

Dynamic Customer Management and the Value of One-to-One Marketing

Romana Khan

McCombs School of Business, University of Texas at Austin, Austin, Texas 78713,
romana.khan@mcombs.utexas.edu

Michael Lewis

Olin Business School, Washington University in St. Louis, St. Louis, Missouri 63130,
michael.lewis@wustl.edu

Vishal Singh

Stern School of Business, New York University, New York, New York 10012,
vsingh@stern.nyu.edu

The concept of one-to-one marketing is intuitively appealing, but there is little research that investigates the value of individual-level marketing relative to segment-level or mass marketing. In this paper, we investigate the financial benefits of and computational challenges involved in one-to-one marketing. The analysis uses data from an online grocery and drug retailer. Like many retailers, this firm uses multiple promotional instruments including discount coupons, free shipping offers, and a loyalty program. We investigate the impact of customizing these promotions on the two most important consumer decisions: the decision to buy from the store and expenditure. Our modeling approach accounts for two sources of heterogeneity in consumers' responsiveness to various marketing mix elements: cross-sectional differences across consumers and temporal differences within consumers based on the purchase cycle. The model parameter estimates are fed into a dynamic programming model that determines the optimal number, sequence, and timing of promotions to maximize retailer profits. A series of policy simulations show that customizing promotions leads to a significant increase in profits relative to the firm's current practice of uniform promotions. However, the effectiveness of various promotions varies because of both cross-sectional differences in consumers as well within consumer heterogeneity due to purchase cycle factors. For instance, we find that free shipping tends to be the preferred instrument for re-acquiring lapsed customers, whereas an across-the-board price cut (via a discount coupon) is the most effective tool for managing the segment of most active customers. Interestingly, we find that customizing based on within-customer temporal heterogeneity contributes more to profitability than exploiting variations across consumers. This is important because the computational burden of implementing the dynamic optimization to account for cross-sectional heterogeneity is far greater than accounting for temporal heterogeneity. Furthermore, targeting promotions based only on timing rather than the nature and magnitude of the offers across consumers alleviates the public relations risks of price discrimination. Implications for marketing managers are also discussed.

Key words: CRM; one-to-one marketing; promotions; dynamic programming; heterogeneity; database marketing

History: Received: December 14, 2007; accepted: February 23, 2009; processed by Scott Neslin. Published online in *Articles in Advance* June 19, 2009.

1. Introduction

Much of the promise of customer relationship management (CRM) is that by understanding individual-level behavior, firms will be able to refine and customize marketing tactics to increasingly fine segments or even to individual customers (Peppers and Rogers 1993). The development of customer databases and communication technologies (Xie and Shugan 2001) has enabled firms to move beyond uniform marketing policies and begin implementing customized marketing strategies. This coincides with a growing body of empirical research focused on the development of individual-level marketing policies (Zhang

and Krishnamurthi 2004, Rust and Verhoef 2005, Lewis 2005b). Although CRM systems have become increasingly prevalent, in many cases firms have been unable to convert the individual-level customer data into profitable marketing policies. This creates significant uncertainty and concern about the economic benefits of these systems (Rigby et al. 2002).

Our objective in this paper is to explore the benefits to a retailer from customizing promotions based on individual-level information. To accomplish this, we estimate a joint model of purchase incidence and expenditure to identify the impact of various promotion incentives on purchase decisions. Our model

accounts for two sources of heterogeneity in responsiveness to promotions: cross-sectional differences across consumers and temporal differences within consumers based on the purchase cycle. The model parameters are used in a dynamic programming model that determines the optimal number, sequence, and timing of promotions to maximize retailer profits. We also conduct a series of policy simulations to show the impact on profitability from accounting for different sources of heterogeneity when determining the optimal policies. Our main contributions are quantification of the benefits from customization in terms of profitability and decomposition of the marginal gains in profitability from accounting for different sources of consumer heterogeneity.

The data for our study come from an online retailer specializing in nonperishable grocery and drugstore items. In addition to customer transaction histories, we observe the firm's marketing promotions, which include discount coupons, free shipping offers, and a reward program. Our data are well-suited to studying multiple promotions policies because these promotional instruments each provide different incentives. The "value" of savings with a percentage-off coupon is determined by the customer who can scale it by increasing the purchase amount while a free shipping promotion removes a fixed transaction cost (Lewis et al. 2006). In contrast, reward programs provide incentives for concentrating purchases over time.

To identify the impact of the promotions on purchase behavior, we estimate a joint model of purchase incidence and expenditure. We model the household purchase decision using a discrete-choice framework with time-varying coefficients (Gupta 1991, Wedel et al. 1995). Conditional on purchase, in-store expenditures are modeled using a semilog specification that has been used extensively in marketing (see Blattberg and Neslin 1990). Our modeling approach accounts for cross-sectional variation in responsiveness to different types of promotions and within-customer temporal heterogeneity because of purchase cycle factors. We use a hierarchical Bayesian estimation approach that allows the full set of model parameters to vary across consumers as a result of both observed (e.g., demographics) and unobserved factors (Allenby and Rossi 1999).

We find distinct effects of each promotion on purchase incidence and expenditures. Although each promotion type increases purchase incidence, the effect on expenditure varies and may be negative. Furthermore, response to the promotions is not static but varies with time since purchase. The heterogeneity in promotion response suggests that the firm can benefit from customizing promotions at the individual level. Cross-sectional heterogeneity in response to

different promotions can be exploited by customizing the particular promotion that a customer will receive. The temporal dimension of heterogeneity can be exploited by the timing of the promotion within a customer's purchase cycle. To determine the optimal sequence and timing of promotions, we develop a dynamic programming model (Lewis 2005b, Simester et al. 2006) of the firm's promotion decision problem. We then conduct a series of policy simulations in which we sequentially account for additional sources of heterogeneity in the dynamic optimization model. This allows us to identify the marginal profits from increasing levels of customization and quantify the gains from accounting for temporal and cross-sectional heterogeneity.

As expected, we find that customized promotions yield large increases in revenue and profits relative to uniform promotion policies. Interestingly, the benefits from accounting for temporal heterogeneity exceed the increases from accounting for cross-sectional heterogeneity. This is important because incorporating cross-sectional heterogeneity is far more computationally intensive because it requires that the dynamic optimization be solved at the individual-customer level. Taking this into consideration, we also investigate the benefits of conducting the dynamic optimization at the segment rather than individual level. We find that this approach significantly reduces the computational burden while retaining significant profitability gains from incorporating both sources of heterogeneity.

Our paper contributes to a growing literature (Zhang and Krishnamurthi 2004, Lewis 2005b) on dynamically customized promotions. Our research is unique in that we use individual-level coefficients to evaluate the benefits of optimizing at the level of individual customers. This allows for a comparison of the benefits of developing true one-to-one policies relative to segment-level policies. Furthermore, our research considers multiple promotion types beyond price discounts. This is an important distinction because these promotions are structurally different and offer different incentives that can alter different elements of consumer behavior.

More generally, our results should be of interest to academicians and practitioners concerned with CRM issues. Previous research in direct marketing has investigated models for selecting prospects (Bodapati and Gupta 2004, Bult and Wansbeek 1995) and measured the effect of loyalty programs (Lewis 2004). Within the larger field of CRM, we add to the optimal contact literature (see Blattberg et al. 2008). The optimal contact literature features dynamic optimization models that determine ideal schedules for catalogs (Simester et al. 2006, Bitran and Mondschein 1996) and models that develop dynamic marketing

policies in circumstances where customers are strategic and forward looking (Gonul and Shi 1998, Lewis 2005a). We contribute to this literature by adopting a multiple-campaign orientation that considers both cross-sectional and temporal heterogeneity in response. Our work is also related to a body of work focused on using customer information to target customers (Heilman et al. 2003, Rossi et al. 1996). Our work adds to this literature by quantifying the value of tracking and acting on both unobserved preference heterogeneity and observed transaction history measures. Finally, because our purpose is to demonstrate the profitability implications for the retailer from customization, our work is also related to research focused on the linkage between marketing actions and financial outcomes (Kamakura et al. 2002, Rust et al. 2000, Gupta et al. 2004).

The remainder of the paper is organized as follows. The next section describes our data and the firm's marketing tactics. Section 3 details the purchase incidence and expenditure model and estimation results. The dynamic programming model and optimization results are discussed in §4. Section 5 concludes the paper with a discussion of managerial implications, limitations to our research, and areas for future inquiry.

2. Customer Data and Firm Marketing Practices

For the empirical analysis, we use data provided by an online retailer specializing in nonperishable grocery and drugstore items. The data set contains records of all customer transaction histories and extensive data related to the firm's marketing policies through the firm's first 14 months of operation. Table 1 provides summary statistics of purchase behavior for the random sample used for estimation. The average order size is \$58, and the typical basket contains more than 20 items. The average interpurchase time is approximately five weeks. For the timing aspects of the analysis, we selected a week as our unit of time based on the firm's marketing practices (promotional pricing, e-mail communications, etc.), which vary on a weekly cycle. Time duration since last purchase, previous purchase amount, and

Table 1 Customer Descriptive Statistics

	Mean	Std. dev.	Minimum	Maximum
Order size (\$)	58.2	34.7	4	391
Number of orders	8.0	6.6	2	67
Interpurchase time	5.3	5.5	1	32
Baby	0.23	0.43	0	1
Child	0.62	0.48	0	1
Pet	0.42	0.49	0	1

Table 2(a) Marketing Mix Descriptive Statistics

Variable	Mean	Std. dev.	Minimum	Maximum
Price	1.92	0.07	1.68	2.1
Banner ad	0.14	0.35	0	1
Coupon	0.21	0.41	0	1
Freeship	0.07	0.25	0	1
Reward	0.16	0.017	0.15	0.25

Note. $n = 56$ weeks.

Table 2(b) Marketing Mix Correlation Matrix

	Price	Banner ad	Coupon	Freeship
Price	1			
Banner ad	0.125 (0.359)	1		
Coupon	−0.105 0.235	−0.089 0.515	1	
Freeship	0.112 0.410	−0.113 0.406	−0.145 0.287	1

Note. Correlation coefficients, $N = 56$ ($p = \text{value } H_0: \rho = 0$).

purchase frequency rate are used in the purchase incidence model to account for possible inventory and attrition effects. The amount of previous purchase and time since previous purchase are also used as covariates in the expenditure model.

Table 2(a) presents summary statistics of the retailer's promotional activities, and Table 2(b) shows the correlations. There is no evidence of significant correlation between the various marketing activities. We define *price* as a summary measure that captures the overall price environment of the store in a given week. The price variable is the weekly log price of the top 100 items in terms of unit sales. In addition to standard short-term, item-specific discounts, the firm used a variety of other promotional instruments including discount coupons, free shipping offers, and a loyalty-based reward program. Each of these possess unique characteristics that may differentially influence individual consumers in terms of incidence and quantity decisions.

The first distinctive promotion involved the distribution of coupons that offered a 10% discount on a purchase made in a specified week. These were delivered by e-mail to the entire customer base rather than targeted to specific individuals. *Coupon* is a binary variable that indicates the availability of a discount coupon during the week. Each coupon could be redeemed over a two-week period. If the coupon was redeemed in the first week, the variable takes a value of zero in the following week for that household. A coupon was available to customers for 10% of the time horizon.

The firm experimented with several different shipping and handling fee policies. The base shipping

schedule charged \$2.99 to ship an order of less than \$50 of merchandise and \$4.99 to ship an order of \$50 or more. *Freeship* is an indicator variable for a free shipping offer during a week. The free shipping offer was available to customers for 7% of the time horizon. The discount coupon and free shipping were nonoverlapping promotions offered at irregular intervals.

The firm also used a loyalty program that rewarded customers based on cumulative expenditures over time. The loyalty program provided customers with 500 frequent flier miles at the airline of their choice when they spent \$1,000 in a 12-month period. Subsequent rewards were offered when consumers reached cumulative spending levels of \$1,500 and \$2,000. To capture the impact of the reward program, we created a measure of closeness to achieving the required expenditure for the reward. *Reward* is computed as $1/\ln(\text{dollar expenditure required for reward})$. This measure is low when the goal is very far and increases as the goal approaches. The average value of this measure over the observation period was 0.16, which corresponds to an average distance from the reward goal of \$471.

The free shipping, the discount coupons, and the loyalty rewards each possess unique characteristics. The base shipping schedule includes elements of nonlinear pricing that may have both positive and negative effects on order-size decisions. The discrete jump in the shipping fee from \$2.99 to \$4.99 at the \$50 threshold may restrain order size by penalizing orders exceeding \$50. However, the base shipping fee also involves a structure that provides quantity discounts. As orders grow larger past the \$50 level, the percentage impact of the shipping surcharge diminishes. The free shipping offer removes these nonlinear pricing structures. Therefore, although free shipping is expected to increase order incidence, it is not clear what the overall effect will be on order size. The discount coupons provide the consumer an opportunity to define the value of the promotional offer. For example, a 10% discount translates to a savings of \$10 on a \$100 order but only \$2.50 on a \$25 order. This type of promotion may therefore increase both incidence and order amount. The loyalty program promotion operates as a dynamic incentive scheme. By operating as a function of cumulative expenditure, the loyalty program can influence behavior over an extended time horizon rather than at just a single purchase occasion.

In the empirical application, we account for the retailer's costs to evaluate the profitability of the retailer's actions and promotion strategies. The retailer's reported average margin across products, approximately 25%, is used to compute profitability based on purchase amount. Note, however, that the margin will vary with the basket contents for each customer,

but this information is not available to us. Ideally, a promotion will shift purchases toward higher margin products, in which case the profitability of promotions will be underestimated.

We also consider costs that are specific to each promotion. The cost of the coupon promotion is the redemption cost. The shipping cost paid by the retailer to the shipper ranges between \$6.57 and \$9.89 depending on the order size. The retailer subsidized the shipping to consumers even when there was no free shipping offer. For the reward offer of 500 frequent flyer miles, we assume a cost of \$10 based on the industry average of \$0.02 cents per mile for partner programs.

In addition to the promotions, the firm used weekly e-mail communications and banner advertising to promote itself. We use information on banner advertisement activity on four major websites: AOL.com, Excite.com, iVillage.com, and Yahoo.com. The variable *banner ad* takes on values between zero and one, indicating whether there was a banner ad placed on one or more of these websites. With this information, we do not know whether the customer was actually exposed to the banner so it enters as a control variable rather than a focal promotion.

There are two key limitations of the data. First, we do not observe information on competitor activities. This is a general problem with almost all data sets in the CRM world. This missing information is economically problematic to the extent that the promotional activities of competitors are correlated with the retailer's promotion activities. This is unlikely to be the case in our application because the promotions are offered by e-mail and the sources of competition are broad. The second limitation is that we do not observe when a customer browses the website but does not make a purchase. This is a general problem with all such data sets including direct marketers and offline retailing. Note that our application is based on all the information that a manager of this firm would observe and use for decision making.

3. Purchase Incidence and Expenditure

3.1. Model

We begin with a description of our joint model of purchase incidence and expenditures. The majority of empirical studies in marketing have used the proportional hazard model to characterize the purchase timing decisions of households, either in continuous time (Jain and Vilcassim 1991, Chintagunta and Haldar 1998) or discrete time (Gupta 1991, Helsén and Schmittlein 1993). For our application, a discrete-time formulation is more appropriate because it allows

us to explicitly account for marketing activity in periods when households do not make a purchase. We use a discrete-choice framework with time-varying coefficients to capture the duration dependence in the consumer's purchase decision and promotion effectiveness.

In each time period, the individual decides whether to make a purchase. Let $U_i(t)$ be the utility for individual i from making a purchase from the store in period t . Assume that

$$U_{it} = X'_{it}\beta_{it} + \varepsilon_{it}, \quad t = 1, \dots, T_i. \quad (1)$$

Here, X_{it} is a vector of time-varying covariates affecting utilities, which can include store-specific time-varying marketing mix variables such as price, the presence of a discount coupon, and the shipping and handling fee schedule, as well as customer-specific variables such as status in a loyalty program and demographics. We define D_{it} as one when $U_{it} > 0$ and zero otherwise.

The model in (1) implies a model for purchase times. Suppose we observe a purchase duration of length t_1 followed by a purchase duration of t_2 . These durations then have likelihood

$$\begin{aligned} & \Pr(T_{i1} = t_1, T_{i2} = t_2 \mid \beta_i^{t_1+t_2}, X_i^{t_1+t_2}) \\ &= \left\{ \prod_{t=1}^{t_1-1} \Pr(D_{it} = 0 \mid X_{it}, \beta_{it}) \right\} \times \Pr(D_{it_1} = 1 \mid X_{it_1}, \beta_{it_1}) \\ & \quad \times \left\{ \prod_{t=t_1+1}^{t_1+t_2-1} \Pr(D_{it} = 0 \mid X_{it}, \beta_{it}) \right\} \\ & \quad \times \Pr(D_{i, t_1+t_2} = 1 \mid X_{i, t_1+t_2}, \beta_{i, t_1+t_2}), \end{aligned} \quad (2)$$

where $\beta_i^{t_1+t_2}$ is the entire time path of coefficients up to T , i.e., $\beta_i^T = \{\beta_i(t)\}_{t=1}^T$; similarly, $x_i^{t_1+t_2}$ is the entire path to T for the covariates $x_i^{t_1+t_2} = \{x_{it}\}_{t=1}^{t_1+t_2}$.

We partition the vector of covariates X'_{it} into X'_{1it} with individual-specific time-varying coefficients β_{1it} and X'_{2it} with individual-specific time-invariant coefficients β_{2i} :

$$U_{it} = X'_{1it}\beta_{1it} + X'_{2it}\beta_{2i} + \varepsilon_{it}. \quad (3)$$

A simple model that allows for duration dependence and heterogeneity in the coefficients β_{1it} is of the form:

$$U_{it} = X'_{1it}\beta_{1i}f(\tau_{it}) + X'_{2it}\beta_{2i} + \varepsilon_{it}, \quad (4)$$

where τ_{it} is time since last purchase, $f(\cdot)$ is a known vector function, β_{1i} is a matrix, and β_{2i} is a vector of individual-specific coefficients. The function $f(\tau_i)$ can be as flexible as needed to capture the duration dependence in the parameters. It is scaled across people using β_{1i} , allowing us to approximate the shape of the household-specific hazard function. We consider

several specifications of $f(\tau_i)$ that are discussed in the next section. Assuming ε_{it} is i.i.d. standard normal, this structure implies that the probability of purchase at time t conditional on last purchase τ_i weeks ago is

$$\Pr(D_{it} = 1 \mid \beta_i, \tau_i) = \Phi(X'_{1it}\beta_{1i}f(\tau_i) + X'_{2it}\beta_{2i}). \quad (5)$$

This is the hazard rate induced by (3) and captures the notion of individual-specific hazard. Although our focus is on capturing duration dependence in the parameters, there are more general approaches to modeling parameter dynamics as in Kim et al. (2005).

To model expenditure conditional on purchase incidence (Boatwright et al. 2006), we use a semi-log specification. In particular, let E_{it} be log expenditures for individual i in time period t (which is 0 unless $U_i(t) > 0$):

$$E_{it} = X'_{e, it}\lambda_i + u_{it}, \quad (6)$$

where $u_{it} \sim N(0, v_i^{-1})$ and $X'_{e, it}$ is a vector of covariates such as basket price, availability of coupon, and shipping and handling charges. Note that ε_{it} and u_{it} are uncorrelated. The coefficients in the expenditure equation (like the purchase time equation) are allowed to be individual specific. Let $\theta_i = (\beta_i, \lambda_i)$ be the full vector of coefficients from the purchase incidence and expenditure equations. We assume that θ_i follows a multivariate normal distribution with a mean vector ΠZ_i and covariance Ω :

$$\theta_i \mid \Pi, Z_i \sim N(\Pi Z_i, \Omega), \quad (7)$$

where Z_i is a vector containing household characteristics (such as demographics) that explain differences in sensitivity as a result of observed factors.

For inference, we use a hierarchical Bayesian approach. In particular, we use a Markov chain Monte Carlo procedure to simulate the posterior distribution of the model parameters and to compute household-level estimates of preferences. Bayesian procedures have become quite standard in the literature and, as discussed in Allenby and Rossi (1999), are well-suited for these models, especially when one is interested in making inference at the individual level. Next, we discuss the criterion for model selection and present the parameter estimates.

3.2. Results

3.2.1. Model Selection: In-Sample Fit and Predictive Performance. The function $f(\tau)$ in Equation (4) captures the underlying duration dependence of the response parameters. To select the specification for $f(\tau)$, we consider over 20 alternative combinations of the vector $(1, \tau, \ln(\tau), \tau^2)'$. We allow $f(\tau)$ to be unique to the intercept, $f_i(\tau)$, and the promotion response parameters, $f_p(\tau)$. Of the 56 weeks of purchase history available, the model was estimated with

Table 3 In-Sample Fit

Model			Purchase incidence			Purchase amount
			Hit rate (%)			
$f_{int}(\tau)$	$f_p(\tau)$	MAD	Overall	Purchase	Nonpurchase	MAD (\$)
(1) Base		0.25	74	16	95	14.78
(2) $1, \tau$		0.40	59	44	64	26.15
(3) $1, \tau, \ln(\tau)$		0.27	64	54	68	16.09
(4) $1, \tau, \tau^2$		0.41	58	36	66	26.20
(5) $1, \tau, \ln(\tau), \tau^2$		0.24	75	4	99	13.67
(6) $1, \tau, \ln(\tau)$	$1, \tau$	0.34	61	42	68	19.21
(7) $1, \tau, \ln(\tau)$	$1, \ln(\tau)$	0.32	69	68	69	20.14
(8) $1, \tau, \ln(\tau)$	$1, \tau, \ln(\tau)$	0.32	66	63	67	20.65
(9) $1, \tau, \ln(\tau)$	$1, \tau, \tau^2$	0.23	76	9	99	13.32
(10) $1, \tau, \ln(\tau)$	$1, \tau, \ln(\tau), \tau^2$	0.25	75	15	95	13.81

the first 50 weeks and the remaining 6 weeks were used for predictive testing. Model selection is based on the in-sample fit and predictive performance measures reported in Tables 3 and 4. For the purchase incidence model, we compute four performance measures: (1) Mean absolute deviation (MAD) is computed as the mean absolute value of the difference between the purchase indicator vector and the predicted probability. (2) Hit rate is computed as the percentage of time that the observed choice has the highest predicted probability (i.e., >0.5). (3) Purchase hit rate is computed as the percentage of time that purchase incidence is correctly predicted. (4) Nonpurchase hit rate is computed as the percentage of time that nonpurchase occasions are correctly predicted. For the purchase amount decision, we compute the MAD by taking the average absolute value of the difference between actual purchase amount and purchase amount predicted by the model.

Tables 3 and 4 report the fit statistics of a selected subset of the alternative specifications considered. A brief description of the specifications shown in Tables 3 and 4 follows. The baseline (1) represents the most restrictive case, where we do not allow for any parameter dynamics so that $f_i(\tau) = f_p(\tau) = (1)$. Models (2)–(5) demonstrate various specifications of $f_i(\tau)$

Table 4 Predictive Performance in Holdout Sample

Model			Purchase incidence			Purchase amount
			Hit rate (%)			
$f_{int}(\tau)$	$f_p(\tau)$	MAD	Overall	Purchase	Nonpurchase	MAD (\$)
(1) Base		0.23	79	13	93	12.33
(2) $1, \tau$		0.44	55	43	57	28.71
(3) $1, \tau, \ln(\tau)$		0.23	61	59	62	12.62
(4) $1, \tau, \tau^2$		0.39	60	35	65	26.02
(5) $1, \tau, \ln(\tau), \tau^2$		0.17	82	1	98	8.91
(6) $1, \tau, \ln(\tau)$	$1, \tau$	0.30	71	30	79	16.59
(7) $1, \tau, \ln(\tau)$	$1, \ln(\tau)$	0.34	66	64	66	20.52
(8) $1, \tau, \ln(\tau)$	$1, \tau, \ln(\tau)$	0.33	67	51	70	18.61
(9) $1, \tau, \ln(\tau)$	$1, \tau, \tau^2$	0.18	81	1	98	9.35
(10) $1, \tau, \ln(\tau)$	$1, \tau, \ln(\tau), \tau^2$	0.21	79	13	98	11.76

while restricting $f_p(\tau) = (1)$. Based on performance measures, specification (3), where $f_i(\tau) = (1, \tau, \ln(\tau))$, best captures the dynamics of the intercept. Models (6)–(10) show alternative specifications of $f_p(\tau)$, where $f_i(\tau) = (1, \tau, \ln(\tau))$. Although we considered many more specifications, we report the performance of a limited set because of space considerations.

First, we discuss the performance measures of models (1)–(5), where the focus is on the specification of $f_i(\tau)$ and we restrict $f_p(\tau) = 1$. The results in Table 3 show that the in-sample performance of the base model (1) and model (5) are best in terms of purchase probability and the amount of MAD as well as overall hit rates. However, a closer examination of the results shows that this is achieved by overpredicting nonpurchases. Because most of the observations are weeks in which individuals make no purchases, it is possible to achieve high hit rates by always predicting no purchases. The nonpurchase hit rate for these models is over 90% but the purchase hit rate is very low. This highlights the importance of looking beyond aggregate performance measures to assess model fit. The performance measures for the holdout data (Table 4) are similar. Model (3) has the better performance both in sample and in the holdout data, reflected in high purchase (59%) and nonpurchase (62%) hit rates and the purchase probability (0.23) and purchase amount (12.62) of MAD. In models (6)–(10), we consider alternative specifications of $f_p(\tau)$ with $f_i(\tau) = (1, \tau, \ln(\tau))$. Models (9) and (10) have the lowest purchase probability and spend MAD and highest overall hit rates. However, once again it is at the expense of the purchase hit rate, which is very low. We select model (7) because it dominates on both purchase (68%) and nonpurchase (69%) hit rates and has reasonable predictive performance in the holdout data. In what follows, we discuss the results based on the estimation of model (7), where $f_i(\tau) = (1, \tau, \ln(\tau))$ and $f_p(\tau) = (1, \ln(\tau))$.¹

3.2.2. Parameter Estimates. The estimation results are shown in Table 5 for the purchase incidence model and Table 6 for the expenditure model. The first two columns report the estimated population mean and standard deviations, followed by the effect of demographic variables. The standard errors are reported in parentheses.

We first discuss general model results and then focus on the impact of promotions. In most cases, the demographic variables have insignificant effects in both equations. The overall variance in response

¹ We use a common f_p function across promotions. This is a limitation because ideally $f_p(\tau)$ would be unique to each promotion, allowing the dynamics to vary by promotion. However, use of a common f_p maintains consistency and facilitates comparison across promotions.

Table 5 Estimation Results: Purchase Incidence Model

Variable	Parameter mean	Std. dev.	Baby	Child	Pet
Intercept	0.649 (0.243)	3.201 (0.092)	0.104 (1.07)	0.972 (1.07)	0.408 (0.98)
τ	−0.094 (0.133)	1.657 (0.019)	−0.852 (0.30)	1.022 (0.27)	0.325 (0.28)
$\ln(\tau)$	0.389 (0.094)	0.968 (0.114)	0.445 (0.18)	−0.159 (0.19)	−0.103 (0.16)
Price	−1.024 (0.149)	1.342 (0.377)	0.015 (0.44)	0.016 (0.43)	0.034 (0.39)
$\ln(\text{previous spend})$	−0.131 (0.035)	0.462 (0.079)	0.053 (0.10)	−0.344 (0.11)	0.017 (0.09)
Banner ad	−0.084 (0.096)	0.504 (0.310)	0.06 (0.09)	0.055 (0.08)	0.022 (0.08)
Coupon	0.145 (0.075)	1.366 (0.080)	−0.352 (0.58)	−0.958 (0.57)	0.075 (0.65)
Coupon * $\ln(\tau)$	−0.058 (0.017)	1.066 (0.235)	−0.191 (0.28)	0.29 (0.29)	−0.04 (0.29)
Coupon * Frequency	0.611 (0.183)	2.592 (0.653)	0.966 (0.94)	−0.303 (1.10)	−0.709 (1.15)
Freeship	−0.288 (0.189)	1.426 (0.020)	−0.008 (0.79)	0.99 (0.73)	−0.488 (0.81)
Freeship * $\ln(\tau)$	0.027 (0.011)	1.260 (0.181)	0.342 (0.36)	−0.515 (0.36)	0.118 (0.35)
Freeship * Frequency	1.225 (0.385)	2.162 (0.645)	−1.014 (1.42)	−0.165 (1.71)	1.981 (1.43)
Reward	0.521 (0.15)	0.250 (0.019)	−0.283 (0.08)	0.158 (0.08)	−0.163 (0.07)
Reward * $\ln(\tau)$	0.261 (0.043)	0.209 (0.007)	0.018 (0.03)	−0.029 (0.03)	0.021 (0.03)
Reward * Frequency	0.325 (0.094)	2.617 (0.734)	0.106 (0.26)	−0.774 (0.26)	−0.404 (0.24)

Note. Standard errors are in parentheses.

parameters (not reported here) accounted for by the observed demographic variables ranges between 2% and 7% in the purchase incidence equation and between 2% and 12% in the expenditure equation. All of the standard deviation parameters, which capture the effect of unobserved sources of heterogeneity, are significant and reflect that most of the heterogeneity across households is due to unobserved factors. This finding of low explanatory power of demographics is consistent with previous research (Gupta and Chintagunta 1994, Rossi et al. 1996).

Weekly price levels affect the purchase incidence and expenditure decision in the expected manner. The price coefficient is negative in both the purchase-incidence (−1.02) and expenditure (−0.55) model so that higher prices decrease purchase incidence and reduce expenditures. The impact of previous expenditures is negative and significant on both purchase incidence (−0.13) and expenditure (−0.28). Because previous expenditure enters as a log term, we can interpret the parameter estimate from the expenditure model as an elasticity, so a 10% increase in expenditure reduces the value of the next purchase by 2.8%. In the

expenditure model, the coefficient on $\ln(\tau)$ (0.018) indicates a positive relationship between expenditure and time since previous purchase.


The banner advertising parameters are insignificant in both the purchase incidence and expenditure equations. This finding adds to the mixed evidence on the effectiveness of banner advertising (Sherman and Deighton 2001, Manchanda et al. 2006). However, we should point out that the effect of banner advertising is more difficult to capture in our study because actual exposure and click throughs on banner advertisements are not observed.

Next, we evaluate the effects of the three types of promotions used by the retailer. The purchase incidence model is structured to measure how the effects of these promotions vary with recency (time since last purchase) and frequency of purchase. The estimated coefficients can also be used to compute the implied hazard—the probability of purchase at time t conditional on the time elapsed since previous purchase. Figure 1 plots the baseline hazard when

Table 6 Estimation Results: Expenditure Model

Variable	Parameter mean	Std. dev.	Baby	Child	Pet
Intercept	5.256 (0.396)	1.706 (0.896)	0.416 (0.33)	−0.235 (0.34)	−0.554 (0.31)
$\ln(\tau)$	0.018 (0.051)	0.149 (0.022)	−0.044 (0.04)	−0.010 (0.08)	0.460 (0.38)
Price	−0.548 (0.182)	0.822 (0.228)	−0.149 (0.15)	0.169 (0.16)	0.198 (0.15)
$\ln(\text{previous spend})$	−0.280 (0.034)	0.241 (0.035)	0.001 (0.03)	−0.024 (0.03)	0.017 (0.03)
Banner ad	0.046 (0.040)	0.284 (0.036)	−0.030 (0.03)	0.007 (0.04)	−0.025 (0.03)
Coupon	0.113 (0.094)	0.404 (0.057)	−0.026 (0.04)	0.099 (0.07)	−0.016 (0.06)
Freeship	−0.227 (0.136)	0.495 (0.086)	0.041 (0.04)	−0.193 (0.09)	0.135 (0.08)
Reward	0.006 (0.020)	0.138 (0.087)	0.016 (0.01)	0.019 (0.01)	0.029 (0.01)

Note. Standard errors are in parentheses.

Figure 1 Impact of Promotions on Purchase Hazard

no promotion is offered and the hazard when each promotion type—coupon, free shipping, and rewards program—is offered.

Coupon. Figure 1 shows that the impact of the coupon on the purchase hazard is most dramatic within eight weeks after a purchase, after which coupon effectiveness decreases. So if a consumer has recently made a purchase, the coupon serves as an effective incentive to induce repurchase but has a limited impact on long-term lapsed customers. The coefficient on the interaction between the coupon and purchase frequency is positive. This indicates that the positive impact of the coupon on purchase incidence is increasing with purchase frequency. For the expenditure equation, we find that the effect of the coupon is positive and significant. After accounting for the discount, the coupon yields an approximately 12% ($\exp(0.113)$) increase in expenditure.

Free Shipping Offer. Figure 1 shows that free shipping is not particularly effective immediately after purchase but becomes more effective as the time since previous purchase increases. In contrast, the coupon was an effective mechanism immediately after purchase and the effectiveness decreased with time elapsed since previous purchase. The shipping and discount promotions likely trigger different responses because their respective structures affect different aspects of consumer decision making. Free shipping removes a transaction cost and eliminates an order size penalty, whereas a 10% discount coupon improves the attractiveness of the individual products for sale. The impact of free shipping is higher among frequent shoppers, similar to the outcome for the coupon.

Also, in contrast to the coupon, the offer of free shipping reduces expenditure by about 20% ($\exp(-0.227)$). This suggests that with the free shipping offer, consumers do not feel the need to order a large basket to justify the shipping cost. That the elimination of shipping fees results in smaller orders suggests that the order size disincentives play a larger role relative to the transaction cost aspect of the shipping fee structure that dissuades small orders.

Reward Program. To determine the impact of the reward program, we assess whether consumers accelerate purchasing or increase expenditures as the cumulative buying goal approaches. Figure 1 shows the implied hazard when consumers are within \$50 of reaching the reward threshold. Although the reward program has a positive impact on the purchase hazard, it is relatively less effective than the other incentives. As time since previous purchase increases, the positive impact on purchase probability diminishes but at a slower rate than the effect of the coupon offer. This results in the reward offer being more effective than the coupon as weeks since previous purchase increase. As with the other promotions, reward program effectiveness increases with purchase frequency. The estimation results do not reveal a significant impact of the reward program on purchase amount. This is not surprising because our formulation is somewhat underspecified. It may also be necessary to more fully capture the dynamics of consumer decision making because consumers can plan to achieve a reward over a multiple-week period.

4. The Value of One-to-One Marketing

Throughout the data collection period, the retailer followed a uniform promotion policy that offered the entire customer base the same promotion at the same time. As discussed in the previous section, there is significant temporal and cross-sectional heterogeneity in response to the different promotions. This provides an opportunity for the retailer to improve performance by customizing both the timing and type of promotion that each individual receives. In this section, we develop a dynamic programming model of the firm's promotion decision problem designed to determine the sequence and timing of promotions to maximize expected profits over some extended time horizon. We conduct a series of policy simulations in which we sequentially account for additional sources of heterogeneity in the dynamic optimization model. This sequence allows us to investigate the marginal profits from customization and to quantify the gains from accounting for temporal and cross-sectional heterogeneity.

4.1. Dynamic Optimization Model

The first step in developing a relationship-oriented promotions policy is the specification of an appropriate dynamic objective function. For the general case, customers are classified into one of $s \in S$ possible states according to their specific transaction history, and the firm is assumed to be able to select marketing actions or controls a from a set of feasible actions A .² Defining the single-period profit function of selecting action a for a customer in state s at time t as $\pi(s_t, a_t)$, the single-period discount factor as α , and the length of the decision horizon as T , the customer management objective function may be written as in Equation (8). This objective function is the discounted sum of single-period profits over a T period horizon:

$$\max_{a_t \in A} E \left\{ \sum_{t=0}^T \alpha^t * \pi(s_t, a_t) \right\}. \quad (8)$$

The customer management goal is to determine the mapping of marketing policies to observed customer behaviors to maximize Equation (8). In our context, this amounts to determining the functional relationship $a(s_t)$ between customer state s_t and marketing promotions a . The set of actions that we

consider includes four options: no promotion, free shipping, discount coupon, and loyalty reward. The loyalty reward promotion is operationalized by setting the reward measure to within \$100 of the reward threshold. Although our focus is on the selection of promotions, the model is adaptable to consider other measures such as overall pricing or customer service level.

The customer management policy, $a(s_t)$, is a rule that links observable measures such as recency and amount of previous purchase with what, if any, promotion should be offered in a given week. The goal is to select the proper promotional action for all possible combinations of previous buying levels, time since last purchase, and previous basket size. We also include past promotional activity as a state variable, which allows for the formulation of policies that limit the number of promotions that a given customer may receive.

The firm's optimization problem is to select the sequence of promotions that maximize profitability over a specified T period time horizon. For a representative customer, 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 s_t . The value function may be given in general equation form as

$$V(s_t) = \max_{a_\kappa \in A} E \left\{ \sum_{\kappa=t}^T \alpha^{\kappa-t} * \pi(s_\kappa, a_\kappa) \right\}. \quad (9)$$

At the level of the individual customer, the value function is equivalent to the maximum expected customer value based on optimal selection of marketing policies.

Thus far, the objective function does not include a direct statement of how customers evolve over time. In Equation (10), the value function is written to make the relationship dynamics more explicit. In this equation $\Pr(s_{t+1} | s_t, a(s_t))$ is the probability that a customer in state s_t transitions to state s_{t+1} given the current state s_t and the firm's action a_t . The transition probabilities are defined in terms of the purchase probability and expenditure models estimated in the previous section:

$$V(s_t, a(s_t)) = E[\pi(s_t, a(s_t))] + \sum_{\kappa=t+1}^T \sum_{s' \in S} \alpha^{\kappa-t} * \pi(s_\kappa, a(s_\kappa)) * \Pr(s_\kappa | s_{\kappa-1}, a(s_{\kappa-1})). \quad (10)$$

The first term on the right side in this expression is the expected one-period profit from applying the selected control to a customer in a given state in the current time period. The expression for the single-period profit $\pi(s_\kappa, a(s_\kappa))$ reflects the expected profit obtained from a customer in a given week. This expression accounts for expected revenue and product costs, the impact of any promotions, and the cost

² It should be noted that our analysis focuses only on the development of customized marketing policies for established customers. The development of policies for new customers or customers with limited transaction histories presents some additional challenges. In the absence of an extensive transaction history, firms may wish to make inferences based on demographic measures and use some type of adaptive or learning-based marketing policy (see Sutton and Barto 1998).

of promotions. For example, if a free shipping promotion is offered, then the firm foregoes any shipping revenues and absorbs the costs to ship the order. The second 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 potential future states (s'_k) of the customer. The outer summation is over future time periods.

The state space for the optimization is mainly composed of the transaction history variables that are used in the customer demand model. Specifically, the state space $s(t)$ includes the customer's time since last purchase (recency), the amount of the last purchase, the number of past purchases, the number of weeks in the system, and cumulative spending.³ In addition, the state space also includes a history of the firm's actions. In this case, we include a state variable that tracks past promotions. This state variable provides a means for constraining the number of promotions that can be offered to a given customer. The choice model defines the transition structure for state space. For instance, the probability that a given customer makes a purchase defines the probabilities for how recency evolves. For example, if $\Pr(\text{buy} | s(t))$ is the probability of purchase at time t with state space $s(t)$, then the probability that the recency level in time $t + 1$ is one is $\Pr(\text{buy})$, whereas the probability that the recency level is one greater than at time t is $1 - \Pr(\text{buy})$.

There are a number of additional decisions that need to be specified to execute the optimization. First, the managerial time horizon of the objective function needs to be specified. For our purposes, we select an eight-week cycle.⁴ This finite horizon allows for the dynamic programming problem to be solved via backward recursion. Second, any constraints that limit the extent to which promotional policies can vary across individuals must also be specified within the dynamic

programming. There are two key reasons for imposing constraints. First, because the model is estimated on historical data, the firm may wish to avoid having the optimization procedure suggest policies that are far outside the range included in the historical record. Second, the use of one-to-one marketing policies has the potential to be controversial. For example, Coca-Cola experienced negative publicity for testing vending machines that varied prices based on the weather (Egan 2001) and Amazon.com came under fire in 2000 when consumers learned they were paying different prices for the same DVDs (Hamilton 2001). For our application, the maximum number of promotions that can be offered is constrained to two. Alternatively, we could require that all customers receive the same number of promotions over the course of the optimization period.

The dynamic optimization is a nontrivial exercise. For example, a state space with 20 levels of cumulative pending, 20 levels of recency, 10 levels of previous order size, and 3 levels of past promotional activity yields 12,000 distinct states. Given an eight-week time horizon, the marketing policy includes 96,000 mappings of promotional actions to different customer states and times. In addition to the state space being relatively large, the more significant computational challenge is that the development of one-to-one marketing policies requires that the optimization be solved for each individual customer. A final complicating factor in implementing the model is that several of the transaction history variables are continuous in nature. To handle this issue, the dynamic programming model is solved for a subset of the points in the state space. The value functions for other points in the state space are computed via interpolation (Keane and Wolpin 1994). We also assume that average prices and banner advertising activity are constant throughout the decision horizon.

4.2. Customization Scenarios

To measure the benefit from customization based on different sources of heterogeneity, we conduct a series of policy simulations. We consider the following four scenarios:⁵

1. Baseline—no customization.
2. Customization based on transaction history; accounts for temporal heterogeneity only.

⁵ We do not include a scenario accounting for only cross-sectional heterogeneity. Not allowing for temporal optimization over the eight-week period meant that in each week the optimal policy was to recommend promotion to a large number of customers. Given the requirement that each customer receive a maximum of two promotions over the eight-week period (to keep the recommended policies in line with current retailer practice), dynamic optimization is required to determine the optimal weeks in which to promote. The case of cross-sectional heterogeneity without dynamic optimization is therefore not feasible.

³ Two elements of the state space require discretization for the model to be implemented. Cumulative spending and the amount of last purchase are by nature continuous variables. The amount of previous purchase is discretized by rounding to the nearest \$25 point. Likewise, cumulative spending is discretized to the nearest \$25.

⁴ The eight-week horizon is intended to provide a sufficient time horizon for the dynamic structure to be meaningful while maintaining computational tractability. The firm typically offers a promotion once every month. The eight-week horizon covers two such monthly promotion cycles and represents a reasonable period over which the firm can plan and manage its promotions. The eight-week horizon is a balance between computational tractability and managerial relevance. It should also be noted that the use of finite time horizon does potentially lead to a concentration of promotions at certain points in time (endgame problems). An alternative approach would be to use a rolling time horizon in the optimization. Lewis (2004) describes how this type of rolling time horizon could be used in the context of a loyalty program.

3. Customization based on transaction history and individual-level preferences; accounts for temporal and cross-sectional heterogeneity at the individual level.

4. Customization based on transaction history and segment-level preferences; accounts for temporal and cross-sectional heterogeneity at the segment-level.

We use a two-stage procedure to evaluate the outcome under each scenario. First, we solve the dynamic programming model at the population, segment, or individual level depending on the scenario. The outputs from the dynamic programming model are policies that map marketing actions to customer states and the associated values of these policies. Customer states are based on recency, previous amount, cumulative buying, and previous number of promotions. The procedure thereby prescribes which promotional instruments, if any, should be used for different customer states. For instance, in the individual-level customization, a mapping of actions to all possible states is generated for each individual. In the second stage, we use the optimal policies from the first step and the individual-level response parameters to run an eight-week simulation for each customer. This allows us to evaluate the expected outcome from following the prescribed optimal policies.

For each scenario, Table 7 shows the expected profits and revenues, the number of purchases, purchase incidence rate, and the average expenditure amount. For comparison, we also show the increase in profits relative to the baseline for each scenario. To understand how customization drives the gains in profitability across scenarios, we decompose the incremental profits into three sources: because of increased purchase incidence, because of increased expenditure, and because of reduced redemption costs. Customization reduces redemption cost primarily in two ways. First, it reduces the likelihood of targeting customers

who will redeem a promotion without changing purchase behavior and contributing to increased profitability. Second, the optimization procedure assigns the most cost-effective policy by considering the revenue gains from a promotion versus the redemption cost.

4.2.1. Baseline—No Customization. Traditional promotion practices typically involve offering the same promotion to all customers in the database simultaneously without reference to individual customer preferences or transaction history. This scenario mirrors the firm's current uniform promotion policy and establishes a baseline relative to which the benefits from customization can be measured. For this scenario, we evaluate all possible combinations of promotion type and timing. We use the individual-level coefficients to run an eight-week simulation under all nine possible combinations of the free shipping, discount coupon, and loyalty promotion and 28 possible combinations of timing. Expected profits were highest with a free shipping offer in week one and a discount coupon in week six. We selected this sequence as it represents the optimal outcome under the retailer's current promotion strategy. The expected eight-week revenue from the sample is \$100,021, generating an expected profit of \$12,987. The purchase incidence rate is 15.6% and average expenditure is \$65.89.

4.2.2. Customization Based on Transaction History. Even in the absence of individual-level response parameters, firm databases provide an opportunity to customize using transaction history data. This approach accounts for temporal heterogeneity because the optimal policy depends on the customer's state, which is determined by purchase cycle factors. It does not account for cross-sectional heterogeneity because all customers in the same state receive the same policy.

Table 7 Dynamic Optimization Results

	(1) Baseline (none)	(2) Transaction history	(3) Transaction history and individual-level preferences	(4) Transaction history and segment-level preferences
Customization based on				
Profit (\$)	12,987	13,994	14,700	14,406
Revenue (\$)	100,021	107,652	112,170	110,039
No. of purchases	1,518	1,590	1,685	1,631
Purchase incidence (%)	15.6	16.3	17.3	16.7
Average amount (\$)	65.8	67.68	66.57	67.47
Increase relative to baseline				
Profit (%)		7.8	13.2	10.9
Percentage of incremental profitability due to				
Incidence		61	83	68
Expenditure		37	9	24
Cost reduction		2	8	8

It is consistent with standard direct marketing techniques, which use recency, frequency, and monetary measures to classify customers and assign promotion policy.

In the first stage, we solve the dynamic programming model using the population-level parameters. This gives us a single set of policies that map customer states to recommended marketing actions. Note that although the dynamic optimization is implemented at the population level, the recommended policies for each week vary across individuals depending on their state. In the second stage, we use the individual-level response parameters to run eight-week simulations for each individual.

With customization based on transaction history only, profitability increases to \$13,994. The 7.8% increase relative to the baseline represents the marginal effect of accounting for temporal heterogeneity. Purchase incidence increases to 16.3% and the average purchase amount increases from \$65.89 to \$67.68. The decomposition of the increase in profitability shows that it is driven primarily by purchase incidence, 61%. Increased expenditure accounts for 37% and cost reduction accounts for only 2% of the profitability increase.

4.2.3. Customization Based on Transaction History and Individual-Level Preferences. The availability of individual-level parameters provides an opportunity to account for both cross-sectional and temporal heterogeneity in response. For this scenario, we solve the dynamic programming model for each customer using the individual-level parameters. This yields individual-specific mappings between customer states and optimal promotion policies. The second stage is the same as in the previous scenario in which we conduct the eight-week simulation. Now the recommended policies are both individual and state specific.

With individual-level customization, profitability increases to \$14,700, which is a 13.2% increase relative to the baseline. Relative to the previous scenario in which only temporal heterogeneity was accounted for, the marginal effect of accounting for cross-sectional heterogeneity is 5.4%. It is interesting that this effect is significantly less than the 7.8% marginal effect of accounting for temporal heterogeneity. It indicates that more than half of the benefits from individual-level customization can be captured by accounting for temporal purchase cycle factors.

Purchase incidence increases to 17.3% and the average expenditure is \$66.57, which is lower but not significantly different than in the previous scenario. This policy increases profitability by inducing purchases that would not have been made otherwise rather than increasing purchase amount. This observation is reflected in a shift of the sources of

the profitability increase. Relative to the previous scenario, the amount accounted for by purchase incidence increases to 83% from 61%, and the amount accounted for by expenditure decreases to 9% from 37%. Overall, the effect of customization is to increase profitability. An observation across scenarios is that finer levels of customization increase profitability primarily by driving purchase incidence rather than increasing expenditure amount. The contribution from purchase incidence becomes increasingly important as the level of customization increases, accounting for over 80% of the incremental profits under individual-level customization.

Given the competitive landscape in most retailing sectors, the gains from customized promotions relative to the baseline are quite significant. However, implementation of this approach is complicated. Estimation of individual-level parameters for a large customer database poses an implementation concern. Although it is possible to recover these parameters using customer demographics and the population-level parameters estimated from a subsample of the database, these estimates will not be as accurate because customer characteristics account for only a small portion of the variation in individual-level parameters. More importantly, the dynamic optimization technique used to develop the promotion policy is fairly complex and computationally burdensome when implemented at the individual level. In environments with even moderately complex state spaces or for firms with millions of customers, the development of true individual-level policies is unlikely to be feasible. For these reasons, we next consider the returns to the retailer from customizing promotional policies at a segment rather than individual level.

4.2.4. Customization Based on Transaction History and Segment-Level Preferences. For this scenario, we use a k -means clustering algorithm to segment the customers into four groups. Note that this clustering is based on individual-level response parameters as well as transaction measures and thus incorporates both temporal and cross-sectional dimensions of customer heterogeneity. The main computational advantage of this scenario is that the dynamic programming model is implemented at a segment-level rather than individual level. The second stage is the same as before except now the recommended policies depend on the customer segment and customer state.

With segment-level customization, the profitability increase relative to the baseline is 11%. The marginal impact on profitability of moving from individual- to segment-level customization is a decrease of 2.3%. Although not insignificant, this loss is relatively small in magnitude. Although individual-level customization yields the highest overall gain in profitability,

the value of the 2% improvement over segment-level customization must be traded off against the effort needed to develop individual-level policies. It should be noted that the segment-level optimization is still based on individual-level information. Customers are clustered based on similarity in their individual-level parameters and transaction history measures. For this reason, the profitability loss from segment-level customization is not very large.

4.3. Type and Sequence of Promotions


The preceding results reveal the potential benefits of exploiting individual-level heterogeneity and transaction history data to set policy. In addition to estimates of profitability, our research also yields marketing policy recommendations that provide a contribution to the limited literature focused on multiple promotions. In this section, we highlight how customers differ in terms of the optimal number, type, and sequence of promotions offered. The policies discussed in this section are based on the individual-level optimization results, i.e., scenario 3 in §4.2.

The complexity of the state space makes the reporting of the complete policies infeasible. The outputted policies do not provide a predetermined multiple-week promotion policy. The policies are developed using an underlying model that captures the probabilistic nature of customer behavior. Therefore, the outputted policies are mappings of (marketing) actions to (customer) states that would be used by continuously updating the customer's current state each week and then retrieving the associated action. Our reporting strategy is to therefore illustrate the policies by reporting summary results from the individual-level policies. We also relate these policies to recency, frequency, monetary (RFM) measures to illustrate how the recommended policies tend to vary across customers with different transaction history profiles.

First we comment on the number of promotions offered to customers. Overall, 57% of customers receive no promotion, 27% receive one promotion, and 16% receive two promotions. Given that standard practice is to offer promotions universally, the percentage of people who receive no promotion is fairly high. However, when one considers the vast differences across the customer base, the number is not surprising.

A key difference between customers who receive one versus two promotions is in the timing of promotions. The differences are illustrated in Figure 2, which plots the probability of receiving a promotion conditional on time since previous purchase for one-promotion and two-promotion customers. For one-promotion customers, the likelihood of receiving a promotion is higher after four weeks have elapsed since the previous purchase. For customers receiving

Figure 2 Probability of Promotion, Conditional on Time Since Previous Purchase and Number of Promotions




two promotions, promotions are prescribed at a much higher rate between two and four weeks after a purchase. Customers receiving only one promotion are responsive to promotions but only after a longer lag since previous purchase. This illustrates that for some customers, it is more profitable to offer a promotion later rather than earlier in the purchase cycle.

Next, we comment on the sequence and type of promotions offered. For customers who receive one promotion, the most frequent promotion is free shipping, offered to 68% of customers. The coupon is offered to 27% customers and the loyalty reward to 6% of the customers. Table 8 shows the sequence of promotions for the 16% of customers who receive two promotions. The most prevalent promotion is the sequence of two free shipping offers, which is offered to 46% of this group. The sequence of two discount coupons is offered to 29% of customers. The interesting finding here is that for the majority of customers, it is optimal to offer the same promotion twice rather than offer a mix of promotions. The coupon and free shipping combination is offered to 19% of customers. The sequence of free shipping followed by the coupon offers a strategy for reacquiring lapsed customers. Free shipping is targeted to customers with a longer time since previous purchase, followed by the discount coupon as this group also has a lower average basket size. Overall, the loyalty reward was offered infrequently. This may be because the reward does not reflect an immediate monetary savings. Interestingly, none of the customers who received this as their

Table 8 Type and Sequence of Recommended Promotions

First promotion	Second promotion (%)		
	Freeship	Coupon	Reward
Freeship	46	15	4
Coupon	4	29	2
Reward	0	0	0

Figure 3 Probability of Promotion Type Conditional on Previous Expenditure

first promotion was selected for a second promotion, reflecting the fact that this segment is very responsive to the reward incentive but not to the discount or free shipping offers.

The primary difference between customers who receive the free shipping versus the coupon offer is in the typical basket size. Figure 3 illustrates the relationship between promotion type and previous order size. It shows the probability of each promotion type being offered for three ranges of order size. For relatively small previous orders, the discount coupon is the most frequently advocated promotion. However, as the order size increases, there is a shift toward offering the free shipping promotion. This is driven by the cost to the firm of each type of promotion relative to the marginal impact of the promotion on purchase behavior. For the discount coupon, consumers can scale the value of the promotion by purchasing a larger quantity. The marginal effect of the coupon is higher for small basket shoppers. Although large basket shoppers are also very responsive to the coupon in terms of purchase incidence, they are not selected for the discount coupon because the marginal impact on their purchase expenditure is significantly lower and the cost of promotion is higher to the firm. Large basket consumers generally use the discount coupon to subsidize their current level of purchase. The optimal promotion therefore tends to be a free shipping offer for large basket customers. The advocacy of the free shipping promotion to larger basket buyers may be particularly effective as it removes a transaction fee that increases as basket size grows. From the firm's perspective, free shipping for large-value customers is preferred over the discount coupon as the cost of the promotion has an upper limit set by the shipper's fixed cost schedule to the retailer.

4.4. Validation of the Dynamic Optimization Policies

To provide validation of the promotion policies recommended by the dynamic optimization, we fol-

Table 9 Validation of Customized Promotion Policies

Policy recommendation	Actual policy: Discount coupon		Actual policy: No promotion	
	(a)	(b)	(c)	(d)
	Percentage of customers	Profit per customer (\$)	Percentage of customers	Profit per customer (\$)
(1) No promotion	85	1.68	86	3.14
(2) Discount coupon	5	3.90	5	2.29
(3) Free shipping	10	2.13	9	2.57

low an approach similar to that used by Zhang and Krishnamurthi (2004). During the period of the hold-out data, the firm offered one promotion—a discount coupon valid for two weeks—to all customers in the database. We use the results from the individual-level dynamic optimization to determine the optimal recommended promotion policy. In Table 9, column a, we divide customers into groups based on the policy recommendation: (1) no promotion to 85%, (2) discount coupon to 5%, and (3) free shipping to 10% of customers. Note that the actual policy during these two weeks was a discount coupon for all customers, so the recommended policy matched the actual policy only for group (2).

In Table 9, column b, we report the actual per customer profit from each group. Customers for whom the discount coupon was the recommended policy yielded the highest per-customer profit of \$3.90. In contrast, per-customer profits are lowest (\$1.68) when the recommended policy is “no promotion” but the customers are offered the discount coupon. Offering a discount coupon when free shipping was the optimal policy also yields reduced average profit (\$2.25). Validation of the policy recommendations comes from noting that per-customer profits are highest for the group whose policy recommendation matches actual practice.

We repeat the validation exercise during a period of no promotional activity. In Table 9, column c, we report the percentage of customers who were recommended each policy: (1) no promotion to 86%, (2) discount coupon to 5%, and (3) free shipping to 9% of customers. In column d, we report the per-customer profit of each group. The per-customer profit for customers where the optimal policy matched actual practice (“No promotion”) was highest at \$3.14. Where the optimal policy was to offer a promotion, average profitability was significantly lower (\$2.29 for discount coupon-recommended customers and \$2.57 for free shipping-recommended customers). The result that per-customer profits are higher when the actual firm policy matches the recommended policy offers validation for the proposed customization approach.

Beyond validation of dynamic programming-based policies, this exercise also provides some further

insights. We find that offering a promotion when none is needed significantly reduces per-customer profit. In Table 9, row 1, the policy recommendation is no promotion. Per-customer profits are \$3.14 when the firm does not offer a promotion, but profits fall to \$1.68 when these customers are offered the discount coupon. In row 2, the policy recommendation is discount coupon. When this policy is implemented, per-customer profits are \$3.90, but profits fall to \$2.29 when customers who should receive a discount coupon receive no promotion. This shows that not offering the discount coupon when it is needed reduces customer profits. On the other hand, when the recommended policy in row 3 is free shipping, offering no promotion yields higher profits (\$2.57) than offering the incorrect promotion (\$2.13). This is important because it illustrates that offering no promotion is better than offering the wrong promotion.

5. Discussion

In this paper, we have examined consumer response to a variety of distinct promotional instruments. We find that the use of customized promotions yields increases in revenue and profitability, which are maximized when both individual preferences and purchase cycle information are used for guiding marketing tactics. The results and the accompanying techniques should be of significant interest to both traditional and online grocery industry practitioners. The grocery industry is currently struggling with a changing competitive dynamic that has led to store closings and reduced profitability. Much of this carnage is the result of Wal-Mart's entry into the grocery business. The current conventional wisdom is that traditional grocers need to exploit information from frequent shopper programs to counteract Wal-Mart's advantages in terms of costs (Singh et al. 2006). For example, Tesco, the leading British retailer, has used its customer-level data to thwart Wal-Mart by designing customized promotion offers (Rohwedder 2006). Our research demonstrates how a grocery retailer can leverage individual-level transaction data to potentially increase profits via customized marketing.

However, it is important to acknowledge the costs associated with implementing such a policy, including warehousing and maintenance of a customer database. The degree of expertise and the amount of data needed increases with the level of customization. For example, to implement customization based on transaction history requires that the firm only track purchase history. To implement customization using individual-level preferences also requires tracking exposure and response to various marketing activities. These data are more easily collected and finely tracked in an online rather than retail store environment. Once these systems have been established,

taking the recommended policies to customers poses further practical concerns. The shift from uniform to customized promotions adds decisions regarding the details of the promotions, followed by issues related to the design and physical delivery of the promotion offer. Here, online stores also have an advantage, particularly in the delivery of the promotion and the ease with which the promotion can be linked back to the retailer's website.

It is also worth noting that although extensive customer databases are increasingly available, we understand that currently relatively few organizations have the internal skills necessary to formulate the type of statistical models we use. Therefore, the linkage of RFM categories to promotional instruments may provide useful guidance to marketing managers. By mapping the various promotional instruments to different RFM profiles, we are able to comment on which type of promotion is generally appropriate for different categories of customers. This analysis is also of interest because the different promotional types provide different types of incentives. There are several customer management issues that may be better understood by studying which promotions are advocated for customers with different transaction histories. For instance, firms may wish to know whether promotions should attempt to change or reinforce previous patterns of purchasing or what type of promotion is more useful for reacquiring lapsed customers (Thomas et al. 2004). For example, our results suggest that the removal of shipping charges that act as a required transaction fee is a more effective reacquisition promotion than a coupon that provides a 10% discount on items purchased. This type of finding suggests a need for future research into why the removal of the transaction fee is more effective in reacquiring customers than the straight discount.

A strength of the research is its consideration of distinct types of promotional instruments, but the data impose several limitations. We are not able to study variations in the value of each promotion. Future researchers may wish to examine the benefits of customizing not only types of promotions but also the value of promotional instruments (Zhang and Krishnamurthi 2004, Lewis 2005b). As the loyalty program was ongoing, the firm did not actually offer the loyalty award as a limited duration promotion, which is likely to evoke a different response from what we were able to capture here. Further research is needed to more accurately assess the response to reward promotions. Another limitation of the data is that actual exposure to promotions is not observed, which may result in an underestimate of promotion effectiveness. A possible approach to correcting this is to include a model of promotion awareness conditional on promotion availability (Erdem et al. 1999).

Other opportunities for future research include the study of different categories. Grocery shopping typically involves short purchase cycles and involves baskets comprised of many relatively low-priced items. Investigations into categories that involve higher-priced products or durable goods may yield interesting results. With lower purchase frequency and higher-priced products, the retailer might shift attention to promotions that result in upselling and cross-selling. Another possibility is to also consider the impact of customer loyalty and experience with the retailer.

Finally, an underappreciated consequence of individual-level marketing is that consumer ill will may result when customers receive different promotional offers. Advocates of one-to-one marketing systems seldom address the possibility that the price discrimination aspects of individual-level marketing systems can harm customer relationships. Although lab studies have shown that consumers find different forms of price discrimination to be unfair (Feinberg et al. 2002), there is little empirical research that documents how perceptions of unfairness influence consumer behavior. Another topic that has not been investigated is how consumers may react when offered promotions that differ in terms of structure. For example, although we expect consumers to object to not receiving a promotion when others do, it is not clear what the reaction would be if one customer received a 10% discount coupon while another was offered free shipping. Our finding that temporal heterogeneity accounts for more than half the gains of customization can be important in this regard, as the ill will generated would be less of an issue if a firm offers similar promotions to all its customers but only manipulates the timing of the offers.

References

- Allenby, G., P. Rossi. 1999. Marketing models of consumer heterogeneity. *J. Econometrics* **89**(1–2) 57–78.
- Bitran, G. R., S. V. Mondschein. 1996. Mailing decisions in the catalog sales industry. *Management Sci.* **42**(9) 1364–1381.
- Blattberg, R., S. Neslin. 1990. *Sales Promotion: Concepts, Methods, and Strategies*. Prentice Hall, Englewood Cliffs, NJ.
- Blattberg, R., B.-D. Kim, S. Neslin. 2008. *Database Marketing*. Springer, New York.
- Boatwright, P., S. Borle, J. Kadane. 2006. A model of the joint distribution of purchase quantity and timing. *J. Amer. Statist. Assoc.* **98**(463) 564–572.
- Bodapati, A., S. Gupta. 2004. A direct approach to predicting discretized response in target marketing. *J. Marketing Res.* **41**(3) 73.
- Bult, J. R., T. Wansbeek. 1995. Optimal selection for direct mail. *Marketing Sci.* **14**(4) 378–395.
- Chintagunta, P., S. Haldar. 1998. Investigating purchase timing behavior in two related product categories. *J. Marketing Res.* **35**(1) 43–53.
- Egan, C. 2001. Vending-machine technology matures, offering branded foods, convenience. *Wall Street Journal* (December 13) B13.
- Erdem, T., M. P. Keane, B. Sun. 1999. Missing price and coupon availability data in scanner panels: Correcting for the self selection bias in choice model parameters. *J. Econometrics* **89** 177–196.
- Feinberg, F., A. Krishna, Z. J. Zhang. 2002. Do we care what others get? A behaviorist approach to targeted promotions. *J. Marketing Res.* **39**(3) 277–291.
- Gönül, F., M. Z. Shi. 1998. Optimal mailing of catalogs: A new methodology using estimable structural dynamic programming models. *Management Sci.* **44**(9) 1249–1262.
- Gupta, S. 1991. Stochastic models of interpurchase timing with time-dependent covariates. *J. Marketing Res.* **28**(1) 1–15.
- Gupta, S., P. K. Chintagunta. 1994. On using demographic variables to determine segment membership in logit mixture models. *J. Marketing Res.* **31**(1) 128–137.
- Gupta, S., D. Lehmann, J. Stuart. 2004. Valuing customers. *J. Marketing Res.* **41**(1) 7–18.
- Hamilton, D. 2001. E-commerce special report: The price isn't right. *Wall Street Journal* (February 12) R8.
- Heilman, C., F. Kaefer, S. Ramenofsky. 2003. Determining the appropriate amount of data for classifying consumers for direct marketing purposes. *J. Interactive Marketing* **17**(3) 5–28.
- Helsen, K., D. C. Schmittlein. 1993. Analyzing duration times in marketing: Evidence for the effectiveness of hazard rate models. *Marketing Sci.* **12**(4) 395–414.
- Jain, D., N. Vilcassim. 1991. Investigating household purchase timing decisions: A conditional hazard function approach. *Marketing Sci.* **10**(1) 1–23.
- Kamakura, W., V. Mittal, F. Rosa, J. Mazzon. 2002. Assessing the service-profit chain. *Marketing Sci.* **21**(3) 294–317.
- Keane, M., K. Wolpin. 1994. The solution and estimation of discrete choice dynamic programming models by simulation and interpolation. *Rev. Econom. Statist.* **76**(4) 648–672.
- Kim, J. G., U. Menzeffricke, F. Feinberg. 2005. Modeling parametric evolution in a random utility framework. *J. Bus. Econom. Statist.* **23**(3) 282–294.
- Lewis, M. 2004. The influence of loyalty programs and short-term promotions on customer retention. *J. Marketing Res.* **41**(3) 281–292.
- Lewis, M. 2005a. Incorporating strategic consumer behavior into customer valuation. (Special Issue on Customer Relationship Management.) *J. Marketing* **69**(4) 230–238.
- Lewis, M. 2005b. Research note: A dynamic programming approach to customer relationship pricing. *Management Sci.* **51**(6) 986–994.
- Lewis, M., V. Singh, S. Fay. 2006. An empirical study of the impact of nonlinear shipping and handling fees on purchase incidence and expenditure decisions. *Marketing Sci.* **25**(1) 51–64.
- Manchanda, P., J. P. Dubé, K. Y. Goh, P. Chintagunta. 2006. The effects of banner advertising on consumer inter-purchase times and expenditures in digital environments. *J. Marketing Res.* **43**(1) 98–108.
- Peppers, D., M. Rogers. 1993. *The One to One Future*. Doubleday, New York.
- Rigby, D., F. Reichheld, P. Scheffer. 2002. Avoid the four perils of CRM. *Harvard Bus. Rev.* **80**(2) 101–109.
- Rohwedder, C. 2006. Stores of knowledge: No. 1 retailer in Britain uses "clubcard" to thwart Wal-Mart. *Wall Street Journal* (June 6) 1.
- Rossi, P. E., R. E. McCulloch, G. M. Allenby. 1996. The value of purchase history data in target marketing. *Marketing Sci.* **15**(4) 321–340.

- Rust, R., P. Verhoef. 2005. Optimizing the marketing interventions mix in intermediate-term CRM. *Marketing Sci.* **24**(3) 477–479.
- Rust, R., V. Zeithaml, K. Lemon. 2000. *Driving Customer Equity*. Free Press, New York.
- Sherman, L., J. Deighton. 2001. Banner advertising: Measuring effectiveness and optimizing placement. *J. Interactive Marketing* **15**(2) 60–64.
- Simester, D. I., P. Sun, J. N. Tsitsiklis. 2006. Dynamic catalog mailing policies. *Management Sci.* **52**(5) 683–696.
- Singh, V. P., K. T. Hansen, R. C. Blattberg. 2006. Market entry and consumer behavior: An investigation of a Wal-Mart supercenter. *Marketing Sci.* **25**(5) 457–476.
- Sutton, R., A. Barto. 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- Thomas, J., R. C. Blattberg, E. J. Fox. 2004. Recapturing lost customers. *J. Marketing Res.* **41**(1) 31–45.
- Wedel, M., W. A. Kamakura, W. S. Desarbo, F. Ter Hofstede. 1995. Implications for asymmetry, nonproportionality, and heterogeneity in brand switching from piecewise exponential mixture hazard models. *J. Marketing Res.* **32**(4) 457–462.
- Xie, J., S. M. Shugan. 2001. Electronic tickets, smart cards, and online prepayments: When and how to advance sell. *Marketing Sci.* **20**(3) 219–243.
- Zhang, J., L. Krishnamurthi. 2004. Customizing promotions in online stores. *Marketing Sci.* **23**(4) 561–578.