

Optimizing marketing planning and budgeting using Markov decision processes: An airline case study

A. Labbi
C. Berrospi

Many companies have no reliable way to determine whether their marketing money has been spent effectively, and their return on investment is often not evaluated in a systematic manner. Thus, a compelling need exists for computational tools that help companies to optimize their marketing strategies. For this purpose, we have developed computational models of customer buying behavior in order to determine and leverage the value generated by a customer within a given time frame. The term “customer value” refers to the revenue generated from a customer’s buying behavior in relation to the costs of marketing campaigns. We describe a new tool, the IBM Customer Equity Lifetime Management Solution (CELM), that helps to determine long-term customer value by means of dynamic programming algorithms in order to identify which marketing actions are the most effective in improving customer loyalty and hence increasing revenue. Simulation of marketing scenarios may be performed in order to assess budget requirements and the expected impact of marketing policies. We present a case study of a pilot program with a leading European airline, and we show how this company optimized its frequent flyer program to reduce its marketing budget and increase customer value.

Introduction

Nearly all industries are vitally interested in making their marketing efforts more accountable to management as well as to shareholders [1]. Marketing expenditures must increasingly be justified in terms of their effectiveness at increasing a company’s revenue. At the same time, the dual role of a marketing department—to increase short-term sales and long-term brand equity—makes it difficult to define precisely which data should be collected in order to measure the performance or success of a given marketing effort. The criteria currently used to measure marketing performance are often associated with specific customers or marketing media, or they may be grouped in categories such as customer awareness and preferences, purchase intentions, customer satisfaction and loyalty, and product shelf space [2]. Unfortunately, this data, although valuable, does not address the financial impact and therefore the ultimate success of a marketing campaign.

In today’s highly competitive environment, marketing departments must account more closely for how effectively they spend their large budgets in terms of the degree to which these expenditures contribute to the profitability, growth, and long-term competitive advantage of an organization. For instance, many studies suggest a positive association between customer satisfaction and financial outcomes such as revenues and return on investment [3–7]. This is especially true for companies whose main performance-measurement criterion is the connection between revenue expansion and cost reduction [8].

In the airline industry, great improvements have been made in managing customer relationships. Airlines collect a large amount of customer data, but instead of leveraging this data to make fine differentiations between various customer groups, most frequent flyer programs (FFPs) consist of only one standard marketing approach

©Copyright 2007 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/07/\$5.00 © 2007 IBM

for all customers. Although the necessary technology is available, airlines continue to use inaccurate models that make rough estimates of customer value, which is defined as the revenue generated from the customer's buying behavior in relation to the costs of marketing campaigns. Moreover, most airlines consider their business and first-class customers—the so-called elite or upper tier—of their FFP to be their most valuable segment. However, most FFPs of today are simplistic schemes that deliver rewards primarily on the basis of the number of return flights or points accrued, regardless of ticket price paid.

In fact, upper-tier travelers are not necessarily the most profitable or even necessarily the most loyal customers. Although they may accumulate the most FFP miles, they may not pay the highest fares, and they are very costly to serve [9]. Therefore, standard marketing efforts directed at such customers may not be optimal in terms of maximizing customer value.

One of the major tasks of marketing departments is to retain customer loyalty, but it is quite challenging to identify the most loyal customers and to calculate the value they create within an airline's FFP. Today, even budget (low-cost) airlines have begun to target business travelers as customers, and this of course has increased price competition more than ever before in the airline industry. The fierce competition to be among the preferred airlines of a large corporation leads some airlines to accept corporate contracts that yield little profit.

Nearly all airlines offer some kind of FFP, including a leading European airline in our case study. As part of its FFP, the airline has conducted numerous marketing campaigns targeting more than 500,000 customers. Every year, customers are exposed to dozens of such campaigns. The campaigns may have different goals, such as tier upgrade, cross- and up-selling, minimizing customer attrition, and introducing new incentives to accrue and redeem FFP points, and the campaigns may be conveyed through various media, such as mailings, in-flight brochures, the Internet, and magazines. Examples of cross-selling include the giving of points when products are purchased in certain shops, or the promotion of holiday packages that include a hotel, entertaining show, and airline flight. Examples of up-selling include the promotion of special discounts for spouses, or promotions that may be pertinent only to business-class flights.

The main business objective of the airline in our case study was to reduce the costs of its FFP while maximizing the lifetime value of its members. To achieve this objective, a team of their marketing managers worked with IBM researchers and consultants to optimize the FFP marketing strategy—that is, to define a business transformation process. The key marketing functions

to be transformed were 1) campaign management, 2) marketing planning, and 3) multichannel communication. This extensive, two-year project was executed in three phases. The purpose of phase 1 was to gain better knowledge of customers by deriving finer loyalty and value metrics and to obtain more homogeneous customer profiles (groups of customers characterized by similar loyalty and value metrics). In phase 2, we attempted to better understand customer behaviors at various phases of the relationship, and to identify more effective ways to leverage each customer contact. In phase 3, we optimized the allocation of marketing resources (the budget) for the FFP by focusing on processes in which cost and revenue can be optimized simultaneously.

Customer lifetime value

Quantitative approaches to the allocation of marketing resources have recently attracted increased research interest, both in the marketing [10–14] and in the data mining and statistics communities [15–19]. It is commonly agreed that marketing efforts should be evaluated in terms of their impact on customer lifetime value (CLV) [7, 20, 21], the long-term value generated by the relationship of a company with a customer. CLV is defined as the sum of the discounted cash flows that a customer generates during the customer's relationship with a company [22].

Several approaches that are useful for estimating CLV using Markov decision process (MDP) techniques can be found in the marketing science literature [18, 23–26]. MDPs include a set of states, actions, and probabilities which are described further in the section on customer dynamics. However, most of these approaches are impeded by certain limitations that are usually very important in marketing practice. These limitations are related mainly to the estimation of accurate MDPs: Most of the models found in the literature assume some *ad hoc* state representation, without providing any theoretical or practical justification for the choice of state definition. Additionally, many current approaches are limited with respect to their scalability; in practice, an evaluation usually requires very large numbers of customer transactions, which are tracked through time and stored in a data warehouse.

Furthermore, some of the marketing literature that deals with optimizing the expected CLV focuses on the allocation of marketing resources [7, 10–12, 25, 27]. However, if marketing investments are to be evaluated from a financial perspective [28, 29], the risk of the investment should be quantified and managed as it is commonly done in financial engineering practice [30, 31]. Our approach attempts to encode the risk information with respect to the expected customer returns given certain marketing actions.

IBM marketing tool

In response to the analytical needs of the marketing industry, IBM Research has developed a solution called the IBM Customer Equity Lifetime Management Solution (CELM). This tool provides a framework to help marketing managers support and optimize their marketing strategies in order to maximize the return on their marketing investments. The CELM technology combines Monte Carlo simulations, portfolio optimizations, and advanced models of MDPs. MDPs are used to model customer dynamics and to find optimal marketing policies that maximize the value generated by a customer over a given time horizon. The CELM technology can also be used to simulate the financial impact of a given marketing policy using Monte Carlo simulations. This allows marketing managers to re-create several targeting scenarios in order to assess budget requirements and to predict the expected impact of a given marketing policy.

In summary, the IBM CELM methodology provides a framework to address practical marketing requirements in which both the customer lifetime value and the associated risks are optimized. Customer lifetime value is modeled using a robust MDP, and a risk-adjusted marketing budget allocation is achieved by using financial engineering techniques for portfolio diversification.

Three-step methodology

CELM involves a three-step methodology that can be summarized as follows and performed in the order suggested here:

- *Customer segmentation:* We introduce advanced value and loyalty metrics and enhance existing customer profiling in order to capture the value, loyalty, and response behavior of customers instead of focusing exclusively on transactions and, for example, mileage-based segmentations in the airline study.
- *Customer dynamics:* We identify customers' different lifecycle phases and dynamics using dynamic programming techniques (MDPs). We estimate customer lifetime value and risk (volatility) over variable time horizons by combining MDP models and Monte Carlo simulations in order to estimate the value-risk profile of customers.
- *Portfolio optimization:* We optimize the planning of campaign sequences for each customer profile in order to avoid saturation and cannibalization, in an effort to maximize the value of customers over a given planning horizon. We optimize marketing budget allocation in order to balance the value-risk tradeoff of the overall portfolio of customers using portfolio diversification techniques. The term *saturation*




Figure 1

Example of a customer purchase time history for the loyalty indicator (*LI*) calculation. To compute the *LI* at a given time t_* , the times between two consecutive purchases accomplished before t_* are averaged and divided by the recency. Each vertical bar represents a purchase by a particular customer.

refers to situations in which marketing campaign effectiveness has decreased significantly because a customer has been subjected to an overwhelming number of campaign appeals. The term *cannibalization* refers to situations in which marketing campaigns directed toward a customer compete for the customer's interest, resulting in a waste of marketing resources.

Customer segmentation

The first step in the CELM methodology requires us to define a set of advanced value and loyalty metrics and select the most suitable features to build a segmentation model. For instance, we have defined a metric to measure loyalty. This loyalty indicator (*LI*) can be described as follows:

$$LI(t_*) = \frac{\text{average}_{i < t_*}(\Delta t_i)}{\text{recency}(t_*)}, \quad (1)$$

which is simply the rate of the average inter-purchase times for the purchase recency. *Recency* is defined as the time between the last purchase time and the current time t_* , as illustrated in **Figure 1**. The average is computed for t_i up through but not including t_* .

LI is used as a loyalty index for a particular customer in a given time frame according to the customer's past purchase behavior, and *LI* is thus a relative loyalty indicator for each customer. As long as $LI > 1$, the customer is in a stage of continuing, or increasing, his or her normal purchase frequency. If $LI < 1$, the purchase frequency slows down and may reveal a process that leads toward defection.

Once the best features have been selected, one must usually discretize the high-dimensional customer feature space into a finite number of states. We propose a list of clustering (partitioning) criteria, which can be divided into two categories: scoring-based clustering and statistics-based clustering. The phrase *feature space*

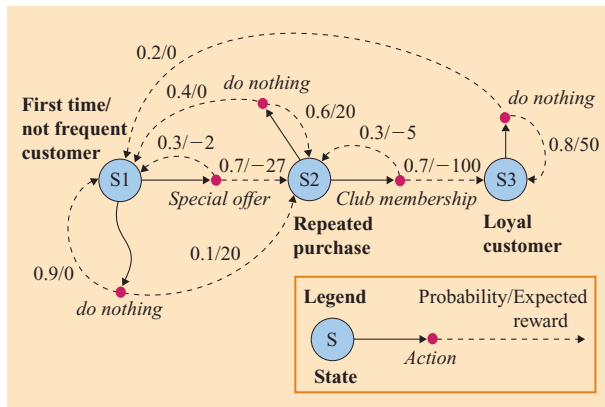


Figure 2

Example of customer dynamics modeling using an MDP. The customers are represented in three states. The effects of marketing actions are modeled by the transition probabilities and the expected values, which are usually referred to as *rewards* in dynamic programming literature. As indicated in the legend, red dots correspond to specific actions.

refers to the collection of metrics associated with every customer, such as demographics (e.g., customer age), transactional volume (e.g., purchase frequency), or marketing data (e.g., campaign response rate).

Scoring-based clustering

Scoring-based clusterings are usually obtained by discretizing recency, frequency, and monetary value separately and then combining the obtained intervals to form clusters, which are given a score. Each clustering criterion can have several parameters. The clusters are defined as follows:

- *RFM(n)*: The RFM method scores the customers according to recency, frequency, and monetary value and defines n intervals of equal size for each of the three metrics. The intervals are then combined to create a total of n^3 clusters. Each cluster is thus defined by an interval of values for recency, frequency, and monetary values.
- *ABC(a, b, c)*: This value scores customers according to value (e.g., the value in the previous three months) and generates three clusters by assigning the first $a\%$ of customers to cluster *A*, the next $b\%$ to cluster *B*, and the remaining $c\%$ to cluster *C*. Customers in clusters *A* and *B* usually account for most of the total generated value.
- *VD(a, b, c)*: This value-defector clustering performs *ABC(a, b, c)* partitioning on value and loyalty characteristics. All possible pairs of one value

ABC cluster and one loyalty *ABC* cluster form the *VD(a, b, c)* clusters (e.g., *AA*, *AB*, *AC*, and so forth).

- *RV(a, b, c)*: Recency-value clustering is similar to *VD* but considers value and recency.

Statistics-based clustering

The following are statistical-clustering-based state definitions that use all available customer characteristics:

- *Trees(n)*: Regression trees [32] are used for supervised clustering. A regression tree is trained to predict the immediate value of each customer. The leaves of the tree correspond to the clusters, which define the states. The parameter n indicates the number of leaves in the final tree. Each leaf represents a *region* in the characteristics space. The characteristics space (also called the feature space) is partitioned to minimize the value prediction error and consequently the variance in each region.
- *SOM(n, m)*: Self-organizing maps [33] allow one to map a high-dimensional characteristics space into a two-dimensional map (an $n \times m$ grid).
- *K-means(n)*: *K*-means clustering [34] derives n clusters, which are defined by the centers minimizing the total within-cluster variance.

Customer dynamics

An MDP consists of a set of states, actions, transition probabilities, and value functions. When an action is applied to a given state, the process moves stochastically to another state and generates a value (e.g., cash flow). The probabilities of moving to a target state (given the original state and the applied action) and the expected values are part of the model specification. In this way, a random sequence of states, actions, and values can be modeled, and the expected cumulative value that is associated with a given state under a specified *policy* (i.e., a mapping from state to actions) can be computed.

We model the customer behavior over time taking into account the marketing actions performed by the company. **Figure 2** depicts an example of an MDP that models the dynamics of customers subject to a set of marketing actions. In the example shown in **Figure 2**, customers are represented as states S1 (low-frequency customers), S2 (repeated-purchase customers), and S3 (loyal customers). According to the model defined in this example, if a special offer is sent to a customer in state S1, there is a 0.7 probability that the customer will respond to the offer. By responding, the customer will move to state S2 and generate a value of -27 (negative value, due to the cost of the promotion, which can be perceived as a long-term investment that would generate a higher value later

on). If customers in state S1 are not targeted, they have a lower probability of moving to state S2, but the company will not allocate marketing budget to promote a special offer.

At each time step, e.g., each month, the company can decide to target customers with the marketing actions allowed in each state. The action *do nothing* is explicitly modeled, since it represents a decision that may have a different effect on the relationship than, for instance, sending a special offer. Once all the states, the transition probabilities, and the expected values are known, it is possible to find the marketing policy that maximizes the expected long-term value generated by the relationship with the customers. For instance, assuming that customers in state S3 are very profitable (their immediate expected profit is $0.8 \times 50 = 40$), it can be appropriate to send a special offer to customers in state S1 and a club membership offer to those in state S2, even if the immediate expected values are negative. The expected value of sending a special offer to customers in state S1 is $-27 \times 0.7 - 2 \times 0.3 = -20.4$, whereas the expected value of sending a club membership offer to customers in state S2 is $-100 \times 0.7 - 5 \times 0.3 = -71.5$. Therefore, a *shortsighted* marketing policy would not send any offer to customers in states S1 and S2. Such a policy would not take into consideration the *long-term value* generated by upgrading customers to state S3. Once a marketing policy is fixed, the future customer dynamics can be simulated for a given time horizon, and the distribution of the future values can be estimated.

We assume that customer historical data is available and that customers are mapped into states using, for instance, any of the segmentation techniques described in the section on customer segmentation. For each customer, the historical data consists of a sequence of events. Each event is defined by a triplet (s, a, r) composed of a state s , an action a , and a value r . The next event is defined by the triplet (s', a', r') , where s' is a potential new state after action a is applied to s and so on. Each customer has an associated sequence of events, defined as an *episode*, that spans a given time window (see **Figure 3** for an example).

In order to completely specify the MDP, we need to define the state space, the action space, the transition probabilities, the expected values, and the decision epochs. Given historical customer data D , the state and action spaces are obtained by considering, respectively, all of the states and all of the actions (as defined above) that appear in D . In order to estimate the transition probabilities, we can simply use the maximum-likelihood estimator

$$p(s'|s, a) = \frac{\#(s'|s, a)}{\#(s, a)}, \quad (2)$$

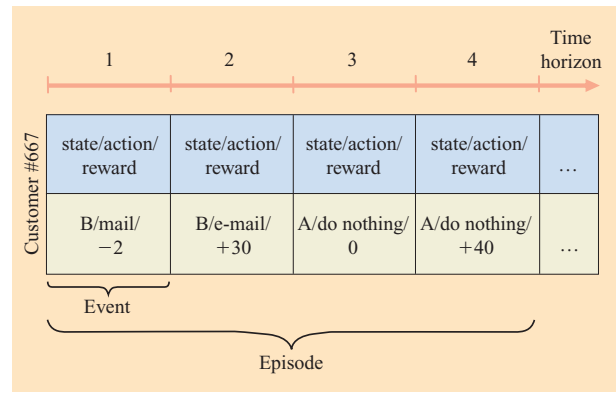


Figure 3

Example of an episode for a given customer. The initial customer state is B . After receiving a mailing, the generated value (reward) is -2 , and the next state is still B . Then an e-mail campaign causes the customer to move to state A , and a value of 30 is generated, and so on.

where $\#(s'|s, a)$ is the total number of transitions from s to s' if action a is applied, and $\#(s, a)$ is the total number of times action a is applied to state s . Because we are estimating the transition probabilities from a limited sample of data, we can assume that the absence of particular transitions does not necessarily imply that the real probabilities are zero or undefined. To address these two issues, we adopt a Bayesian approach incorporating a prior model of transition probabilities $\hat{p}_{s'|s,a}$ into Equation (1). This leads to the following estimator [35] for the posterior probability

$$p(s'|s, a) = \frac{\#(s'|s, a) + m_1 \hat{p}_{s'|s,a}}{\#(s, a) + m_1}. \quad (3)$$

It is easy to verify that, when $\#(s, a) = 0$, the posterior and prior probabilities are the same. The value of m_1 can be interpreted as the number of instances, following the prior probability, that are injected into the data set D . The variable m_1 acts therefore as a weight defining the relative importance of the prior probability with respect to the probability estimated from the data.

Two possibilities exist for modeling the prior transition probability $\hat{p}_{s'|s,a}$: a) adopt a *state-driven* approach, emphasizing the role of the origin state, or b) adopt an *action-driven* approach, emphasizing the role of the action. In the state-driven case, the prior transition probability is modeled as follows:

$$\hat{p}_{s'|s,a} = p(s'|s) = \frac{\#(s'|s) + m_2 \hat{p}_{s'|s}}{\#(s) + m_2}, \quad (4)$$

where $\#(s)$ is the number of times state s appears in the data set D and $\#(s'|s)$ is the number of times a transition

from state s to state s' is observed. Like m_1 , the variable m_2 acts as a weight that defines the relative importance of the prior probability with respect to the probability estimated from the data. Finally, the nested prior probability $\hat{p}_{s'|s}$ is estimated as follows:

$$\hat{p}_{s'|s} = p(s') = \frac{\#(s') + m_3 \hat{p}}{\sum_{s \in S} \#(s) + m_3}. \quad (5)$$

In the action-driven approach, the prior probability is modeled as follows:

$$\hat{p}_{s'|s,a} = p(s'|a) = \frac{\#(s'|a) + m_2 \hat{p}_{s'|a}}{\#(a) + m_2}, \quad (6)$$

where $\#(a)$ is the number of times action a appears in the data set D , and $\#(s'|a)$ is the number of times that a transition to state s' is due to the application of action a . The expected value $r(s, a, s')$, if action a is applied to state s and a transition to state s' occurs, can be estimated as follows:

$$r(s, a, s') = \frac{\sum_{(s,a,s') \in D} r_{(s,a,s')}}{\#(s, a, s')}, \quad (7)$$

where $r_{(s,a,s')}$ is the value observed in the data when action a is applied to state s and a transition to state s' occurs. If $\#(s, a, s')$ is zero or null, because action a has never been applied to the transition s to s' , we can estimate the expected value considering either a state-driven approach or an action-driven approach. The state-driven estimate is

$$r(s, a, s') = r(s, s') = \frac{\sum_{(s,s') \in D} r_{(s,a,s')}}{\#(s, s')}, \quad (8)$$

whereas the action-driven estimate is

$$r(s, a, s') = r(a) = \frac{\sum_{(a) \in D} r_{(s,a,s')}}{\#(a)}. \quad (9)$$

Finally, we need to define the decision epochs or stages, T , which are the points of time at which an action can effectively be applied to a state. We determine the stages according to the action planning of a specific marketing strategy.

Portfolio optimization

In the previous section we showed how to model an MDP and simulate a targeting policy over a given time horizon. So far, we have imposed no constraints on the total cost implied by such a policy. However, in practice, marketing plans are usually subject to budgeting constraints.

Because some policies cannot be applied to all customers, one has to determine the optimal number of customers

per state to be targeted at each decision time. To answer this question, we adopt the classical *mean-variance* formulation framework for portfolio optimization. The following definitions are used in order to formulate the optimization problem:

- α_{ast} is the number of customers to target in state s , at time t , using action a . We refer to these as the target or decision variables in the optimization problem.
- r_{ast} is the average value generated by customers in state s , at time t , when targeted with action a .
- σ_{ast}^2 is the variance of the value generated by customers in state s , at time t , when targeted with action a .
- c_{ast} is the average cost per customer in state s , at time t , when targeted with action a .

The mean and variance of the distribution of the value generated by an action a at time t can then be written as

$$R(a, t) = \sum_s \alpha_{ast} r_{ast} \quad \text{and} \quad \sigma^2(a, t) = \sum_s \alpha_{ast}^2 \sigma_{ast}^2. \quad (10)$$

The above definition assumes independence of values generated by different states at each decision time. This assumption can be motivated by the fact that customers are individuals whose responses to a targeted marketing action are independent of one another.

The budget-constrained optimization problem is essentially the finding of the optimal target variables α_{ast} that maximize the expected value $R(a, t)$ while minimizing the variance $\sigma^2(a, t)$ over a and t . The expected targeting cost is constrained to be below a user-defined budget B , and is represented as a linear constraint: $\sum_{a,s,t} \alpha_{ast} \leq B$.

Additional constraints related to the total size of the customer set are considered. These constraints can be divided into two constraint categories, as described below.

Customer dynamics constraints

These constraints represent the fact that the total number of customers across all states is constant. Any customer moving out of a state at time t should be in another state at time $t + 1$. For such a purpose, the set of possible states should include future prospects and inactive customers in order to account for new acquisitions and defections. These constraints are represented as

$$\sum_a \alpha_{as(t+1)} = \sum_{a,s'} \alpha_{as't} P(s|s', a) \quad \text{for } t > 0. \quad (11)$$

The term $\alpha_{as't} P(s|s', a)$ represents the number of customers expected to be in state s at time $t + 1$ after receiving action a and starting from state s' . The sum of these numbers over all actions and all previous states s'

should be identical to the number of customers in state s at time $t + 1$, which is equal to $\sum_a \alpha_{as(t+1)}$.

Initial state constraint

The previous set of inequalities would be ill-defined if there were no initial condition on the number of customers per state. Therefore, the number of customers in state s at time $t = 0$ is a constant in the optimization problem. This is simply measured by the initial distribution of customers across different states. This set of constraints is represented as

$$\sum_a \alpha_{as0} = |S_s|, \quad (12)$$

where $|S_s|$ is the number of customers in state s at time $t = 0$.

Formulation of optimization problem

Given the definitions and constraints mentioned above, we can formulate the following optimization problem:

$$\begin{aligned} \text{Maximize}_{a,s,t} \quad & \sum_{a,t} R(a, t) - \lambda \sigma^2(a, t) \\ & = \sum_{a,s,t} \alpha_{ast} r_{ast} - \lambda \sum_{a,s,t} \alpha_{ast}^2 \sigma_{ast}^2, \end{aligned}$$

subject to the following constraints expressed in equation form:

- $\sum_{a,s,t} \alpha_{ast} \leq B$
(Budgeting constraint: B is the total available budget);
- For all s , $\sum_a \alpha_{as0} = |S_s|$
(Initial state constraints: $|S_s|$ is the cardinality of the state s at time $t = 0$);
- For all s and for $t > 0$, $\sum_a \alpha_{as(t+1)} = \sum_{a,s'} \alpha_{as't} P(s|s', a)$
(Customer dynamics constraints).

The parameter λ controls the tradeoff between the mean (expected) and the variance of the value generated by a policy. Ranging between 0 and 1, the parameter λ models the degree of risk aversion in the optimization problem. A high value of λ penalizes the variance of the value generated by a policy (referred to as a conservative strategy), whereas a low value of λ favors the maximization of the expected generated value (referred to as a greedy strategy).

The solution of this problem provides a natural way to define a constrained marketing policy. Let us denote by α_{ast}^* the optimal value for the target variable. The optimal policy π^* is defined as $\pi^*(s, a, t) = (\alpha_{ast}^* / \sum_a \alpha_{ast}^*)$, where $\pi^*(s, a, t)$ is the fraction of customers in state s at time t to

target with marketing action a . This policy π^* can be compared with the historical policy π by replacing the optimal target variable α_{ast}^* with some value derived from historical data. The historical policy can be derived in different ways. For instance, if the marketing policy is assumed to be stationary, π can be estimated as follows:

$$\pi(a|s) = \frac{\#(a|s) + m \hat{\pi}_{a|s}}{\#(s) + m}, \quad (13)$$

where $\#(a|s)$ is the number of events (i.e., transactions) with state s and action a , and $\#(s)$ is the number of events with state s . The quantity $\hat{\pi}_{a|s}$ is the prior probability of applying action a to state s and can be estimated with Equation (14):

$$\hat{\pi}_{a|s} = p(a) = \frac{\#(a) + 1}{\sum_{a \in A} \#(a) + |A|}, \quad (14)$$

where $\#(a)$ is the total number of actions of type a in all of the events, and $|A|$ is the number of available actions.

It is important to note that the optimization provides only the optimal number of customers to be targeted in each state at each time. The customers who will actually be targeted in each state are selected using some scoring mechanism that can be driven by any type of criteria (e.g., customer demographics and probability of responding).

CELM technology

The IBM Zurich Research Laboratory and the IBM Center for Business Optimization implemented the CELM solution as a Java** application. The solution consists of the following main components:

- *CELM data model*: This component allows the selection of the customer historical data by connecting to various data sources in which the behavioral, financial, and demographic characteristics of customers are stored. Moreover, this component computes derived customer metrics such as loyalty indices, value, recency, and frequency.
- *State definition models*: This component discretizes a high-dimensional customer characteristics space into a finite number of customer states. A list exists that contains proposed partitioning algorithms, including statistical clustering algorithms. Moreover, the system allows us to import other partitioning criteria, which can be defined by the end user. If the user has existing state definitions, this step can be omitted.
- *Customer dynamics modeling*: Using some state definition, which can be derived in the previous step, an MDP is estimated that models the customer dynamics. By fixing a time horizon and using Monte Carlo simulations, the customer lifetime

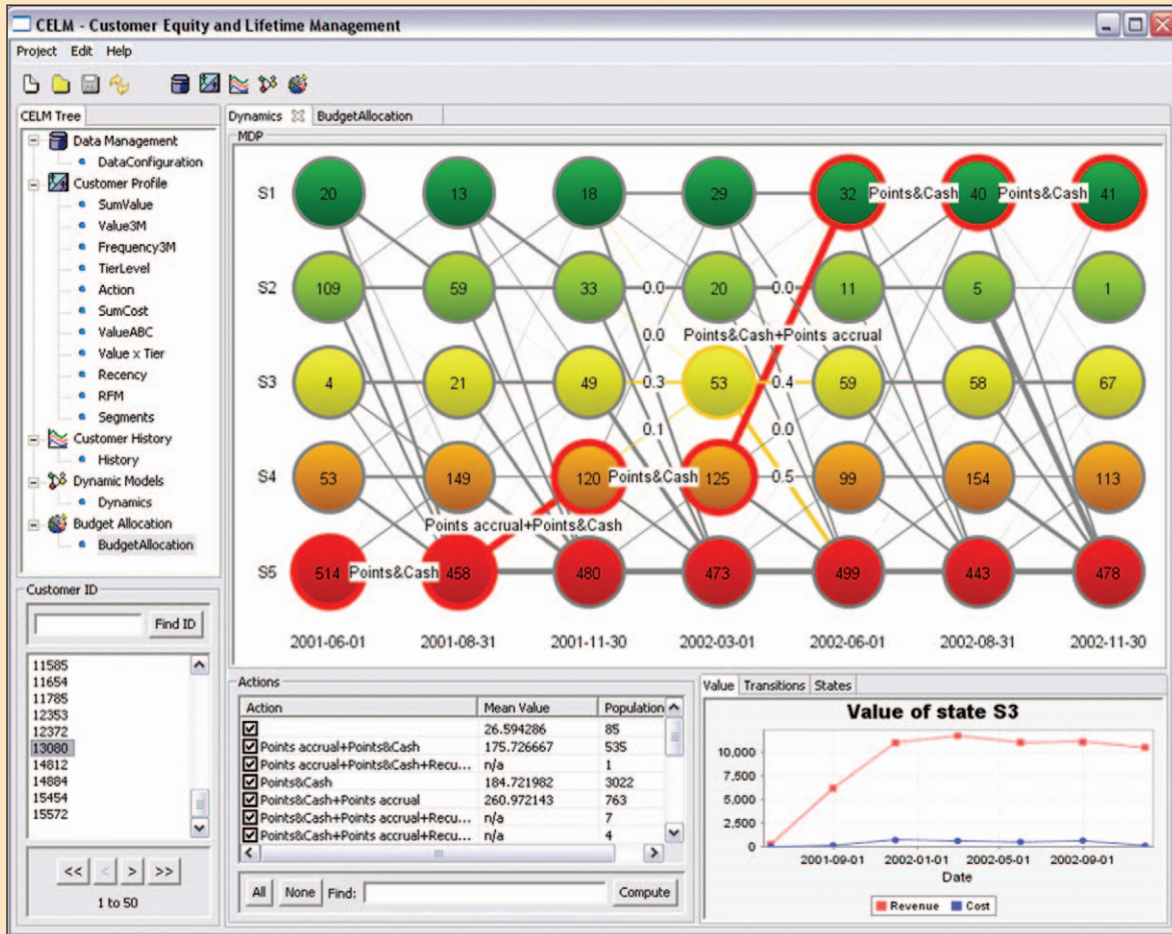


Figure 4

MDP of customer dynamics in CELM. The thick red line shows the trajectory followed by customer #13080 (selected in the left panel). The trajectory shows the states (colored circles), through which the customer has been moving at different time stages (labeled directly beneath the red circles), together with the marketing actions (labels between any two circles) that targeted such a customer.

value generated by the historical marketing policy is computed. Moreover, by using dynamic programming, the optimal marketing policy and the optimal customer lifetime value are estimated.

Figure 4 shows the CELM user interface with a graphical representation of a customer dynamics model.

- *Marketing budget allocation:* Using the MDP estimated in the previous step, a time horizon (e.g., 12 months), and a marketing policy, the financial profile of each customer state is derived by means of Monte Carlo simulations. The optimal customer portfolio that maximizes expected future value while minimizing risk (variance) is then derived by using the optimization framework described in the previous

section. The amount of marketing budget to invest, and the actions to be used, are derived for each customer state.

Business impact

Our pilot project with a leading European airline identified several strategic marketing processes to be optimized. For example, let us consider the customers in state S3, a subset of the upper-tier members who are highly loyal and are referred to as *Gold* and *Platinum* members. Figure 5 shows the distribution of marketing campaigns for customers in that state during a given time period. The pie chart on the left corresponds to the campaign distribution as observed in historical data, whereas the chart on the right shows the optimal

campaign distribution estimated by CELM. The optimal policy suggests sending no campaign material to about 60% of the customers in state S3, whereas the historical policy would have indicated omitting only 25% of customers. Reducing the frequency of such campaigns (e.g., *Points&Cash*, defined in the caption for Figure 5) results in significant cost savings and avoids customer saturation, because high-tier members of most loyalty programs tend to receive large amounts of campaign material. Despite their high flight frequency, these customers have low response rates to several kinds of campaigns, including so-called premium campaigns, which can be much more costly than standard campaigns. However, these customers respond strongly to collecting bonus points. Therefore, the optimal marketing policy produced by CELM targets these customers with a combination of *cost-awareness* marketing activities (omitting certain campaigns that have proved ineffective with this group), cross-selling campaigns such as *points accrual* (defined in the caption for Figure 5), and other programs such as all-inclusive travel packages. This combination allows these customers to earn and redeem large numbers of bonus points. The lifetime value of customers in state S3 is therefore optimized by minimizing targeting costs and maximizing cross-selling revenue.

With the use of CELM, the airline reported a significant impact on the planning and costs of its marketing campaigns and on the allocation of its marketing resources. The airline used CELM data to derive highly homogeneous FFP customer profiles to which it can most effectively direct its various marketing actions. CELM gave the marketing managers of airlines vital support in moving from a mileage-based to a value-based management of its frequent flyers. Consequently, the airline reported that marketing costs were reduced by more than 20% and response rates were improved by up to 10%. These excellent results convinced the airline to integrate CELM into the loyalty system that it is currently developing. Also as a result of this successful pilot project, CELM was awarded the 2005 INFORMS Marketing Society Practice Prize at the 27th Marketing Science Conference held at the Goizueta Business School (Emory University, Atlanta, Georgia).

Conclusion

In this paper, we have described our marketing optimization tool, the IBM Customer Equity Lifetime Management Solution (CELM), and shown how it was applied to optimize a leading European airline's customer relationship management processes within the company's frequent flyer program. The benefits provided by the CELM solution apply to segmentation,

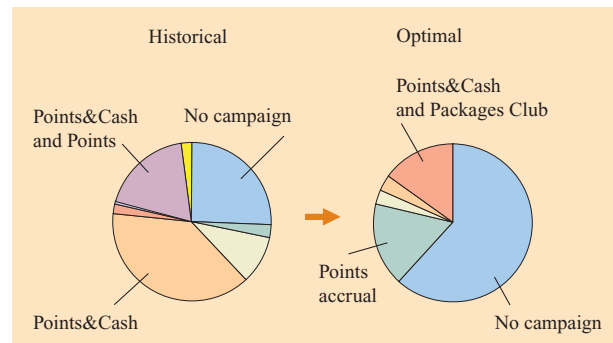


Figure 5

Comparison between historical and optimal marketing action distribution for state S3. The composite term “Points&Cash” refers to promotions on flights that combine cash and loyalty points. “Points” refers to promotions on flights using loyalty points only. “Packages Club” refers to travel packages that include the flight and an entertainment event (e.g., theater) or an overnight stay in a hotel. “Points accrual” refers to campaigns that increase the point balance—example, the offering of extra points if a car is rented.

customer dynamics, portfolio optimization and service differentiation in the following ways.

Segmentation

CELM helps companies define and/or redefine customer segments on the basis of specific value and loyalty metrics, in addition to traditional demographics and transactional metrics such as recency and frequency. The result is a collection of value–loyalty states that characterize key phases of the customer lifecycle.

Customer dynamics

Using CELM, enterprises can identify customer transitions (the probability of moving to a higher/lower value state); estimate customer lifetime value and risk (volatility); predict the impact of future marketing actions on customer lifetime value; identify optimal future marketing policies; and assess historical and optimized marketing policies. It is important to note that CELM estimates the optimal budgeting and scheduling of marketing plans over variable time horizons. This is in contrast to traditional campaign management systems, which manage single campaigns one at a time and usually do not take into account the so-called portfolio effect, i.e., the mutual dependencies and impact of various campaigns.

Portfolio optimization/service differentiation

CELM helps companies allocate their marketing resources in order to maximize return on investment and

minimize risk (uncertainty). Whereas this approach has been used primarily to plan and schedule marketing campaigns, the customer insights gained allow companies to better understand and manage their customer relationships. Through optimal diversification of marketing actions and targeted customer profiles, CELM helps marketing managers optimize the business performance of a company by improving marketing efficiency and effectiveness, optimizing budgets and resources for each customer segment over variable time frames, and maximizing the value/risk ratio across the customer portfolio.

Acknowledgments

We gratefully acknowledge the collaboration of Giuliano Tirenni and André Elisseeff, who were involved in all phases of the CELM solution, and Daniel Ramage, for his contribution to the software development. We also thank our colleagues Chris Rospenda, Kirsti Lindfors, and Tarja Ruuska from IBM Finland, and Caryn Bloom, Catherine Burton, and Tobin Cook from the IBM Center for Business Optimization.

*Trademark, service mark, or registered trademark of International Business Machines Corporation in the United States, other countries, or both.

**Trademark, service mark, or registered trademark of Sun Microsystems, Inc. in the United States, other countries, or both.

References

1. J. Nail, *Where Is Marketing Measurement Headed?*, Forrester Research, Inc., Cambridge, MA, January 2005; see <http://www.forrester.com/Research/Document/Excerpt/0,7211,33836,00.html>.
2. R. Srivastava and D. J. Reibstein, "Metrics for Linking Marketing to Financial Performance," *Marketing Science Institute Special Report*, Marketing Science Institute, Cambridge, MA, Winter 2005, pp. 85–109.
3. E. W. Anderson, C. Fornell, and D. R. Lehmann, "Customer Satisfaction, Market Share and Profitability: Findings from Sweden," *J. Marketing* **58**, No. 3, 53–66 (1994).
4. E. W. Anderson, C. Fornell, and R. T. Rust, "Customer Satisfaction, Productivity, and Profitability: Differences Between Goods and Services," *Marketing Sci.* **16**, No. 2, 129–145 (1997).
5. R. N. Bolton, "A Dynamic Model of the Duration of the Customer's Relationship with a Continuous Service Provider: The Role of Satisfaction," *Marketing Sci.* **17**, No. 1, 45–65 (1998).
6. R. T. Rust, A. J. Zahorik, and T. L. Keiningham, "Return on Quality (ROQ): Making Service Quality Financially Accountable," *J. Marketing* **59**, 58–70 (1995).
7. R. T. Rust, V. Zeithaml, and K. Lemon, *Driving Customer Equity: How Customer Lifetime Value Is Reshaping Corporate Strategy*, Free Press, New York, 2000.
8. V. Mittal, E. W. Anderson, A. Sayrak, and P. Tadikamalla, "Dual Emphasis and the Long-Term Financial Impact of Customer Satisfaction," *Marketing Sci.* **24**, No. 4, 544–555 (2005).
9. S. M. Shugan, "Brand Loyalty Programs: Are They Shams?" *Marketing Sci.* **24**, No. 2, 185–381 (2005).
10. G. L. Lilien and A. Rangaswamy, *Marketing Engineering*, Prentice Hall, Upper Saddle River, NJ, 2003.
11. S. Gupta, D. R. Lehmann, and J. A. Stuart, "Valuing Customers," Paper No. 03-08, HBS Marketing Research, Boston, MA, 2003.
12. R. T. Rust, K. Lemon, and V. Zeithaml, "Return on Marketing: Using Customer Equity to Focus Marketing Strategy," *J. Marketing* **68**, 109–127 (2004).
13. R. T. Rust and P. C. Verhoef, "Optimizing the Marketing Interventions Mix in Intermediate-Term CRM," *Marketing Sci.* **24**, No. 3, 477–489 (2005).
14. G. Tirenni, *Allocation of Marketing Resources to Optimize Customer Equity*, Ph.D. Thesis, University of St. Gallen, Switzerland, 2005.
15. S. Rosset, E. Neumann, U. Eick, N. Vatnik, and S. Idan, "Lifetime Value Models for Decision Support," *Data Mining & Knowledge Discovery J.* **7**, No. 3, 321–339 (2003).
16. J. H. Drew, D. R. Mani, A. L. Betz, and P. Datta, "Targeting Customers with Statistical and Data-Mining Techniques," *J. Service Res.* **3**, No. 3, 205–219 (2001).
17. K. Gelbrich and R. Nakhaeizadeh, "Value Miner: A Data Mining Environment for the Calculation of the Customer Lifetime Value with Application to the Automotive Industry," *Proceedings of the Eleventh European Conference on Machine Learning*, Springer-Verlag, New York, 2000, pp. 154–161.
18. E. Pednault, N. Abe, B. Zadrozny, H. Wang, W. Fan, and C. Apte, "Sequential Cost-Sensitive Decision Making with Reinforcement Learning," *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Edmonton, Alberta, Canada, 2002, pp. 259–268.
19. G. Tirenni, A. Labbi, A. Elisseeff, and C. Berrospi, "Efficient Allocation of Marketing Resources Using Dynamic Programming," *Proceedings of the Fifth SIAM International Conference on Data Mining*, Newport Beach, CA, 2005, pp. 581–585.
20. D. Jain and S. S. Singh, "Customer Lifetime Value Research in Marketing: A Review and Future Directions," *J. Interactive Marketing* **16**, No. 2, 34–46 (2002).
21. R. C. Blattberg, G. Getz, and J. S. Thomas, *Customer Equity: Building and Managing Relationships as Valuable Assets*, Harvard Business School Press, Boston, MA, 2001.
22. P. D. Berger and N. I. Nasr, "Customer Lifetime Value: Marketing Models and Applications," *J. Interactive Marketing* **12**, No. 1, 17–30 (1998).
23. G. R. Bitran and S. V. Mondschein, "Mailing Decisions in the Catalog Sales Industry," *Manage. Sci.* **42**, No. 9, 1364–1381 (1996).
24. F. Gönül and M. Z. Shi, "Optimal Mailing of Catalogs: A New Methodology Using Estimable Structural Dynamic Programming Models," *Manage. Sci.*, **44**, No. 9, 1249–1262 (1998).
25. P. Pfeifer and R. Carraway, "Modeling Customer Relationships as Markov Chains," *J. Interactive Marketing* **14**, No. 2, 43–55 (2000).
26. W. K. Ching, M. K. Ng, K. K. Wong, and E. Altman, "Customer Lifetime Value: Stochastic Optimization Approach," *J. Oper. Res. Soc.* **55**, 860–868 (2004).
27. S. Gupta and D. R. Lehmann, "Customers as Assets," *J. Interactive Marketing* **17**, No. 1, 9–24 (2003).
28. P. Doyle, "Value-Based Marketing," *J. Strategic Marketing* **8**, No. 4, 299–311 (2000).
29. R. K. Srivastava, T. A. Shervani, and L. Fahey, "Market-Based Assets and Shareholder Value: A Framework for Analysis," *J. Marketing* **62**, No. 1, 2–18 (1998).
30. D. G. Luenberger, *Investment Science*, Oxford University Press, New York, 1997.
31. R. Brealey and S. Myers, *Principles of Corporate Finance*, McGraw-Hill, New York, 1996.
32. L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen, *Classification and Regression Trees*, Wadsworth International Group, Belmont, CA, 1984.

33. T. Kohonen, *Self-Organizing Maps*, Springer, New York, 2006.
34. T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning*, Springer-Verlag, New York, 2003.
35. T. Mitchell, *Machine Learning*, McGraw-Hill, New York, 1997.

Received September 18, 2006; accepted for publication October 14, 2006; Internet publication May 17, 2007

Abdel Labbi *IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (abl@zurich.ibm.com)*. Dr. Labbi is CRM Research Leader at the IBM Zurich Research Laboratory. His research, teaching, and consulting activities in the last 15 years have been focused on quantitative approaches to the optimization of various business processes, especially in the CRM and risk management areas for several industries. He has published more than 40 papers in international business and scientific journals and conferences and holds several patents, including a patent for the IBM Customer Equity Lifetime Management Solution (CELM), which was awarded the 2005 Best Practice Prize of the INFORMS Society for Marketing Science (ISMS).

Cesar Berrospi *IBM Research Division, Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (ceb@zurich.ibm.com)*. Mr. Berrospi is a researcher at the IBM Zurich Research Laboratory. He has also worked as a consultant at Accenture in Barcelona, Spain. He holds an M.S. degree in mathematics from the Technical University of Catalonia, Spain, and an M.S. degree in mathematical engineering from Ecole Polytechnique, France, and Ecole Polytechnique Fédérale de Lausanne, Switzerland. Mr. Berrospi has received an IBM Research Division Award and the 2005 INFORMS Society for Marketing Science Practice Prize. His research interests and publications are primarily in the areas of data mining, customer segmentation, campaign management, and marketing portfolio optimization.