as prices and on individual-level factors such as inventories. A myopic decision maker would select from the options 1, ..., J and consider the relative benefits associated with each option, denoted as R_j. The standard assumption is that consumers select the option that yields the optimal benefits for the purchase occasion. However, a loyalty program may provide an incentive for customers to view weekly transactions as a sequence of related decisions. If rewards are earned over time, a multiple-period objective function is appropriate for the customer's decision problem. The objective function of a dynamically oriented customer[1] is

$$(1) \qquad \max E \left[\sum_{t=1}^{T} \alpha^{t-1} \sum_{j=1}^{J} R_j(t) d_j(t) | S(t) \right],$$

where $R_j(t)$ is the single-period reward associated with option j at time t, $d_j(t)$ is an indicator variable that is set equal to 1 if option j is chosen at time t and is set equal to zero otherwise, and α is a one-period discount factor. An important element of this formulation is that the reward functions $R_j(t)$ are conditional on the state of the environment, $S(t)$. The state space, $S(t)$, is a vector of information about the environment that is relevant to the customer's forward-looking optimization problem. The state space may consist of marketing-mix elements, such as the pricing environment in a given week, and customer-specific information, such as cumulative purchases.

The customer's decision problem involves selecting the level of buying in each period to maximize the expected discounted utility for the remainder of the time horizon.[2] The value function, V, is the maximum value of discounted expected utility over the horizon.

$$(2) \qquad V[S(t)] = \max_{d_j(t)} E \left[\sum_{\tau=t}^{T} \alpha^{\tau-t} \sum_{j=1}^{J} R_j(\tau) d_j(\tau) | S(\tau) \right].$$

The value function is the utility that can be achieved over the time horizon if the customer selects the optimal sequence of quantities. The value function may also be

written in terms of alternative specific expected value functions as $V[S(t)] = \max_{j \in J} V_j[S(t)]$, where the alternative specific value function $V_j[S(t)]$ is the expected value of the customer selecting option j at time t and then selecting optimal actions subsequently. Furthermore, the alternative specific value functions satisfy the Bellman (1957) equation:

$$(3) \qquad V_j[S(t)] = E[R_j(t)|S(t)] + \alpha E\{V[S(t+1)]|S(t), d_j(t) = 1\},$$

$$\text{when } t = 1, ..., T - 1,$$

and

$$V_j[S(T)] = E[R_j(T)|S(t)], \text{ at } t = T.$$

The form of the alternative specific value functions implies that decisions are based on both the immediate utility provided by an option and the expected total future utility from subsequent purchases and loyalty rewards. A key detail is that the expected future benefits from period t + 1 forward may depend on the option selected. In many loyalty programs, an important element of the state space, which evolves according to customer actions, is the level of cumulative spending. In general, the evolution of the state space is a function of both the current environment, $S(t)$, and the choices made by the customer, $d_j(t)$, over time.

The discount factor, α, defines the degree to which consumers are forward-looking or dynamically oriented. A positive α ($0 < \alpha < 1$) implies that customers consider future benefits when making current decisions. If α is equal to zero, customers completely discount future benefits. In the context of a loyalty program, a positive α means that customers partially base current purchasing decisions on the expectation of earning a reward. When α is equal to zero, customers do not consider future benefits.

Customer Choices and State Description

As I previously stated, the setting for the empirical analysis is an online retailer that specializes in nonperishable grocery and drugstore items. I model customer choices in terms of four purchase levels during weekly periods: no order, small order, medium order, and large order. Small orders are ones in which merchandise costs less than $50. Medium orders are ones in which merchandise costs at least $50 but less than $75. Large orders cost at least $75. The options available in each week are indexed by j, where

j = no: no purchase in a given week,
j = sm: purchase of a small basket of goods,
j = med: purchase of a medium basket of goods, and
j = lrg: purchase of a large basket of goods.

The order options are defined to be mutually exclusive so that $\Sigma d_j(t) = 1$. This discretization is not arbitrary but is based on the shipping fees charged by the firm.

I use weekly purchasing decisions to update a state variable that tracks cumulative expenditures Cum_{it} for individual i at time t. This variable is maintained by rounding the actual cumulative expenditures to the nearest $33 increment.[3] For purposes of the customer's expectations of the future state, the variable is updated according to Equation

---

[1]The model is developed for a representative customer. Given that response to this type of program is likely to vary across the population, I estimate a finite mixture model, which is detailed in the "Model Estimation" section.

[2]The loyalty program under study includes a time constraint. Specifically, the program is based on customers' annual purchasing, and points are not carried over into future annual cycles. This is an important element of the program's structure because it dictates the formulation of the dynamic program used to model the customer's optimization problem.

---

[3]These increments roughly reflect the sizes of the average orders observed in each category.

4.[4] For example, if a customer has $500 in cumulative expenditures at a given time t and makes a small purchase, the customer is assumed to behave as if the future value of Cum_{it+1} is $533.

$$(4) \qquad Cum_{it+1} = Cum_{it} \sum_{j=1}^{4} \left[(100/3)(j-1)d_j(t) \right].$$

I use the cumulative buying variables to determine when a loyalty-based reward is earned. A reward variable, L_{it}, is set to one if a purchase causes a reward threshold to be crossed and is set to zero otherwise. This occurs when a customer whose previous level of cumulative purchases is less than a reward threshold makes a purchase that brings cumulative purchasing over the next reward threshold. For the program under study, the reward thresholds are set at $1,000, $1,500, and $2,000 worth of annual purchases. If a customer with cumulative purchases of $925 makes a purchase worth $75 or more, L_{it} equals one, but L_{it} equals zero if only a $50 purchase is made. Rewards are paid in blocks of 500 frequent-flyer miles for the $1,000 and $1,500 thresholds and 1000 miles for the $2,000 level. Therefore, when the $2,000 threshold is crossed, the L_{it} variable equals two, to reflect that the reward earned is double the value of the previous rewards.

I do not treat the cumulative expenditure variables as a direct component of customers' reward functions. Rather, I define variables that account for the potential benefits of being an experienced customer. My strategy is to use indicator variables to define broad categories of previous activity. I define two indicator variables (FM_1 and FM_2) that define classes of previous purchasing activity for a given customer as follows: $FM_1 = 1$ if $Cum_{it} \geq \$300$, and $FM_1 = 0$ otherwise, and $FM_2 = 1$ if $Cum_{it} \geq \$500$, and $FM_2 = 0$ otherwise. The variables are intended to capture dynamic effects beyond the loyalty program. For example, as cumulative purchasing increases in terms of number and size of transactions, it is likely that customers gain a measure of expertise that effectively reduces transaction costs.

Another element of the state space is the amount of time since a customer's previous order. This measure is intended to control for possible inventory or customer attrition effects and is adjusted each period as follows: Dur_{it} = weeks since a purchase was made; if $j \in$ {small, medium, large} (i.e., a purchase is made), then $Dur_{it+1} = 1$; and if $j =$ no, then $Dur_{it+1} = Dur_{it+1}$. I subsequently define the actual implementation of the duration variable, which is capped at nine weeks through a series of categorical variables.

Before the empirical analysis, the role of the duration variable is not well defined. It may capture inventory-related decision making if the natural buying cycle is greater than the weekly cycle used in the model. However, if the probability of observing a purchase decreases as the time since the previous purchase increases, the duration variable will capture customer attrition. The duration categories are indexed by k, where k = 1 indicates a purchase

within the previous two weeks, k = 2 indicates between three and five weeks since the previous purchase, k = 3 indicates between six and eight weeks since the previous purchase, and k = 4 indicates no purchase in the previous nine weeks. The level of duration for a representative consumer at time t is denoted by a series of indicator variables, $Rcn(k)_{i,t}$, that are defined as mutually exclusive, such that $\Sigma_{k=1}^{4} Rcn(k)_{i,t} = 1$.[5]

The Rcn and FM variables are consistent with the recency, frequency, and monetary (RFM) value variables that are popular in direct-marketing response models. The RFM variables reflect important individual-level differences and have been used in previous dynamic programming models of consumer behavior (Gönül 1999; Gönül and Shi 1998).

An important marketing-mix element for direct retailers is shipping and handling fees. A unique element of the data is that the online retailer in this study experimented with various shipping fee structures during data collection. The shipping fee structure in place at a given time is important, because it provides incentives and penalties for various order sizes and thus may affect both order incidence and distribution of order sizes. Specifically, graduated shipping fee structures impose an element of nonlinear pricing on buying decisions. For example, a shipping fee structure used in the data is a schedule that charges $6.97 in fees for orders less than $75 and provides free shipping for orders greater than $75. This type of policy provides an incentive for customers to increase order sizes. Customers may find themselves in situations in which the total order cost (i.e., the sum of merchandise and shipping fees) is less for a basket with a higher merchandise total, because the marginal cost of an incremental item (that causes the merchandise order to cross a threshold) can be negative. Similarly, a schedule that involves fees that increase as order size increases provides an incentive for customers to purchase smaller baskets of goods. The cost of shipping each order size under each shipping fee schedule is provided in Table 1.

Shipping variables, $SH(h)_t$, indicate the shipping fee structure used at a given time, where h = 1 indicates a mixed/decreasing fee schedule, h = 2 indicates increasing fee schedule (high), h = 3 indicates increasing fee schedule (low), and h = 4 indicates free shipping for all order sizes. For example, $SH(4)_t = 1$ indicates that the firm used a free shipping policy at time t. The shipping and handling fee structures are mutually exclusive, such that $\Sigma SH(h)_t = 1$.

The online retailer under study carries in excess of 10,000 items. Therefore, rather than use specific prices in the model, I use a summary measure that describes the pricing environment. This price variable is defined as the average price of the 50 top-selling items during the data collection period in a given week and is calculated as $P_t = (\Sigma_{item=1}^{50} P_{item,t})/50$. The price variable accounts for the relative attractiveness or competitiveness of the firm's pricing in a given week. An additional price-related factor is e-mail

---

[4]For purposes of the loyalty rewards, this variable is updated exactly on the basis of previous levels of buying and the quantity purchased. I use the discretization to approximate the customer's expectations of the next period state.

[5]Recency also could be constructed as a continuous variable. The use of a series of categorical variables is a simplifying assumption that helps maintain computational tractability. In addition, I based the selection of the break points on patterns of incidence observed in the data. For example, purchase incidence drops from 9.5% when duration is eight weeks to 5% when it is nine weeks; after nine weeks, the rate stabilizes near 3%.

promotions. The firm's practice has been to send occasional e-mail coupons to existing customers for a 10% discount on purchases during a given week. The variable C_{it} is set equal to one if a 10% discount coupon is sent in week t and is set equal to zero otherwise.

In addition to the role of current-period prices, expectations of future prices may also be salient. Recall that in Equation 3, the value function involves an expectation of the future value of program participation. When expected future prices are high, customers expect to pay a significant premium to obtain the points needed for a reward. Expectations of future prices have been considered by several marketing researchers (Kalwani et al. 1990; Narasimhan 1989) and have been estimated with dynamic programming models (Gönül 1999).

Expectations of future prices are modeled as a function of current prices with an ordered logit structure. Specifically, expectations of the future pricing environment are modeled by means of a discrete three-point distribution that enables customers to anticipate an average, low (one standard deviation below the average), or high (one standard deviation above the average) price in the next period. The probabilities are estimated with the following equations, where P_t is the current-period price, and θ_P, θ_{Low}, and θ_{High} are parameters to be estimated:

$$(5)\quad Pr(P_{t+1} = P_{Low}) = \frac{\exp(\theta_{Low} + \theta_P P_t)}{[1 + \exp(\theta_{Low} + \theta_P P_t)]},$$

$$(6)\quad Pr(P_{t+1} = P_{ave}) = \frac{\exp(\theta_{High} + \theta_{Low} + \theta_P P_t)}{[1 + \exp(\theta_{High} + \theta_{Low} + \theta_P P_t)]}$$
$$- \frac{\exp(\theta_{Low} + \theta_P P_t)}{[1 + \exp(\theta_{Low} + \theta_P P_t)]},$$

and

$$(7)\quad Pr(P_{t+1} = P_{High}) = 1 - \frac{\exp(\theta_{High} + \theta_{Low} + \theta_P P_t)}{[1 + \exp(\theta_{High} + \theta_{Low} + \theta_P P_t)]}.$$

The complete set of state and control variables is summarized in Table 2, which also includes the variables constructed from the state variables and the levels the variables may take.

Utility Functions

The single-period utilities associated with each of the J options are given subsequently, where $R_{j,i}(t)$ corresponds to the utility to individual i of option j at time t:

$$(8)\quad R_{no,i}(t) = \beta_{no} + \varepsilon_{no,it},$$

Table 1
SHIPPING AND HANDLING FEE STRUCTURES

Shipping and Handling Policy	Description	Small Basket: $0–$50	Medium Basket: $50–$75	Large Basket: Greater Than $75
1	Mixed/decreasing	$4.99	$6.97	$0
2	Increasing (high)	4.99	6.97	8.95
3	Increasing (low)	2.99	4.99	4.99
4	Free	0	0	0

Table 2
STATE AND CONTROL VARIABLES

Variable	Description	Levels
State Variables and Related Covariates		
Cum	Cumulative purchasing	$33 dollar increments to $3,033 (91 levels)
FM(f)	Indicates how to classify customers according to levels of cumulative spending	Not a state variable; constructed from the cumulative spending variable
L	Indicates whether a loyalty reward is earned in period t	Binary variable; not a state variable but a function of cumulative expenditures and customers' actions at time t
Dur	Time (in weeks) since the previous purchase	Ranges from 1 to 9 weeks (9 levels)
Rcn(k)	Indicator for recency categories	Not a state variable (4 levels)
P	Price	Continuous for the current period
C	E-mail coupon	Binary variable (2 levels)
SH(h)	Shipping and handling schedule	Four different shipping and handling fee schedules used during data collection (4 levels)
T	Individual-specific time (in weeks) relative to the program's finite (52-week) horizon (more formally, should be denoted as t_i)	$5 < t \le 52$ (Weeks 1–5 are initialization period)
$E(P_{t+1})$	Expectations of future prices	Function of current-period prices
Control Variable		
$d_j(t)$	Purchase quantity decision	4 levels (no buy, small, medium, and large)

(9) $R_{sm,i}(t) = \beta_{sm} + \beta_{P,sm}P_{it} + \beta_{C,sm}C_{it} + \sum_{h=1}^{H} \beta_{h,sm}SH(h)_{it}$

$$+ \sum_{k=1}^{K} \beta_k Rcn(k)_{it} + \sum_{f=1}^{2} \beta_{f,sm}FM(f)_{it} + \beta_L L_{it}$$

$$+ \varepsilon_{sm,it},$$

(10) $R_{med,i}(t) = \beta_{med} + \beta_{P,med}P_{it} + \beta_{C,med}C_{it}$

$$+ \sum_{h=1}^{H} \beta_{h,med}SH(h)_{it} + \sum_{k=1}^{K} \beta_k Rcn(k)_{it}$$

$$+ \sum_{f=1}^{2} \beta_{f,med}FM(f)_{it} + \beta_L L_{it} + \varepsilon_{med,it},$$

and

(11) $R_{lrg,i}(t) = \beta_{lrg} + \beta_{P,lrg}P_{it} + \beta_{C,lrg}C_{it} + \sum_{h=1}^{H} \beta_{h,lrg}SH(h)_{it}$

$$+ \sum_{k=1}^{K} \beta_k Rcn(k)_{it} + \sum_{f=1}^{2} \beta_{f,lrg}FM(f)_{it} + \beta_L L_{it}$$

$$+ \varepsilon_{lrg,it}.$$

Equation 8 represents the reward associated with the no-purchase option, and Equations 9–11 represent the rewards associated with small, medium, and large orders. For example, in Equation 9, the utility of a small order is a function of β_{sm}, which is the constant associated with a small order, the price index, and whether an e-mail coupon was available at time t. The shipping charges enter through the summation of $\beta_{h,sm}$ and $SH(h)_{it}$. For example, $\beta_{1,sm}$ is the coefficient for a small order (option j = 2) when the shipping schedule is of Type 1 (h = 1). The second summation term, which involves β_k and the duration indicator variables $Rcn(k)_{it}$, captures the effects of time since the previous purchase. The third summation term involves the cumulative purchasing variables, $FM(f)_{it}$. In the term for the loyalty indicator, β_L is the coefficient that describes the benefit of a loyalty reward. The ε_{jit} terms correspond to the random component of utility. In addition, for identification purposes, the constant of the no-purchase option is normalized to zero ($\beta_{no} = 0$).

MODEL ESTIMATION

The estimation of the parameters of the reward functions is somewhat analogous to the approach used in standard discrete-choice models. As in standard choice models, the estimated probability of an observed choice is based on a comparison of the deterministic portion of the utilities of each feasible alternative. The major difference with the dynamic programming model is that the utility associated with each alternative involves both an expression for current-period utility and an expectation of conditional future benefits. Thus, for dynamic programming models, the estimation of choice probabilities is based on a comparison of the alternative specific value functions given in Equation 3.

In each period, it is assumed that customers select the alternative that maximizes utility over the remaining decision horizon. The probability that a customer selects an alternative at time t is defined in terms of the alternative specific value functions. Under the assumption that the random terms are distributed extreme-value i.i.d., it is possible to obtain the following closed-form expression for the choice probabilities (Rust 1994):

(12) $Pr[d_j(t) = 1|S(t)] = \dfrac{\exp\{v_j[S(t)]\}}{\displaystyle\sum_{j' \in \{no, sm, med, lrg\}} \exp\{v_{j'}[S(t)]\}}.$

In this expression, v_j represents the deterministic portion of the alternative specific value function. Therefore, model estimation requires the repeated solution of a dynamic programming problem to calculate the value functions. Under the assumption of a homogeneous population, the log-likelihood function to be maximized consists of the sum of the logarithms of the choice probabilities defined in Equation 12.

The estimation procedure involves nesting an algorithm to solve the dynamic programming problem within a maximum-likelihood routine. At each iteration, the dynamic programming problem is solved by the use of the current estimate of the parameter vector associated with the reward functions. The alternative specific value functions are used to compute the log-likelihood for the current iteration. The maximum-likelihood algorithm then updates the parameter vector, and the process is repeated until convergence.

Because the loyalty program under study has a finite horizon, I solve the dynamic programming model used to generate the choice probabilities with backward recursion. Given the annual time horizon of the program and the assumption of weekly decisions, the decision horizon involves 52 periods. The recursion procedure is executed as follows: For a set of parameter values, I compute the alternative specific value functions for the final period (T = 52) for all combinations of the state space S:

(13) $V_j[S(52)] = E[R_j(52)|S(52)].$

I then use the values for $V_j(52)$ to calculate the value functions for t = 51. Specifically, the alternate specific value function for an arbitrary state at t = 51 is the expected reward during that period plus the discounted expected reward of an optimal action in the next period:

(14) $V_j[S(51)] = E[R_j(51)|S(51)] + \alpha \sum_{s}^{S_{pot}} Pr[s|S(51), j]V[s(52)],$

where S_{pot} is the set of potential states that can be reached in the next period, given the current state. The process continues until the value functions are computed for all periods. I then use the value functions to calculate the probability of each observed choice in the data using Equation 12 (for more details on the general estimation approach, see Rust 1994).

As stated, the model applies to a population with homogeneous preferences for the firm's services and the offered rewards. This is potentially problematic for a model of response to a loyalty program. A more realistic assumption

is that the loyalty program alters the behavior of only some fraction of the customer population. To extend the model to account for variability in customer preferences, I use a latent class approach (Kamakura and Russell 1989). This approach assumes that the population consists of M types, where π_m is the proportion of type m in the population. For the finite-mixture approach, the sample likelihood is

$$(15) \qquad \prod_{i=1}^{N} \sum_{m=1}^{M} Pr(d_{1i}^m, d_{2i}^m, ..., d_{Ti}^m | type = m)\pi_m,$$

where

$$Pr[d_i^m(1), d_i^m(2), ..., d_n^m(T)] = \prod_{t=1}^{T} Pr[d(t)|S(t), type = m].$$

The use of a finite-mixture model to account for customer heterogeneity significantly increases the computational burden, because the dynamic optimization problem must be repeatedly solved for each type in the population.

Because the direct estimation of the discount factor α tends to be difficult, this parameter is fixed before the estimation of the other model parameters. The value of α used in the estimation procedure plays a key role in the analysis. By setting α equal to zero, the model is estimated under an assumption that customers are myopic and do not consider the value of future rewards in current decisions. If α is set to a high value (i.e., close to one), the model assumes that customers value potential future benefits. The empirical strategy is to estimate both types of models. A comparison of model fits then provides evidence of whether customers have a dynamic response to the loyalty-based rewards.

EMPIRICAL IMPLEMENTATION

The complete data set from the Internet retailer includes transaction histories for more than 30,000 customers and spans a 13-month period from 1997 to 1998. For estimation purposes, I selected a small subsample of the population. Because the primary focus is the effectiveness of the reward program, the empirical analysis focuses on customers with meaningful transaction histories. The selection criterion is that customers make at least two purchases during their first 5 weeks in the database and that there are at least 30 weeks of observations. The final sample consists of 1058 customers who made an average of 10.45 purchases per year.

I estimated models for various single- and multiple-segment formulations and for different assumptions regarding the discount rate. Fit statistics and assumptions regarding segment discount rates for single-, two-, and three-segment models are provided in Table 3. Table 4 presents the estimated coefficients for a two-segment mixture model in which both segments possess a dynamic orientation.[6] The two-segment dynamic model is the best fitting of all the two-segment models, according to the Bayesian information criterion (BIC) fit measure, and is the basis for subsequent analysis. In addition to providing a better fit than the competing two-segment models, the two-segment dynamic model is superior to the homogeneous population

Table 3
SUMMARY OF MODEL FITS

Model	Number of Parameters	BIC
Single Segment		
Myopic	28	35,419.5
Dynamic	31	34,996.8
Two Segments		
Two myopic	60	34,388.0
Two dynamic	66	33,996.7
One dynamic and one myopic	63	34,230.5
Three Segments		
Three dynamic	95	34,003.9

models. The addition of a third segment marginally improves fit but yields several nonintuitive parameters.

Estimation results provided in Table 4 indicate two fairly different segments. In broad terms, Segment 1 purchases less frequently and favors larger order sizes, whereas Segment 2 purchases more frequently and favors smaller orders. To generate descriptions of the two segments, Table 5 reports results of simulation studies based on each segment's parameter values and a marketing policy similar to that employed by the firm. Segment 1 places an average of 8.7 orders per year and spends more than $670 annually. Segment 2 spends an average of $586 dollars per year and places approximately 14 orders. In terms of the loyalty rewards, approximately 15% of the Segment 1 population earns at least one reward, whereas only approximately 1% of Segment 2 earns a loyalty reward. Table 5 also shows how the two segments combine to an overall population and how this simulated population compares with the actual sample.[7]

An important difference in the segments is the levels of dispersion. Although the two segments have fairly similar mean levels of annual spending, there is a significant difference in terms of variability. Segment 1 tends to be more volatile, which is why it contains the majority of award recipients. Figure 1 shows the distribution of annual purchasing for both segments and illustrates the differences in the dispersion of the two groups. Although the BIC measures help justify the incorporation of dynamic factors and heterogeneity controls, measures that show how these model enhancements lead to improved managerial forecasts are also important. Because I am interested in predicting long-term behavior, I define measures that describe how well each specification can replicate the actual distribution of cumulative purchasing. To construct the measures, I use each model to simulate the distribution of annual spending for 100,000 customers. I then compare the simulated distributions with the actual percentages of the population in terms of absolute deviations (Table 6). This provides an indication of how well the specifications capture the variability in the population. To show how well the specifications capture the influence of the loyalty program, I also include the percentage of customers who are predicted to earn a loyalty reward. In general, the dynamic specifications help better account for the spikes caused by the reward pro-

---

[6]To improve readability, Table 4 does not report coefficients that are not significant. However, when relevant, I mention such coefficients in the text. Full results are available on request.

[7]I use simulation studies rather than posterior probabilities to describe the segments in order to be consistent with the simulation studies described in the "Policy Experiments" section.

Table 4

TWO-SEGMENT MODEL RESULTS

	Segment 1		Segment 2	
	Coefficient	Standard Error	Coefficient	Standard Error
Small-Order Parameters				
Intercept: β_{sm}	N.S.		N.S.	
Shipping policy 1: $\beta_{h=1,sm}$	−1.02**	.15	−.32**	.07
Shipping policy 2: $\beta_{h=2,sm}$	−.61*	.26	−.28*	.12
Shipping policy 4: $\beta_{h=4,sm}$	N.S.		N.S.	
Price: $\beta_{P,sm}$	−1.90*	.89	−1.38*	.74
E-coupon: $\beta_{C,sm}$	N.S.		.36**	.11
FM$_1$: $\beta_{f=1,sm}$	N.S.		.31*	.13
FM$_2$: $\beta_{f=2,sm}$	N.S.		.48**	.08
Medium-Order Parameters				
Intercept: β_{med}	N.S.		N.S.	
Shipping policy 1: $\beta_{h=1,med}$	−.95**	.17	−.63**	.24
Shipping policy 2: $\beta_{h=2,med}$	−.62*	.33	−.45*	.24
Shipping policy 4: $\beta_{h=4,med}$	N.S.		N.S.	
Price: $\beta_{P,med}$	−2.13*	1.14	−2.20*	1.18
E-coupon: $\beta_{C,med}$	N.S.		.19*	.11
FM$_1$: $\beta_{f=1,med}$	.94**	.18	1.10**	.34
FM$_2$: $\beta_{f=2,med}$	.91**	.13	1.12**	.18
Large-Order Parameters				
Intercept: β_{lrg}	N.S.		N.S.	
Shipping policy 1: $\beta_{h=1,lrg}$	.40**	.05	1.02**	.26
Shipping policy 2: $\beta_{h=2,lrg}$	N.S.		N.S.	
Shipping policy 4: $\beta_{h=4,lrg}$	.19*	.09	N.S.	
Price: $\beta_{P,lrg}$	−2.49**	.76	−2.88**	1.15
E-coupon: $\beta_{C,lrg}$	.33**	.12	.35*	.17
FM$_1$: $\beta_{f=1,lrg}$	1.59**	.23	1.29*	.55
FM$_2$: $\beta_{f=2,lrg}$	1.90**	.15	1.42**	.29
Common Parameters				
Reward: β_L	.75**	.27	N.S.	
Recency level 2: $\beta_{k=2}$	N.S.		−.46**	.07
Recency level 3: $\beta_{k=3}$	−.73**	.08	−1.01**	.11
Recency level 4: $\beta_{k=4}$	−1.73**	.19	−1.84**	.15
Price Expectation Model				
Intercept 1: θ_{Low}	N.S.		N.S.	
Intercept 2: θ_{High}	N.S.		N.S.	
Price effect: θ_P	−.65*	.35	N.S.	
Segment size: θ_{seg1}	.89**	.11		

*$p < .10$.
**$p < .01$.
Notes: I estimated the segment size parameter using a logit formulation such that $\pi_1 = \exp(\theta_{seg1})/[1 + \exp(\theta_{seg1})]$. N.S. = not significant.

Table 5

SEGMENT DESCRIPTIONS

	Segment 1	Segment 2	Simulated Population	Actual Sample
Percentage of population	71%	29%		
Order incidence rate	17%	27%	19%	20.1%
Annual orders	8.7	14	10.3	10.45
Annual spending	$672	$586	$647	$637.28
Average order size	$ 77	$ 42	$ 63	$ 60.98
Percentage earning reward	15%	1%	11%	19.7%

gram, and the heterogeneity controls help better reflect the variability in the population.

The segments are substantially different in terms of the loyalty reward effects. Segment 1 yields a parameter of .75, whereas the parameter for Segment 2 is just .22. In addition, the t-statistic for Segment 2's response to the loyalty program is just 1.3. To some extent, this is not surprising, because the rewards are rarely relevant to members of Segment 2. The estimated reward parameters from the dynamic

and myopic homogeneous models are also notable.[8] In the dynamic model, the reward coefficient is positive and significant (.77), whereas the static model yields a smaller positive estimate that is only marginally significant (.56). The static model yields a lower estimate for the reward parameter because the specification does not consider the ability of

---

[8]Full estimation results for the models are available on request.

Figure 1
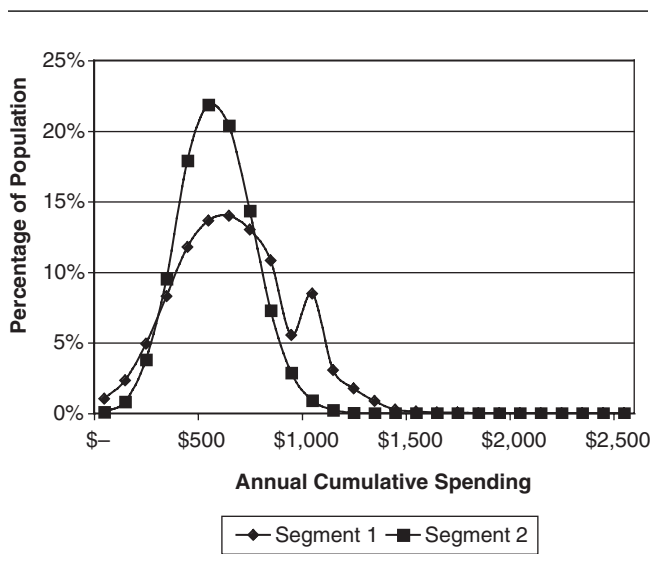COMPARISON OF SIMULATED POPULATIONS

customers to delay rewards. For example, a customer who has sufficient cumulative purchases for a reward with an appropriate purchase in the next period may decide to wait or to make a smaller purchase. The reward may seem not to be valued, but the customer may actually be planning to claim the reward in the future.

In addition to the loyalty rewards, the results indicate the importance of several other factors. Shipping fees have significant effects on both order incidence and order size.[9] The policy (h = 1) that provides free shipping for large orders has a significant, positive impact on the percentage of large orders for both segments. Conversely, this policy results in significantly fewer small and medium-sized orders. The estimation results also suggest that Policy 2 (increasing/high) results in significantly fewer small and medium-sized orders than does Policy 3 (increasing/low). Notably, the free shipping promotion (Policy 4) has only a minor effect compared with the base policy (Policy 3) for the population under study. In terms of segment comparisons, Segment 1 tends to favor large orders, and Segment 2 favors small orders.

The price coefficients are negative and significant for each order size and reveal important segment differences. Segment 1's price coefficients are –1.89 for the small

Table 6
ALTERNATIVE FIT MEASURES

	Two-Segment (Dynamic)	Homogeneous (Dynamic)	Homogeneous (Myopic)
Absolute deviations	65	84	90
Percentage earning rewards	10.7%	9.8%	.6%

category, –2.12 for the medium category, and –2.48 for the large category. In contrast, Segment 2's coefficients are –1.38, –2.20, and –2.88, respectively. These patterns are as I expected, because higher prices have a larger absolute impact on larger orders. In terms of segment differences, the results suggest that higher prices tend to shift Segment 2 to smaller order sizes at a greater rate than Segment 1. Response to the e-mail coupons is similar for both segments. The e-mail coupon coefficients are of the expected sign and are significant for the large and small categories.[10]

Expectations about the future pricing environment are similar for both segments. The coefficients related to the expectations suggest that an increase in current-period prices creates an expectation of higher prices in future periods. This result suggests that high prices can deter current-period ordering and negatively influence retention. This is important for the operation of a loyalty program, because expectations of high prices can reduce the attractiveness of program participation. However, this conclusion should be viewed with some caution, because the coefficients for the expectation terms do not yield significant t-statistics. The terms are included in the model because they significantly improve overall fit.[11]

The time duration parameters are similar for both segments. These variables indicate that the probability of a customer placing an order decreases as the time since the previous order increases. A difference is that the effect of recency accelerates more quickly for Segment 2. The coefficients for the constructed cumulative purchasing variables, FM_1 and FM_2, are also of the expected sign and structure. For both segments, I observe the same pattern of positive signs for all terms and increasing parameters for the larger sizes. The results suggest that customers with higher levels of cumulative purchasing buy more frequently and in larger quantities.

In addition, although time (t) and cumulative purchasing (Cum) are not elements of the reward functions, they are meaningful elements of the state space. The importance of these factors for Segment 1 is illustrated in Figure 2, which gives the probability of a purchase as a function of time and cumulative spending and illustrates how the pull of the loyalty program varies by proximity to a reward threshold and time constraints for three levels of cumulative spending ($500, $700, and $900).[12] For the customer with $900 worth of cumulative spending, the probabilities rise continuously as the expiration time approaches; at the $700 level, probabilities first rise and then decline. As time approaches approximately 35 weeks, the probabilities increase for the $700 level and exceed the probabilities for the $900 level. However, as the remaining time becomes short, the probabilities for this group plunge and are close to the probabilities for the $500 level by the end of the horizon. This occurs

---

[9]Note that order-size intercepts are defined in terms of the third shipping fee schedule. A chi-square test of restricting the shipping coefficients to zero results in a t-statistic of 465.9 (18 degrees of freedom), which is significant at $p < .001$.

[10]Note that for several of the pricing parameters (e.g., shipping, e-mail coupon), the coefficients for the different order sizes often are not dramatically different.

[11]Restricting these terms to $P_{low} = 1/3$, $P_{ave} = 1/3$, and $P_{high} = 1/3$ yields a model with a log-likelihood of –16,628.31. A model comparison test results in a chi-square statistic of 37, which is significant at $p < .001$.

[12]Probabilities are calculated with a fixed state for all variables with the exception of time and cumulative buying.

Figure 2

PURCHASE PROBABILITIES AS A FUNCTION OF TIME AND
CUMULATIVE BUYING (SEGMENT 1)

- ·-·-- $500 Cumulative spending
- -▲-- $700 Cumulative spending
- -■- $900 Cumulative spending
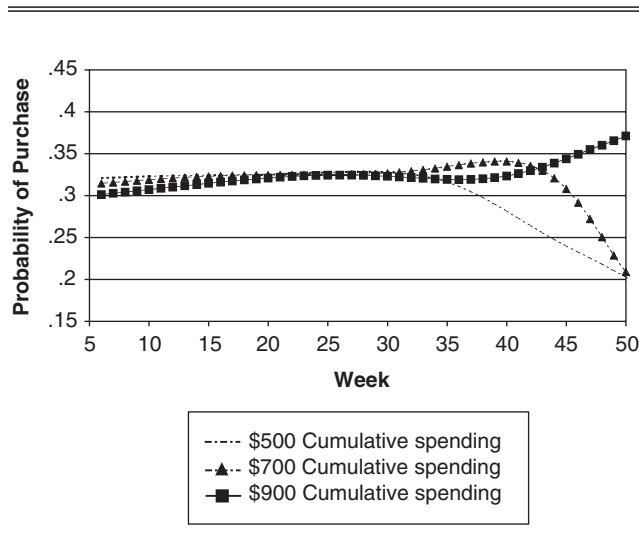
because the likelihood of earning a reward is similar for the two levels ($500 and $700) when only 1 or 2 weeks remain.

POLICY EXPERIMENTS

An issue associated with dynamic programming models of behavior is that their complexity makes it difficult to interpret model coefficients. To understand the model's implications more fully, the results may be evaluated with simulation or policy experiments. This section examines the effects of changing or eliminating the loyalty program. I conducted the policy experiments using the estimated parameters to simulate customer behavior over an extended time period. I evaluated the consequences of eliminating the program by comparing a simulation that uses the full dynamic finite-mixture model structure with a simulation that removes the incentives associated with cumulative buying.[13]

I performed the simulations using the estimated coefficients to compute the probabilities of each alternative (no purchase, small, medium, or large orders) for each possible customer state. I then simulated consumer purchases using a random-number generator. For the purpose of the simulations, a small order is valued at $33, a medium order at $67, and a large order at $100. I then determined customers' future states on the basis of the variable definitions and the laws of motions described in the "Model Estimation" section. The results reported in this section are based on simulations of 100,000 customers who make 51 sequential decisions using a marketing policy (e.g., prices, e-mail coupons) that is similar to that employed by the firm.

Table 7 presents summarized results from a simulation experiment that eliminates the loyalty program. The simulation suggests that removal of the loyalty program will decrease the purchase incidence rate from 19.7% to 19.2% per week. This change translates to a predicted drop in revenue of approximately $13 per customer.

[13]Specifically, I removed the loyalty program by setting the coefficients (β_L) associated with the rewards equal to zero.

Figure 3 illustrates the effects of the loyalty program in more detail. The series coded as "Loyalty Program" shows the predicted distribution of annual ordering for the population when loyalty awards are available, and the series coded as "No Loyalty Program" shows the prediction without loyalty rewards. With loyalty awards, a spike occurs at an annual level of purchasing of slightly more than $1,000. This corresponds to customers who make sufficient purchases to earn an award. In contrast, there is no incentive to reach purchasing targets without the loyalty program, so the model predicts a single modal distribution.

In addition to the policy experiment of eliminating the reward program, many other analyses are possible. Policy experiments provide a means for evaluating the long-term impact of short-term promotions such as e-mail coupons. For a point of comparison, Table 8 reports the results of an experiment that evaluates the effect of an incremental e-mail coupon. This experiment suggests that an additional e-mail coupon has a slightly smaller overall impact than the loyalty program. A notable contrast between the experiments is that whereas the loyalty program substantially affects only one segment, the e-mail coupon leads to more buying by both segments.

The e-mail coupon experiment also highlights the model's decision-support capabilities in that it captures interactions between short-term and dynamic promotions. An additional coupon stimulates demand in two ways. First, it effectively lowers prices and thus directly increases demand in the current period, but a more subtle benefit

Table 7

LOYALTY PROGRAM IMPACT

	Scenario 1: Loyalty Program	Scenario 2: No Loyalty Rewards
Purchase incidence rate	19.7%	19.2%
Mean annual orders per customer	10.25	9.98
Average customer revenue	$647	$634

Figure 3

POLICY EXPERIMENT: COMPARISON OF ANNUAL
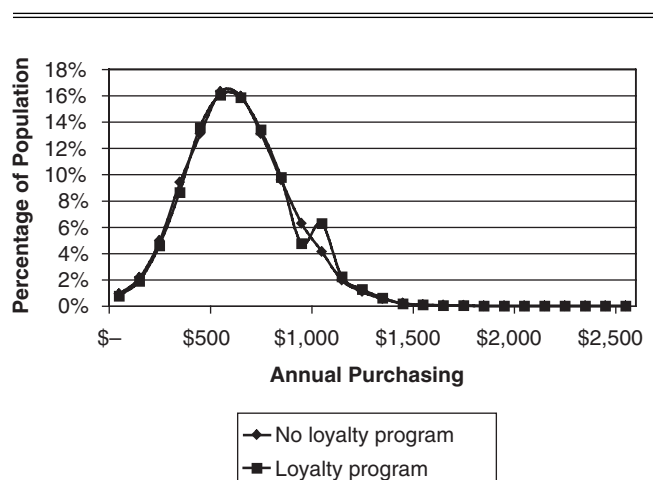PURCHASING WITH AND WITHOUT A LOYALTY PROGRAM

- ◆ No loyalty program
- ■ Loyalty program

Table 8

POLICY EXPERIMENT 2: ADDITIONAL E-MAIL COUPON

	Scenario 1: Base Case	Scenario 2: Additional E-mail Coupon
Purchase incidence rate	19.7%	19.9%
Mean annual orders per customer	10.25	10.37
Average customer revenue	$647	$657

occurs in an interaction between the short-term coupon promotion and the long-term loyalty promotion. Second, by inducing customers to purchase, the e-mail coupon has the effect of increasing a customer's cumulative buying or investment in the reward program. As the customer moves closer to attaining a reward, the pull of the loyalty program is increased, which stimulates additional purchases in future periods.[14]

LIMITATIONS AND FURTHER RESEARCH

Although loyalty programs are a common marketing instrument, empirical examinations of their effectiveness are limited (Drèze and Hoch 1998; Sharp and Sharp 1997). This is largely because the programs are difficult to evaluate with standard techniques, because the rewards for cumulative purchasing can influence behavior over an extended period. Although the dynamic programming method overcomes this difficulty, there are opportunities for further research.

Methodologically, there are opportunities to develop techniques for programs with structures different from the one evaluated in this article. A salient characteristic of the program under study is that it has a strictly defined finite horizon. The selection of a program's time horizon is an important element of loyalty program design. In contrast to the program under study, many programs have rolling or even indefinite time horizons.[15] Moreover, programs may be designed as hybrids so that some elements have strict finite time horizons, and other elements do not include expirations. For example, airline programs often determine a traveler's classification (e.g., gold, silver) on the basis of miles flown in a calendar year, but the miles earned from travel may expire after a few years or never. Reward programs with these types of time horizons can be evaluated with dynamic programming models, but they require adjustments to the solution approach and additional data.

Programs with rolling or indefinite time horizons may be evaluated by modeling the customer's objective function as an infinite-horizon dynamic programming problem. Infinite-horizon problems can be solved with techniques such as value or policy iteration (Bertsekas 1996). Although such techniques can be more computationally intensive than the backward recursion used in the finite case, they are still feasible for fairly large problems.

The main complication of modeling a program with a rolling time horizon is that it is necessary to retain the customer's purchase history for the length of the specified time horizon. With a rolling structure, the customer's relevant cumulative purchases may change from period to period for two reasons. First, as mentioned previously, the customer's investment in the program may increase in response to current-period buying. Second, with a rolling time horizon, it is possible that investment in the program decreases through the expiration of previously earned points. From the perspective of the firm, customer transitions between levels of cumulative purchasing could include both a probabilistic element related to current-period buying decisions and a deterministic element related to the number of points that are about to expire.

A potential difficulty with this type of program is the maintaining of computational tractability. The issue is the degree to which previous purchase history must be maintained. Theoretically, it is necessary to maintain the customer's complete transaction history as elements of the customer's state space. However, in practical terms, it is likely that most of these data could be approximated or summarized, particularly because it is doubtful that customers themselves have perfect memory of their transaction histories.

Related to the issue of computational tractability are questions about potential extensions and refinements of the model. For example, it may be useful to consider the interaction between a customer's most recent purchase and recency. This would potentially help disentangle inventory-based inactivity from attrition-based inactivity. The difficulty with this extension is that it requires the augmentation of the state space to include the customer's previous order size. Although this is a minor extension to the state space (one variable with three levels), the end result is a tripling of the total state space. Another potential avenue of study is interactions of the marketing variables with the customer's level of cumulative purchasing. For example, although the model accounts for the effects of time and cumulative purchasing, it is conceivable that these factors interact with marketing-mix elements, such as price.

The modeling of shipping fees as categorical rather than dollar values also merits discussion. There are potentially significant benefits to using the dollar values of shipping fees as covariates in the model. With dollar values, analyses that quantify the effects of changes to such fees could be performed. An understanding of the elasticities associated with shipping fees is important, but it is left to further research for two reasons. First, the development of an appropriate specification may not be straightforward, because the choice of a given order size may be a function of multiple shipping fees. As a case in point, consider the customer response to the free shipping policy and the policy offering free shipping for large orders. Although both charge the same price for a large order, they result in different average order sizes.[16] Modeling of the impact of shipping fees requires a specification that includes direct and relative shipping cost terms. Second, the use of categorical variables helps maintain computational feasibility. In the

[14]The results of an experiment that simultaneously removes the loyalty program and adds an incremental e-mail promotion suggest that the average annual purchasing increases from $634 to $642.

[15]By "rolling," I mean programs that involve points that expire according to the date that each point is acquired.

[16]The average order size in the free shipping policy is $44; in Policy 1, the average size is $65.

specification herein, the shipping fee schedules are represented as four distinct states, and the use of dollar values would require three distinct variables with three levels.

Another area for further research pertains to the effectiveness of alternative reward types. A limitation of the current data set is that rewards are limited to a single type. With appropriate data, it may be possible to measure the relative effectiveness of different rewards or the benefits of offering reward choices.

CONCLUSION

This article presents a dynamic programming model of customer response to a loyalty program and other marketing tactics. The model measures the influence of rewards by considering customers' sequences of purchases as a solution to a dynamic optimization problem. A primary strength of the approach is that the dynamic programming framework provides a rich structure for modeling customer behavior. The model includes both the influence of previous behavior in terms of cumulative purchases and forward-looking factors, such as expectations of future prices and loyalty rewards. The inclusion of forward-looking behaviors is computationally expensive but intuitively appealing for modeling customer response to a loyalty program. Given the prevalence of loyalty programs, this is a salient topic for researchers and practitioners.

The model estimates the effects of e-mail coupons, pricing changes, shipping fees, and the loyalty program. The sign and statistical significance of the parameter associated with the reward in the dynamic model is evidence that the loyalty program effectively increases repeat-purchase rates. Further evidence of the program's effectiveness is the relative fit of the dynamic model compared with a static formulation in which customers do not consider future benefits. The results should provide solace to advocates of loyalty programs.

The model also provides a platform for conducting what-if analyses of customer retention. Simulation studies can be used to study the relative power of each marketing instrument to increase long- and short-term purchasing. Such studies are useful because the estimation results indicate that multiple instruments can stimulate repeat buying. Furthermore, given appropriate cost information, it is a simple extension to use these experiments to estimate the profitability of alternative policies for customer relationship management.

REFERENCES

Bellman, Richard (1957), *Dynamic Programming*. Princeton, NJ: Princeton University Press.

Bertsekas, Dimitri (1996), *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific.

Boulding, William, Ajay Kalra, Richard Staelin, and Valarie Zeithaml (1993), "A Dynamic Process Model of Service Quality: From Expectations to Behavioral Intentions," *Journal of Marketing Research*, 30 (February), 7–27.

Bult, Jan Roelf and Tom Wansbeek (1995), "Optimal Selection for Direct Mail," *Marketing Science*, 14 (4), 378–95.

Deighton, John (2000), "Frequency Programs in Service Industries," in *Handbook of Services Marketing and Management*, T. Swartz and D. Iacobucci, eds. London: Sage Publications, 401–408.

Dowling, Grahame R. and Mark Uncles (1997), "Do Customer Loyalty Programs Really Work?" *Sloan Management Review*, 38 (Summer), 71–82.

Drèze, Xavier and Stephen J. Hoch (1998), "Exploiting the Installed Base Using Cross-Merchandising and Category Destination Programs," *International Journal of Research in Marketing*, 15 (5), 459–71.

Eckstein, Zvi and Kenneth Wolpin (1989), "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models," *Journal of Human Resources*, 24 (4), 562–98.

Erdem, Tülin, Susumu Imai, and Michael Keane (2003), "Brand and Quantity Choice Dynamics Under Price Uncertainty," *Quantitative Marketing and Economics*, 1 (1), 5–64.

——— and Michael Keane (1996), "Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Markets," *Marketing Science*, 15 (1), 1–20.

Gönül, Füsun (1999), "Estimating Price Expectations in the OTC Medicine Market: An Application of Dynamic Stochastic Discrete Choice Models to Scanner Panel Data," *Journal of Econometrics*, 89 (1–2), 41–56.

——— and Mengze Shi (1998), "Optimal Mailing of Catalogs: A New Methodology Using Estimable Structural Dynamic Programming," *Management Science*, 44 (9), 1249–62.

——— and Kannan Srinivasan (1996), "Estimating the Impact of Consumer Expectations of Coupons on Purchase Behavior: A Dynamic Structural Model," *Marketing Science*, 15 (3), 262–79.

Kalwani, Manohar U., Chi Kin Yim, Heikki J. Rinne, and Yoshi Sugita (1990), "A Price Expectations Model of Consumer Brand Choice," *Journal of Marketing Research*, 27 (August), 251–62.

Kamakura, Wagner and Gary J. Russell (1989), "A Probalistic Choice Model for Market Segmentation and Elasticity Structure," *Journal of Marketing Research*, 26 (November), 379–90.

Kim, Byung-Do, Mengze Shi, and Kannan Srinivasan (2001), "Reward Programs and Tacit Collusion," *Marketing Science*, 20 (2), 99–120.

Kivetz, Ran and Itamar Simonson (2002), "Earning the Right to Indulge: Effort as a Determinant of Customer Preferences Toward Frequency Program Rewards," *Journal of Marketing Research*, 39 (May), 155–70.

Kopalle, Praveen and Scott A. Neslin (2003), "The Economic Viability of Frequency Reward Programs in a Strategic Competitive Environment," *Review of Marketing Science*, (accessed April 16, 2004), [available at http://www.bepress.com/romsjournal/vol1/iss1/art1].

Lemon, Katherine, Tiffany Barnett White, and Russell Winer (2002), "Dynamic Customer Relationship Management," *Journal of Marketing*, 66 (January), 1–14.

Narasimhan, C. (1989), "Incorporating Consumer Price Expectations into Diffusion Models," *Marketing Science*, 8 (4), 343–57.

O'Brien, Louise and Charles Jones (1995), "Do Rewards Really Create Loyalty?" *Harvard Business Review*, 73 (May–June), 75–82.

Rust, John (1994), "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics*, Vol. 4, D. McFadden and R. Engle, eds. Amsterdam: North-Holland, 3081–3143.

Sharp, Byron and Anne Sharp (1997), "Loyalty Programs and Their Impact on Repeat-Purchase Loyalty Patterns," *International Journal of Research in Marketing*, 14 (5), 473–86.

Soman, Dilip (1998), "The Illusion of Delayed Incentives: Evaluating Future Effort-Money Transactions," *Journal of Marketing Research*, 35 (November), 427–37.

Sun, Baohong, Scott A. Neslin, and Kannan Srinivasan (2004), "Measuring the Impact of Promotions on Brand Switching When Consumers Are Forward Looking," *Journal of Marketing Research*, 40 (November), 389–405.

Zhang, Z. John, Aradhna Krishna, and Sanjay K. Dhar (2000), "The Optimal Choice of Promotional Vehicles," *Management Science*, 46 (3), 348–62.