J. of the Acad. Mark. Sci. (2013) 41:320–337
DOI 10.1007/s11747-012-0312-7

ORIGINAL EMPIRICAL RESEARCH

Overcoming the "recency trap" in customer relationship management

Scott A. Neslin · Gail Ayala Taylor ·
Kimberly D. Grantham · Kimberly R. McNeil

Received: 15 December 2011 / Accepted: 3 July 2012 / Published online: 4 August 2012

Abstract Purchase likelihood typically declines as the length of time since the customer's previous purchase ("recency") increases. As a result, firms face a "recency trap," whereby recency increases for customers who do not purchase in a given period, making it even less likely they will purchase in the next period. Eventually the customer is effectively lost to the firm. We develop and illustrate a modeling approach to target a firm's marketing efforts, keeping in mind the customer's recency state. This requires an empirical model that predicts purchase likelihood as a function of recency and marketing, and a dynamic optimization that prescribes the most profitable way to target customers. In our application we find that customers' purchase likelihoods as well as response to marketing depend on recency. These results are used to show that the targeting of email and direct mail should depend on the customer's recency and that the optimal decision policy enables the average high recency customer, who currently is virtually worthless to the firm, to become profitable.

Keywords Customer relationship management · Customer lifetime value · Optimization · Migration model · Customer recency

The authors are grateful to Brett Gordon for his help and guidance in implementing the value iteration method used in this work. We are also grateful for comments made by Olivier Toubia and students in his marketing models seminar. We are indebted to Rong Guo for preparing the data used in the analysis. We thank an anonymous online food service company for providing the data used in this research. Finally, we thank the three anonymous reviewers for their valuable comments and suggestions.

S. A. Neslin (✉) · G. A. Taylor
Tuck School of Business, Dartmouth College,
100 Tuck Drive,
Hanover, NH 03755, USA
e-mail: scott.neslin@dartmouth.edu

G. A. Taylor
e-mail: Gail.A.Taylor@Dartmouth.edu

K. D. Grantham
Terry College of Business, University of Georgia,
310 Herty Drive,
Athens, GA 30602, USA
e-mail: kdg@UGA.edu

K. R. McNeil
School of Business and Economics, Quiester Craig Hall,
North Carolina A&T State University,
Greensboro, NC 27411, USA
e-mail: krmcneil@NCAT.edu

Introduction

Customer relationship management (CRM) gives executives the mindset that customers are "manageable strategic assets of the firm" (Reimann et al. 2010, p. 329). CRM provides firms with the opportunity to contact the right customer at the right time through the right marketing medium. The challenge, however, is in the *dynamics*. The marketing decisions we make today may affect what we wish to do tomorrow. For example, due to saturation, the customer who received a direct mail piece last week may be less receptive to a direct mail piece this week. This suggests that a firm must manage the timing of its direct mail efforts in order to obtain optimal results. Customers are always changing over time. Not taking into account these dynamics can result in mistargeting and mistiming of marketing actions.

An important dynamic is recency—how long it has been since the customer's previous purchase. Recency, along with its cousins, frequency and monetary value, is one of the pillars of CRM (Blattberg et al. 2008; Hughes 1996). Studies show recency is highly correlated with customer purchase and directly relates to customer lifetime value

(CLV) through the "migration model" of CLV (Berger and Nasr 1998; Blattberg et al. 2008; Pfeifer and Carraway 2000). A common finding is that higher recency[1] is associated with lower purchase likelihood (Bitran and Mondschein 1996; Bult and Wansbeek 1995; Fader et al. 2005; Rhee and McIntyre 2008). As a result, firms face a "recency trap": when customers do not purchase in a given period, this increases their recency, which makes it less likely they will purchase in the next period, which in turn increases their recency, making them even less likely to purchase in the period after that, etc. The result is that the customer drifts away from the firm and the lifetime value of the customer decreases.

The recency trap is highly relevant to managers attempting to allocate marketing resources effectively across customers. Confronted with the recency trap, should the firm turn up its marketing efforts for high recency customers, or give up and let them lapse into oblivion? If the firm turns up its marketing efforts for high recency customers, is it wasting money and even desensitizing the customer to future efforts? On top of this, the firm has multiple marketing instruments at its disposal. Which ones should it use and when?

The purpose of this paper is to devise and demonstrate a procedure that prescribes how marketing efforts should be targeted to which customers at which time in order to manage the recency trap. We estimate a purchase model that that depends on recency, marketing, their interaction, and other marketing dynamics including carryover and saturation. We then derive the optimal decision policy for two marketing instruments, email and direct mail promotions.

We aim to contribute to the burgeoning literature on what Blattberg et al. (2008) call "optimal contact models." The theme of these models is that CLV is something to be *managed*, not merely *measured*. They rely on two components: (1) a customer response model, quantifying customers' purchase likelihood as a function of marketing, and in our case, a detailed array of recency effects, and (2) a dynamic optimization that prescribes the actions the firm should take at a given time for a given customer in order to maximize CLV. Our paper demonstrates how an optimal contact model can be used to effectively manage the recency trap.

We apply our approach to a meal preparation service provider whose key marketing tools are email and direct mail promotions. Our customer response model shows that recency is related negatively to purchase probability, setting up the recency trap. We also find that marketing interacts with recency and is subject to carryover and saturation. These effects differ for email and direct mail. Direct mail has more carryover and interacts positively with recency. Email is subject to saturation effects although it has zero distribution cost. Our optimization balances these considerations while accounting for customer migration between recency states. Our application suggests four findings regarding the firm's average customer: (1) the firm should increase marketing efforts as recency increases, (2) the firm currently is underutilizing both email and direct mail, (3) the firm should allocate more budget to direct mail than email, and (4) the implementation of our procedure would increase CLV by a predicted \$175–\$200, with particularly important increases for the average customer with high recency.

We proceed to review the literature in more detail. Then we discuss a detailed illustration of the recency trap, followed by a description of our response model and optimization. Next we describe the data for our application, and finally, the application itself. We close with a discussion of implications for researchers and practitioners.

Literature review

The concept of recency

Blattberg et al. (2008, p. 324) define recency as "the last time the customer purchased from the company . . . the elapsed time . . . since the last purchase." The origin of this concept dates to the 1960s. As Blattberg et al. discuss, recency can predict whether the customer will respond to a current marketing effort; response is a function of how recently the customer has purchased. Recency is commonly calculated with respect to purchasing from "the firm," and the concern with the recency trap is that the customer is drifting away from the firm. However, one can define recency at any level, for example item or brand as well as firm. The customer may drift away from General Motors' Chevrolet brand but drift into General Motors' Buick. In that case, the recency trap is bad news for Chevrolet's management but good news for Buick's management and arguably good news for General Motors as a whole. The model we develop in this paper could be applied to that situation—we would define recency with respect to each of the firm's brands. However, we take the perspective of the general manager responsible for overall firm profits. We therefore define recency at the level of the firm, in our case, a food preparation service company.

Customer recency and customer lifetime value

Several researchers have included recency in models that predict customer behavior. Bult and Wansbeek (1995), Bitran and Mondschein (1996), Fader et al. (2005), and Rhee and McIntyre (2008) find a negative association between recency and purchase likelihood. These findings

[1] Note that given the definition of recency as time since previous purchase, "higher recency" or "increased recency" means a longer time period has elapsed since the customer last purchased.

reinforce the common belief that "consistently, the most recent buyers out-perform all others" (Miglautsch 2002, p. 319) and that "many direct marketers believe that the negative relationship is a law" (Blattberg et al. 2008, p. 325).

Blattberg et al. (2008) note that the relationship between recency and purchase likelihood may differ by category. For example, Khan et al. (2009) find for an online grocery retailer that the relationship is positive at first, peaks at about 4 weeks, and then declines. This still begets a recency trap, because once customers have not purchased in 4 weeks, they tend not to purchase and their recency increases. Some researchers have found, within the range of their data, a positive relationship between recency and purchase, e.g., Gönül and Shi (1998) and Gönül et al. (2000) for a durable goods cataloger, and Van den Poel and Leunis (1998) for financial services. These findings may be due to a long purchase cycle; with a long enough data history, high recency would presumably result in lower purchase likelihood. For example, a customer may replace a television every 5 years, so purchase likelihood might increase as a function of recency up to 5 years. If those 5 years pass by and the customer has not purchased from the company, it is likely the customer has purchased from a competitor, and hence the probability of purchasing from the focal company declines with higher recency.

In summary, while there are exceptions, the common finding is that higher recency means lower purchase likelihood. There is empirical evidence and it is reasonable to believe that even if the relationship is not negative at first, it becomes negative in the long run. This begets the recency trap. Our procedure does not require the negative relationship; it is applicable for any relationship between recency and purchase. We focus on the negative relationship and the resultant recency trap because the negative relationship appears to be most common.

A key breakthrough in recency research was incorporating the relationship between recency and purchase likelihood into calculations of customer lifetime value (Berger and Nasr 1998; Pfeifer and Carraway 2000). This work views CLV as a Markov chain with recency serving as a "state variable," i.e., a variable that changes over time and influences the customer's likelihood of purchasing. Customers move from one recency state to another depending on whether they purchase or not. If the customer purchases, he or she is placed in recency state 1, meaning "just purchased." If the customer purchases in period 1 but not in period 2, the customer transitions to recency state 2, meaning that at the very end of period 2 (or the outset of period 3), the customer last purchased two periods ago, in period 1. Berger and Nasr show the details for calculating CLV using this framework, and Pfeifer and Carraway provide general formulas using matrix algebra.

Customer response to email and direct mail

Many predictive models find that email and direct mail affect purchase likelihood. The evidence regarding email is more recent and less definitive. An important paper by Drèze and Bonfrer (2008) finds that the *scheduling* of email solicitations affects customer retention as well as the customer's tendency to open and click on an email message. This is perhaps related to the traditional effects found with regard to advertising, namely *carryover* and *saturation*. Carryover means that a marketing activity in period t has an impact on purchase likelihood in period $t+1$. This may be due to the customer remembering the message for more than one period, or due to a delay between the reception of the message and the opportunity to act upon it. Pauwels and Neslin (2008) find evidence of carryover. Saturation means that marketing in period t reduces the impact of marketing in period $t+1$. For example, it could be that no new information is provided by the firm's marketing message in period $t+1$ because the customer just received a message in period t. Ansari et al. (2008) and Rhee and McIntyre (2008) find evidence of saturation.

A counter-argument for saturation would be a "foot-in-the-door" effect (Freedman and Fraser 1966) in which consumers who have complied with an initial request are more likely to comply with a subsequent request. In the context of communication, this would suggest that attending to an initial message would make the consumer more likely to attend to a subsequent message. This would produce an increased impact of future communication. Whether a saturation or foot-in-the-door effect is operative is an empirical question, which we will measure in our analysis. Note also that carryover and saturation pertain to the *customer response* to the communication—they are therefore a property of the customer; i.e., customer response is saturated, not the communication *per se*.

An extreme form of saturation, *supersaturation*, has been conjectured (e.g., Leeflang et al. 2000, p. 68), whereby high levels of marketing in previous periods result in current period marketing actually *decreasing* purchase likelihood. This could be due to customer irritation (Van Diepen et al. 2009) or information overload; after a surfeit of emails, additional emails merely encourage the customer to collect them in his/her inbox and ignore them all. As a result, the email–purchase relationship becomes negative. Van Diepen et al. (2009) look for supersaturation and don't find it. However, early field experiments by Ackoff and Emshoff (1975) find evidence of supersaturation, and more recent work by Naik and Piersma (2002) finds that cumulative marketing expenditures related negatively to customer goodwill. Ray and Sawyer (1971) find some evidence of supersaturation.

The concepts of carryover, saturation, and supersaturation can be applied at the level of a specific form of communication (e.g., email) or across all forms of communication (totaled across email and direct mail). We apply these concepts to the form of communication for two reasons. First, it has been shown that these effects can differ by form of communication. For example, Ray and Sawyer (1971) find different saturation effects for "grabber" versus "non-grabber" advertising copy. Second, our managerial interest is in specifying whether the firm should send an email or direct mail (or both or neither) to customers over time. It is therefore important to understand how customer response to both forms of communication might differ.

Optimal contact models

One of the most exciting areas of CRM research is optimal contact models (Blattberg et al. 2008). Optimal contact models determine what marketing efforts should be expended on which customers at what time. They integrate dynamic phenomena such as recency and carryover into a prescription for the optimal marketing policy. The first optimal contact models were applied to sales force management (Beswick 1977; Montgomery et al. 1971; Zoltners and Sinha 1980). In CRM, the pioneering work was by Bitran and Mondschein (1996), followed by important contributions from Gönül and Shi (1998), Gönül et al. (2000), Elsner et al. (2003, 2004), Rust and Verhoef (2005), Simester et al. (2006), Venkatesan et al. (2007) and Khan et al. (2009).

Many of these papers focus on catalog mailings. The catalog industry is a major innovator in CRM, so this is not surprising. Rust and Verhoef (2005), Venkatesan et al. (2007), and Khan et al. (2009) focus on multiple marketing activities. Rust and Verhoef consider direct mail and a customer relationship magazine; Venkatesan et al. consider sales calls and direct marketing; Khan et al. consider discount coupons, loyalty rewards, and free shipping. Consideration of multiple marketing instruments is important because it is realistic and makes the analysis more relevant.

As noted earlier, the basic components of an optimal contact model are (1) a customer response model, i.e., a predictive model of how customers respond to marketing, and (2) a method for optimization. Previous papers have used a variety of response models, including hazard models (Khan et al. 2009), RFM categorizations (Bitran and Mondschein 1996), and decision trees (Simester et al. 2006). The optimization usually employs dynamic programming. Dynamic programming is necessary because "forward looking" is crucial—the actions the firm takes with the customer today may influence what actions it may want to take in the future. Dynamic programming methods include infinite horizon (Simester et al.), rolling horizon (Neslin et al. 2009), and finite horizon approaches (Khan et al.). Khan et al. note each has advantages and disadvantages. Finite horizon optimization can run into end-game issues, distorting marketing efforts at the end of the time horizon (T) because there are no explicit costs or benefits in time T+1. On the other hand, infinite horizon methods can be computationally cumbersome. We use an infinite horizon dynamic program solved using value iteration, which is relatively simple to program.

Contribution of this paper

In summary, previous research suggests there is a recency trap. The contribution of this paper is to call attention to this phenomenon and demonstrate how to target the right marketing to the right customers at the right time to maximize customer lifetime value when the recency trap is operative. Our approach is to focus on this managerial issue and to develop an optimal contact model that is "complete on the important issues" (Little 1970) yet relatively simple and managerially relevant. Our paper is unique in its focus on the recency trap and how to manage it. This is manifested in our customer response model through the inclusion of a main effect of recency, interactions between recency and marketing, and carryover and saturation effects of marketing. We show how our dynamic optimization explicitly takes into account the customer's current recency state, as well as her or his future recency state depending on which marketing action is taken.

Our paper extends important research in this area. Two closely related papers are Rust and Verhoef (2005) and Khan et al. (2009), because they both deal with multiple marketing instruments. Compared to Rust and Verhoef, we emphasize the role of recency, consider interactions between email and recency as well as saturation and carryover effects, and perform an infinite horizon optimization. Compared to Khan et al., we consider saturation and carryover effects and perform an infinite horizon optimization. Overall, because of our focus on the recency trap, we model recency, carryover, and saturation effects more thoroughly.

Modeling framework

Our modeling framework consists of three elements: (1) quantification of the role of recency in determining customer purchase, (2) a logistic customer response model that focuses heavily on recency, and (3) a dynamic programming optimization that recognizes recency as an important characterization of the customer at any point in time. These three elements enable us to formulate a model that determines the targeting, timing, and total quantity of marketing efforts, as well as the relative allocation of funds spent on different marketing efforts (in this case, email and direct mail).

The role of recency and the recency trap

Figure 1 highlights the key phenomenon at work: recency is negatively associated with purchase likelihood. Figure 1 is calculated using actual raw data from our application and shows that the effect is particularly pronounced; customers who purchased in the previous period (recency state 1) have a 23.1% chance of repurchasing in the current period, whereas customers who last purchased 5 months ago (recency state 5) have only a 4.6% chance of purchasing.

The ramification of the recency/purchase relationship for CLV is shown vividly in Table 1. Table 1 uses the migration model of CLV to calculate the probability a customer acquired in period 1 will be in various recency states at all future points in time (see Blattberg et al. 2008). For example, by the end of period 7, there is a 9% chance that a customer acquired in period 1 will be in recency state 5, i.e., the customer's last purchase was five periods ago. The "Recency 1" column in Table 1 is most crucial because it shows the probability that the customer will purchase each period. The numbers in the top row of Table 1 govern these calculations and are taken from the probabilities shown in Fig. 1. They represent the conditional probability the customer will purchase in the current period, given his or her recency state *S* (*ProbPurch (S)*). In Table 1, the customer migrates to state 1 (just purchased) with probability *ProbPurchase(S)*. However, with probability 1 – *ProbPurchase(S)*, the customer migrates to a higher recency state, *S*+1, creating the recency trap.[2]

Table 1 shows how the recency trap plays out. The dominant tendency for the newly acquired customer is to make an initial purchase and then not purchase for several periods, sliding to recency state ≥20. Sometimes the customer in a high recency state makes a purchase. For example, even a customer who has not purchased in 13 periods has a 1.2% chance of purchasing in the current period. But clearly the company is losing hold of the customer.

The recency trap is somewhat related to the concept of customer defection or churn (Tokman et al. 2007; Blattberg et al. 2008, chapter 24). Churn is most applicable to a subscription service, where the customer formally ceases his or her relationship with the firm. This is a very relevant issue in industries such as publishing, telecom, and financial services. The problem then becomes one of predicting churn (Neslin et al. 2006) and winning back customers who have churned (Tokman et al.). The recency trap is more applicable to a non-subscription service, where the customer does not inform the company he or she has defected, but we infer the customer is drifting away by noting her or his probability of purchasing is decreasing. We can think of the high-recency customer as "inert," i.e., the probability of purchase is very low, but it does not decline exactly to zero. There is always a chance the high-recency customer is just on hiatus and will resume purchasing as personal circumstances not measured by the data change. However, the chance of this happening dwindles with the passage of time, unless perhaps the firm takes a proactive stance and tries to market to high recency customers.

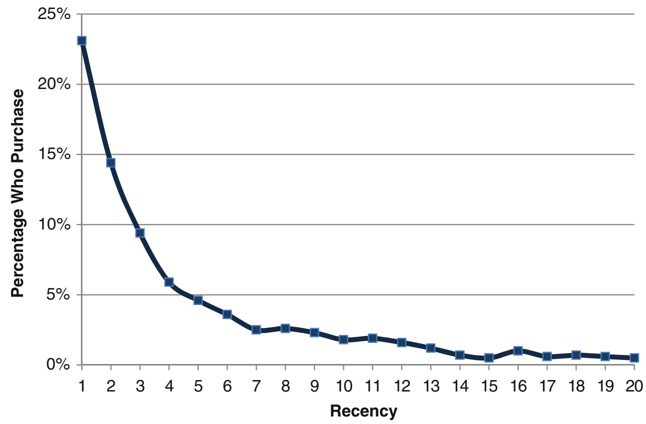

Fig. 1 Purchase likelihood vs. recency calculated directly from the data. Descriptive statistics based on 61,952 customer/month observations

Indeed, Fig. 1 and Table 1 capture the managerial problem in vivid terms: to devise a profitable targeted marketing strategy that will arrest the drifting away of a newly acquired customer. This strategy will depend on customer purchase probabilities, since they drive the recency trap. We will now estimate these probabilities *as a function of marketing*.

Logistic response model of purchase likelihood

The logistic customer response model that has been used in numerous applications (e.g., see Neslin et al. 2006). The dependent variable is $\mathrm{Purchase}_{it}$, a dummy variable equal to 1 if customer *i* purchases in period *t*; 0 if not.

We include the following explanatory variables for predicting this variable:

- Email_t and Dmail_t: the marketing efforts expended by the firm in period *t*, in our case, either email or direct mail offers.

[2] A customer's recency state is assigned at the end of the period. As in all CLV models, since we start the calculation from the time the customer makes her or his first purchase, we know the customer is in recency state 1 at the end of period 1. So the probabilities in Table 1 start off with a probability of 1 (the customer buys in period 1) by definition. The customer then has a 0.23 probability of purchase in period 2, because ProbPurchase(1) = 0.23. The subsequent purchase probabilities, and hence the states, are determined by the migration probabilities at the top of Table 1.

Table 1 The recency trap: the progression of the customer to higher and less profitable recency states*

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	≥ 20
Prob(Purchase | Recency) =	0.231	0.144	0.094	0.059	0.046	0.036	0.025	0.026	0.023	0.018	0.019	0.016	0.012	0.007	0.005	0.010	0.006	0.007	0.006	0.005
Recency State (Periods since last purchase):	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	≥ 20
Period: 1	1.00																			
2	0.23	0.77																		
3	0.16	0.18	0.66																	
4	0.12	0.12	0.16	0.60																
5	0.10	0.10	0.11	0.14	0.56															
6	0.08	0.07	0.08	0.10	0.13	0.54														
7	0.07	0.06	0.06	0.07	0.09	0.12	0.52													
8	0.06	0.05	0.05	0.06	0.07	0.09	0.12	0.51												
9	0.05	0.04	0.05	0.05	0.05	0.07	0.08	0.12	0.49											
10	0.05	0.04	0.04	0.04	0.05	0.04	0.06	0.08	0.11	0.48										
11	0.04	0.04	0.03	0.03	0.04	0.04	0.05	0.06	0.08	0.11	0.47									
12	0.04	0.03	0.03	0.03	0.03	0.04	0.04	0.05	0.06	0.08	0.11	0.46								
13	0.04	0.03	0.03	0.03	0.03	0.03	0.04	0.04	0.05	0.06	0.08	0.11	0.46							
14	0.03	0.03	0.03	0.02	0.03	0.03	0.03	0.03	0.04	0.05	0.06	0.07	0.11	0.45						
15	0.03	0.02	0.02	0.02	0.02	0.03	0.03	0.03	0.03	0.04	0.05	0.06	0.07	0.10	0.45					
16	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.03	0.04	0.04	0.06	0.07	0.10	0.45				
17	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.03	0.04	0.04	0.06	0.07	0.10	0.44			
18	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.04	0.04	0.06	0.07	0.10	0.44		
19	0.02	0.02	0.02	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.04	0.04	0.06	0.07	0.10	0.44	
20	0.02	0.02	0.01	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.04	0.04	0.06	0.07	0.10	0.43
21	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.04	0.04	0.05	0.07	0.53
22	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.04	0.04	0.05	0.60
23	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.04	0.04	0.65
24	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.04	0.69
25	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.72
….	….	….	….	….	….	….	….	….	….	….	….	….	….	….	….	….	….	….	….	….
50	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.89
51	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.89
52	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.89

*Recency state represents the number of periods since the previous purchase. The customer is acquired in period 1. Cell entries represent the probability the customer will be in each state in each time period. Recency column 1 represents the probability a customer will purchase in each period

- $Recency_{it}$: the recency state of customer i in period t, i.e., how long it has been since the customer's previous purchase.
- $Recency_{it}^2$ and $Recency_{it}^3$: the possibility that the relationship between recency and purchase is non-linear, beyond the inherent nonlinearity included in a logistic regression. We allow for a third-order effect because a second-order model may not be flexible enough to capture a monotonically declining relationship between recency and purchase probability. If $Recency_{it}^3$ is not needed, its coefficient will be insignificant. By including $Recency_{it}^3$, we gain added flexibility and the data will tell us if this is not necessary.
- $Email_t \times Recency_{it}$: the interaction between email and recency; a significant coefficient means that customers in different recency states respond differently to email offers.
- $Dmail_t \times Recency_{it}$: the interaction between direct mail offers and recency; a significant coefficient means that customers in different recency states respond differently to direct mail offers.
- $Email_t \times Recency_{it}^2$, $Email_t \times Recency_{it}^3$, $Dmail_t \times Recency_{it}^2$, and $Dmail_t \times Recency_{it}^3$: these interactions capture the possibility that the relationship between marketing efforts and customer response may vary non-linearly as a function of recency. For example, it may be that customers in the middle recency states (e.g., 5–10) respond more readily to marketing solicitations, compared to customers in low or high recency states.
- $Email_{t-1}$ and $Dmail_{t-1}$: carryover effects of marketing. That is, an offer received in period t-1 may have an impact on purchasing in period t.
- $Email_t \times Email_{t-1}$ and $Dmail_t \times Dmail_{t-1}$: potential saturation effects. For example, a negative coefficient for $Email_t \times Email_{t-1}$ means that large email efforts in the previous period render the email efforts in the current period less effective. It is possible of course that the coefficient could be positive, which would represent synergistic effects of prolonged campaigns.
- $Month_t$: the month pertaining to the particular customer observation (January, February, etc.). We use a dummy variable for each month, since month is the unit of observation.
- $First\_Amt_i$: to control for inherent cross-customer differences in preference for the firm. It equals the amount the customer spent on the first purchase when he or she was acquired. We expect the coefficient for this variable to be positive, because customers who start off by making a large purchase are probably very sure they like the product and are therefore likely to purchase on an ongoing basis (see Fader et al. 2007). Given Fader et al.'s results (pp. 61 (Fig. 1) and 64), this variable is associated with both frequency and monetary value (the F and M in RFM).

Collecting these variables (plus a column of ones for the intercept) into a k+1 vector X_{it}, where k is the number of explanatory variables described above, yields the following logistic regression model:

$$Prob(Purchase_{it} = 1) = \frac{1}{1 + exp^{-X_{it}\beta}} \quad (1)$$

where β is a k+1 vector representing the impact of each of the k explanatory variables on probability of purchase.

We use linear, squared, and cubed terms for recency to flexibly capture the potentially nonlinear relationship between recency and purchase. We expect a monotonically declining relationship since recency may indicate growing dissatisfaction with the firm or better offerings from competition, but there is no guarantee this will happen in an orderly fashion; hence the need for a highly nonlinear function. For example, Khan et al. (2009) use linear and log terms in their model. Since the Taylor series expansion for any continuous function is expressed in terms of polynomials, we use polynomials to be as flexible as possible. We believe this flexibility is important because the goal of the paper is to illustrate and address the recency trap, so it makes sense to model this important phenomenon in as much detail as possible.

Our focus is on the recency trap and therefore we need to model recency in depth. One could argue that variables such as frequency, monetary value, and duration might be added to the model. As mentioned above, First_Amt captures at least some of these effects. Including them explicitly would increase the complexity of the analysis significantly. This would be manifested in the response model but more importantly in the optimization. As we will describe shortly, we have 20 recency states, which combined with the other variables yields 216,000 total states. This is a large but manageable number. Including frequency, monetary value, and duration would create three more state variables. If we included just 10 states for each of these, we would have a total of $216{,}000 \times 10^3 = 216{,}000{,}000$ states, clearly too many for reasonable optimization. We would have to resort to interpolation methods (e.g., Keane and Wolpin 1994). Khan et al. (2009) do this and it suits their purposes. However, our purpose is to study recency in detail, so we chose not to use approximations.

Optimization

The context for our application is a meal preparation service. Once we have estimated Eq. 1, we know how the service's customers respond to marketing. We then need to derive a decision policy that will maximize the CLV of the service's customers. The lifetime value of customer i (CLV_i) can be expressed as:

$$CLV_i = \sum\nolimits_{t=0}^{\infty} \pi_{it}(D|S) \times \delta^t \tag{2}$$

where:

- $\pi_{it}(D|S)$ = Profit contributed by customer i in period t, given the customer is in state S in that period and marketing decision D is made with respect to that customer. The decision D in our case will be how much emailing and direct mailing effort the service will expend on that customer. The "state variables" that describe the customer are those that affect current profitability and change over time. In our case, recency, previous email/direct mail efforts, and month are the state variables. The notation "(D|S)" should be read as "the marketing decision the service will make for customers who are in state S."
- δ = discount factor, e.g., 0.995 on a monthly basis means that profits achieved 1 year from the present are worth 94% (0.995^{12}) of what they are worth today.

Equation 2 emphasizes that CLV is not a static number—it is an objective to be managed through marketing efforts. These efforts should depend on the recency state of the customer. The challenge is to find the decision policy (D|S) that maximizes CLV, taking into account that current actions may place the customer in a different state in the next period, which affects our optimal decision in that next period. For example, if the customer of the service has recently been saturated with emails regarding special discounts, etc., it may be optimal not to email the customer in the current period but instead to wait until the next period when the customer will pay attention to a new email. This must be balanced against the lower purchase probability that results from letting the customer lapse to a higher recency state.

We use dynamic programming to determine the optimal policy (D|S). This requires computation of the value function, $V_{it}(S)$. The value function has an intuitive interpretation in our context—it represents the lifetime value the service can expect from a customer i who is in state S at time t, given that it uses the *optimal* policy (D|S). The value function in period t can be expressed as the expected profit the service obtains from the decision that maximizes the current period profit of the customer, *plus* what the service henceforth expects to gain (on a discounted basis) from the customer, given the decision made in period t (Judd 1998):

$$V_{it}(S) = \max\nolimits_{D|S}\left\{E(\pi_{it}(D|S)) + \delta E\left(V_{it+1}\left(S'\right)\right)\right\} \tag{3}$$

Equation 3 presents the customer management viewpoint that we should maximize current period expected profits in light of the future profits we expect to result from actions we take in the current period. Note that in the current period the customer is in state S, but depending on the actions we take, the customer will be in a different state, S', in the next period.

In our case, the expected future profits, $E(V_{it+1}(S'))$, take on a simple form, because the customer either will purchase or not purchase. In particular:

$$\begin{aligned} E(V_{it+1}(S')) = \delta[&ProbPurch(S)_t \times E(V_{it+1}(S' = 1) \\ &+(1 - ProbPurch(S)_t) \times E(V_{it+1}(S' = S + 1)] \end{aligned} \tag{4}$$

$ProbPurch(S)_t$ is calculated using the logistic purchase model and will depend on the customer's current state as well as the marketing actions taken. The expected future value functions depend on which state,S', the customer transitions to in the next period. For example, consider the case of the customer who currently is in state S=5. If the customer purchases in period t, the customer will be in recency state 1 in period t+1, so S'=1. However, if the customer doesn't purchase in period t, the customer shifts to recency state S'= 6, because now it has been six periods since the customer purchased.

Equations 3 and 4 show that we take the action that maximizes current period profit *plus* expected future profits. Future profits depend on whether the customer buys (and moves to recency state 1), which happens with probability $ProbPunch(S)_t$, or does not buy (and migrates to recency state S+1), which happens with probability (1-$ProbPurch(S')_t$). The model manages the recency trap by explicitly considering what actions should be taken now, taking into account what happens if the customer does not purchase and slides to a higher recency state.[3]

We have not yet specified the profit function $E(\pi_{it}(D|S))$. This function involves application-specific costs, etc., and so we describe it fully in the Application section

Data

The data for our application are from a meal preparation service. Customers log on to the company's website and order the meals they will assemble during their visit to the service establishment or the pre-assembled meals they wish to pick up. Ordering is primarily done online, allowing for the easy collection of customer-level data. The data span 25 months, October 2006 through November 2008. We have data on 4,071 customers who made an initial purchase.

[3] Note that if as in Table 1, the highest recency state is ≥20, once ≥20 becomes that customer's current state and he or she does not buy, he or she "migrates" to state ≥20, since we are collapsing states 20, 21, 22, etc. into one state, ≥20. That is, S+1 is state ≥20 if the current state is ≥20.

These customers made a total of 4,442 additional purchases, an average of one additional purchase per customer (consistent with the data in Table 1). It is important to note that there is adequate variation in the number of additional purchases. Although consistent with the recency trap identified in Fig. 1, 62% of customers made no more purchases after the initial purchase, 16% made one additional purchase, 8% made two additional purchases, and the remaining 14% made three or more additional purchases, over the 25 periods of the data. We thus have ample variation to estimate the relationship between recency and purchase, while the recency trap ("one and done") shown in Table 1 is still an obvious characteristic of the data.

Customers are acquired at different times; on average we observe the customer for 15.2179 months. This means in total we have $4071 \times 15.2179 = 61{,}952$ customer-month observations available for the logistic customer response model.

The chief marketing instruments used by the firm were email and direct mail promotions. These promotions offered a discount on purchased merchandise during certain periods of time. For example, an email could alert the customer that a promotion was in effect during a specified three-week period. We therefore created monthly email and direct mail variables so that a month-long promotion would assume a value of one. A value of 0.75 associated with an email means that the email announced a promotion that was available for 3 weeks. This procedure yielded monthly email and direct mail variables representing how many months worth of promotion were announced by these communications. The average email variable was 0.67 (see Table 2), meaning the average email-communicated promotion was in effect for a little less than 3 weeks during the month it was announced. The values for the email variable ranged from 0 to 2.5, while the values of the direct mail variable ranged from 0 to 2.9, with an average of 1.52. Values greater than one are possible because there may have been overlapping emails or direct mail in a given month. Table 2 describes these and other variables used in the model.

Application

Logistic regression of purchase probability

We estimate Eq. 1 in stages, adding variables to demonstrate the impact of email and direct mail, and particularly the role of recency. Table 3 shows the results.

The base model includes just recency, monthly dummies and the First_Amt control variable. The recency variables—linear, squared, and cubed terms—are all highly significant as expected given Fig. 1. The First_Amt variable is highly significant, and six of the 11 monthly dummies are significant at the 5% level.

Model 2 adds the basic marketing variables: (1) the current period effect (Email and Dmail), (2) the lagged effect (Lagged_Email and Lagged_Dmail), and (3) the interactions between the main and lagged variables, measuring saturation. The addition of these variables "costs" 6° of freedom, but the likelihood ratio test shown in the bottom three lines of Table 3 finds that the contribution to fit is statistically significant. The results suggest that carryover and saturation effects are present. Carryover is particularly strong for direct mail; saturation is present for both email and direct mail, significant at the 0.039 level for email, albeit only marginally significant for direct mail. Overall the key finding is that the classic marketing effects—current period, carryover, and saturation—are apparent in the data and add to overall fit.

Model 3 adds interactions between marketing and recency. The likelihood ratio test is significant at the 0.014 level, indicating that these interactions add to fit. The interaction is

Table 2 Key variable definitions and descriptive statistics*

Variable	Description	Mean	Std. Dev.	Min.	Max
$Recency_{ht}$	Periods since last purchase by household *h*, with "1" signifying the purchase was made in month *t*.	7.659	5.621	1	20
$First\_Amt_h$	Amount spent on first purchase by household *h*.	116.264	58.068	−110**	480
$Email_t$	Level of Email marketing activity in month *t*, scaled so that one email campaign lasting one month would be scored as 1. The variable therefore can be interpreted as number of months worth of email campaigning in month *t*.	0.672	0.689	0	2.548
Dmail	Level of direct marketing activity in month *t*, scaled so that one direct mail campaign lasting one month would be scored as 1. The variable therefore can be interpreted as number of months worth of direct mail campaigning in month *t*.	1.524	0.805	0	2.933
$Purchase_{ht}$	= 1 if household *h* purchased in month *t*; 0 otherwise.	0.069	0.253	0	1

*Based on $n = 61{,}952$ household-week observations

**There was one customer outlier with a negative value for First_Amt. The rest of the values were above zero. We decided to leave this customer in the data, although this had virtually no influence on the results

Table 3 Logistic regression results

Variable	Base Model		Model 2		Model 3		Model 4		Model 5	
	Coef	P-val	Coef	P-val	Coef	P-val	Coef	P-val	Coef	P-val
Intercept	−1.105	<.001	−1.774	<.001	−1.720	<.001	−1.626	<.001	−1.623	<.001
Recency	−0.714	<.001	−0.715	<.001	−0.741	<.001	−0.790	<.001	−0.791	<.001
$Recency^2$	0.050	<.001	0.049	<.001	0.049	<.001	0.052	<.001	0.053	<.001
$Recency^3$	−0.00134	<.001	−0.00130	<.001	−0.00131	<.001	−0.00123	<.001	−0.00131	0.002
Email			0.375	0.007	0.357	0.012	0.364	0.014	0.453	0.0004
Lagged_Email			0.095	0.441	0.092	0.455	0.086	0.487	0.085	0.491
Email × Lagged_Email			−0.247	0.039	−0.242	0.043	−0.221	0.067	−0.219	0.068
Email × Recency					−0.003	0.717	−0.026	0.270	−0.107	0.045
Email × $Recency^2$							0.00198	0.206	0.016	0.057
Email × $Recency^3$									−0.00056	0.091
Dmail			0.140	0.191	0.105	0.335	0.034	0.731	−0.0046	0.970
Lagged_Dmail			0.375	0.001	0.394	0.0003	0.391	0.0003	0.384	0.0004
Dmail × Lagged_Dmail			0.090	0.116	−0.105	0.071	−0.104	0.072	−0.103	0.076
Dmail × Recency					0.020	0.011	0.065	0.001	0.102	0.020
Dmail × $Recency^2$							−0.004	0.016	−0.103	0.134
Dmail × $Recency^3$									0.0003	0.291
First_Amt	0.00349	<.001	0.00343	<.001	0.00344	<.001	0.00344	<.001	0.00344	<.001
Jun	0.103	0.208	0.067	0.461	0.077	0.385	0.085	0.353	0.087	0.338
Jul	0.006	0.947	−0.195	0.125	−0.163	0.201	−0.146	0.252	−0.142	0.267
Aug	0.239	0.003	0.223	0.205	0.242	0.176	0.230	0.197	0.229	0.201
Sep	0.548	<.001	0.495	0.002	0.504	0.002	0.489	0.003	0.488	0.003
Oct	0.183	0.021	−0.096	0.469	−0.074	0.579	−0.079	0.554	−0.072	0.590
Nov	−0.101	0.222	−0.167	0.092	−0.162	0.103	−0.169	0.088	−0.166	0.094
Dec	−0.120	0.188	−0.210	0.055	−0.191	0.080	−0.179	0.102	−0.184	0.094
Jan	−0.113	0.211	0.006	0.955	0.002	0.985	−0.007	0.950	−0.011	0.923
Feb	0.333	0.007	0.536	<.001	0.534	<.001	0.524	<.001	0.520	<.001
Mar	0.243	0.002	0.221	0.025	0.244	0.014	0.256	0.010	0.254	0.011
Apr	−0.093	0.263	0.140	0.230	0.130	0269	0.124	0.291	0.122	0.301
N	61,952		61,952		61,952		61,952		61,952	
-2log_likelihood	25925.793		25887.240		25878.750		25872.857		25869.933	
Incremental Log_LL			38.553		8.490		5.893		2.924	
Incremental P-value			<0.001		0.014		0.053		0.232	
AIC	25957.793		25931.240		25926.750		25924.857		25925.933	
BIC	26102.339		26129.990		26143.568		26159.744		26178.888	

not significant for email but is significant (p-value = 0.011) for direct mail. The positive sign means that as customers transition to higher recency states, they become more receptive to direct mail.

Model 4 adds interactions between marketing and recency-squared. The likelihood ratio test here is significant at the 0.053 level. Model 5 adds interactions between marketing and recency-cubed. This model clearly does not improve fit—the likelihood ratio test has a p-value of 0.232.

In choosing which model to use, we decided to reach a middle ground between fit and simplicity. Model 5 has the highest log likelihood, but that is a mathematical truism since it has the most parameters. Table 3 also includes AIC and BIC statistics which penalize a model for added parameters—BIC is known to penalize more than AIC. Model 4 has the best AIC, while Model 1 has the best BIC. Model 1 omits too many important phenomena, both managerially and statistically, given the likelihood ratio test. The choice is between Models 3 and 4. Model 4 has a better AIC, but Model 3 has a better BIC, plus the likelihood ratio test clearly establishes the statistical significance of the variables it adds, while the significance of the variables Model 4 adds is marginal. We therefore decided to use Model 3 for our optimization. One could argue for Model 4, but we decided to be conservative and retain the simpler model.

Figures 2, 3 and 4, based on Model 3, provide graphical illustrations of the effects quantified by the logistic regression. Figure 2 graphs probability of purchase as a function of recency, using the coefficients for recency, recency^2, and recency^3 in Model 3. As expected, the shape of the graph is very similar to that shown in Fig. 1, calculated from actual data. This says that the results in Fig. 1 were not due to a confound with other variables.

Figure 3 illustrates the interaction between marketing and recency. From Table 3, the interaction between recency and

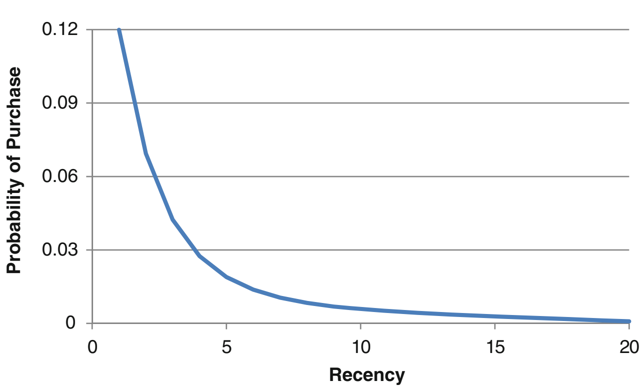

Fig. 2 Purchase likelihood as a function of recency calculated from the model. Calculation assumes no marketing effort, i.e., Email and Dmail = 0, and lagged Email and Dmail = 0; Month = 0. The shape of the curve is unaffected by changes in these assumptions

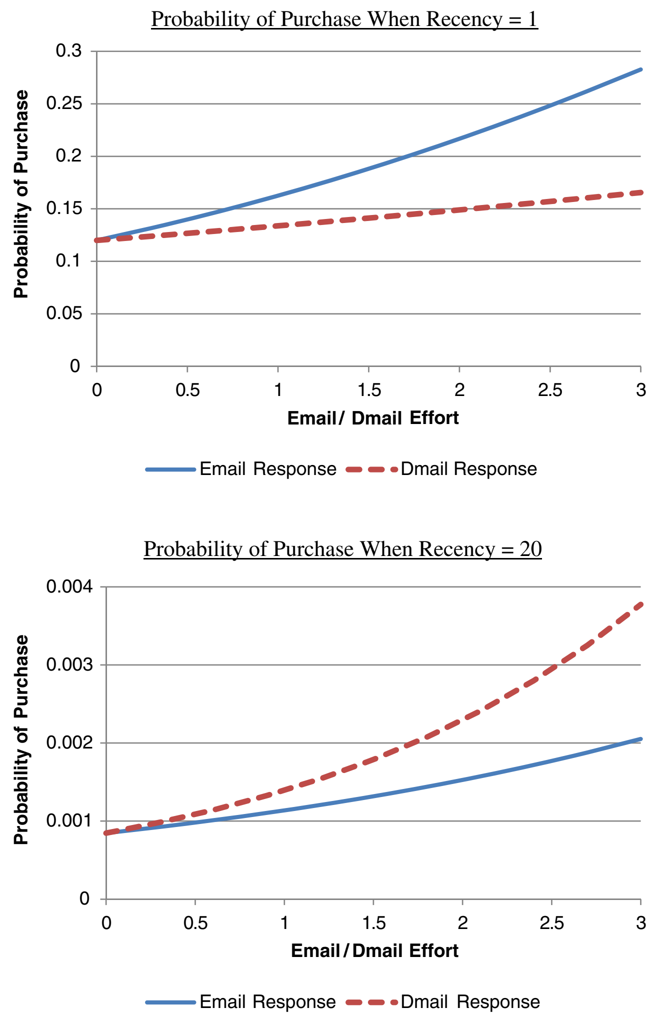

Fig. 3 Purchase likelihood as a function of email and dmail for different recency states

marketing is statistically significant and positive for direct mail. This means that direct mail response becomes more pronounced for higher recency states. This is illustrated in Fig. 3, which compares response to email and direct mail for customers in different recency states. When recency equals 1, the dotted line, representing direct mail response, is positive but has smaller slope than the solid email line. When recency equals 20, the direct mail slope is noticeably steeper than the email line. In terms of the coefficients in Table 3, when recency equals 20, the email response slope is $0.357–20\times0.003=0.297$, while the direct mail response slope is $0.105+20\times0.020=0.505$. When recency equals 1, the response slope for email is $0.357–1\times0.003=0.354$, while for direct mail it is $0.105+1\times0.020=0.125$.

Figure 4 demonstrates saturation effects. These are driven by the negative email × lagged_email and dmail × lagged_dmail coefficients in Table 3. This means that the slope of purchase probability as a function of

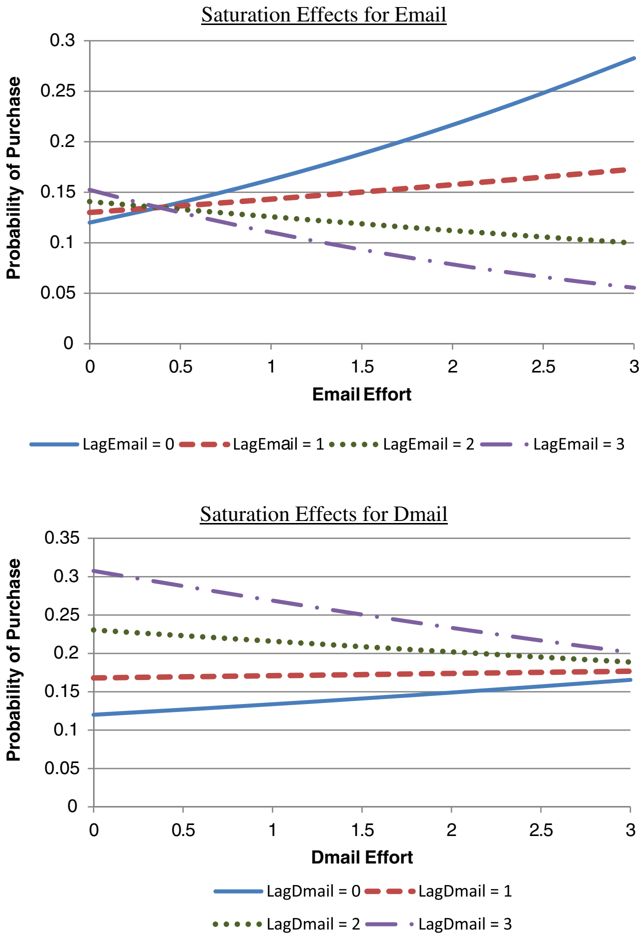

Fig. 4 Purchase likelihood as a function of Email and Dmail depending on previous Email and Dmail – illustrating saturation effects

marketing decreases if a large level of marketing has been employed in the previous period. These saturation effects are similar to those found by Ansari et al. (2008) as well as Dreze et al. (2009) and represent a "cost" to marketing beyond distribution or discount costs. The slope of purchase probability versus email or direct mail gets smaller but still positive when lagged email/direct mail equals 1. But at lagged email/direct mail equals 2 or 3, the slope actually becomes negative, suggesting supersaturation. This effect is particularly strong for email. Apparently, when the company is emailing heavily, the customer becomes so frustrated with the company, or so overloaded with emails, that continued high levels of emailing actually backfire, making the customer less likely to purchase.

We conducted three robustness checks for Model 3. First is the possibility that the negative effect of recency on purchase is becoming less severe over time as customers become intrinsically more loyal to the firm. We tested this by interacting a time period variable (numbered 1 through 25) with the recency variables, and found these interactions were not significant. This could be due to the relatively short duration of our data (25 months; a little more than 2 years). With a longer time series, we might have observed such interactions. Second, we tested the predictive ability of Model 3 on a holdout sample. To do this we divided customers into two equally sized calibration and validation samples. We estimated the model on the calibration sample and tested it on the validation sample. We calculate the ROC curve often used in database marketing (Blattberg et al. 2008, pp. 315–317) on the calibration and validation samples. The ROC curve plots "false positives" (i.e., predicting a purchase although there was none) versus "true positives" (i.e., predicting a purchase when in fact there is one) for various probability cut-off values for predicting purchase versus no purchase. The ROC curve is shown in Fig. 5. The ideal ROC curve bends more severely toward the top left, i.e., there are cut-off values that yield very few false negatives but many true positives. The 45° line represents random prediction. Figure 5 shows two things: (1) the ROC curves for both calibration and validation samples lie noticeably above the 45° line, but more importantly, (2) there is little discernible degradation in predictive performance in moving from the calibration to validation samples. This is encouraging because our model has many predictors and there is ample multicollinearity because of the squared and cubic terms, coupled with interactions.

Third, as a final check for model stability, we estimated the model separately for the calibration and validation samples and compared the coefficients. All coefficients were within two standard errors of each other, using the average standard error for the two samples. There was some instability for the email × recency and email × emaillagged variables (these were 1.95 standard errors apart) so

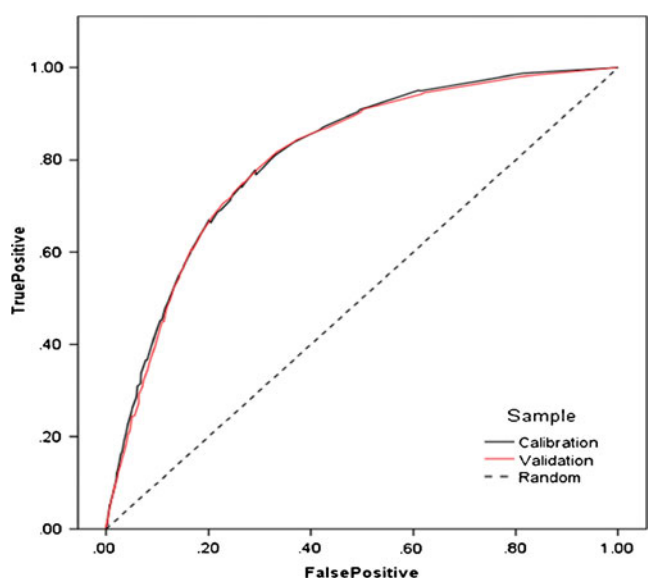

Fig. 5 ROC curve for Model 3–Calibration vs. Validation

multicollinearity is probably playing a role here. However, the most important function of the model is to predict accurately, because this is used to calculate the predicted probabilities that feed the value function. The strong ROC performance attests that while multicollinearity might understandably be causing some instability in the coefficients, overall predictive validity is not hampered.

In summary, our logistic regression contains (1) a pronounced impact of recency, (2) significant current period, carryover, and saturation effects, and (3) interactions between marketing and recency. We now can appreciate the complexity of the task at hand. For example, the saturation effects present for both email and direct mail suggest that "pulsing" may be optimal, in that if we employ a lot of marketing when the customer is in state S, we will be less apt to use marketing in state *S′*, the state that follows depending on whether the customer buys or not. However, direct mail has particularly high carryover, plus it interacts positively with recency, so this may bode for steadily increasing levels of direct mail. Of course there is also the negative main effect of recency to contend with, i.e., the recency trap. How these factors balance out to achieve the optimal policy will be demonstrated in the next section.

Optimization

State variables We use Eqs. 3 and 4 to calculate the optimal policy function, (D|S), using the method of "value iteration" (Judd 1998, pp. 412–413). Value iteration solves for the optimal policy that will not depend on the time period *per se*, but only on the state variables that describe the customer. In our application, we have four state variables:

- *Recency* If the customer is in recency state *r*, the customer moves to recency state 1 if he or she purchases, or state *min*{*r*+1, *Maxrecency}* if he or she does not purchase. That is, if the customer has not purchased for 5 months and does not purchase in the current period, the customer now has not purchased in 6 months so is in recency state 6. While in theory recency could increase indefinitely, for tractability and to ensure not working outside the range of the data, we put a cap on recency, called "Maxrecency." We use Maxrecency = 20. Once the customer gets to recency state 20 and does not purchase, we consider the customer still in recency state 20 (state ≥20 in terms of Table 1).
- *Month*: There are 12 months in the year. Table 3 shows that month influences purchase probability, and obviously changes from period to period.
- *Lagged_Email*: Table 3 shows carryover effects of email, and this variable will change over time, depending on the level of emailing in the previous period. Therefore it is a state variable. Technically, it is a continuous state variable. However, states need to be defined discretely in order to solve the dynamic program. The maximum value for monthly email was close to 3; the minimum was obviously zero. We divided this variable into 30 equal increments (i.e., 0, 0.1, 0.2, etc. up to 3.0). This means that in any period, the customer could be in one of 30 possible lagged_email states.
- *Lagged_Dmail*: Table 3 shows carryover effects of direct mail, and as for Lagged_Email, we created 30 lagged_direct_mail states.

In summary, recency, month, and two lagged marketing variables describe the customer at any point in time. There are 20 recency states, 12 months, and two 30-level lagged marketing states, so the total number of states is 20×12×30×30=216,000. This means we have 216,000 value functions, each representing the lifetime value of a customer who starts in state *S* and is marketed to optimally according to the decision rule, (D|S), derived from value iteration.

Since we have 216,000 states, it is impossible to show the transition matrix in full. However, we describe how we calculate transition probabilities. The customer's state is described by recency, month, lagged email, and lagged direct mail. The current state is denoted by *S* and the transitioned-to state by *S′* (Eq. 3). Assume it is January, the customer last purchased 5 months ago, received one email and zero direct mails in January, and zero of both in December. Then *S* = {Recency = 5, Month = January, Lagged Email = 0, Lagged Direct Mail = 0}. Assume we calculate that the probability of purchase equals 0.15. Then with probability 0.15, the customer transitions to state *S″* = {Recency = 1, Month = February, Lagged Email = 1, Lagged Direct Mail = 0} with probability 0.15, and *S″* = {Recency = 6, Month = February, Lagged Email = 1, Lagged Direct Mail = 0} with probability 1−0.15 = 0.85. In general, recency re-sets to one if the customer purchases, or increases by one if the customer does not purchase.[4] Month will always advance 1 month, and the lagged email and direct mail variables depend directly on the emailing and direct mailing the customer received in the current month.

Value iteration method Value iteration is an iterative approach. Each of the 216,000 value functions is approximated at each iteration. The procedure terminates when each value function changes by some small tolerance level, in this case we used $0.00001. The procedure was programmed in C and required approximately 1,600 iterations and 15 hours to converge. The program is available from the authors. The

[4] Note the exception that since we have 20 recency states, if the customer is in recency state ≥ 20 and doesn't purchase, he or she remains in state ≥20.

procedure is quite straightforward and is outlined in the Appendix.

Profit function Implementing the algorithm above requires specification of the current period profit function. For our application that function is expressed as:

$$\begin{aligned} E[\pi_{it}(S)] = M \times ProbPurch(Email_{it}, Dmail_{it}, S) \\ - DISTE \times Email_{it} - DISTD \times Dmail_{it} \\ - DISCE \\ \times ProbPurch(Email_{it}, Dmail_{it}, S) \\ \times min\{Email_{it}, 1\} - DISCD \\ \times ProbPurch(Email_{it} Dmail_{it}, S) \\ \times \min\{\mathrm{Dmail}_{\mathrm{it}}, 1\} \end{aligned} \quad (5)$$

where:

$\pi_{it}(S)$	Net profit contributed by customer *i* in time *t*.
M	Gross profit contribution if customer *i* makes a purchase.
$Email_{it}$	Level of emailing targeted at customer *i* in time *t*.
$Dmail_{it}$	Level of direct mailing targeted at customer *i* in time *t*.
$ProbPurch(Email_{it}, Dmail_{it}, S)$	Probability customer *i* purchases in time *t* if customer is in state *S* at that time and receives marketing equal to $Email_{it}$ and $Dmail_{it}$.
DISTE	Distribution cost per unit of emailing effort.
DISTD	Distribution cost per unit of direct mailing effort.
DISCE	Average price discount when customer buys under an email promotion.
DISCD	Average price discount when customer buys under a direct mail promotion.

The first term in Eq. 5 represents the expected positive contribution, equal to the average contribution (*M*) multiplied times the probability the customer makes a purchase. This probability depends on what state the customer is in, plus the level of emailing and direct mailing the customer receives. Information provided by the firm in this application suggested that *M*=\$71.93. Purchase probability was calculated using the logistic regression model. Note that to calculate purchase probability, we need the value of First_Amt, which varies across customers. We used the average value, \$116 (Table 2), so the optimization results and optimal policy pertain to the average customer.

The next two terms in Eq. 5 represent distribution costs, e.g., sending a direct mail or email. Information provided by the firm was that *DISTE*=\$0 and *DISTD*=\$0.40. The final two terms reflect the expected price discount when the customer responds to an email or a direct mail. Calculations using customer purchase records suggested *DISCE*=\$10.70 while *DISCD*=\$6.16. The use of the "min" function in the final two terms reflects the empirical fact that no customers purchased more than once a month in the data. This assures that the customer never can gain more than *DISCD* when purchasing under a direct mail promotion.

Optimization results Figure 6 shows the optimal email and direct mail policies for the average customer as functions of recency and compares them with the company's current policy. Recall we have 216,000 possible states. To assess the relationship between email/direct mail policies and recency, we conduct four regressions: one for each of the two marketing instruments (email/direct mail) and one for both the optimal and current policies. For the optimal policy, the dependent variable is the optimal level of email/direct mail. For the current policy, we use the current data, at the customer/time level, and use the actual level of email/direct mail used by the meal preparation establishment. The explanatory variables in both cases are the state variables: recency (19 dummy variables), month (11 dummies), lagged email, and lagged direct mail (scaled from 0 to 3 in increments of 0.1). Figure 6 displays the estimated recency state dummies.

Figure 6 leads to the following conclusions for the average customer:

- Optimal levels of direct mail are higher than optimal levels of email. This makes sense in that (1) emailing has higher saturation effects (Fig. 4), (2) direct mail has much stronger carryover effects, and (3) emailing yields larger discounts off regular price and so is more costly.
- Optimal levels of both email and direct mail generally increase with recency. This reinforces the theme that marketing should be used to arrest the progression of the customer to higher recency states (Table 3, Figs. 1 and 2).
- We see some signs of "pulsing" in the optimal email policy. For recency levels 13–18, high levels of email when the customer is in state *r* are followed by low levels of email if the customer does not purchase and therefore progresses to state *r*+1. This is probably due to the saturation effects shown in Fig. 4. If the customer does not purchase and moves to state *r*+1, it becomes unprofitable to follow up with additional emailing which will just be ignored due to saturation. It is better to wait to see if the customer drifts further, to state *r*+2, and then

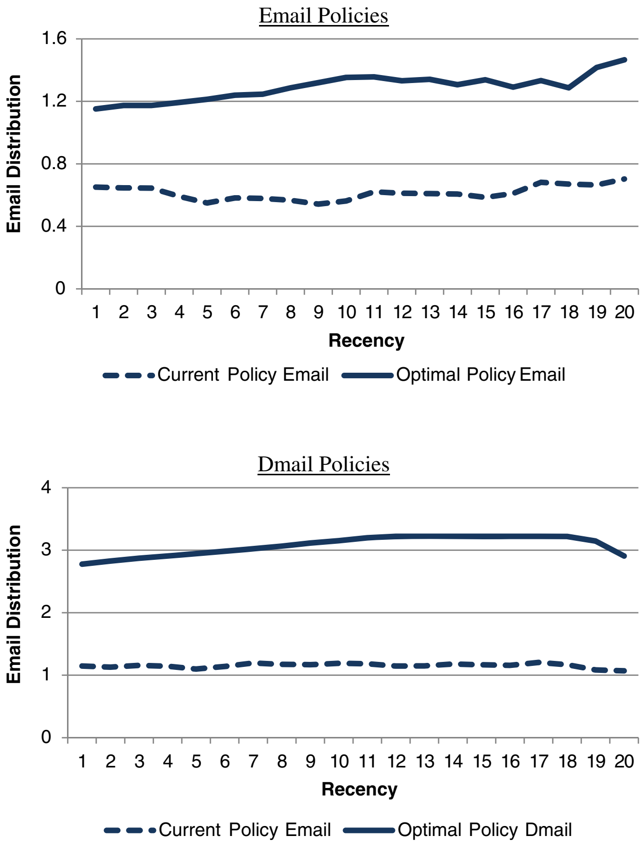

Fig. 6 Optimal and actual Email/Dmail policies as function of recency

expand emailing when the customer is more receptive to it.

- When the customer is in state 20, direct mailing falls off while emailing increases. We interpret this to be the result of strong carryover for direct mail. Part of the attraction in direct mail is the carryover effect, which means that high levels of direct mail ensure the customer will be more likely to purchase in the next period even if the customer does not purchase and drifts to a higher recency state. However, when the customer gets to state 20, that additional insurance benefit no longer is in play, since if the customer does not purchase when in state 20, he or she stays in state 20.
- The optimal policy suggests the firm should be spending more on both email and direct mail, compared to their current policy. One might assert it is obvious to increase email because email has zero distribution cost. But email has two costs: First is the discount it advertises, which is expended on baseline sales as well as any incremental sales stimulated by the email. Second is the saturation effects of email (Fig. 4) that show that too much emailing in period t decreases its effectiveness in period t+1. So the cost of emailing in period t is that it makes emails less effective in period t+1.
- The company's current policy is not to target based on recency (the relationships between recency and email/direct mail distribution are basically flat).[5]

Figure 7 shows a revealing picture of what could be gained by following the optimal policy. It displays CLV for the average customer (i.e., the customer with First_Amt = $116), as a function of recency. For the optimal policy, these are the average value functions after controlling for our other state variables (see footnote to Fig. 7). For the current policy, these values were calculated using a simulation of the current policy over the lifetime of the average customer. As can be seen, CLV decreases markedly as a function of recency—even with the optimal policy, a high recency customer is not as profitable as a low recency customer. But the difference between current practice CLV and optimal CLV is clear, on the order of $150–$200 per customer. The results for high recency customers are particularly salient: these customers are currently virtually worthless to the firm, but our optimization suggests that with proper marketing, they would be worth roughly $150–$175. Together with Fig. 6, this suggests the firm is now giving up too soon on these customers.

Summary and avenues for future research

We have developed and demonstrated an approach to managing the recency trap that maximizes customer lifetime value. The approach consists of three key elements: (1) focus on customer recency and the related customer migration model of CLV, (2) estimation of a customer-level marketing response function that includes several recency phenomena as well as marketing carryover and saturation, (3) use of a dynamic program that utilizes the estimated response function to derive a customer-specific optimal policy for utilizing two marketing tools—in this case, email and direct marketing. The optimization explicitly accounts for the recency trap by determining what marketing decision to make given the customer's current recency state, taking into account the ramifications for which recency state the customer is likely to transition to next period. This is shown explicitly in Eq. 3 and 4.

Our paper can be seen as an advocacy for recency and the migration model of CLV, but recency is not the only

[5] In fact, discussion with the firm's management suggested that the company was not currently targeting email or direct mail efforts in any way, i.e., they were not using previous purchase, etc., to target marketing. If they had been, this would have created an endogeneity that we would have had to handle in our estimation of the logistic customer response function (see Rhee and McIntyre 2008).

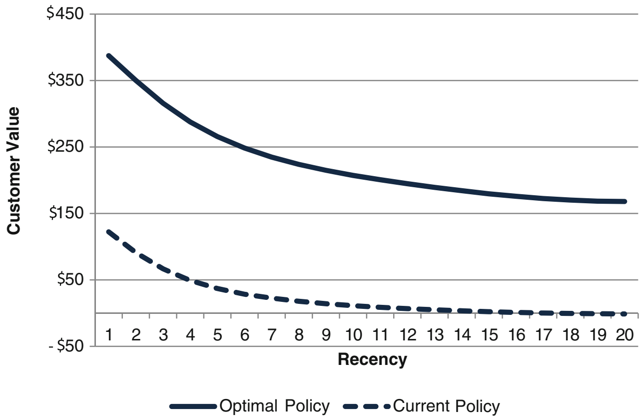

Fig. 7 Customer lifetime value: optimal vs. current policies. Graph is based on regression of state-specific value functions vs. recency, month, and LastEmail/Dmail. Graphed numbers use month=0, Last Email=0, and LastDmail=0 as base cases. Changing these bases would change the level of the graphs slightly but the general trends and difference between optimal and current policy would remain roughly the same

phenomenon to be factored into optimal customer-targeted marketing programs. Marketing carryover and saturation play a crucial role. The need to keep track of these variables increased the complexity of the optimization—as not only recency but recent marketing efforts also became state variables—but our application shows that incorporating these factors is feasible.

Our application serves as an interesting case study. This company was truly falling victim to the recency trap, as shown in Table 1. Their marketing program for the average customer was underfunded and did not expend the *additional* efforts needed as the customer moved to higher recency states. Our prescribed policy calls for increasing efforts before the customer drifts away. Note this is a function of the particular response function we estimated and for the average customer for whom we ran the simulation. One could imagine, for example, a different application where higher recency groups might become significantly *less* responsive to marketing, whereby beyond a point, when recency becomes too high, it no longer becomes worth it and the firm lets the customer drift away.

The implications of our work for researchers are: (1) Recency and the migration model of CLV are key tools that merit increased attention in customer management models. (2) Recency and marketing response can interact, reinforcing Khan et al. (2009). This inclusion of this interaction can be very important in prescribing the optimal marketing policy. (3) In fact, several response phenomena—interactions, carryover, and saturation—all need to be factored into an optimal targeted marketing policy.

The implications of our work for managers are: (1) Recency and migration model diagnostics shown in Table 1 and Fig. 1 should constantly be monitored by firms. It is possible that in a given circumstance companies will not be at the mercy of the recency trap. But the accumulated evidence, including this paper, suggests this is a key phenomenon. (2) The tools to derive an optimal CLV marketing policy are feasible for practical implementation. Logistic regression is very well known to companies, and the value iteration solution of a dynamic program can be easily programmed. (3) The optimization derives recommendations for targeted one-to-one marketing. But the approach also contributes important strategic guidance—in this case, increase marketing efforts and use recency as a crucial criterion for targeting marketing efforts. (4) Optimization can have a large impact on CLV. Our results suggest that in this case, customer value would increase by hundreds of dollars per customer and customers who heretofore were worth virtually $0 to the company could be converted to customers worth roughly $150 on average. (5) Finally, this work reinforces the emerging view that customer lifetime value is something to be *managed*, not merely *measured*. A key challenge is to derive a set of marketing policies that will maximize CLV.

While we believe this paper has covered and addressed several key issues in managing customer lifetime value, there are many opportunities ahead. Purchase expenditure could be included in the approach. In our case, we used an average customer contribution in our profit function. However, contribution could be influenced by marketing. and customers could migrate among different contribution states. Omitting this simplifies the analysis so we can focus on the recency trap. To include contribution would require two “doable” steps: (1) Contribution would be included as a state variable, so we would have Recency = 1; Contribution = $100, Recency = 1; Contribution = $200, etc. This obviously would add many states and undoubtedly we would have had to decrease the number of recency states. (2) We would need to add a customer “Expenditure” equation to the model (see Khan et al. 2009). Again, doing this would not obliterate the recency trap effect; it would just make it less detailed than we wanted to focus on in our research. Having said this, a promising area for future research would be to examine whether for example increased recency is associated with lower expenditures as well. While the logistic regression model includes an implicit interaction between email and direct mail, we did not model this explicitly in order to keep the model as simple as possible. This would be feasible within our framework because it would not expand the state space required to solve the dynamic program.

We have taken the perspective of the firm as the profit center of interest. However, if we think of the brand level or the product line level within a firm, one brand’s recency trap could be another brand’s recency gain. Companies may be particularly interested in this because the profit margins associated with one brand may be higher than that for others. This could be handled within our framework by defining recency with respect to a particular brand. The response model would be altered, perhaps to a nested logit where the model would predict whether a customer would purchase from the firm, and if so, which brand, as a function of brand recency variables. This would increase the number of states and hence the complexity of the analysis, but would be an interesting area for future research.

Our response model includes many sources of customer heterogeneity, for example the First_Amt variable that differs across customers, and differences in response to marketing related to recency. As in any model-building effort, it is possible there are other sources of observed heterogeneity we might have included in the model (e.g., purchase frequency) but which we did not in the interest of model simplicity. In addition, there may be unobservable sources of heterogeneity that could be incorporated using for example random effects, latent segments, or random parameters. These approaches identify additional sources of heterogeneity among consumers, almost always improve model fit, and can prevent the detection of spurious dynamic effects. However, these models require assumptions, for example that random effects must be uncorrelated with variables included in the model, which are difficult to verify. They also are more difficult to implement. For example, a random effects model with individual-level unobserved sources of heterogeneity may be estimated on 4,000 customers but the company may have 4,000,000 customers and wish to apply the model to them. It is difficult to infer the individual-specific coefficients for these 4,000,000 customers. Khan et al. (2009) incorporate unobserved heterogeneity at the segment level. However, in applying the model, they would still have to infer segment membership for the rest of the firm's customer base. Additionally, our demonstration of the model was for the average customer defined by the average of our First_Amt variable. It is possible the optimal policy would differ for customers with different values for this variable.

In summary, there are trade-offs in formulating the customer response model involving how much and which forms of customer heterogeneity to include. More heterogeneity, both observed and unobserved, is good because it can uncover real differences among customers, improve fit, and prevent spuriously estimated dynamic effects. However, including more heterogeneity has costs in terms of model complexity, implementation and required assumptions. We took a middle ground, focusing on observed sources of heterogeneity particularly related to recency (e.g., email × recency interactions) since recency is the key phenomenon in our study. An interesting avenue for future research would be to include as many sources of heterogeneity, both observed and unobserved, as possible and examine how addressing the recency trap may differ across customers.

Our work is highly suggestive of the gains to be had by managing customer recency effectively. However, the efficacy of the approach should be demonstrated in a field test, which would provide convincing evidence. We indeed encourage future researchers to undertake these important improvements over our current paper.

Appendix

Value iteration algorithm used to derive the optimal policy (D*|S)

1. Let $V(S)^w$ = the value function at iteration w for a customer who is in state S.
2. We have two decision variables—email effort and direct mail effort. These each range from 0 to 3 (3 is the maximum value of these variables in the data, and we wanted to stay within the range of the data). We divide each of these variables into 30 increments of 0.1. Therefore, the policy function (D|S) is a set of two values for each state, consisting of one of 30 possible email decisions and one of 30 possible direct mail decisions.
3. Find initial values, $V(S)^0$. We did this by computing the short-term profit function, $\pi_{it}(S)$ for each of the 900 possible email/direct mail combinations, for each state, and taking the maximum as the initial value, $V(S)^0$.
4. For each state the maximum value of $V(S)^w = \max_D \left\{\pi(S) + \delta(ProbPurch \times V\left(S'\right)^{w-1} + (1 - ProbPurch) \times V\left(S''\right)^{w-1})\right\}$ by trying all 900 possible combinations of email and direct mail. Call the combination that produces this maximum $(D^*|S)$.
5. Test whether $V(S)^w - V(S)^{w-1} < 0.00001$ for each state S. If this holds, the process has converged and the current value of $V(S)^w$ is the value function for customers in state S, and the most recently used combination of email and dmail is the optimal policy function $(D^*|S)$. If the condition does not hold for all states S, set $V(S)^{w-1} = V(S)^w$ and proceed back to step 4 for another iteration. Note that we have updated the value function because now in step 4, the new value functions we created on the left side of the equation will be on the right side of the equation.

References

Ackoff, R. L., & Emshoff, J. R. (1975). Advertising research at Anheuser-Busch, Inc (1963–68). *Sloan Management Review, 16* (2), 1–15.

Ansari, A., Mela, C. F., & Neslin, S. A. (2008). Customer channel migration. *Journal of Marketing Research, 45*(1), 60–76.

Berger, P. D., & Nasr, N. I. (1998). Customer lifetime value: marketing models and applications. *Journal of Interactive Marketing (John Wiley & Sons), 12*(1), 17–30.

Beswick, C. A. (1977). Allocating selling effort via dynamic programming. *Management Science, 23*(7), 667–678.

Bitran, G. R., & Mondschein, S. V. (1996). Mailing decisions in the catalog sales industry. *Management Science, 42*(9), 1364–1381.

Blattberg, R. C., Byung-Do, K., & Neslin, S. A. (2008). *Database marketing analyzing and managing customers*. New York: Springer.

Bult, J. R., & Wansbeek, T. (1995). Optimal selection for direct mail. *Marketing Science, 14*(4), 378–394.

Dreze, X., & Bonfrer, A. (2008). An empirical investigation of the impact of communication timing on customer equity. *Journal of Interactive Marketing, 22*(1), 36–50.

Elsner, R., Krafft, M., & Huchzermeier, A. (2003). Optimizing rhenania's mail-order business through dynamic multilevel modeling (DMLM). *Interfaces, 33*(1), 50–66.

Elsner, R., Krafft, M., & Huchzermeier, A. (2004). The 2003 ISMS practice prize winner-optimizing rhenania's direct marketing business through dynamic multilevel modeling (DMLM) in a multicatalog-brand environment. *Marketing Science, 23*(2), 192–206.

Fader, P. S., Hardie, B. G. S., & Lee, K. L. (2005). RFM and CLV: using iso-value curves for customer base analysis. *Journal of Marketing Research, 42*(4), 415–430.

Fader, P. S., Hardie, B. G. S., & Jerath, K. (2007). "Estimating CLV using aggregated data: the tuscan lifestyles case revisited. *Journal of Interactive Marketing, 21*(3), 55–71.

Freedman, J. L., & Fraser, S. C. (1966). compliance without pressure: the foot-in-the-door technique. *Journal of Personality and Social Psychology, 4*(2), 195–202.

Gönül, F., & Shi, M. Z. (1998). Optimal mailing of catalogs: a new methodology using estimable structural dynamic programming models. *Management Science, 44*(9), 1249–1262.

Gönül, F., Kim, B.-D., & Shi, M. (2000). Mailing smarter to catalog customers. *Journal of Interactive Marketing, 14*(2), 2–16.

Hughes, A. M. (1996). *The complete database marketer, revised*. New York: McGraw-Hill.

Judd, K. L. (1998). *Numerical methods in economics*. Cambridge: MIT Press.

Keane, M. P., & Wolpin, K. I. (1994). The solution and estimation of discrete choice dynamic programming models by simulation and interpolation. *Review of Economic Statistics, 76*(4), 648–672.

Khan, R., Lewis, M., & Singh, V. (2009). Dynamic customer management and the value of one-to-one marketing. *Marketing Science, 28*(6), 1063–1079.

Leeflang, P. S. H., Wittink, D. R., Wedel, M., & Naert, P. A. (2000). *Building models for marketing decisions*. Boston: Kluwer.

Little, J. D. C. (1970). Models and managers-concept of a decision calculus. *Management Science Series B-Application, 16*(8), B466–B485.

Miglautsch, J. (2002). Application of RFM principles: what to do with 1-1-1 customers? *Journal of Database Marketing, 9*(4), 319.

Montgomery, D. B., Silk, A. J., & Zaragoza, C. E. (1971). A multiple-product sales force allocation model. *Management Science, 18*(4), P3–P24.

Naik, P. A., & Piersma, N. (2002). *Understanding the role of marketing communications in direct marketing*. Rotterdam: Econometric Institute.

Neslin, S. A., Gupta, S., Kamakura, W., Junxiang, L., & Mason, C. H. (2006). Defection detection: measuring and understanding the predictive accuracy of customer churn models. *Journal of Marketing Research, 43*(2), 204–211.

Neslin, S. A., Novak, T. P., Baker, K. R., & Hoffman, D. L. (2009). An optimal contact model for maximizing online panel response rates. *Management Science, 55*(5), 727–737.

Pauwels, K., & Neslin, S. A. (2008). *"Building with bricks and mortar: the revenue impact of opening physical stores in multichannel environment." working paper*. Hanover: Tuck School of Business, Dartmouth College.

Pfeifer, P. E., & Carraway, R. L. (2000). Modeling customer relationships as markov chains. *Journal of Interactive Marketing (John Wiley & Sons), 14*(2), 43–55.

Ray, M. L., & Sawyer, A. G. (1971). Repetition in media models: a laboratory technique. *Journal of Marketing Research, 8*(1), 20–29.

Reimann, M., Schilke, O., & Thomas, J. S. (2010). Customer relationship management and firm performance: the mediating role of business strategy. *Journal of the Academy of Marketing Science, 38*(3), 326–346.

Rhee, S., & McIntyre, S. (2008). Including the effects of prior and recent contact effort in a customer scoring model for database marketing. *Journal of the Academy of Marketing Science, 36*(4), 538–551.

Rust, R. T., & Verhoef, P. C. (2005). Optimizing the marketing interventions mix in intermediate-term CRM. *Marketing Science, 24* (3), 477–489.

Simester, D. I., Sun, P., & Tsitsiklis, J. N. (2006). Dynamic catalog mailing policies. *Management Science, 52*(5), 683–696.

Tokman, M., Davis, L. M., & Lemon, K. N. (2007). The WOW factor: creating value through win-back offers to reacquire lost customers. *Journal of Retailing, 83*(1), 47–64.

Van den Poel, D., & Leunis, J. (1998). Database marketing modeling for financial services using hazard rate models. *International Review of Retail, Distribution & Consumer Research, 8*(2), 243–257.

Van Diepen, M., Donkers, B., & Franses, P. H. (2009). "Does Irritation Induced by Charitable Direct Mailings Reduce Donations? *International Journal of Research in Marketing, 26*(3), 180–188.

Venkatesan, R., Kumar, V., & Bohling, T. (2007). Optimal customer relationship management using bayesian decision theory: an application for customer selection. *Journal of Marketing Research, 44*(4), 579–594.

Zoltners, A. A., & Sinha, P. (1980). Integer programming models for sales resource allocation. *Management Science, 26*(3), 242–260.