

Determining the Direct Mailing Frequency with Dynamic Stochastic Programming

Nanda Piersma¹

Jedid-Jah Jonker²

Econometric Institute Report EI2000-34/A

Abstract

Both in business to business and in consumer markets direct mailings are an important means of communication with individual customers. This paper studies the mailing frequency problem that addresses the issue of how often to send a mailing to an individual customer in order to establish a profitable long-term relationship rather than targeting profitable groups of customers at every new mailing instance. The mailing frequency is optimized using long-term objectives but restricts the decisions to the number of mailings to send to the individual over consecutive finite planning periods. A stochastic dynamic programming model that is formulated for this problem is easy to solve for many applications in direct mailing. A particular implementation of the model will provide the direct mailer with controls to stimulate desired response behavior of their customers. The model is calibrated for a large non-profit organization and shows that very large improvements can be achieved by approaching the mailing strategy with the mailing frequency problem, both in the number of mailings to send and in the profits resulting from the responses.

Keywords: Direct Marketing, Stochastic Dynamic Programming

¹ Corresponding author, Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands, e-mail:piersma@few.eur.nl

² Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands, e-mail:jjjonker@few.eur.nl

1. Introduction

In 1995, 12% of total sales on business-to business markets and 5% of total sales on consumer markets in the United States were generated by direct marketing (DMA, 1995). Direct marketing corporations frequently send mailings to their (potential) customers. For each mailing the corporation selects addresses that are likely to respond to the mailing. Based on the historical mailings and its responses this selection process is repeated at every mailing instance. Being part of the marketing activities of a firm, direct marketing has become a serious research topic in the last 10 years in the marketing research literature. Typically mathematical models for direct marketing aid the address selection process for a single mailing based on the probability to respond (Roberts and Berger, 1999). The models provide a mailing strategy by explaining response behavior from previous behavior, social demographic variables and other available information. This information is used to predict the probability to respond to the next mailing and addresses are selected for the new mailing based on their probability to respond. Stochastic models that determine a mailing strategy based on the response behavior of the customers include binary choice models (Bult and Wansbeek, 1995), latent class models (DeSarbo and Ramaswamy, 1994) and neural networks (Levin and Zahavi, 1996).

For a direct marketing firm sending a mailing is not a one time event, but part of a flow of mailings sent over a longer period of time. Making a selection for only one period neglects the dynamics in the response to a mailing: the direct mailer wants to establish a long-term relationship with the customer and the decision to send a mailing to a person today influences the probability this person will respond to the next mailing. In this paper we study the mailing frequency pattern for individual customers and not the address selection for a single mailing. The focus of our model is on the *optimal frequency of mailings* to individual customers in order to maximize the total long-term expected revenue of their combined responses. Throughout this paper we will refer to the mailing frequency problem as the decision on the mailing pattern for an individual over an infinite horizon and with the objective to establish a profitable long-term customer relationship. The combination of mailings and responses comprises a Mailing and

Response pattern (further denoted by MR pattern). Single mailing models decide on the individuals that will receive a mailing at a certain point in time. When the single mailing policy is evaluated over a number of mailing occurrences usually one observes an erratic MR pattern. A mailing policy that is based on frequency of mailings rather than address selection will result into a MR pattern that is more regular in terms of intermediate time between successive mailings and responses. A well chosen mailing frequency policy will result into a regular mailing pattern that will increase customer profitability and will reduce wasteful mailings. Since it is more expensive to obtain new customers than to maintain a relationship with current customers (Peppers and Rogers, 1993) an optimal frequency policy will also be cost efficient. Although the importance of long-term customer relationships has long been recognized, maximizing mailing policies on the basis of long-term customer profitability has not been included in the quantitative models that aid the direct mailing decisions until recently.

We will model the mailing frequency problem for each individual as a Markov decision chain with the objective to maximize customer profitability over an infinite time horizon. We define finite planning periods within the infinite horizon that each consist of multiple mailing occurrences. The mailing frequency problem decides on the number of mailings to send over each planning period and records the response behavior. The new planning period considers new mailing occurrences immediately following those in the previous planning period. The state space defines the MR pattern of the customer in a planning period. The mailing frequency decision is based on the state in the previous planning period and the probability to move to a new state based on the decision made at this time point. Our main contribution is the development, estimation and implementation of a stochastic dynamic programming model to maximize long-term customer profitability with the notion of a shifting planning horizon. We can use this model to determine long-term mailing frequency policies with very small computational effort.

Two other references also apply a Markov decision model for direct mailing decisions. Bitran and Mondschein (1996) determine the optimal mailing policy over multiple mailing occurrences using stochastic dynamic programming. In their model they consider

the mailing frequency decision in combination with inventory control, with applications to the catalog industry. However they consider the mailing frequency problem where multiple mailings can only result into one response. The decision is therefore how many mailings will trigger the response of a customer, where we allow multiple responses to multiple mailings. Also the decision to send a number of mailings at the time occurrence is based only on the probability to respond and does not consider the size of the response. Gönül and Shi (1998) determine an optimal mailing policy over an infinite planning horizon using a Markov decision model. For each individual customer a Markov decision model is formulated that decides whether to send a mailing in each future mailing occurrence based on customers that maximize utility and direct mailers that maximize profits. In each (discrete) time period the decision to mail a specific customer is based on its probability to respond and the expected future value of this customer (based on its utility). A drawback of their model is the increase of the size of state space over time that in combination with the infinite time horizon leads to lengthy computational efforts. As they remark, the inclusion of the monetary value in the state definition will increase their state space even more, making the model infeasible for practical applications.

We define the states of the customer such that they include mailing pattern, number of response and the size of the responses. With this definition the direct mailer can identify preferable states in terms of customer profitability and mailing intensity. Customers that will respond to one mailing with a high response are usually preferred over customers that do not respond or that have a small response to multiple mailings. By assigning weights to the different MR patterns and optimizing the long-term weighted profit of a MR pattern the direct mailer can determine the mailing frequency that will have the highest probability to observe customers in preferred states, thus showing preferred response behavior. We will show that the mailing frequency policy that takes preference into account is competitive to mailing frequency policies based on more traditional long-term objectives such as customer life-time value. Also the short-term effects show that the mailing pattern will converge to the profitable steady state within a small number of steps. Thus a simple decision tool that identifies preferable MR patterns in combination

with the assignment of weights to the MR patterns will enable the direct mailer to find very profitable mailing frequency patterns.

The outline of this paper is as follows. In Section 2 we present the Markov decision model and the preferred states setup. In Section 3 we calibrate the model for a Dutch charitable organization. In Section 4 we apply the model to a mixture of loyal (current) and new customers, and model when old customers become inactive and when new customers will enter. In the final section we will discuss the marketing implications and define a simple decision support system for the mailing frequency problem.

2. Modeling the mailing frequency problem

The timing of the mailings within each planning period is modelled by mailshots, i.e. specific dates when a mailing is sent to the customers. For each period the decision is how many mailings each customer will receive. If the decision is to send a specific customer three mailings within a four mailshot period then we assume that the most profitable mailshots are known and will be used. Most direct mailers consider a list of loyal customers (houselist) and a list of potential customers (rental list). For the customers from the houselist we determine the number of mailshots within the next planning period where this customer will receive a mailing. For the customers from the rental list we determine the initial number of mailings that each customer in this group will receive in the next period.

A state $s=s[t]$ records the MR pattern of customer i in the current period t . A state is represented by a three dimensional vector $s[t]=(s_1[t], s_2[t], s_3[t])$. The first element $s_1[t]$ reflects the number of mailings that are sent in period t . The second element $s_2[t]$ reflects the total amount donated by the customer over the mail-shots in period t and $s_3[t]$ holds the number of mailings that the customer responded to in period t . Given a state $s[t]$ of customer i at the end of period t , the direct mailer decides on the number of mailings that this customer will receive in period $t+1$. This action a is taken at the beginning of period

$t+1$. The states are thus measured at the end of t , and the decision on the number of mailings for the next period is made at the beginning of $t+1$. The next state is then observed at the end of $t+1$. We assume that the direct mailer will take the same action whenever the customer is found in the same state. The action is also assumed to be stationary, that is, it is independent of t . The collection of actions corresponding to all the states is called a policy $R = \{a_s, s \in \hat{S}\}$.

For a customer in a given state $s[t]$ we define an immediate reward (profit) $r_s(a)$, if action a is taken in state s . This reward is the total amount spent by this customer in this period given the number of mailings sent by the direct mailer. A reward thus depends on the state of the customer and the action taken by the direct mailer. The net reward is the reward corrected for the costs of the mailings. The response of a customer is described by the immediate reward obtained in the current period and the number of mailings that the customer responded to. Depending on the action a taken for a customer that is found in a state $s[t]$, there is a transition probability $p_{su}(a)$ to the state $u[t+1]$ of the customer in the next period. This probability thus depends on the state $s[t]$ in the current period and the action a taken by the direct mailer. The Markov decision model is based on the assumption that the transition probability depends only on the current period, the so-called Markov property.

The states at the end of each period t are closely related with a segmentation of the customers. By observing the state of every customer in the house list, these customers are divided over segments defined by each of the states. We assume that the segments show homogeneous behavior with respect to response behavior (both the probability to respond and size of the response). This assumption relates to the assumption that the probability to observe a MR pattern in a planning period is the same every time that a customer is observed in a certain state at the beginning of the period and the same action is taken.

2.1 Optimizing Customer Lifetime Value

If one would apply a different policy every year it is not clear which policy is responsible for an increase or decrease in revenues. Only by applying the same policy for a number of years and comparing that policy with another one that has been applied for a number of years one can distinguish these policies in quality. But in practice there is not enough data for the evaluation of all possible policies. A theoretical measure for quality of a mailing policy is the long-run average performance. This measure reflects the performance of a mailing policy if this policy would be applied over an infinite number of periods. The measure can be seen as an honest comparison for the effectiveness of different policies. A well known performance measures applied to customer relationships is the discounted future revenues, i.e. the customer lifetime value. With the discount factor α , the objective function based on expected total discounted return, given that the starting state s_0 at period 0, is given by

$$\max_{R:\{a,a \in A(s), s \in S\}} E \left[\sum_{n=0}^{\infty} r_{s_n}(a) \alpha^n \mid s_0 \right]$$

where action a is applied at each planning instance where the customer is observed in state s_n .

For bounded rewards the optimal policy A^* exists for the mailing frequency problem with this objective and can be found for instance using an LP formulation of the optimization problem:

$$\begin{aligned} \min \quad & \sum_{s \in S} u_s \\ \text{subject to} \quad & \\ u_s \geq & r_{sa} + \alpha \sum_v p_{sv}(a) u_v \quad \forall s, a \\ u_s \geq & 0 \quad \forall s \end{aligned}$$

2.2 Optimizing preferred states

Besides maximizing profits, a direct mailer may have other objectives. Another goal may be to minimize the number of mailings, or maximize the number of profitable customers

that are being approached. The Markov model allows the direct mailer to evaluate the performance of alternative criteria. In this section we analyze the effect of maximizing the number of individuals in so-called preferred states. We illustrate the usefulness of this strategy by the following example: consider the state (3,50,1) where the customer responds only once to three mailings with a response of 50. If one could have the same response with one mailing (state (1,50,1)) then the net reward will be higher. However, if an extra mailing would trigger an extra response (say of size 50) resulting into state (4,100,2), the net reward will increased. Clearly the states (1,50,1) and (4,100,2) are preferred over the state (3,50,1) with respect to net reward. However some states are not as easily distinguished. State (1,10,1) has smaller reward than state (3,14,1) but sends less mailings, resulting in the same net reward. When budget restrictions are in use, state (3,14,1) may not be preferred because of the higher cost, but otherwise the mailer may prefer 3 mailings in order to enhance visibility. (see also Bitran and Mondschein who address the problem of how many mailings will trigger a response).

The direct mailer can prefer to minimize the number of mailings, to maximize response size or any combination of these two. To give the direct mailer the control over the multiple objectives every state is assigned a weight that reflects the relative preference of this state.

This mailing policy is then based on the long-run average weighted probability to observe customers in certain states. When state s is given weight w_s then the objective can be expressed by

$$\max_{R: \{a: a \in A(s), s \in S\}} \lim_{n \rightarrow \infty} \frac{E\left[\sum_{j=0}^n w_{s(n)} r_{s(n)}(a) \mid s(0) = i\right]}{n+1}$$

where $r_{s(n)}(a)$ is the net reward when action a is taken in state $s(n)$.

If the weight of a state is equal to one in each state then the objective function will become the standard total profit criterium for stochastic dynamic programming (Ross

1983, Ch 4.). The existence of the optimal policy is satisfied in our model for every nonnegative weight.

The model can be is solved by the linear programming representation:

$$\begin{aligned}
 & \max \sum_s \sum_a r_{sa} w_s x_{sa} \\
 & \text{subject to} \\
 & \sum_a x_{sa} = \sum_v \sum_a x_{sa} p_{sv}(a) \quad \forall s \\
 & \sum_s \sum_a x_{sa} = 1 \\
 & x_{sa} \geq 0 \quad \forall s, a
 \end{aligned}$$

In the next section we will compare the performance of Markov decision model with the preferred states criterion to the customer lifetime value for a charitable organization in the health care sector.

3. Calibration of the model

Our model is calibrated with data from a large charitable organization in the health care sector whose identity cannot be disclosed. The data consists of the complete mailing and response history from February 1994 till December 1999. There are approximately 600.000 customers in the data set. For each customer there is a record with personal information (postal code, registration number, house number), customer information (when active, how was the customer approached, current status, date inactive, reason inactive etc), and mailing information (date of each mailing, date of each response, size of the response). The organization uses planning periods of one year. At the beginning of each year it determines who of the active donors will receive 0, 1, 2, 3 or 4 mailings. New donors are also actively recruited each year. This makes the composition of the database dynamic: there are those who enter and there are those who become inactive.

The organization sends at most 4 mailings per year, so the response of the customers is limited to at most 4 reactions per year. A careful comparison of the donations showed that the amount donated per year can be aggregated in five intervals [0,10], [10,25], [25,50], [50,100], [100,+) (amounts in the local currency). There are thus exactly 55 states. In general it is not hard to see that with m possible donation intervals and at most n mailings in a planning period the state space contains exactly $m * n * (n+1) / 2 + n + 1$ states. To see this observe that for every donation interval one can observe one response to one mailing, one or two responses to two mailings up to n possible responses to n mailings: $n(n+1)/2$. Also, with any number of mailings ($n+1$) sent there is a possibility of no response. The associated 55x55 transition probabilities for each action can also be estimated directly from the data.

The mailing policy that is derived from the historical data is given in the Appendix (Table 1) together with a description of the states, the average giftsize for each state (with standard deviation), the number of observations for each state and the current action. We will refer to the historical mailing policy as the current policy. In Section 2 we describe the maximizing customer lifetime value and the preferred states criterion. For a test setup

of this approach we divided the states into three equally sized groups. The most favorable states are defined as the ones where an individual made a cumulative donation of at least 50 guilders in the previous year and responded to more than 50% of the mailings received in the previous year. The moderately favorable states are defined as those remaining states where an individual made a cumulative donation of at least 20 guilders in the previous year and responded to at least 33% of the mailings received in the previous year. The remaining states are classified as unfavorable. The weights that are assigned to the states are 3000 for the most preferred states, 20 for the moderately preferred and 0,1 for the less preferred states. These weights ensured that the relative profit (profit multiplied by weight) for each state and action resulted into a ordering of states such that the most preferred states have the highest relative profit.

We first compare this model to the current mailing policy of the charity fund on the basis of a number of long-run and short-run performance criteria. For the long-run comparison the steady-state probability is determined of observing a customer in a certain state. The policies are applied to a fictitious number of customers who are divided over the states according to these steady-state probabilities. By applying the mailing strategy that resulted from the optimization of the model under consideration the performance of the policy is measured by

1. Total number of mailings send to the customers as a results of the mailing policy.
2. Total number of responses observed based on long-run probability to move to a new state in the next year, where the number of responses is part of the state description.
3. Total revenue observed based on long-run average donation in each state.

Four other measures are also recorded, they are derived from the measure 1, 2 and 3.

4. Total revenue observed based on long-run average donation in each state and corrected for the cost of the mailings (the cost of a single mailing is set to 1).
5. Net revenue per mailing, i.e. total net revenue divided by the total number of mailings sent.
6. Net revenue per response, i.e. total net revenue divided by the total number of responses

7. Response percentage, i.e. the total number of responses divided by the total number of mailings sent.

We also compare the preferred states model to the objective to maximize total discounted average profit described in Section 2.

3.1 Long-run performance

The current policy can be characterized by sending almost all customers either 3 or 4 mailings and to send new and inactive customers one mailing. The long-term profitability of this policy is compared to the discounted net profit model and the preferred states model.

Table 2: Long-term performance

	Number of mailings per individual	Number of response per individual	Revenue per individual	Net Revenue per individual	Revenue per mail	Revenue per response	Response Percentage
Current policy	3.31	0.89	23.87	20.56	7.22	26.93	26.8
Preferred states	2.43	1.05	51.50	49.08	21.21	49.20	43.1
Discounted net profit (a=0.8)	2.60	1.08	51.87	49.27	19.95	48.00	41.6

Table 2 shows that optimizing by means of the mailing frequency problem results into a much more profitable mailing policy than the currently used mailing strategy. Comparing the results of the optimizing policies with the policy currently used by the organization, we see a significant improvement for both models. Less mailings are sent and the response is higher, hence the mailings are used more efficiently. Total (average) net revenue is more than doubled and higher revenues are generated by sending less mailings. The preferred states model slightly outperforms the discounted net profit model, as it sends less mailings for approximately the same net revenue per individual.

We allow for the decision not to send a mailing when a state is observed. The customers in this state are thus not profitable enough in the long run to justify sending any mailing in the next year, and are put on the inactive mailing list. Surprisingly more than 50% of

the states are considered to be inefficient and do not receive a mailing in the next year. These states include customers that did respond to every mailing that they received, but who donated only small amounts. In the long run, it is thus more profitable to spend the budget on customers that donate higher amounts of money. The customer life time value criterion results into a comparable policy for the discount factor 0.8 and 0.9. However for a small discount factor the number of mailings send increases significantly and the less profitable groups also receive (a number of) mailings. Figure 3 shows the net profit and the number of mailings per individual for discount factor 0.1 up to 0.9. It shows that putting more emphasize on the short-term profit results into frequent mailings, where incorporating the long-run relationship with the customer results into fewer mailings, with higher efficiency.

==INSERT FIGURE 3==

3.2 Short-term performance

For the short-term performance we simulate the mailing process over a number of planning periods. The customers migrate over the states on the basis of their one-step transition probability under the specified policy. The response behavior and collected revenue is measured in each period and compared to the long-run performance measures of the considered mailing policy. The short-term results are calculated using as a starting point that an individual has an equal probability of being in any of the states. A starting population of 55.000 individuals is divided equally among the states. Although the starting situation is not representative for the steady-state behavior, it does show how fast the mailing policy will result in steady-state response behavior.

Tables 4a – 4c contain the short-term performance for the different decision rules applied to the charitable organization. The budget is assumed to be unlimited, but with the standard cost of 1 per mailing, the number of mailings send per customer is a measure of the costs of sending the mailings.

Table 4a: Current decision rule

Year	Number of mailings per individual	Number of response per individual	Revenue per individual	Net revenue per individual	Revenue per mail	Revenue per response	Response percentage
1	3.36	1.22	31.36	28.00	9.33	25.64	36.4
2	3.30	1.07	27.76	24.46	8.42	25.91	32.5
3	3.27	0.99	25.67	22.40	7.86	26.02	30.2
5	3.25	0.91	23.92	20.67	7.37	26.32	28.0
10	3.26	0.88	23.47	20.21	7.20	26.82	26.9
20	3.28	0.88	23.66	20.38	7.21	26.91	26.8
L.T.	3.31	0.89	23.87	20.56	7.22	26.93	26.8

Table 4b: Preferred states decision rule

Year	Number of mailings per individual	Number of response per individual	Revenue per individual	Net revenue per individual	Revenue per mail	Revenue per response	Respos percentage
1	2.96	1.27	57.94	54.97	19.55	45.57	42.9
2	2.10	0.93	46.38	44.28	22.08	50.02	44.1
3	2.67	1.15	55.34	52.67	20.73	48.24	43.0
5	2.54	1.09	53.30	50.76	20.99	48.80	43.0
10	2.41	1.04	51.24	48.83	21.25	49.26	43.2
20	2.43	1.05	51.52	49.09	21.21	49.19	43.1
L.T.	2.43	1.05	51.50	49.08	21.21	49.20	43.1

Table 4c: Discounted net revenue maximizing decision rule

Year	Number of mailings per individual	Number of response per individual	Revenue per individual	Net revenue per individual	Revenue per mail	Revenue per response	Respos percentage
1	3.10	1.28	57.64	54.55	18.62	44.87	41.5
2	2.30	0.97	47.83	45.53	20.81	49.37	42.1
3	2.75	1.14	53.68	50.93	19.52	47.08	41.5
5	2.64	1.10	52.43	49.78	19.82	47.74	41.5
10	2.60	1.08	51.84	49.24	19.96	48.01	41.6
20	2.60	1.08	51.87	49.27	19.95	48.00	41.6
L.T.	2.60	1.08	51.87	49.27	19.95	48.00	41.6

The results show that the convergence to the long-term performance is very fast for all the models. Only during the first 5 years there is some fluctuation in the response. The response percentages and the profits are higher in the first 5 years and then converge to their steady state values, showing that implementing the strategies should yield an immediate effect.

3.3 Inclusion of a rental list.

In the previous results only the house list is considered. The customers on the rental list enter the model through state (0,0,0), that is, we assume that these customers did not receive a mailing in the previous planning period. The policy will consider the customers on the house list that are inactive in the previous period (state (0,0,0)) in the same way as customers on the rental list. Therefore, the decision not to send a mailing to a customer in the next planning period, will move this customer to an inactive list, and the customer can be seen to leave the system. At the beginning of every planning period the inactive list is combined with a rental list and if needed cleaned up. The remaining customers will enter the system again through the state (0,0,0).

Table 5a: Preferred states decision rule

Year	Number of mailings per individual	Number of response per individual	Revenue per individual	Net revenue per individual	Revenue per mail	Revenue per response	Response percentage
1	2.98	1.29	56.22	53.24	18.89	43.59	43.3
2	2.26	0.95	45.27	43.01	20.01	47.59	42.0
3	2.75	1.15	54.61	51.86	19.86	47.50	41.8
5	2.64	1.09	52.67	50.03	19.95	48.25	41.3
10	2.51	1.03	50.26	47.75	20.03	48.72	41.1
20	2.53	1.04	50.66	48.13	20.01	48.66	41.1
L.T.	2.53	1.04	50.64	48.11	20.02	48.66	41.1

Table 5b: Discounted net revenue maximizing decision rule

Year	Number of mailings per individual	Number of response per individual	Revenue per individual	Net revenue per individual	Revenue per mail	Revenue per response	Response percentage
1	3.29	1.29	57.61	54.32	17.51	44.81	39.1
2	2.98	1.05	50.30	47.31	16.87	47.82	35.3
3	3.13	1.07	50.50	47.37	16.15	47.03	34.3
5	3.13	1.04	48.96	45.83	15.65	47.02	33.3
10	3.14	1.03	48.43	45.29	15.45	46.91	32.9
20	3.14	1.03	48.41	45.27	15.44	46.91	32.9
L.T.	3.14	1.03	48.41	45.27	15.44	46.91	32.9

Again we tested the model using a simulation of the mailing process that starts with a population of 55,000 customers equally divided over the states. Using the combined house list and the rental list and applying the same policy to both groups will decrease the response percentage after a number of years for all the models. The customers on the

rental list require more mailings to trigger a response and the probability to larger donations is smaller. Response percentages are lower as is to be expected, as the number of mailings sent increases. Also the convergence of the performance is of the same order as the models without the rental list.

4. Concluding remarks

The mailing frequency problem is a new approach to direct mailing that takes the dynamics in the mailing process into account. Contrary to other dynamic models the dimensions of the problem are bounded by the concept of the finite planning periods. A small planning period will decrease the state space, and enables the direct mailer to make frequent decisions on the number of mailings sent. Notice that a planning period consisting of one mailshot relates to the traditional single mailing decision problem. But longer planning periods can also be modeled within this framework. A longer planning period enables the direct mailer to make long-run predictions on mailing costs and response behavior.

The goal of the direct marketing firm is to maximize profits. Sending all the customers in the database a mailing can be a profitable strategy when the costs of sending a mailing are low. However, even when the costs are low, sending all the customers a mailing is usually not a preferred strategy. A direct mailing company will want to minimize the "waste" or non-response from a cost minimizing perspective but also from a customer perspective: sending unwanted mailings can harm the long term relationship with a customer as it can lead to irritation towards the firm. This results in a trade-off in the number of mailings that will be sent. The mailing frequency problem models this trade-off by deciding on the mailing pattern over multiple mailing instances and it uses a long term expected profit as a criterion. The customer relationships can therefore be modeled long term and on an individual level. The profitability of this model is clearly exhibited by Figure 2, where the discount factor is varied between 0.1 and 0.9. Incorporating the long term profitability will boost the profits.

The mailing frequency problem can be solved at individual level, but is also closely related to segmentation of the customers. Each state of the Markov model is related to a segment of the customers in the database. If a customer is observed to be in a certain state, then the customer is assumed to belong to the segment defined by this state. The transition to another state in the next planning period thus relates to another segmentation

of the customers. The Markov decision models can easily be extended to include a budget constraint (see Gönül and Shi (1998) and Bitran and Mondschein (1996)).

The dataset that is used in the calibration of the model has the specific property that each response can be traced to the mailing that triggered the response. The customers fill out a form authorizing the direct mailer to withdraw an specified amount from their account. This forms contains a code that identifies the exact mailshot when the customer was approached. However, the Markov decision process with finite planning periods also applies to the case where responses cannot be related to the triggering mailing.

The Markov decision model for the mailing frequency model can be used as the basis of a decision support system (DSS) where the direct mailer has control over the following parameters:

1. The length of the planning periods.
1. The maximum number of mailings that are send in a planning period.
2. The costs of a single mailing to an individual customer.
3. The preferred MR patterns.

The results for a DSS model based on three classes of preferred states show that even a simple DSS for the mailing frequency problem can be of good use for practitioners. The concept of sending customers on the house list frequent mailings according to a consistent policy that is applied to finite planning horizons is a very promising concept compared to single mailing address selection. Applications to other organizations for instance in charity fundraising or catalog sales require minor adjustments of the model.

References

Bitran, G.R. and S.V. Mondschein, “*Mailing Decisions in the Catalog Sales Industry*” , Management Science, Vol 42, 1364-1381, 1996.

Bult, J.R. and T. Wansbeek, “*Optimal Selection for Direct Mail*”, Marketing Science, Vol 14, 378-394, 1995.

DeSarbo, W.S. and V. Ramaswamy, “*CRISP: Customer Response Based Iterative Segmentation Procedures for Response Modeling in Direct Marketing*”, Journal of Direct Marketing, Vol 8, 7-20, 1994.

Direct Marketing Association, “*Economic Impact: U.S. Direct Marketing Today*”, DMA Inc., New York, 1995.

Gönül, F. and M.Z. Shi, “*Optimal Mailing of Catalogs: A New Methodology Using Estimable Structural Dynamic Programming Models*” , Management Science, Vol 44, 1249-1262, 1998.

Levin, N. and J. Zahavi, “*Segmentation Analysis with Managerial Judgment*”, Journal of Direct Marketing, Vol 10, 28-37, 1996.

Pepper, D. and M. Rogers, “*Share of Customer, Not Share of Market*”, in: *The One to One Future*, Currency Doubleday, New York, 19-50, 1993.

Roberts, M.L. and P.D. Berger, “*Direct Marketing Management*”, Englewood Cliffs, New Jersey: Prentice Hall, Inc, 1989.

Ross, S., *“Introduction to Stochastic Dynamic Programming”*, Probability and Mathematical Statistics, A Series of Monographs and Textbooks, Academic Press, New York, 1983

Table 1 Statistics Charity fund, current policy and weights of preferred states model

STATE	gift size	#mailings	#responses	average giftsize	#observations	action
0	0	0	0	0(0)	152749	1
1	0	1	0	0(0)	219461	1
2	0	2	0	0(0)	290569	2
3	0	3	0	0(0)	97594	3
4	0	4	0	0(0)	35195	3
5	10	1	1	4.73 (1.2)	16251	2
6	10	2	1	4.74 (1.2)	14646	3
7	10	3	1	4.85 (1.1)	12778	3
8	10	4	1	4.99 (1.2)	2200	3
9	20	1	1	10.77 (1.8)	44343	3
10	20	2	1	10.80 (1.8)	38833	3
11	20	3	1	11.9 (1.6)	47880	3
12	20	4	1	11.5 (2.1)	22599	4
13	50	1	1	25.09 (3.3)	14060	4
14	50	2	1	24.73 (2.8)	16765	3
15	50	3	1	24.35 (2.9)	17274	4
16	50	4	1	25.12 (3.2)	28366	4
17	100	1	1	51.11 (4.4)	1966	4
18	100	2	1	51.54 (5.9)	2061	4
19	100	3	1	50.95 (4.1)	1959	4
20	100	4	1	51.72 (6.3)	5781	4
21	>100	1	1	168.88 (201)	699	4
22	>100	2	1	162.70 (293.9)	759	4
23	>100	3	1	156.81 (299.4)	703	4
24	>100	4	1	200.93 (548.7)	2682	4
25	10	2	2	5.16 (1.6)	1748	3
26	10	3	2	5.21 (1.5)	1574	3
27	10	4	2	5.3 (1.6)	345	3
28	20	2	2	11.62 (2.3)	5431	3
29	20	3	2	11.62 (2.3)	7812	3
30	20	4	2	12.5 (2.3)	3105	4
31	50	2	2	24.45 (6.2)	10198	4
32	50	3	2	23.36 (4.3)	27868	4
33	50	4	2	26.67 (7.6)	25411	4
34	100	2	2	53.46 (8.3)	2534	4
35	100	3	2	52.71 (6.6)	5229	4
36	100	4	2	53.69 (8.6)	15538	4
37	>100	2	2	141.38 (97.3)	438	4
38	>100	3	2	141.99 (122.1)	790	4
39	>100	4	2	157.17 (171.5)	3437	4
40	10	3	3	6.63 (1.4)	996	4
41	10	4	3	6.82 (1.7)	361	4
42	20	3	3	14.93 (1)	3795	4
43	20	4	3	15.24 (1.5)	2093	4
44	50	3	3	31.24 (5.9)	16314	4
45	50	4	3	33.25 (6.2)	18224	4
46	100	3	3	67.95 (10.2)	3739	4
47	100	4	3	69.74 (9.8)	13116	4

Figure 3 discounted net profit

