

“Adaptive” Learning and “Proactive” Customer Relationship Management

Baohong Sun¹
Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA15213
Tel: 412-268-6903
Fax: 412-268-7357
Email: bsun@andrew.cmu.edu

Shibo Li
Kelley School of Business
Indiana University
1309 E. 10th Street
Bloomington, IN 47405
Phone: 812-855-9015
Fax: 812-855-6440
Email: shili@indiana.edu

Catherine Zhou
Accenture
Spear Street Tower
Suite 4200, One Market
San Francisco, California 94105
Tel: 1 (415) 537.5000
Fax: 1 (415) 537.5037
Email: iwifwhy999@yahoo.com

January 2006

¹ Baohong Sun is Associate Professor of Marketing at Carnegie Mellon University. Shibo Li is Assistant Professor of Marketing at Indiana University. Catherine Zhou is Manager of the Customer Insight Group at Accenture.

Adaptive Learning and Proactive Customer Relationship Management

Abstract

CRM is about introducing the right product to the right customer at the right time through the right channel to satisfy the customer's evolving demands. Ideally, it should follow the development of each individual customer and develop integrated multi-segment, multi-stage, and multi-channel CRM decisions in order to maximize the total customer lifetime profit. However, most existing CRM practice and academic research focuses on methods to select the most profitable customers for a scheduled CRM intervention. This campaign-centric approach deviates from the goal of customer-centric CRM.

In this article, we discuss the two-step procedure ("adaptive" learning and "proactive" CRM decisions) and three-components for customer-centric CRM, *adaptive learning* (of customer individual preference), *forward-looking* (into future marketing consequences of current CRM interventions), and *optimization* (to optimally balance cost and benefit). We then formulate CRM interventions as solutions to a stochastic dynamic programming problem under demand uncertainty in which the company learns about the evolution of customer demand as well as the dynamic effect of its marketing interventions, and make optimal CRM decisions to balance off the cost of interventions and the long-term payoff with the goal of maximizing each customer's "long-term" profit. The framework allows us to integrate all the inter- and state-dependent factors that drive the CRM decisions and results in inter-temporally related path of CRM solutions that are consistent with customer-centric CRM. Finally, we choose two examples to demonstrate the input, output, and benefit of "adaptive" learning and "proactive" CRM.

The proposed solution meets the recent trends of companies seeking real-time solutions for integrating database and CRM decisions that are empowered by the advancement of technology.

Keywords: Customer-centric CRM; CRM intervention; Adaptive Learning; Forward-looking; Optimization; Proactive CRM; Dynamic Programming Approach; Dynamic Control Problem.

1. Introduction

The tools and technologies of data warehousing, data mining, and other customer relationship management (CRM) techniques provide greater opportunities for today's companies to establish and sustain long-term relationships with their customers than ever before (Winer 2001). The ultimate goal is to transform these relationships into greater profitability by improving the effectiveness of CRM programs, increasing customer loyalty and purchase probability, and lowering the cost of serving, thereby increasing profitability. Realizing the increasing importance of customer orientation, companies from all types of industries, ranging from manufacturing to information, are exploring relationship building as a promising means of differentiation, competition, and revenue growth opportunities (Sawhney, Balasubramanian and Krishnan 2004). In addition, contemporary practice of CRM has been integrated into every step of the marketing process - telemarketing, advertising, transaction, service, and survey. Furthermore, the traditional process of mass marketing is being challenged by the new approach of interactive marketing (Blattberg and Deighton 1991, Haeckel 1998) or one-to-one marketing (Peppers, Rogers, and Dorf 1999). Companies focus on the depth of each customer's needs and endeavor to establish a long-term relationship with each customer.

Most of the current CRM practices are campaign-centric in the sense that they focus on methods to select the most desirable customers for a scheduled CRM intervention. However, by definition, CRM should be customer-centric. It requires the company to develop the right interventions for the right customer at the right time through the right channel to meet the customer's need. To achieve this, the company needs to develop detailed customer knowledge, follow the development of each individual customer, adopt CRM interventions that are relevant to the status and preference of each individual customer. These steps help build a stronger one-on-one relationship. The resulting CRM solutions should be an integrated sequence of multi-segment, multi-stage, and multi-channel CRM decisions with the goal of maximizing the total customer lifetime profit.

Complete CRM decisions should follow a two step process: (1) learning of customer insights and (2) determining and executing the best CRM action. There are three important components in CRM decisions that are needed to support the two step process. The first component is *adaptive learning* to extract hidden predictive information from large databases to

identify valuable customers, learn about their preference, predict future behaviors, and estimate customer value. The second component is *forward-looking* into future marketing consequences of current CRM interventions, which is crucial for making a choice among alternative CRM intervention decisions. The third component is *optimization* from which an integrated sequence of CRM solutions can be explicitly derived by integrating all the driving factors and optimally balancing them between their short-term costs and long-term benefits.

[Insert Table 1 About Here]

There is an abundance of research on CRM in the marketing literature. The recent articles by Winer (2001), Rust and Chung (2005), and Kamakura et al. (2005) give excellent reviews of existing marketing models of service and customer relationship management. In Table 1, we classify recent research on database marketing and CRM based on whether the three components are modeled and whether the resulting CRM decisions are integrated multi-segment, multi-stage, and multi-channel solutions. There is research on building consumer models to predict product-next-to-be-purchased and lifetime value (Edwards and Allenby 2003; Li, Sun and Wilcox 2005; Kamakura, Ramaswami and Srivastava 1991; Kamakura, Wedel, de Rosa and Mazzon 2003; Reinartz and Kumar 2000; Reinartz and Kumar 2003) or compare customer satisfaction and loyalty online vs. offline (Shankar, Smith, and Rangaswamy 2003). Assuming the company's CRM decisions are given, this stream of research focuses on developing customer response models to profile the heterogeneity of customers. The CRM decisions are implicitly discussed (not explicitly derived) based on the parameters estimated from the consumer model. Recently, several studies treat companies as decision makers and show that it is important to take into account the future consequences of current marketing mix decisions. This stream of research attempted to solve for a company's optimal decisions such as mailing of catalog and relationship pricing in order to improve the company's profits (Bult and Wansbeek 1995; Bitran and Mondschein 1996; Gönül and Shi 1998; Elsner, Krafft and Huchzerneier 2004; Lewis 2005). For example, Venkatesan and Kumar (2004) develop a dynamic framework that enables managers to improve customer relationship proactively through marketing contacts across various channels and to maximize customer lifetime value. Assuming a myopic company, Rust and Verhoef (2005) derive the optimal marketing interventions mix in intermediate-term CRM. Although an optimization

component is allowed, this stream of research either lacks the adaptive learning component or misses the forward-looking components that are crucial for integrated multi-segment, multi-stage, and multi-channel solutions.

From the above discussion, we see that the existing literature on CRM can be improved by either adding the execution step or integrating the three components of complete CRM solutions. In this paper, we aim to contribute to the literature by formally illustrating the ideas of adaptive learning and proactive CRM that are recently observed in scattered academic discussions (e.g. the idea of proactive marketing contacts in Venkatesan and Kumar 2004). Specifically, we discuss the concepts of *adaptive learning*, *forward-looking*, and *optimization* and point out the importance of incorporating them to transform CRM decisions from being campaign-centric to being customer-centric. We then formulate CRM interventions as solutions to a stochastic dynamic programming problem under demand uncertainty in which the company learns about the evolution of customer demand, the dynamic effect of its marketing interventions, the heterogeneity of customer preferences, the cost of interventions, and the long-term payoff, and make optimal CRM decisions to maximize “long-term” profit of each customer. This framework allows us to integrate all the inter- and state-dependent factors that drive the CRM decisions and results in inter-temporally related path of CRM solutions that are consistent with customer-centric CRM. We then use two examples to demonstrate the input, output, and benefit of “adaptive” learning and “proactive” CRM. Finally, we briefly discuss the possible automated real-time implementation of our proposed CRM solution, that is enabled by recent technology advancement (Rust and Chung 2005; Winer 2001).

2. Adaptive Learning and Proactive CRM

2.1. Adaptive Learning

For CRM to be customer-centric, the first necessary step is to learn about the evolving needs and preferences of individual customer. Based on accrued information on customer history as well as the feedback obtained from the last executed CRM decision, the company should be able to continuously learn and improve the accuracy of its knowledge on each individual. To be more specific, the ideal learning should have the following properties: (1) the accrued information is used to continuously to update the company’s knowledge of the customer’s preferences; (2) the

company's strategic decision is adapted according to the updated knowledge; and as a result, (3) the company can revise its belief in the next period based on the successful and unsuccessful interactions with the customer. We term this type of learning as "adaptive learning."² Adaptive learning offers the company the opportunity to learn about customer preferences and adapt its strategies in a real-time fashion. It is an important class of learning algorithms in a stochastic environment.

With the improvement of the accuracy of the company's knowledge of customer preferences, customers are better served and may be more likely to stay with the company. In addition, by immediately adapting its CRM interventions to the most updated knowledge on customer, the company can improve the effectiveness of its interactions with its customers. Thus, by following the foot steps of a customer's development, adaptive learning is more aligned with the idea of "relationship learning." Should the business environment change such as intensifying competition and increasing customer sensitivity to price or service quality, the adaptive learning ensures the incorporation of the new nature of customer response.

What need to be learned adaptively?

First, one purpose of developing relationship is for the company to cross-sell additional products and services to existing customers. To achieve this, companies need to learn about the development of individual customer demand. As demonstrated by Kamakura, Ramaswami and Srivastava (1991) and Li, Sun and Wilcox (2005), customers' demand for various products is governed by a latent and evolving demand state or maturity, which develops over time with change of life-stages, accumulation of consumption experience, available financial resources, learning of a particular product, etc. The evolving—but latent—demand maturity represents an individual customer's readiness for a particular product at a certain time. It is an important predictor for products that are most likely to be purchased at a certain time by a particular customer. Accurate knowledge on the development of each individual customer's demand is crucial for improving the targetability and effectiveness of cross-selling campaigns.

Second, it is important to learn the status a customer is in and the differential roles of CRM interventions at different stages of customer development. In addition to generating instantaneous

² A similar idea has been adopted in conjoint analysis to reveal customer preferences as demonstrated by Toubia, Simester, Hauser and Dahan (2004).

sales as for the frequently purchased products, CRM campaigns serve an additional function of interacting with customer development of maturity, educating customers, and cultivating customers' needs. It is important to understand the status of a customer in order to know when to send educational campaigns and when to send promotional campaigns.³ In educating the customers, the company can provide information on which products meet customers' needs even before they know it themselves. For example, a company could use data mining to characterize the behavior surrounding retirement, and send out informational campaigns to those who are about to retire to prepare them for the products in the near future.

Third, it is important to predict customer lifetime value because it relates directly to customer revenue, cost of acquisition, and customer profitability. The term "customer lifetime value" refers to the potential revenue obtained from a customer during his or her relationship with a company. Companies invest in customer relationships in order to increase long-term customer revenue in (at least) three ways: (1) increase their use or purchases of products they already have; (2) sell them more or higher-margin products; and (3) keep the customers for a longer period of time. In addition, managers need to track and compute the cost of acquiring each customer and then relate this cost to the profits the customer produces over his or her lifetime. Knowing future profitability serves two functions: to calculate future discounted value of existing customers without intervention (as documented by customer lifetime value analysis in most existing marketing literature) and to estimate the impact of current CRM intervention on future customer value (which we will elaborate more in the session on proactive CRM).

Fourth, customers may have different preference for communication channels as shown by Kumar and Venkatesan (2005). Since the 1990s, companies have been racing to add 24-hour call centers, direct mail, email, fax, and webpage as new service channels. Among them, automated service channels such as Internet access, voice-recognition phone systems, and transaction kiosks are given special emphasis to encourage self-service. Each customer may have a different preference in every step of their decision process, such as information search, purchase, transaction, and post-purchase service. (e.g., Ansari, Mela and Neslin 2005 and Sullivan and

³ For example, Li, Sun and Montgomery (2005) develop a stochastic dynamic optimization framework with hidden Markov process to capture the evolution of customer financial maturity. They recognize the educational role of cross-selling campaigns by allowing them to affect the development of financial maturity in addition to triggering immediate purchase as sales promotions. Applying to cross-selling financial products, they provide empirical evidence that these campaigns educate customers about what products can satisfy their financial needs or even how to advance their financial needs.

Thomas 2004) How do we successfully blend the various functions of multiple communication channels? How do we steer customers to their most preferred channels? How do we direct the self-sufficient customers to self-learning channels? Given the high cost of customer communication, it is important to learn the unobservable customer channel preference and to determine the optimal allocation of resources over multiple channels in order to improve the effectiveness and efficiency of channel mix strategies for CRM.

Fifth, it has been shown that customer sensitivities to companies' marketing variables such as price and quality change over time (Chen, Sun and Singh 2006). This is particularly true for managing long-term dynamic relationships. In a long-term relationship, customers change their sensitivities to explanatory variables such as product offerings, prices, promotions, and channels. It is important to learn these changes in sensitivities in order to offer more customized products, targeted deals, and discounts.

Difference between “passive learning” and “adaptive learning”

In campaign-centric CRM, various segment approaches such as latent class models are adopted to identify desirable customers. These segmentation models help econometricians to build a snapshot profile of a customer's heterogeneous sensitivities to marketing variables. We term this type of learning as “passive learning” because segmentation is based on pooled historical data and inference are made in an *ad hoc* fashion. Taking cross-selling as an example, most existing models are consumer models aimed at predicting the next-to-be-purchased product instead of company decision models solving for the best product to be cross-sold (Edwards and Allenby 2003). The underlying assumption is that customers with similar demographics should own similar products. Products to be cross-sold are recommended based on the purchasing behavior of all other customers. Similarly, the book recommendation system on Amazon.com is based on a giant matrix characterizing the correlations among millions of products. The recommendation is based on cross-product comparison. In either case, the path of an individual's purchase history is ignored and the interaction of CRM recommendations with customer development is not taken into account.

Adaptive learning differs from passive learning in the following ways. First, adaptive learning refers to the case of a company being defined as a decision maker and actively gaining

knowledge about customers instead of the econometricians learning about customer heterogeneity, as in the passive learning. Models with passive learning do not treat companies as CRM decision makers. Second, adaptive learning gains customer knowledge based on the development path. On the contrary, passive learning relies more on a snapshot cross-customer comparison (e.g. the latent class models separate customers based on the similarity of customer sensitivities to explanatory variables) or cross-product comparisons. Third, adaptive learning occurs is learning about each individual customer in real-time fashion. This is as opposed to passive learning which happens after the event, and the gained knowledge on customers is assumed to hold for a different sample or the same sample at a different time. Fourth, the company continuously updates its beliefs on customer preferences according to the feedback obtained from the last executed decision. This allows the company to take into account the effect of each CRM interventions on the path of customer development. However, the knowledge gained from passive learning is static and, when applied across the sample, is subject to the endogeneity bias caused by ignoring the interdependence between the company's CRM interventions and customer reaction (Rust and Chung 2005).

2.2 Proactive Customer Relationship Management

While gaining advanced customer knowledge provides the foundation for CRM decisions, developing and executing CRM programs to act upon customer insights is the ultimate step to successful CRM. Determining the best action that follows the learning of customer knowledge is where the hard work lies.

Forward-looking into the future consequences of current CRM interventions

In relationship management, each CRM intervention changes future customer behavior such as customer retention, purchase, and profitability. When making decisions on which CRM intervention to choose, the company needs to take into account the dynamic effect of each considered CRM intervention on future profit. To achieve this, the company needs to be treated as a forward-looking decision maker which forms expectations on future customer behavior, predicts consequences associated with each possible CRM intervention, calculate profit implications of

current CRM interventions, and chooses the CRM intervention that maximizes the sum of discounted future profits. The choice of the best intervention is made by comparing the sum of current and future profits achieved by alternative CRM interventions. Being able to forward look into future marketing consequences allows the company to make dynamic decisions to act upon long-term marketing consequences.

Note that this forward-looking component is different from most of the existing literature on customer lifetime value (LTV) analysis. LTV analysis calculates the static discounted present value of customer profit without taking into account the future consequences of CRM interventions. The goal of this analysis is to treat the value as another segmentation variable to differentiate profitable customers from unprofitable ones to guide targeting strategies. It is again a campaign-centric approach. As pointed out by Rust and Chung (2005) and Rust and Verhoef (2005), this approach is subject to the endogeneity problem that a company's intervention changes customers' future purchase probabilities.

Optimization

To make the optimal choice among alternative CRM interventions, the company needs to integrate its knowledge on the heterogeneity of customer preferences and the evolution of customer demand maturity, predict the dynamic consequences of CRM interventions, estimate the cost of acquisition and long-term financial payoff, and trade off short-term intervention cost and long-term financial payoff. The above listed factors drive CRM decisions and are also driven by each CRM interventions. Also, the idea of customer-centric CRM requires all the CRM interventions to be inter-temporally related over time. This is a complicated problem that can not be easily handled by traditional consumer models or regression. Given the nature of the problem, treating CRM decisions as solutions to a stochastic optimization problem under demand uncertainty may be the best or even most parsimoniously way to obtain the answer. The demand uncertainty comes from latent customer demand and preference, especially unknown future customer behavior, which can be reduced by adaptive learning.

3. An Integrated Framework

Tools of adaptive learning

Adaptive learning is more like an idea than a technique. The idea of adaptive learning is similar to Bayesian updating that accruing customer information is used to update the company's knowledge on customers. Thus, all the consumer learning models existing in the marketing literature can be modified to be useful tools.

The data mining and machine learning can also serve as base for adaptive learning. These techniques combine algorithms developed in computer science with statistical techniques to automatically search for patterns in large data sets. In the last decade, its commercial use has caught on in a big way due to the availability of large volumes of data, dramatic improvement of computing power, and greatly increased industrial competition. Nowadays many software providers offer data-mining tools: Angoss, BusinessMiner, Data Distilleries, IBM, Megaputer, Quadstone, Urban Science, Visual Insights, SAS, and SPSS. However, data mining techniques are still passive learning unless they are modified to have the three characteristics specified in the definition of adaptive learning.

Framework of Proactive CRM

Methodologically, we formulate CRM interventions decisions as solutions to a stochastic dynamic control problem under demand uncertainty with built-in customer reactions (as represented by $E[Q_{ij\tau} | INFO_{ij\tau}, Z_{ij\tau}]$). As we discuss before, the adaptive learning, forward-looking, and optimization components allow us to derive an integrated sequence of multi-step, multi-segment and multi-channel CRM interventions decisions about *when* to contact *which* customer with *what* product or content using which communication channel (*how*).

$$(1) \quad \underset{Z_{ijk\tau}}{Max} \quad \sum_{i=1}^I \sum_{\tau=t}^T \delta^{\tau-t} \sum_{j=1}^J E[Q_{ij\tau} (PRICE_{ij\tau} - MC_{ij\tau}) | INFO_{ij\tau}, Z_{ij\tau}]$$

$Z_{ijk\tau}$ = Marketing intervention k such as price or cross-selling campaign

δ = Time-discounting variable

$E[.]$ = Expected profit

$Q_{ij\tau}$ = Quantity of product j being purchased by customer i at time τ including no-purchase

$PRICE_{ij\tau}$ = Price paid by customer i for product j at time τ

$MC_{ij\tau}$ = Marginal cost of providing product j at time $t\tau$ for customer i

$INFO_{ij\tau}$ = Information available to the company about customer i for product j at time τ

[Insert Figure 1 About Here]

In Figure 1, we set up the proactive CRM problem: At the beginning of time t , the company observes information set $INFO_{it}$ consisting of past history and demographic variables. Based on the accrued information, the company updates its belief on a customer according to a pre-specified adaptive learning rule. Given this updated knowledge on a customer, the company calculates the total expected cost and revenue resulting from each possible CRM decision for all time periods from time t and further. With expected profit being defined by the above equation, the company chooses the best CRM decision that maximizes the expected long-term profit.

By solving the dynamic programming problem, we can obtain a sequence of optimal CRM intervention decisions. The derived CRM solutions possess the following properties. First, the solution is *customized* because adaptive learning allows the company to continuously learn about the heterogeneous preferences of each customer and make customized offers to better match customer preferences. Second, the decisions are *forward-looking* because the company takes into account the dynamic consequences of current CRM intervention decisions. Third, the marketing intervention is also *interactive* because it follows each stage of customer demand maturity and intervenes with the most appropriate marketing tool. Fourth, it is also *dynamic* because the company is allowed to sacrifice short-term profit by incurring higher cost to offer better service to a potential profitable customer in order to increase her chance of staying and for future profit. Fifth, the framework allows the company to be *experimental* and learn about customer preference by sampling. For example, Amazon.com can experimentally recommend a few books to a sample of buyers and measure their responses for the purpose of learning about their preference and evaluating current recommendation system. At a small and immediate cost, the experimentation allows the company to gain better knowledge on customers and improve its recommendation system (this is the so called “bandit problem”). Finally, the framework results in a sequence of *consistent* actions over time. The adjustment of the CRM interventions over time stems from the improvement of the accuracy on customer knowledge.

Given the above stated features, we term the CRM solution derived from the above framework as “proactive” CRM. This is because our framework integrates the three components of customer-centric CRM which allows the company to actively follow the development of each customer, forecast the effect of today’s marketing intervention on future profitability, and take preventative and interactive actions to improve the customer’s experience and relationship with the company. Our framework is more consistent with the ideal CRM intervention models described in the review paper by Rust and Chung (2005) that “(1) have marketing intervention levels personalized for each customer; 2) consider the effect of multiple marketing interventions; 3) maximize customer lifetime value, at least up to an arbitrary time horizon; and 4) fully address the endogeneity issue.”

The proposed framework is different from existing studies on CRM in the following ways. First, the existing literature on CRM focuses on modeling customer decision process with *ad hoc* segmentation approaches assuming company’s decisions are given and verbally discuss implications on company’s CRM intervention decisions. Allowing the customer’s decisions to be affected by the company’s decisions, the framework of proactive CRM treats the company as a decision maker and explicitly derives its optimal decisions. Second, most current research emphasizes developing better approaches to model customer heterogeneity, in which segmentation is based on the pooled historical data and inference is made in an *ad hoc* fashion. We formulate the idea of adaptive learning in which accruing customer information is adopted and integrated into the company’s period decisions. As a result, the company continuously updates its belief on customer preferences according to the feedback obtained from the last executed decision. Third, we treat the company as a forward-looking decision maker which takes into account the long-term profit implications of customer attrition when making a current decision. The future consequences are built into the derived optimal decisions. This is different from most of the existing literature on customer lifetime value analysis that calculates the net present value of customers’ future profit and treats the value as another segmentation variable to guide targeting strategies. Fourth, since the customer decision model and the company’s optimization model are solved simultaneously, this framework allows customer reaction and company’s CRM interventions decisions to be inter-dependent. This mitigates the endogeneity problem in existing research on CRM.

There are a few things worth mentioning regarding applying the proposed framework to real business problems. First, each of the components of the proposed general framework needs to be modified to make it more practical. For example, how many future periods a company should consider depends on the nature of the business environment and the company's goal. Suppose the average customer tenure is about 3 years, the cross-selling bank may only need to look at a three year period. Also, the real-time learning can be more practically defined as daily, monthly, or quarterly. Second, the rapid updates of customer information leads to an "explosion" in the number of states and control variables. This problem hinders estimation and is appropriately called the "curse of dimensionality." In such environments, heuristics such as interpolation method developed by Keane and Wolpin (1994) that closely resemble the optimal solution and are faster to obtain, are perhaps a more practical solution. Third, even for companies that can not implement the proposed framework, it is useful for them to conceptually frame their CRM problems in a similar way and/or derive some heuristic rules that resemble optimal solution. For example, from the proposed framework, we can derive some statistical properties that describe how the optimal CRM decisions are driven by important customer or company specific variables.

4. Two Illustration Examples

[Insert Table 2 About Here]

We choose two particular cases, allocation of service calls and cross-selling campaigns, to demonstrate the ideas and value of adaptive learning and proactive CRM. In Table 2, we summarize the formulation of the two CRM problems with emphasis on the three key components.

4.1 Call allocation

The problem

With the fast development of off-shore call centers, companies start to show increasing concern with the placement of a key CRM asset in the hands of a third-party provider. How can we leverage the strength of off-shore centers to improve the service effectiveness without

incurring significant costs? This requires the company to develop a CRM algorithm that allows for learning about the heterogeneous preference of customers as well as the comparative advantage of call centers, calculates the trade-off between operating efficiency and marketing effectiveness, and allocates customers to call centers according to their preferences.

The solution

We treat service duration as a measure of operation efficiency, as well as a determinant of customer retention or marketing effectiveness. With service duration, customer satisfaction and retention decision models characterizing customer response to the firm's interventions, the company's optimal allocation decision is formulated as a dynamic control problem in which the company learns about the heterogeneous preferences of customers as well as the comparative advantages of off-shore centers in an adaptive fashion, balances the trade-offs between short-term cost benefit and long-term customer reactions, and makes optimal allocation decisions that best match customer preferences for service duration and maximize long-term profit.

Applying our framework to a call allocation history data, we analyze the relationship among service allocation, service duration, retention, and profit. Based on the parameter estimates, we run simulations to derive the optimal call allocation decisions. We demonstrate how the operating decisions are driven by marketing consequences, as well as the dynamic, customized, and state-dependent nature of the derived allocations strategies.

The results

Our proposed framework allows the company to make customized and dynamic decisions that save costs, increase customer retention, and improve profit. Compared with the currently adopted "skill-based" routing (i.e. the rule of routing the customer to the agent with the lowest average service duration), the optimal allocation decisions derived from our framework (1) improve customer retention by 8%; (2) reduce average service costs by 6%; and (3) enhance total profits by 15% because of the growing relationships. Thus, service effectiveness can be improved without incurring a significant cost. The proposed solution is different from conventional ways of improving customer retention by incurring more costs to increase service quality.

4.2 Cross-selling Campaigns

The Problem

Despite the increasing investment in cross-selling efforts, companies find that million-dollar marketing campaigns often fail to generate the responses necessary to create revenue or even cover the cost of the campaign. Cross-selling companies are challenged by how to improve the effectiveness of cross-selling campaign in a cost efficient way. Managers may be left with many puzzling questions: How to design the most relevant cross-selling campaign that is tailored to each customer's evolving needs and preference? How to improve the average response rate of a cross-selling campaign? In order to address these questions, it is important to understand the role of a cross-selling campaign, how it interacts with the development of customer demand, and the trade-off between short-term campaign costs and long-term profit gains.

The Solution

Treating customer purchase decision as customer response model with hidden Markov process to capture the development of their financial maturity, we allow the company to adopt a Bayesian updating rule to update its perceived likelihood of a customer belonging to a type (i.e. hand-holding or self-sufficient) classifying their intrinsic preference for service duration based on observed past service duration and resulting retention. We formulate cross-selling campaigns as a stochastic dynamic programming problem with adaptive learning for the company whose goal is to develop a sequence of optimal cross-selling campaigns strategies in order to maximize the long-term profit of its existing customers. The proposed framework integrates the development of the latent financial needs, the different functions of various cross-selling campaigns at different customer financial stages and heterogeneity of customer preferences with the goal of maximizing long-term value of each individual customer.

Applying our model to cross-selling campaigns and purchase data provided by a national bank, we parameterize customer demand based on estimation results and then run simulations to demonstrate the dynamic and state-dependent nature of the optimal cross-selling campaign decisions and derive structural properties of the customized and dynamic cross-selling campaign strategies. We demonstrate that the company's decisions are a multi-step, multi-segment and multi-channel cross-selling campaign process about when to target which customer with what product using the best campaign channel (how).

The Results

We demonstrate that the instantaneous effectiveness of a cross-selling campaign is improved from 5.6% to 11.2%. We also calculate the indirect educational effect of a cross-selling campaign to be 63%. We show that the ROI of a cross-selling campaign is improved by 40.8% and demonstrate the increase of long-term profit when the company shifts its cross-selling strategy from campaign-centric to customer-centric.

[Insert Figure 2A and Figure 2B about Here]

In Figure 2A and Figure 2B, we use sample data to demonstrate how the customer knowledge obtained from adaptive learning and the optimal CRM action derived from the proactive CRM can be translated to customer scoring and recommended CRM interventions. Take the service allocation as an example, when a customer calls in, her background information is traced and shown on the operator's screen. The results of adaptive learning algorithm show that this customer is a hand-holding customer who prefers longer service duration. Given this updated information on customer type, the expected service durations and probability of retention are calculated using the customer response models of duration and retention. The total profit is also calculated as the customer lifetime value similarly defined by equation (1). Comparing the total profit of on-shore and off-shore routing, off-shore routing is recommended as the company's action.⁴

5. Brief Discussions on Automated Implementation of Proposed Solutions

In summary, with adaptive learning, forward-looking, and optimization, our approach is more akin to the decision support system of CRM, in which the company obtains more information about the customer in order to make dynamic and customized decisions with the goal of maximizing long-term profit. To implement the solutions, it requires the company to be able to have immediate access to customer database, to learn about the customer's intrinsic preference, to

⁴ Interested readers can refer to Sun and Li (2005) and Li, Sun and Montgomery (2005) for more detailed information on these two examples.

solve the dynamic programming problem to obtain the optimal allocation decision, and to update its belief based on the successful and unsuccessful interactions. All these need to be done within seconds, which is impossible for a human being. However, today's technology provides companies with the ability to instantly retrieve real-time customer information, automatically analyze customer insights, respond directly to customer requests, and provide the customer with a highly customized intervention decision. It is possible for companies to develop software-based *automated decision support systems* to accomplish these sophisticated decision processes. This is especially true for internet and direct marketing industry for which it is easier to implement adaptive learning and proactive CRM. Automation reduces the risk of human error by automatically preparing and modeling thousands of input variables.

Take the call center as an example. As a minimum, the company's call center has both a CRM system ready to record customer call history and ACD (automated call distribution) systems in place to automatically allocate service calls. CRM can be integrated into the ACD system. When customers call in, they are required to provide their account numbers. With immediate access to its CRM system, the company has an integrated view of every customer's call history and all other related information up to the current period. According to the pre-specified rule of learning, the company can update its knowledge of customer preferences. With future marketing consequences in mind, the company solves for optimal CRM decisions according to the proposed framework which maximize customer long-term profit. The decision can then be automatically implemented by the ACD system. As for the cross-selling campaign example, there are campaign execution and tracking software that allow users to develop and deliver targeted messages in a test-and-learn environment.

6. Conclusion

Recent technological developments have opened massive possibilities to advance and automate CRM decisions. This transforms the CRM system from a data collection and storage technology into service excellence and revenue growth opportunities (Winer 2001). Ideal CRM interventions should be customer-centric by following the development of each individual customer and making multi-segment, multi-stage, and multi-channel CRM interventions with the goal of maximizing the total customer profit. However, most recent CRM practices are still conveniently campaign-

centric, in which the focus is on selecting the most profitable customers for a scheduled CRM intervention.

In this article, we discuss the two-step procedure (“adaptive” learning and “proactive” CRM decisions) and three-components for customer-centric CRM. We formulate CRM interventions as solutions to a stochastic dynamic programming problem under demand uncertainty in which the company needs to learn about the evolution of customer demand, the dynamic effect of its marketing interventions, the heterogeneity of customer preferences, the cost of interventions, and the long-term payoff, with the goal of maximizing the “long-term” profit of each customer. We also use two examples to demonstrate the input, output, and benefit of “adaptive” learning and “proactive” CRM. The proposed solution meets the recent trend of companies seeking real-time solutions for integrating database and CRM decisions that are empowered by the advancement of technology.

As Shugan (2003) pointed out, “Extraordinary increases in computational speed allow sellers to use more sophisticated tools to quickly analyze traditional databases and to continuously improve targeting strategies.” The fast development of on-line and direct marketing industries creates enormous opportunity for adaptive learning and proactive CRM. This calls for substantial research to develop statistical algorithms to measure customer insights and develop optimization routines as decision-support systems to automate the implementations of marketing decisions to better manage customer relationships. For companies with difficulties automating the implementation of CRM decisions, the statistical properties derived from the proposed framework provide descriptive guidance for managers to adjust their CRM decisions.

Reference

- Ansari, Asim, Carl Mela, and Scott A. Neslin (2005), "Customer Channel Migration," Working Paper, Tuck School of Business, Dartmouth College, January 2005
- Berry, Michael and Gordon Linoff (1997), *Data Mining Techniques: Techniques: For Marketing, Sales, and Customer Support*, John Wiley and Sons, Inc.
- Berry, Leonard L. (1995), "Relationship Marketing In Services: Growing Interest, Emerging Perspectives," *Journal of the Academy of Marketing Science*. 12(winter), 17-30.
- Bitran, Gabriel R., and Susana V. Mondschein (1996), "Mailing Decisions in the Catalog Sales Industry," *Management Science*, 42(9), 1364-1381.
- Blattberg, Robert C., and John Deighton (1991), "Interactive Marketing: Exploiting the Age of Addressability," *Sloan Management Review*, (fall), 5-14.
- Bolton, Ruth N. (1998), "A Dynamic Model of the Duration of the Customers' Relationship with a Continuous Service Provider," *Marketing Science*, 17(1), 45-65.
- Boulding, W., A. Kalra, R. Staelin, and V. A. Zeithaml (1993), "A Dynamic Process Model of Service Quality: from expectations to behavioral intentions," *Journal of Marketing Research*, 30, 7 – 27.
- Bult, Jan-Roelf, and Ton Wansbeek (1995), "Optimal Selection fro Direct Mail," *Marketing Science*, 14(4), 378-394.
- Chen, Tao, Baohong Sun and Vishal Singh (2006), "Did Marlboro Man Fell off His Horse?," working paper, Carnegie Mellon University.
- Edwards, Y. D., and G. Allenby (2003), "Multivariate Analysis of Multiple Response Data," *Journal of Marketing Research*, 40(3), 321-334.
- Elsner, Ralf, Manfred Krafft, and Arnd Huchzerneier (2004), "Optimizing Rhenania's Direct Marketing Business through Dynamic Multi-level Modeling (DMLM) in a Multicatalog-brand Environment," *Marketing Science*, 23(2), 196-206.
- Gans, N., Koole, G., Mandelbaum, A. Telephone Call Centers: Tutorial, Review and Research Prospects. Invited review paper by *Manufacturing and Service Operations Management (M&SOM)*, 5 (2), 2003.
- Gönül, Füsün, and Meng Ze Shi (1998), "Optimal Mailing of Catalogs: A New Methodology Using Estimable Structural Dynamic Programming Models," *Management Science*, 44(9), 1249-1262.
- Haeckel, Stephan H. (1998), "About the Nature and Future of Interactive Marketing," *Journal of Interactive Marketing*, 12(1), 63-71.
- Kamakura Wagner et al, "Choice Models and Customer Relationship Managements – Summary Paper for the Sixth Choice Symposium," forthcoming at *Marketing Letters*.
- Kamakura W., S. Ramaswami, and R. Srivastava (1991), "Applying Latent Trait Analysis in the Evaluation of Prospects for Cross-Selling of Financial Services," *International Journal of Research in Marketing*, 8: 329-349.

- Kamakura, Wagner A, Michel Wedel, Fernando de Rosa and Jose A. Mazzon (2003), "Cross-selling through Database Marketing: A Mixed Data Factor Analyzer for Data Augmentation and Prediction," *International Journal of Research in Marketing* 20, 45-65.
- Keane, Michael P., and Kenneth Wolpin (1994), "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence," *The Review of Economics and Statistics*, 76(4), 648-673.
- Kumar, Piyush, Manohar U. Kalwani, Maqbool Dada (1997), "The Impact of Waiting Time Guarantees on customers' Waiting Experiences." *Marketing Science*. 16(4), 295-314.
- Kumar, V., and Rajkumar Venkatesan (2005), "Who are the Multichannel Shoppers and How do they Perform?: Correlates of Multichannel Shopping Behavior," *Journal of Interactive Marketing*, 19(2), 44-62.
- Lemon, Katherine N., Tiffany B. White, and Russell S. Winer (2002), "Dynamic Customer Relationship Management: Incorporating Future Considerations into the Service Retention Decision," *Journal of Marketing*, 66(1), 1-14.
- Lewis, Michael (2005), "A Dynamic Programming Approach to Customer Relationship Pricing," *Management Science*, 51(6), 981-994.
- Li, Shibo, Alan Montgomery and Baohong Sun, "Introducing What Financial Product to Which Customer at What Time -An Empirical Analysis of Customized and Dynamic Cross-selling Campaigns," working paper, Carnegie Mellon University.
- Li, Shibo, Baohong Sun and Ronald Wilcox (2005), "Cross-selling Sequentially Ordered Products: An Application to Consumer Banking Services," *Journal of Marketing Research*, XLII, May, 233-239.
- Montgomery, Alan L. and Kannan Srinivasan (2003), "Learning About Customers Without Asking", Nirmal Pal and Arvind Rangawamy (eds.), *The Power of One - Leverage Value from Personalization Technologies*, eBRC Press, Penn State University.
- Montgomery, Alan, and Kannan Srinivasan (2003), "Learning about Customers without Asking", *The Power of One*, 122-143.
- Peppers, Don, Martha Rogers, and Bob Dorf (1999), "Is Your Company Ready for One-to-One Marketing," *Harvard Business Review*, 77(1), 151-160.
- Reinartz, Werner J., and V. Kumar (2000) "On the Profitability of Long-Life Customers in a No contractual Setting: An Empirical Investigation and Implications for Marketing," *Journal of Marketing*, 64(4), 17-35.
- Reinartz, Werner J., and V. Kumar (2003) "The Impact of Customer Relationship Characteristics on Profitable Lifetime Duration," *Journal of Marketing*, 67(1), 77-99.
- Rust, Roland T and Tuck Siong Chung (2005), "Marketing Models of Service and Relationships," forthcoming, *Marketing Science*.
- Rust, Roland T., Katherine N. Lemon and Valarie A. Zeithaml (2004), "Return on Marketing: Using Customer Equity to Focus Marketing Strategy," *Journal of Marketing*, 68 (1), 109-127.
- Rust, Roland T., and Peter C. Verhoef (2005), "Optimizing the Marketing Interventions Mix in Intermediate-Term CRM," *Marketing Science*, 24(3), 477-489.

- Rust, Roland T., Valarie A. Zeithaml, and Katherine N. Lemon (2000), *Driving Customer Equity: How Customer Lifetime Value Is Reshaping Corporate Strategy*. The Free Press, New York..
- Sawhney, Mohanbir, Sridhar Balasubramanian and Vish Krishnan (2004), "Creating Growth With Services", *MIT Sloan Management Review*, 45(2), Winter.
- Shankar, Venkatesh, Amy K. Smith, and Arvind Rangaswamy (2003), "Customer Satisfaction and Loyalty in Online and Offline Environment," *International Journal of Research in Marketing*, 20(2), 153-175.
- Shugan, Steven M. (2003), "Editorial: Defining Interesting Research Problems," *Marketing Science*, 22(1), 1-15.
- Sullivan UY and Jackie Thomas (2004), "Customer Migration: An Empirical Investigation Across Multiple Channels", Working Paper, Northwestern University.
- Sun, Baohong and Shibo Li (2005), "Improving Effectiveness of Customer Service in a Cost-Efficient Way - Empirical Investigation of Service Allocation Decisions with Out-Sourced Centers," Working Paper, Carnegie Mellon University.
- Sun, Baohong, Yacheng Sun and Shibo Li (2005), "Purchase to Secure Future Consumption ----- An Empirical Investigation of Consumer Choice of Service Plan of Online DVD Rental," working paper, Carnegie Mellon University.
- Taylor, Gail Ayala and Scott A. Neslin (2004), "The Current and Future Sales Impact of a Retail Frequency Reward Program," working paper, Dartmouth University.
- Toubia, Olivier, Duncan I. Simester, John R. Hauser, and Ely Dahan (2003), "Fast Polyhedral Adaptive Conjoint Estimation," *Marketing Science*, 22(3), 273-303.
- Toubia, Olivier, John R. Hauser, and Duncan I. Simester (2004), "Polyhedral Methods for Adaptive Choice-Based Conjoint Analysis," *Journal of Marketing Research*, 41(1), 116-131.
- Venkatesan, Rajkumar, and V. Kumar (2004), "A Customer Lifetime Value Framework for Customer Selection and Resource Allocation Strategy," *Journal of Marketing*, 68(October), 105-125.
- Witten, Ian H. and Eibe Frank (2005), "Data Mining: Practical Machine Learning Tools and Techniques." Elsevier.
- Winer, Russell (2001), "A Framework for Customer Relationship Management," *California Management Review*, 2001, 43 (Summer), 89-105.
- Xie, Jinhong and Steven Shugan (2001), "Electronic Tickets, Smart Cards, and Online Prepayments: When and How to Advance Sell," *Marketing Science*, Vol. 20 (Summer) 219-243.

Table 1. Selected Literature on Customer Relationship Management

Papers	CRM Decision	Consumer Model	Company Decision			multi-state, segment, channel, solution
			Adaptive Learning	Forward-looking	Optimization	
Edwards and Allenby 2002; Li, Sun and Wilcox 2005; Kamakura, Ramaswami and Srivastava 1991; Kamakura, Wedel, de Rosa and Mazzon 2002;	Cross-selling	x				
Shankar, Smith, and Rangaswamy 2003;	Customer satisfaction and loyalty online vs. offline	x				
Reinartz and Kumar 2000; Reinartz and Kumar 2003;	Profitability of long-life customer and lifetime duration	x		x		
Toubia, Simester, Hauser and Dahan 2003; Toubia, Hauser and Simester 2004;	Conjoint analysis to reveal customer preference	x	x			
Bult and Wansbeek 1995;	Direct mailing list	x			x	
Bolton 1998;	Customer duration	x	x			
Rust, Zeithamal and Lemon 2000; Rust, Lemon and Zeithaml 2004	Customer lifetime value	x		x		
Kumar and Venkatesan (2005)	Channel selection	x				
Bitran and Mondschein 1996; Gönül and Shi 1998; Elsner et al. (2004)	Catalogue mailing	x		x	x	
Lewis 2005;	Relationship pricing	x		x	x	
Venkatesan and Kumar (2004); Rust and Verhoef 2005	Marketing allocation			x	x	x
Proposed Framework	general application	x	x	x	x	x

Table 2. Applications of “Adaptive” Learning and “Proactive” CRM

	Allocation of Service Calls	Cross-selling Campaign
CRM Decisions	allocating customers to call centers	when to introduce what financial product to which customer using what channel
What to be Learned?	customer preference for on-shore and off-shore centers; comparative advantages of centers	financial maturity; state of financial maturity; preference for channel
Input of Learning	History on call durations, customer retention	purchase history of various financial products; demographics; sampling and experimentation
Output of adaptive learning	Learn about the sensitivity to service duration	development of financial maturity; preference for campaign channel
Forward looking Component	Future probabilities of customer retention and profit implications	future probabilities of purchasing other products and profit implications
Optimization problem	tradeoff the short term saving of operating cost (by assigning off-shore centers) and long-term marketing consequence of alienating customers	tradeoff the short-term costs of education and future benefit of higher future cross-selling opportunities
Improvement over passive learning	(1) Improve customer retention by 8%; (2) reduce average service costs by 6%; and (3) enhance total profits by 15% because of the growing relationships.	Improve the instantaneous effectiveness of a cross-selling campaign from 5.6% to 11.2%. Find that the indirect educational effect of cross-selling campaign is about 63%. Improve ROI of cross-selling campaign by 40.8%.

Figure 1. Timeline of Company Decision Process

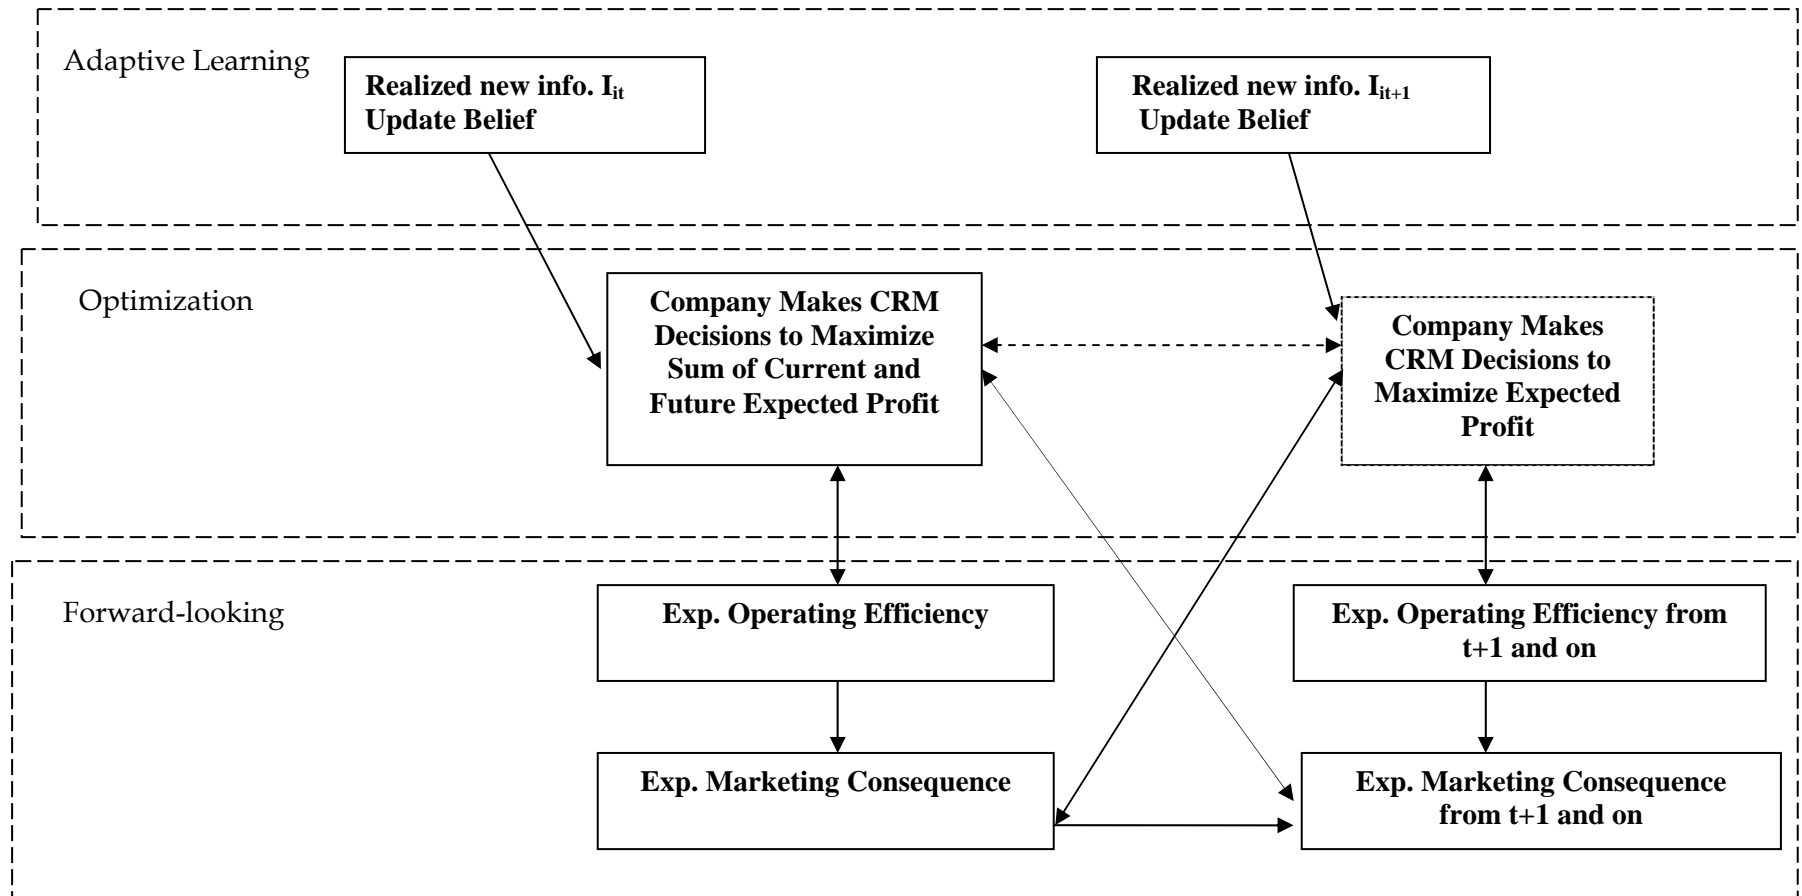


Figure 2A. A Sample Output of Call Allocation Decisions

Customer ID: 123456		
Gender: female		
Education: high school		
Type of call-in Question: software installation		
Results of Customer Knowledge from “Adaptive” Learning		
Customer Type: hand-holding customer who prefer longer service duration		
	On-shore	Off-shore
Expected service duration for this call	14 min	12 min
Expected probability of retention after this call	0.75	0.72
Expected service cost for this call	\$11	\$7
Expected lifetime profit	\$297	\$341
Recommended Decision Resulting from “Proactive” CRM Framework		
Action: allocate to off-shore center		

Figure 2B. A Sample Output of Cross-selling Campaign

Customer ID: 654321		
Gender: male		
Education: college		
Current ownership: checking, saving, money market		
Previous cross-selling contacts: email 6 times, telephone call 5 times		
Results of Customer Knowledge from “Adaptive” Learning		
Financial maturity: stage 3 (investment and risk covering)		
Duration in current financial state: 13 months		
Estimated months to switch state: 5 months		
Channel preference: email		
	email	telephone
Expected probability of immediate response	0.23	0.11
Expected cost for this campaign	\$0.57	\$1.54
Expected lifetime profit	\$1000	\$900
Recommended Decision Resulting from “Proactive” CRM Framework		
Next product to introduce:		
life insurance, brokerage, annuity (in the order of purchase likelihood)		
When to contact:		
1 Educational campaigns within the next 3 months followed by		
2 promotional campaigns		
How to contact:		
email or direct mail		