MANAGEMENT SCIENCE
Vol. 51, No. 6, June 2005, pp. 986–994
ISSN 0025-1909 | EISSN 1526-5501 | 05 | 5106 | 0986

DOI 10.1287/mnsc.1050.0373

Research Note: A Dynamic Programming Approach to Customer Relationship Pricing

Michael Lewis
Department of Marketing, University of Florida, Bryan Hall, Gainesville, Florida 32611, mike.lewis@cba.ufl.edu

The practice of offering discounts to prospective customers represents a rudimentary form of using transaction history measures to customize the marketing mix. Furthermore, the proliferation of powerful customer relationship management (CRM) systems is providing the data and the communications channels necessary to extend this type of pricing strategy into true dynamic marketing policies that adjust pricing as customer relationships evolve. In this paper, we describe a dynamic programming–based approach to creating optimal relationship pricing policies. The methodology has two main components. The first component is a latent class logit model that is used to model customer buying behavior. The second component is a dynamic optimization procedure that computes profit-maximizing price paths. The methodology is illustrated using subscriber data provided by a large metropolitan newspaper.

The empirical results provide support for the common managerial practice of offering discounts to new customers. However, in contrast to current practice, the results suggest the use of a series of decreasing discounts based on the length of customer tenure rather than a single steep discount for first-time purchasers. The dynamic programming (DP) methodology also represents an important approach to calculating customer value (CV). Specifically, the DP framework allows the calculation of CV to be an explicit function of marketing policies and customer status. As such, this method for calculating CV accounts for the value of managerial flexibility and improves upon existing methods that do not model revenue and attrition rates as functions of marketing variables.

Key words: pricing research; customer relationship management; customer valuation
History: Accepted by Jagmohan S. Raju, marketing; received February 20, 2004. This paper was with the author 4 months for 2 revisions.

1. Introduction

A desire to customize the marketing mix offered to individual customers is a prime motivation for firms to develop customer relationship management (CRM) systems. The promise of personalized marketing is that firms will be able to optimize the value of customer assets by leveraging individual-level information. However, CRM efforts often fail to yield a positive return on investment, and surveys have found CRM to have a low managerial satisfaction rate (Rigby et al. 2002). A common impediment to relationship management is a lack of techniques for using individual-level data to create dynamic customer-marketing policies. In this paper, we present a dynamic programming–based method for customer relationship pricing.

Our approach to relationship pricing involves three main components. First, a system for classifying customers based on observed transaction histories is developed. Second, a demand model is used to understand the link between pricing and migration between classifications. The third element is a dynamic optimization procedure designed to compute the pricing structure that optimally manages transitions between customer classes. For our empirical implementation, the choice model results suggest that price sensitivity decreases with customer tenure for the vast majority of customers, and acquisition discounts should be based on time since attrition and on the customer's previous activity with the firm. The optimization results suggest that the firm can significantly enhance profitability by shifting from using a single steep discount to acquire customers to offering a series of decreasing discounts as customer tenure grows.

Methodologically, our work is especially related to several papers that utilize dynamic programming (DP) models to create optimal catalog-mailing policies (Gonul and Shi 1998, Bitran and Mondschein 1995).[1] In these models, customers are classified according to transaction history measures. These classification systems, combined with purchasing rates for each class, provide a means of modeling customer migration as a Markov chain. The firm's optimization problem

[1] There is also a body of work in the marketing literature that estimates DP models of consumer behavior (Sun et al. 2004, Erdem et al. 2003). In contrast, we use DP to develop policy rather than to model decision making.

is to maximize customer value (CV) by determining which transaction history classes should receive mailings. Gonul and Shi (1998) include an additional element by performing the firm-level optimization under an assumption that customers also act as dynamic optimizers.

In addition to addressing a different application, dynamic pricing, this paper complements the literature in two respects. First, rather than focus on mailing decisions, our emphasis is on how to price based on past behavior. This represents a generalization of the catalog-mailing models where the firm's decisions are binary mailing decisions. Additionally, the inclusion of price in the model means that revenue and profits are estimated directly rather than assumed. This is important because the absence of price and promotion in the catalog models means that these models assume fixed managerial policies beyond the mailing decision. Second, the paper makes an effort to account for customer heterogeneity. While this is a relatively minor refinement, it is an important factor when determining policy as a function of transaction history.

The results are relevant to the dynamic-pricing, customer-valuation, and customer-base analysis literatures. In contrast to pricing over the product life cycle (see Rao 1993), we are concerned with pricing over customer relationship stages. Similar to the dynamic pricing of products, the customer-level dynamic-pricing problem involves balancing current revenues with the future value of a customer relationship. Also, because the investments in customer relationships involve promotional discount, the results are relevant to the promotions literature (Neslin 2002).

The underlying objective of maximizing long-term CV means that the method is related to work on customer valuation (Berger and Nasr 1998). The standard approaches to estimating customer equity and customer asset value assume revenue and retention rates that are either fixed or vary purely as a function of time (Berger and Nasr 1998, Gupta et al. 2004). These implicit assumptions mean that existing customer valuation methods do not account for managerial flexibility and therefore provide downwardly biased CV estimates. By simultaneously addressing customer valuation and customer management, the technique removes a source of bias in customer valuation and directly links CV to marketing policy.

Our results are also relevant to the customer-base analysis literature concerned with the relationship between transaction history and customer behavior. Reichheld and Teal (1996) and Reinartz and Kumar (2000) study how behavior changes over the customer life cycle. We add to this literature with further empirical results and provide an extension by showing how behavioral differences between customer segments can be used to enhance profitability.

This paper is organized as follows. Section 2 introduces the idea of treating customer management as a dynamic optimization problem. Section 3 describes our empirical context and develops the customer response model. Section 4 details the estimation results for the demand model. Section 5 describes the optimization procedure and reports the optimal pricing policies. Section 6 discusses managerial issues, substantive findings, and model limitations.

2. Modeling Customer Management

Our empirical analysis is performed using customer data from a major metropolitan newspaper. Specifically, we study the problem of using price to manage the acquisition and retention of newspaper subscribers. Subscribers are a prime example of customers who are best viewed as assets, as firms frequently make investments in terms of discounts to acquire a customer who will provide a stable stream of future revenues. The subscriber management problem is to determine the series of prices and promotions that maximize the contribution provided by customer assets. For prospective customers, the question is what levels of discounts, if any, should be offered. For active subscribers, the marketing question is what price should be offered when subscriptions expire. The marketing decisions are complicated because the maximization of future profits requires balancing immediate revenues and long-term retention.

The data for our empirical analysis involves customer histories for 1,326 active customers and prospects and includes a monthly record of current-period purchasing activity, subscription type, and billing information such as price paid. The marketing data include information about pricing and promotional activity. The average price paid during the data-collection period was approximately \$2.40 per week and ranged from \$1.75 to \$3.00 per week. Table 1 shows the average and the range of prices paid by customers with varying degrees of customer tenure. Prices charged to new customers are generally lower because a significant fraction of new customers are acquired using introductory subscription discounts.

Table 2 illustrates the behavioral differences of lapsed and active customers as a function of months since attrition or months of consecutive buying. For

Table 1 Pricing by Tenure

Months of tenure	Average price (\$)	Min price (\$)	Max price (\$)
1 to 3	2.26	1.75	3.00
4 to 6	2.55	1.75	3.00
7 to 12	2.57	1.95	3.00
13 to 18	2.70	1.95	3.00

Table 2 Historical Rates

Lapsed customers			Active customers		
Months lapsed	Reacquisition rate (%)	Solicitation rate (%)	Months tenure	Retention rate (%)	Renewal rate (%)
1 to 6	5.02	8.5	1 to 6	81	37
7 to 12	3.25	11.3	7 to 12	93	58
13 plus	1.18	13.7	13 plus	96	61

lapsed customers, Column 2 illustrates how the probability of reacquiring a customer decreases as time lapsed increases. Column 3 provides information on the firm's reacquisition efforts. The probability that a lapsed customer is solicited ranges from 8.5% per month for recently lapsed customers to 13.7% for former subscribers who have not purchased in a year or more.

The second half of the table illustrates the relationship between time as a subscriber and likelihood of future loyalty. The column labeled "Retention rate" gives the probability that a customer continues a subscription until expiration. For example, customers who have been active for over a year cancel subscriptions prior to expiration about 4% of the time. The column labeled "Renewal rate" indicates that customers with over a year of tenure renew subscriptions at expiration just over 60% of the time.

This cursory examination reveals significant relationships between transaction history and the likelihood of future buying. The data suggest that consumer behavior evolves as time as an active buyer or as a nonactive prospect increases. This evolving behavior, combined with a managerial desire to consider long-term CV, makes subscriber management an application addressable via DP techniques. Such models can accommodate the stochastic evolution of consumer preferences and long-run managerial planning horizons.

2.1. Objective Function

The first step in developing a relationship pricing policy is specifying a suitable forward-looking objective function. For the general case, customers are classified into one of $x \in X$ possible states based on transaction history, and the firm is assumed to be able to select marketing actions or controls, d, from a set of feasible actions, D. Defining the single-period reward function as $\pi(x_t, d_t)$, a single-period discount factor as α, and the decision-horizon length as T, the objective function may be written as

$$\max_{d_t \in D} \mathrm{E}\left\{\sum_{t=0}^{T} \alpha^t * \pi(x_t, d_t)\right\}.$$

This objective function is the discounted sum of single-period rewards over a T-period horizon.

The customer management goal is to determine the mapping of marketing policies to observed customer states $d(x_t)$ that maximize customer value. In our context, this involves determining the functional relationship $P(x_t)$ between customer state and price. Specifically, the firm's optimization problem is to select the sequence of prices and associated marketing actions that maximizes profitability over some time horizon. The solution to the dynamic optimization problem is defined in terms of the value function, denoted as V. By definition, the value function is the greatest feasible expected payoff from time t forward, given that the customer's state at time t is x_t. At the level of the individual customer, the value function is equivalent to the maximum expected CV based on optimal selection of marketing policies:

$$\mathrm{CV}(x_0) = \max_{P_t \in D} \mathrm{E}\left\{\sum_{t=0}^{T} \alpha^t * \pi(x_t, P(x_t))\right\}. \tag{1}$$

Thus far, the objective function does not include a direct statement of how customers evolve over time. In Equation (2), optimal CV is rewritten to make the relationship dynamics more explicit. In this equation, $\Pr(x_{t+1} \mid x_t, P(x_t))$ is the probability a customer in state x_t transitions to state x_{t+1} given the price specified by $P(x_t)$:

$$\begin{aligned}\mathrm{CV}(x_t, P(x_t)) &= \mathrm{E}[\pi(x_t, P(x_t))] \\ &\quad + \sum_{\tau=t+1}^{T} \sum_{x' \in X} \alpha^{\tau} * \pi(x'_{\tau}, P(x_{\tau})) * \Pr(x'_{\tau+1} \mid x_{\tau}, P(x_{\tau})).\end{aligned} \tag{2}$$

The first term on the right side is the expected one-period profit from applying the selected control to a customer in a given state in the current time period. The term involving the double summations represents the sum of the discounted expected profits for all future periods. This term includes the transition probabilities through the inner summation over the customer's potential future states (x'). The outer summation is over future time periods.

The expression thus incorporates the probabilistic evolution of consumer transaction histories. Further, given that from the firm's perspective the consumer behaves nondeterministically, the dynamic optimization problem is to identify the pricing strategy that optimally controls the probabilities of transitions from one transaction history state to another.

2.2. Single-Period Revenue

We begin constructing the single-period contribution from a customer by defining a binary variable Y_{it} equal to 1 to indicate a purchase by customer i at time t, and equal to 0 if no purchase is made. The one-period profit is the probability of purchase, $\Pr(Y_{it}=1)$, multiplied by the total of all revenue sources less product and marketing costs. The single-period contribution π is the sum of the offered price P_t and a nonprice revenue component M, less the cost to serve a customer C_{cts} and any marketing costs C_d. Further, the probability of purchase $\Pr(Y_{it}=1 \mid x_t, d_t)$ is conditional on the customer's state and the marketing mix:

$$\pi(x_t, d_t) = \Pr(Y_t = 1 \mid x_t, d_t) * (P_t + M - C_{cts}) - C_d. \quad (3)$$

In our application, cost to serve a customer is a function of factors ranging from editorial costs to cost of materials like paper and ink. The nonprice component of revenue M merits some explanation. A customer base often has value beyond the revenue due to merchandise sales. For instance, in the newspaper industry, advertising revenue may be three times as large as circulation revenue (Vogel 1998). For the empirical work that follows, we use a value for the quantity of M less C_{cts} suggested by the firm. The last term, C_d, represents costs specific to marketing actions taken, such as the cost of soliciting a prospective customer.

3. Transition Model and Customer State Space

Transaction History Measures. The key to developing a customer classification system is identifying observable data for categorizing customers into homogeneous groups. Transaction histories provide information regarding customers' propensities for loyalty or repeat purchasing. Also, using transaction history measures as state variables means state-to-state transitions are defined by customer buying decisions. Specifically, we use measures that reflect current tenure as a subscriber, F_{it}, or alternatively, the amount of time since a subscription has lapsed, L_{it}. These variables are adjusted based on purchase activity each month according to the Markov transition processes defined in Equations (4a) and (4b). Individual i's time as a lapsed subscriber, L_{it}, increases each month that a purchase is not made ($Y_{it}=0$) and is set to 0 if a purchase is made ($Y_{it}=1$):

$$L_{it+1} = \begin{cases} 0 & \text{if } Y_{it}=1, \\ L_{it}+1 & \text{if } Y_{it}=0. \end{cases} \quad (4a)$$

The equation of motion for the purchasing tenure variable F_{it} is similar in structure. Individual i's time as an active subscriber, F_{it}, is incremented each month that a purchase is made and is reset to 0 if a purchase is not made:

$$F_{it+1} = \begin{cases} 0 & \text{if } Y_{it}=0, \\ F_{it}+1 & \text{if } Y_{it}=1. \end{cases} \quad (4b)$$

We use the time-lapsed measure directly as an explanatory variable, whereas the purchase-tenure measure is used to construct indicator variables that approximately correspond to the number of times a customer has renewed a subscription. The F_{it} measure is used to create binary variables that define four mutually exclusive categories of active customers. Active customers are classified as either NEW (1 month $\le F_{it} <$ 3 months), REP1 (3 months $\le F_{it} <$ 6 months), REP2 (6 months $\le F_{it} <$ 12 months), or LOYAL ($F_{it} \ge$ 12 months).

In addition to the customer's current spell, information regarding total previous duration of purchasing is also maintained. The logic for using this information is that customers who have purchased for extended time periods in the past but who are currently lapsed may behave differently than new prospects. Cumulative months of purchasing for customer i at time t is denoted CUM_{it}. The CUM_{it} measure is reset to zero when a lapsed customer's time since last purchase as a lapsed customer exceeds 12 months. This means that any customer who has not purchased in over a year is classified the same as a prospect that has not previously purchased. The cumulative buying measure is used to define dummy variables that account for the potentially different probabilities of reacquiring lapsed customers with diverse prior buying histories. Two binary variables are defined. For pure prospects and customers who have been lapsed for more than a year, the variable PREV0_{it} is set equal to one. For lapsed customers with less than six months of prior buying, the variable PREV6_{it} is set equal to one.

Marketing Variables. The key marketing-mix variables are price, P_{it}, and solicitation, S_{it}. The price variable is the weekly price available to customer i during month t. Solicitation is a binary variable that is set to 1 if a solicitation is made to individual i in month t. For lapsed customers, solicitation is directly linked to available price. In the absence of a solicitation, lapsed customers are assumed to have access only to the regular price.

Marketing and Transaction History Interactions. In addition to weekly prices, we also include a term to account for the effect of price increases. The previous price, P_{it-1}, is used to compute the percentage level of a price increase relative to the current offered price. This variable is intended to capture reference price effects. Also, because prices can rise only when subscriptions expire, the price increase variable is applicable only at times of renewal. Further, because the rejection of a previous solicitation may be a meaningful element of data, we define a binary variable to indicate a given customer has been previously solicited (PSOL = 1). For current customers, subscription expiration status is a binary variable equal to 1

Table 3 Summary of Covariates

Covariates and categories	Definition
Current transaction history measures	
Months lapsed: L	Months of consecutive nonpurchasing
Months lapsed squared: L^2	Months lapsed squared
Months of subscribing: F	Months of consecutive purchasing; used to define NEW, REP1, REP2, and LOYAL classes of active customers
Previous transaction history	
PREV0	Binary variable to indicate that an individual is a prospect
PREV6	Binary variable to indicate that an individual has previously subscribed for less than six months
Marketing variables	
Price: P	Weekly price available
Promotional solicitation: SOL	Binary variable that indicates a lapsed subscriber receiving a promotional solicitation
Marketing/transaction history interactions	
Price increase: PINC	Percentage change in prices from period $t-1$ to t
Multiple solicitation: PSOL	Binary variable to indicate a previous solicitation
Subscription expiration: EXSUB	Binary variable to indicate the final month of a subscription

if the customer's subscription expires in the current period.[2] The complete set of explanatory variables is summarized in Table 3.

3.1. Transition Model

The customer classification system is based largely on transaction history measures. This structure facilitates the determination of transition probabilities because state-to-state transitions are based on buying and marketing decisions. The transition probabilities can therefore be modeled via a binary logit model. As before, Y_{it} is a binary variable equal to 1 if a purchase is made in period t, and to 0 otherwise. Y is defined in terms of a measure of latent utility Y^*, which is a function of Z_{it}, a matrix of all state variables and marketing actions that influence customer i's decision to buy in period t and a vector of parameters that reflect the impact of the Z variables, β:

$$Y_{it} = \begin{cases} 1 & \text{if } Y_{it}^* \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$
$$\text{where } Y_{it}^* = \beta Z_{it} + \varepsilon_{it} \quad \text{or} \quad Y_{it}^* = v_{it} + \varepsilon_{it}.$$

Assuming the error term, ε_{it}, follows the extreme value distribution, the probability of a purchase by customer i in month t is $\Pr(Y_{it} = 1) = \exp(v_{it})/(1 + \exp(v_{it}))$.

If the individual is currently a subscriber, the deterministic portion of utility, v_{it}, is

$$v_{it,\,c} = \beta_c + \beta_{c,\,p} * f(P_{it}, F_{it}) + \beta_{c,\,\text{exsub}} * \text{EXSUB} + \beta_{c,\,\text{pinc}} * \text{PINC}. \tag{5}$$

Similarly, the expression for a nonactive individual is

$$v_{it,\,l} = \beta_l + \beta_{l,\,p} f(P_{it}, L_{it}) + \beta_s S + \beta_{\text{psol}} S * \text{PSOL} + \beta_{l,\,\text{prev0}} \text{PREV0}_{it} + \beta_{l,\,\text{prev6}} \text{PREV6}_{it}. \tag{6}$$

Note that we include an undefined function of price and the duration variables (L_{it} and F_{it}). This function incorporates state dependence by defining the relationship between transaction history and price sensitivity. To account for state dependence for current customers, we employ a step function that approximates the number of times a customer has renewed a subscription to define the relationship between transaction history measures and price sensitivity:

$$f(P_{it}, F_{it}) = P_{it} * \left(\beta_P + \beta_{\text{REP1}} \text{REP1} + \beta_{\text{REP2}} \text{REP2} + \beta_{\text{LOY}} \text{LOY}\right). \tag{7}$$

For lapsed customers, we use a quadratic formulation that implies a continuous relationship between months lapsed and price sensitivity. The quadratic term is intended to account for possible nonlinear evolution of price sensitivity:

$$f(P_{it}, L_{it}) = P_{it} * \left(\beta_P + \beta_{P,\,L} L_{it} + \beta_{P,\,L_{it}^2} L_{it}^2\right). \tag{8}$$

The model thus far is structured to account for possible state dependence. To account for heterogeneity in individual preferences, we estimate models that treat the population as a mixture of unobserved types (Kamakura and Russell 1989). This approach assumes that the population consists of M distinct types who differ in their preferences and responsiveness to marketing activities. In this formulation, a vector of parameters is estimated for each type in the population, and the likelihood function is a finite mixture or weighted average of the type-specific likelihoods. For a sample of N individuals, each making t_n choices, the likelihood function under an assumption of M types is

$$\prod_{n=1}^{N} \sum_{m=1}^{M} \Pr\left(Y_{1n}^m, Y_{2n}^m, \ldots, Y_{t_n n}^m \mid \text{type} = m\right) * \gamma_m, \tag{9}$$

where γ_m is the proportion of type m in the population. Additionally, segment sizes γ_m are determined via the following expression, where the ϕs are parameters to be estimated:

$$\gamma_m = \frac{\exp(\phi_m)}{\sum_{i \in M} \exp(\phi_i)}. \tag{10}$$

[2] Because subscriptions may be cancelled, the existence of a subscription represents a transaction cost associated with the nonpurchase option rather than a commitment to buy.

Table 4 Parameter Estimates

	Segment 1		Segment 2	
	Estimate	Std. error	Estimate	Std. error
Lapsed customers				
Intercept	−1.63***	0.40	−0.73*	0.41
Price	−2.18***	0.47	−2.62***	0.45
Price × L	−0.30***	0.10	0.52	0.10
Price × L^2	0.01**	0.005	−0.017***	0.004
Solicitation	1.43***	0.36	1.97***	0.38
Solicitation × PSOL	−0.03	0.22	−0.73	0.49
PREV0	−3.33***	0.46	−3.50***	0.41
PREV6	−1.37***	0.15	−0.81***	0.22
Active customers				
Intercept	4.86***	0.22	4.03***	0.57
Price	−2.44***	0.25	−2.06***	0.60
Price × REP1	0.03	0.14	−0.27	0.24
Price × REP2	0.64***	0.17	0.47*	0.23
Price × LOYAL	1.32***	0.15	1.57***	0.57
Expiration	−1.56***	0.16	−1.09***	0.27
Price increase %	−3.82***	0.33	−2.52***	0.83
Segment size	2.66***	0.21		
Observations	35,135			
Log-likelihood	−3,566.2			

*Significant at the 0.10 level, **significant at the 0.05 level, *** significant at the 0.01 level.

4. Estimation Results

Overall, the estimation results suggest that it is important to include terms that capture the evolution of price sensitivity and to control for unobserved heterogeneity. Estimation results for a two-segment solution using the specification detailed above are summarized in Table 4. The two-segment solution yields a BIC measure of 7,456.9 versus 7,612.6 and 7,476.2 for a one- and a three-segment specification, respectively.

The segment-size parameter indicates that the population is comprised of one large segment and a smaller secondary segment. The dominant segment comprises 93%[3] of the population and exhibits the expected pattern of price sensitivities that diminish (increase) with time as a customer (time lapsed). For the first segment, all parameters are of the expected sign and, with the exception of the multiple solicitation term and the REP1 price interaction, all are significant. For the second segment, the estimates reveal a fundamentally different population whose price sensitivity decreases as time lapsed increases. In some respects, this result is not surprising because we might reasonably expect some fraction of the population to be switchers. In this context, a switcher is a customer who alternates between active and nonactive status. Because this type of customer tends to go back and forth, we would expect negative purchase feedback effects.

[3] The segment size parameter of 2.66 listed in Table 4 is used to compute the segment size of 93% via Equation (10).

For the first segment, the parameters evolve as expected. For noncustomers, price sensitivity increases with time lapsed. For current customers, price sensitivity diminishes with time as a buyer. However, there is no evidence to indicate NEW subscribers ($1 \leq F < 3$) are more price sensitive than subscribers renewing for the first time ($3 \leq F < 6$). As customer tenure increases and base price sensitivity decreases, reference price effects become increasingly salient. However, while the coefficients for the price increase term are large, it should be noted that these terms are multiplied by the percentage of a price increase. The results are salient because, unlike in previous studies (Reichheld and Teal 1996), they are generated via a model that accounts for unobserved preference heterogeneity. We also find that previous cumulative buying is an important factor, as lapsed customers with large amounts of previous buying are significantly more likely to be reacquired.

5. Dynamic Optimization Results

In this section, the preceding demand models are used to create optimal marketing policies. For the optimization, the estimation results in Table 4 are used to compute transition probabilities. Beyond the probabilities, the optimization routine requires the definition of feasible managerial controls and the time horizon for the decision problem. While price is often treated as a continuous variable, the numerical solution to the DP problem requires that price be treated as a discrete variable. Based on expressed managerial preference for certain types of price points, the optimization is restricted to price points ending in five-cent increments ranging from \$1.65 to \$3.25 per week. The optimization problem is formulated as an infinite-horizon model and is solved using successive approximations and a 10% annual discount rate. However, the customer valuations estimates are for only the next 36 months of activity. A finite horizon is used for the valuations to alleviate potential concerns about unforeseen changes to the competitive environment.

To account for the two segments in terms of state-to-state transitions, the objective function in Equation (2) is generalized to reflect the multiple population types identified by the latent class model. Because each type has a different utility function, and thus different purchase probabilities, transition rates must be computed for all possible future states for each type. Policies are created under an assumption that the firm does not make probabilistic judgments about each customer's unobserved type.

5.1. Price Paths

Optimal price policies are displayed in Figure 1 under an assumption of an average subscription length of three months. Recommended prices are presented for

Figure 1 Optimal Price Paths

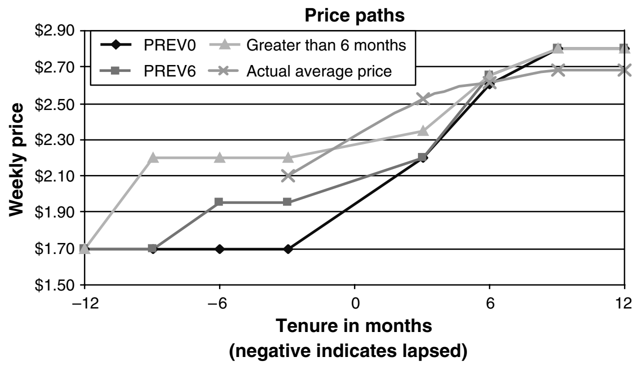

Note. These price paths are created assuming costless solicitations. If solicitations are assumed to cost $2 per call, the policy is not to solicit (and therefore not offer discounts to) customers who have lapsed more than a few months.

customer states ranging from 12 months since subscription lapse to 12 months or more of consecutive purchase. Also, Figure 1 includes paths for each of the three previous cumulative buying classifications (PREV0, PREV6, and "greater than six months"). For reference, the figure includes the actual average prices charged to customers with different purchase tenures.

The display of optimal pricing policies is complicated by the problem's dimensionality. Prices are a function of current time of tenure or time since subscription lapse, total previous purchasing duration, and previous price paid. The price paths displayed assume that customers pay the highest previous price specified and are acquired via the advocated acquisition policies. For example, at the first renewal point it is assumed that customers are paying the price specified for recently lapsed customers ($L_{it} < 3$) rather than the price offered to long-term lapsed customers.

For pure prospects ($\mathrm{PREV0}_{it} = 1$), the prices range from $1.70 for customers who have not purchased in more than a year (or who have been in the database for a year without ever having purchased), $2.20 for customers at first renewal, $2.60 for the second renewal, and a full price of $2.80. The price paths for the other two categories of previous purchasing involve shallower discounts to lapsed subscribers. For customers who have previously subscribed for between one and six months, the reacquisition price is $1.95 per week for recently lapsed individuals and $1.70 per week for those who have not purchased in six months. For customers with greater than six months of cumulative subscribing, the reacquisition price is $2.20. The three paths converge to the full price of $2.80 by the customer's third renewal.

The preceding results are generated under an assumption that the firm employs an extended planning horizon. Alternatively, the firm could use a shorter or a transactional objective function. Under a transactional objective, the firm's goal is to maximize the contribution provided by the current subscription. This amounts to discounting revenue streams beyond the current subscription period to zero.

When the optimization is performed using a finite three-month time horizon, it yields a policy of relatively high prices. The price path begins at $2.30 for the least promising segment and reaches a full price of $3.15. Compared to the infinite-horizon policy, the suggested policy recommends prices 35% higher for the worst prospects ($2.30 versus $1.70), 25% higher for recently acquired customers ($2.75 versus $2.20), and 12.5% higher for long-term subscribers ($3.15 versus $2.80).

5.2. Customer Valuation

In addition to determining dynamic marketing policies, a key feature of the DP approach is that it provides an improved method for calculating customer value by allowing attrition and revenue to be functions of the marketing mix. Table 5 provides customer value estimates for a variety of transaction histories. The column, labeled "Optimal," lists the estimated values when the optimal pricing strategy detailed in Figure 1 is used. The estimates range from $0.60 for a customer who has not subscribed in a year or more to $46.71 for a customer who has purchased for at

Table 5 Customer Lifetime Value Estimates

Population	Time as subscriber	Optimal (Figure 1)	3 price ($1.80, $2.40, $2.80)	AVG. 2 price* ($2.20, $2.80)	1 price ($2.80)
37%	Lapsed 12 months	$0.60	$0.53	$0.35	$0.12
16%	Lapsed 9 months	$2.47	$1.98	$1.02	$0.42
19%	Lapsed 3 months	$4.29	$4.01	$2.25	$0.83
13%	Tenure 3 months	$30.63	$30.51	$29.17	$31.96
6%	Tenure 9 months	$41.46	$39.43	$37.77	$36.54
10%	Tenure 12 months	$46.71	$45.83	$42.78	$42.47
	Average CV	$12.28	$11.89	$10.77	$10.50
	versus AVG 2 price schedule*	13.9%	10.3%	—	–2.6%

* The estimates of CV are point estimates. The establishment of confidence intervals is not straightforward, because the estimates are a function of the complete set of parameters.

least a year. These estimates are forward looking and, in the case of a positive tenure ($F_{it} \geq 1$), assume that customers are at the point of subscription expiration. Also, the estimates are a function of both the customer's state and the price offered. For a long-term lapsed customer, the offered price is $1.70 per week and the CV is $0.60 while the CV of $46.71 for a long-term loyal customer is based on an offered price of $2.80.

The cost of using nonoptimal policies is easily evaluated using constrained versions of the optimization model. For the policies reported thus far, the optimal policy is not limited in the number of discount levels allowed. The optimal policy given in Figure 1 calls for three discount levels ranging from about 40% for the least promising customer class to a 20-cent weekly discount for customers who have previously purchased for six months.

Table 5 illustrates the effects of increasingly restricted policies on CV. The column labeled "3 price" is the optimal pricing schedule when the optimization model is restricted to two price increases ($1.80, $2.40, and $2.80 per week). The two-price policy corresponds to the modal price points actually used, and the one-price policy is simply the full price advocated by the optimization model. As the degree of restriction increases, the effect on long-term value becomes particularly severe for lapsed customers. For customers who have not purchased in a year, the restriction to three prices reduces CV by 11.7% while restriction to a single price yields a CV of just $0.12. For current customers, the impact of more restrictive policies is also significant. Under the optimal policy, the CV of long-term subscribers is $46.71 compared to $45.83 under the three-price policy, $42.78 for the two-price policy ($2.20 and $2.80), and $42.47 for the single-price policy ($2.80).

The fact that the table includes instances where it appears the optimal policy does not strictly dominate the restricted policies merits explanation. For example, the estimated CV for a three-month subscriber is $31.96 under the single-price policy and only $30.63 under the optimal policy. This difference occurs because the more restrictive policy charges a higher price to lapsed customers, and the table does not reflect the previous price dimension of the optimal policy. For the single-price policy, the table gives the CV for a customer acquired at a price of $2.80, while the table lists the CV of a customer acquired at $1.70 who at the completion of the initial subscription is offered a renewal price of $2.80. The optimization procedure does provide the price to be charged and the accompanying CVs for a customer acquired at a price of $2.80, but this is not reported in the table. The CV estimates and price paths for the optimal policy are reflective of results of all customers being acquired using the recommended policies. From another perspective, the higher CVs associated with more restrictive policies would be achieved at the cost of acquiring fewer customers.

Assuming that the data sample is representative of the overall population, it is possible to compare the various policies' impact on the firm's customer equity. The first column in Table 5 gives the proportion of each transaction history type in the sample at the start of the data-collection period. Using the two-price policy as a base case, the optimal policy improves the value of the overall customer equity by 13.9%, the three-price policy leads to an increase in equity of 10%, and a restriction to a single price results in a 2.6% decline. Note, however, that if the population were comprised only of prospective and lapsed customers, the one-price policy would reduce customer equity by more than 50%.[4]

6. Discussion

In this paper, we compute relationship pricing strategies that optimize CV. The structure of the recommended pricing strategy should be of particular interest to firms that manage subscriber relationships. For our application to newspaper subscriptions, the procedure recommends the use of multiple discount levels. This conclusion contrasts with the current practice of using single steep, short-term discounts and then charging all customers the same regular price. However, while deviating from current practice, our findings are consistent with magazine industry reports that a series of diminishing discounts is more effective than a single steep discount (Barboza 1993, Freedman 1997).

The analysis is also relevant to the ongoing debate about the relationship between customer loyalty and profitability. In contrast to Dowling and Uncles (1997) and Reinartz and Kumar (2000), who argue that the link between length of purchasing and profits is small, our results are consistent with Reichheld and Teal (1996) and suggest purchase feedback effects exist and can be exploited to increase profitability. These different findings may be due to the categories studied. In contrast to the categories studied by Dowling and Uncles, the firm under study has direct relationships with customers and can customize prices.

Finally, a potential limiting factor for this type of individual-level marketing is the possible adverse reaction of loyal customers to the offering of discounts to prospects. Our proposed policies involve extensive use of this type of price discrimination. Feinberg et al. (2002) find that loyal customers react negatively when

[4] A sensitivity analysis of the relationship between CV and a five-cent incremental price increase across all relationship stages is predicted to decrease customer equity by 1.8%.

prospects are offered discounts. The intuitiveness of this laboratory result contrasted with the prevalence of price discrimination in practice leaves this as an open and important research issue. Another issue that merits further research is the effect of using a series of decreasing discounts. While researchers have found that introductory discounts can have adverse effects on repeat purchasing (see Neslin 2002), there is little work that examines psychological response to a sequence of diminishing discounts.

References

Barboza, Steven. 1993. The price is right. *Folio* **22**(8) 58–59, 118–119.

Berger, Paul, Nada Nasr. 1998. Customer lifetime value: Marketing models and applications. *J. Interactive Marketing* **12**(1) 17–30.

Bitran, Gabriel R., Susana V. Mondschein. 1995. Mailing decisions in the catalog sales industry. *Management Sci.* **42**(9) 1364–1381.

Dowling, Grahame R., Mark Uncles. 1997. Do customer loyalty programs really work? *Sloan Management Rev.* **38**(Summer) 71–82.

Erdem, Tulin, Susuma Imai, Michael Keane. 2003. Brand and quantity choice dynamics under price uncertainty. *Quant. Marketing Econom.* **1**(1) 5–64.

Feinberg, F., Aradhna Krishna, Z. John Zhang. 2002. Do we care what others get? A behaviorist approach to targeted promotions. *J. Marketing Res.* **39**(3) 277–291.

Freedman, Eric. 1997. Price-raising tricks of the trade. *Folio* **26**(8) 54.

Gonul, Fusun, Meng Ze Shi. 1998. Optimal mailing of catalogs: A methodology using estimable structural dynamic programming models. *Management Sci.* **44**(9) 1249–1262.

Gupta, Sunil, Donald Lehmann, Jennifer Stuart. 2004. Valuing customers. *J. Marketing Res.* **41**(1) 7–18.

Kamakura, Wagner, Gary J. Russell. 1989. A probabilistic choice model for market segmentation and elasticity structure. *J. Marketing Res.* **26**(November) 379–390.

Neslin, Scott. 2002. *Sales Promotion*. Marketing Science Institute, Cambridge, MA.

Rao, V. 1993. Pricing models in marketing. J. Eliashberg, G. Lilien, eds. *Handbooks in OR & MS*, Vol. 5. Elsevier Science, North Holland, The Netherlands, 517–549.

Reichheld, F., T. Teal. 1996. *The Loyalty Effect*. Harvard Business School Press, Boston, MA.

Reinartz, W., V. Kumar. 2000. On the profitability of long-life customers in a noncontractual setting: An empirical investigation and implications for marketing. *J. Marketing* **64**(October) 17–35.

Rigby, D., F. Reichheld, P. Schefter. 2002. Avoid the four perils of CRM. *Harvard Bus. Rev.* **80**(2) 101–109.

Sun, B., S. Neslin, K. Srinivasan. 2004. Estimating brand switching elasticities more accurately by accounting for rational consumer purchase strategies. *J. Marketing Res.* **40**(4) 389–405.

Vogel, Harold. 1998. *Entertainment Industry Economics*. Cambridge University Press, Cambridge, UK.