

Exploiting Response Models - Optimizing Cross-Sell and Up-Sell Opportunities in Banking

Andrew Storey

Scotiabank

1 Richmond W., 2nd Floor

Toronto, Ontario, Canada M5H 1Y8

416.933.2436

andrew.storey@scotiabank.com

Marc-david Cohen

SAS Institute Inc.

SAS Campus Drive

Cary, NC 27513

919.531.7630

marc.cohen@sas.com

ABSTRACT

The banking industry regularly mounts campaigns to improve customer value by offering new products to existing customers. In recent years this approach has gained significant momentum because of the increasing availability of customer data and the improved analysis capabilities in data mining. Typically, response models based on historical data are used to estimate the probability of a customer purchasing an additional product and the expected return from that additional purchase. Even with these computational improvements and accurate models of customer behavior, the problem of efficiently using marketing resources to maximize the return on marketing investment is a challenge. This problem is compounded because of the capability to launch multiple campaigns through several distribution channels over multiple time periods. The combination of alternatives creates a complicated array of possible actions. This paper presents a solution that answers the question of what products, if any, to offer to each customer in a way that maximizes the marketing return on investment. The solution is an improvement over the usual approach of picking the customers that have the largest expected value for a particular product because it is a global maximization from the viewpoint of the bank and allows for the effective implementation of business constraints across customers and business units. The approach accounts for limited resources, multiple sequential campaigns, and other business constraints. Furthermore, the solution provides insight into the cost of these constraints, in terms of decreased profits, and thus is an effective tool for both tactical campaign execution and strategic planning.

Keywords

Database marketing, Cross-selling, Up-selling, Profit Optimization, Assignment Problems, Constrained Optimization, Response Models.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGKDD '02, July 23-26, 2002, Edmonton, Alberta, Canada.

Copyright 2002 ACM 1-58113-567-X/02/0007...\$5.00.

1. INTRODUCTION

The new mantra of database marketing in banking and financial services is “the right product to the right customer at the right time”. However, a practical and effective implementation of this goal is not easy to accomplish. What makes this particularly difficult is that companies have multiple products and operate under a complex set of business constraints. Choosing which products to offer to which customers in order to maximize the marketing return on investment and meet the business constraints is enormously complex. This paper outlines a framework for solving this problem and presents an example using data from Scotiabank.

Scotiabank is one of North America’s premier financial institutions; it is comprised of Domestic Banking, Wealth Management, International Banking and Scotia Capital groups. The Domestic Bank employs more than 23,000 people and has over 6 million customers. The Wealth Management Group incorporates key personal investment and advisory activities within the Scotiabank Group. Scotiabank is the most international of all Canadian banks, its International Banking Group has more than 21,000 employees and provides retail banking services in over 50 countries. The Scotia Capital Group provides corporate and investment banking on a global basis. Because of its breadth, Scotiabank is able to offer a full suite of financial products to its clients.

Scotiabank has made a deliberate effort to become a customer-focused institution, as opposed to a vertical product driven company. The bank’s formally stated goal is “to be the best at helping customers become financially better off by providing relevant solutions to their unique needs”. A direct consequence of this goal is that marketing campaigns are multiple product campaigns as opposed to single product campaigns. This transforms the data mining and campaign targeting process from a fairly simple application of individual response models into a significantly more complex problem of choosing which product, if any, to offer to which customer and through which channel. The benefit is that campaigns are more customer-focused than in the past.

1.1 Business Problem

The database marketing community has changed significantly over the last several years. In the past, database marketers applied business rules to target customers directly. Examples include; targeting customers solely on their product gaps or on marketers’ business intuition. Marketers have also applied RFM

type analysis where general *recency*, *frequency*, and *monetary* measurements as well as product gaps are used to target customers for specific offers. The current approach, which has widespread use, relies on predictive response models to target customers for offers. These models accurately estimate the probability that a customer will respond to a *specific* offer and can significantly increase the response rate to a product offering. However, simply knowing a customer's probability of responding to a particular offer is not enough when a company has several products to promote and other business constraints to consider in its marketing planning.

Marketing departments also face the problem of knowing which *product* to offer to a *customer*, not just which customer to offer a product. In practice, many ad hoc rules are used. Prioritization rules based on response rates or estimated expected profitability measures have been used; business rules to prioritize products that can be marketed are sometimes used; and product response models to select customers for a particular campaign are also used. One approach that is easily implemented but, for reasons outlined later, may not produce optimal customer contact plans relies on a measure of expected offer profitability (the estimated probability of response multiplied by the profit given customer response less direct costs) to choose which products to offer customers. However, a shortcoming of this approach is its inability to effectively handle complex constraints on the customer contact plan.

1.2 Business Constraints

Database marketing departments face several types of business constraints. Typically, there are

- restrictions on the minimum and maximum number of product offers that can be made in a campaign,
- requirements on minimum expected profit from product offers,
- limits on channel capacity,
- limits on funding available for the campaign, and
- campaign return-on-investment hurdle rates that must be met.

These are a sample of the constraints that marketing departments must meet when executing a campaign. Ad hoc approaches are also typically used in an attempt to meet these constraints.

The opportunity costs of the business constraints are generally not known. Constraints are usually negotiated between marketing, product lines and delivery channel management. If the cost of a constraint was known, then the company could choose to tighten or to relax the constraint by removing or adding more resources. For example, channel capacity could be increased if it were known that there was a significant return on the investment by doing so. Knowledge of the opportunity costs could help evaluate these management decisions. Applications of this will be discussed in the "STRATEGIC USAGES" section of this article.

Ultimately, the database marketer needs a concrete framework to effectively act on "the right product to the right customer at the right time" mantra. The approach we take is to transform the database-marketing problem into an optimization problem that is designed to generate the maximum incremental profit from a limited amount of resources subject to the necessary business constraints. This paper will describe an actionable framework that will satisfy this business problem.

2. SOLUTION FRAMEWORK

It is helpful in understanding the solution framework to understand the data that are available for marketing campaign planning. Understanding the data will help make the problem more concrete.

2.1 Data

We assume that there has been a thorough analysis of historical marketing campaigns and that accurate response probability models exist for all products in question. The result of these data mining exercises is a data set that contains an expected profit for each product for each customer, where the expected profit is derived from the customer specific probability of response and the profit generated given a customer response. Needless to say, these data sets can be rather large. It is not unusual to have over 5 million customer records in such a data set. Let's assume that there are 10 products and 1 million customers, and that for each customer and product we have an estimate of the expected profit given that each customer is offered each product.

2.2 Ideal Approach

The ideal approach to solving this problem is to model it as a specialized type of assignment problem. This type of problem is an integer program. It can be unambiguously expressed with a mathematical formulation. Let $x_{ij} = 1$ if customer i is offered product j , and 0 if not; let r_{ij} the expected profit of offering customer i product j ; let c_{ij} the cost of offering customer i product j ; let R be the corporate hurdle rate. Then, a very simplified version of the problem can be expressed as finding the x_{ij} that satisfy

$$\text{Max } \sum_{ij} x_{ij} r_{ij}$$

Subject to:

$$\sum_{ij} x_{ij} c_{ij} \leq \text{Campaign budget}$$

$$\sum_i x_{ij} \geq \text{Minimum offers of product } j$$

$$\sum_{ij} x_{ij} r_{ij} \geq \sum_{ij} x_{ij} c_{ij} (1 + R) \text{ Corporate hurdle}$$

$$x_{ij} \in \{0, 1\}$$

This formulation captures only the bare elements of the problem. It does not account for multiple campaigns composed of different products, multiple channels, and channel capacity constraints just to name a few possibilities. However, the model can easily be extended to cover most typical business constraints encountered in practice, but the basic formulation remains the same. It is important to note that this ideal formulation is difficult to solve because of its scale. For 1 million customers and 10 products there are 10-million integer variables x_{ij} , this yields $2^{10,000,000}$ possible customer-offer combinations. Using standard optimization methods a problem of this size can, in principle, result in a branch and cut tree of as many nodes. Because of this problems of this size are extremely difficult to solve, so we propose an alternative approach. While not providing a strictly optimal solution, the alternative approach does provide an approximately optimal solution that in preliminary studies has shown to be a good approximation.

2.3 Practical Approach

Although it is not practical to solve problems formulated in this ideal way, it is possible to approximate the ideal formulation and arrive at a formulation that is practical to solve. There are numerous ways to approach this approximation; one approach is to sample from the customer base and use that sample as representative for the optimization. Another approach (and the one that we take) is to aggregate customers based on the coefficients r_{ij} in the ideal formulation. Aggregation can be considered natural in this setting particularly when we understand that much of the data is consistent and estimated. For example, the cost data c_{ij} are most likely to be consistent across customers for a given product. Similarly, the estimated expected profit r_{ij} is most likely the result of data mining techniques such as predictive response models. The implementation of this framework is loosely coupled to the chosen form of the predictive response models. As long as the customer/offer specific response rate is represented as a probability, the proposed framework can handle it. Scotiabank uses standard, accepted statistical and data mining approaches to obtain these estimates.

The aggregation process we use involves conversion of the raw data into a form that can be used *naturally* in a linear programming optimization model. The key is to cluster the raw data r_{ij} and use the clusters as the aggregate. Unlike the usual use of clustering, the purpose here is not the identification of customer segments or to differentiate groups of customers, but to aggregate customers into similar groups. This is an important distinction to keep in mind since clustering is most frequently used to distinguish, not to aggregate. If the clusters are internally consistent, then the cluster centroids can be used as representative of the data for all the customers within a single cluster.

This aggregation enables the problem to be reformulated as a linear program so that rather than assigning offers to individual customers, as the ideal integer program does, the program identifies proportions within each cluster for each product offer. This can be accomplished with similar constraints to those of the ideal formulation. Moreover, the linear program is much smaller and much easier to solve. Note however, the solution may require that multiple products are offered to proportions of customers within a single cluster. When that happens, a new problem is defined that is a simple assignment problem at the level of the cluster, where multiple offers are to be assigned within the cluster, and it is relatively easy to solve.

It should also be noted that preliminary experiments solving the approximation with varying numbers clusters indicate that as the number of clusters increases the value of the objective quickly rises then slowly converges to the integer relaxed solution. Further study is needed to identify a good number of clusters that work in a “typical” setting.

2.3.1 Formulation

Consider the following variables defining raw data as input into the solution algorithm. Let y_{ij} be the number of customers in cluster i that are offered product j ; let r'_{ij} be the estimated expected profit given that customer in cluster i is offered of product j ; let c'_{ij} be the cost of offering a customer in cluster i product j ; let R be the corporate hurdle rate. Then, a very simplified version of the problem can be expressed as finding the y_{ij} that satisfy

$$\text{Max } \sum_{ij} y_{ij} r'_{ij}$$

Subject to:

$$\begin{aligned} \sum_j y_{ij} &\leq \text{Number customers in cluster } i \\ \sum_{ij} y_{ij} c'_{ij} &\leq \text{Campaign budget} \\ \sum_i y_{ij} &\geq \text{Minimum offers of product } j \\ \sum_{ij} y_{ij} r'_{ij} &\geq \sum_{ij} y_{ij} c'_{ij} (1 + R) \text{ Corporate hurdle rate} \\ y_{ij} &\geq 0 \end{aligned}$$

Once the y_{ij} that satisfy the formulation are found, the optimal proportions that they give must be applied to the customers within the specific clusters. For example, suppose that y_{ij} is the total number of customers in cluster i . Then, every customer in that cluster should be offered product j . Alternatively, suppose that for a given i , $y_{ij} > 0$ and $y_{ij'} > 0$ for $j \neq j'$. Then, y_{ij} of customers in cluster i must be offered product j and $y_{ij'}$ of customers in cluster i must be offered product j' . The optimal way to do that is to solve a simple assignment problem using the estimated expected profit r_{ij} for the individual customers and not the clusters. It is important to note that some of the constraints may be violated as a result of solving this assignment problem particularly if the cluster centroids used in the linear program formulation are involved in a tight constraint and not consistent within the cluster.

3. A TACTICAL EXAMPLE

We demonstrate this approach with data from Scotiabank and using existing procedures within the SAS system to implement the formulation described above. The details of the SAS code will not be given.

Eleven unique offers were to be considered: five investment, three lending and three day-to-day banking offers. The investment offers included GICs, mutual funds, Registered Education Savings Program (RESP) and two unique discount brokerage offers. The lending offers included a mortgage and two credit card offers. The day-to-day banking offers included one of two Scotia online banking service offers and a deposit account acquisition. The term campaign is used here to imply one large pro-active customer contact campaign that it comprised of eleven distinct offers, it can be thought of as eleven single product campaigns that are being offered at generally the same time to a non-overlapping set of customers. For the purposes of this paper the detailed product offer descriptions have been suppressed. Approximately 2.5 million customers were included in the potential universe for the campaign.

Ultimately, the goal of marketing campaigns is to produce a positive return on investment for the company that exceeds the corporate investment hurdle rate. Although the timeframe upon which this investment should be measured may be debatable, the goal is fundamental to the bank. To achieve this specific objective, the bank can execute marketing campaigns that are designed to maximize the expected incremental profit through making one of several offers to some of its customers, or potential customers.

3.1 Response Models

The expected incremental profit of a specific offer to a customer is an estimate based on response models and detailed product profitability calculations. These response models are used to estimate *the probability that a customer will accept a specific offer*. Scotiabank's data warehouse has detailed account level profitability calculations for all of its products. This profitability information is used to estimate the near term incremental profit *given that the customer accepts the specific offer*. Once a specific offer is made to a customer there are two possible outcomes: the customer can accept or reject the offer. Using the offer specific response models the probability of both states is known for each customer. The incremental profit for both states is also known; it is zero if the customer rejects the offer and the mean near-term profitability for new accounts of the specific type if the offer is accepted. With this information, the expected incremental profit of the offer can be calculated for each customer/offer combination. The cost of making each offer is also known and is largely dependent on the channel through which the offer is made.

3.2 Channels

Scotiabank has several distribution channels through which campaigns can be executed. The main channels for direct marketing are direct mail, retail branch centres and call centres. For this example we assume that leads sent to the branch officers and call centres are follow-ups from a direct mail piece and that offers designated as direct mail are direct mail only. The use of the branch and call centres for follow-ups has been shown to have a positive effect on the probability of response to the offer when compared to direct mail alone. Of course, the lead delivery costs vary with the channel used. In this example we have used costs per lead of \$3.00, \$1.50 and \$1.00 for the branch, call centre and direct mail only channels respectively.

3.3 Business Constraints

Several practical issues surround the campaign execution process that affects the customer/offer selection process, for this application to be acceptable for implementation these business constraints must be maintained. The following business rules have been translated into constraints that can be applied to the optimization model:

- Campaign costs cannot exceed \$1 million.
- The campaign must have a return on investment of at least the corporate hurdle rate. In this example we have used 20%, which is not necessarily the bank's actual corporate hurdle rate.
- The branch and call centre channels have a certain capacity constraint for timely processing of campaign generated leads. In the example, the call centre can accept up to 500,000 leads, the branch can handle up to 250,000 leads and direct mail is unlimited.
- Product offer minimums are also required to satisfy internal bank objectives. For the purposes of this example we set all offer minimums to 20,000 with two exceptions. The RESP offer, which has an extremely limited eligible universe, had a lower bound of 2,500 and one of the Scotia online offers had a lower bound of 5,000.
- Cannot offer products to customers who already have that product at Scotiabank.

- The standard marketing exclusions, such as credit risk or do not solicit, must also be strictly adhered to.

3.4 Solution

The estimates for customer/offer expected incremental profit, costs and business constraints serve as inputs to the profit optimization phase of the campaign design. The profit optimization phase is independent of the construction of these inputs. This means that as response models, profit estimates or costs are refined as long as the results are represented in the same manner, the optimization phase will be able to accept them as inputs. This property is important as the bank is constantly testing and refining these inputs as the marketplace is ever changing.

3.4.1 Results

The result of this algorithm is an allocation, of a specific offer, or no offer, to each customer. Also output is the associated expected incremental profit by customer making that offer. This solution is a SAS data set that has a customer identifier, the expected return, offer and channel designation. The full data set is 2.5 million records; the table below shows the first 25 records.

Customer Id	Product	Channel	Return
00182723			.
00200688			.
00032937			.
00722119			.
02137391			.
00992639			.
00060721			.
00483601	Offer 2	Direct Mail	0.0005
01164964			.
00025469			.
01008244			.
00179891			.
00410488	Offer 10	Direct Mail	3.1852
01484008			.
00184804			.
03350118			.
00983111			.
00387834	Offer 5	Call Center	13.0782
01100914			.
01507075			.
01559899			.
00309931			.
00657640			.
01095694			.
02075404			.

Figure 1. Sample of the solution dataset.

To better understand the solution, it is useful to look at several charts that summarize the solution and a report that is produced by the algorithm.

3.4.2 Offers

The Offer Frequencies by Channel chart, **Figure 2**, provides a graphical representation of the distribution of offers by channel and is useful for understanding the solution. This figure shows that few of the contacts have a branch follow-up treatment, some have call centre follow-ups and most have just the direct mail treatment. Offer 10 has the largest quantity of contacts and is spread across the direct mail and call centre channels. The Constraint Report, **Figure 3**, provides significantly more insight into the nature of the derived solution.

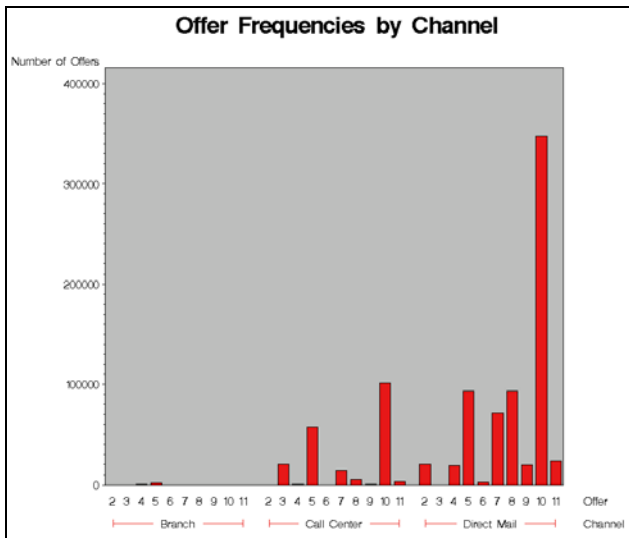


Figure 2. Offer Frequencies.

3.4.3 Business Constraints

The Constraint Report summarizes the constraints applied to the problem as well as outlines the chosen level and marginal costs associated with each constraint.

Constraint Name	Lower Target	Level	Upper Target	Marginal Value
Branch Capacity		2180.00	250000	0.00000
Call Center Capacity		202258.00	500000	0.00000
Campaign Cost		1000000.00	1000000	1.52605
Offer 10	20000	448325.00		0.00000
Offer 11	5000	7046.00		0.00000
Offer 12	20000	20000.00		-1.12965
Offer 2	20000	20000.00		-1.50959
Offer 3	20000	20000.00		-0.58726
Offer 4	20000	20000.00		0.00000
Offer 5	20000	152468.00		0.00000
Offer 6	2500	2500.00		-1.98643
Offer 7	20000	85665.00		0.00000
Offer 8	20000	98507.00		0.00000
Offer 9	20000	20000.00		-1.47605
Objective		3583503.25		0.00000

Figure 3. Constraint Report.

- The last line of the Constraint Report in **Figure 3** shows that the objective function, expected profit, was maximized at \$3.58 million from an expenditure of \$1 million. This results in a 258% return on investment for the campaign.
- The first two lines in **Figure 3**, *Branch Capacity* and *Call Center Capacity*, summarize the results of the branch and call centre capacity constraints respectively. The branch capacity for follow-up contacts was limited to 250,000 and the call centre to 500,000. In the solution, only 2,180 contacts were assigned to the branch, and 202,258 to the call centre, for follow-up calls. This low quantity of follow-up contacts, at either the branch or the call centre, is due to the conservative estimate of the increased response rates resulting from the follow-up and the significantly higher cost, \$3.00 for branch and \$1.50 for call centre, as compared to direct mail only, \$1.00.

- The third constraint, *Campaign Cost*, limits the total costs for the campaign to \$1 million. This constraint is in fact tight, meaning that the optimal solution was restricted by the condition. The marginal value of the constraint is \$1.53, this means that an additional \$1 spent on the campaign would result in a \$1.53 increase in expected profit.
- The *offer* constraints show the lower bounds for each of the specific offers. Notice that 5 of the 11 products were limited by their lower bounds. In the solution, *Offer 9* was only made to 20,000 people. If that constraint were to be *decreased* by one unit, e.g. to 19,999, then the objective function, expected profit, would be *increased* by \$1.48 – so the cost, on expected profit, of this business constraint is clear.

3.4.4 Profitability

The estimated Profitability by Channel report, **Figure 4**, clearly reveals the quality of leads that are sent to the respective channels. The branch follow-up leads have a significantly higher expected incremental profit than for call centre follow-up leads or direct mail alone. The call centre is also sent leads that are more profitable than direct mail alone. A few comments about the way that the channel effects were modeled are necessary to fully understand this result.

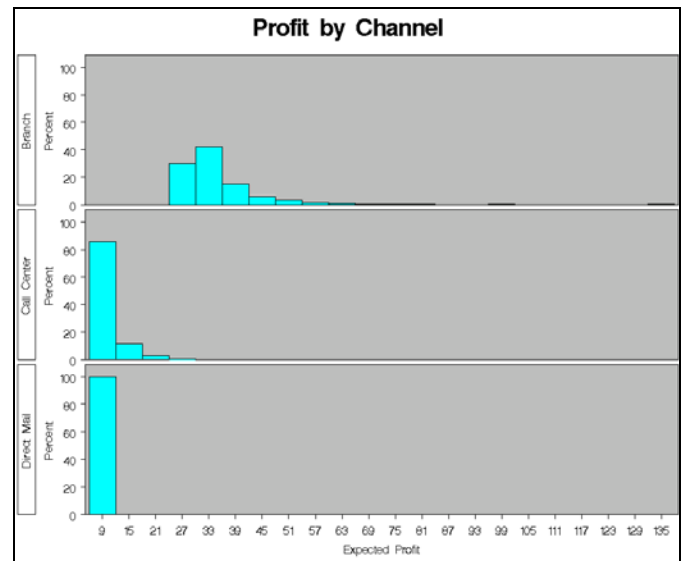


Figure 4. Profitability Report.

The differential effects of the various channels enter the response models as a *main effect*. A call centre follow-up treatment increases the probability of response as compared to no call centre follow-up treatment. The branch follow-up treatment has the same directional effect as the call centre, although larger. As such, everything else the same, the expected profit from making an identical offer to a customer with a branch or call centre follow-up is greater than without the follow-up. As such we would expect to see a higher expected rate of return when applying the additional follow-up treatments, although not at this magnitude. There is some other factor driving the higher rate of return, in fact, it is the channel selection by the optimization routine.

Recall that the channel costs are fixed at \$1.00, \$1.50 and \$3.00 for direct mail only, call centre follow-up and branch follow-up respectively. The incremental expected profit enters the equation through an increase in response rate. For a branch follow-up to be more profitable than an offer made without a branch follow-up it must be the case that:

$$\begin{aligned} & \text{Prob}_{\text{Branch}} \cdot \text{Profit} - \text{Cost}_{\text{Branch}} > \text{Prob}_{\text{DM}} \cdot \text{Profit} - \text{Cost}_{\text{DM}} \\ \Rightarrow & \text{Profit} \cdot (\text{Prob}_{\text{Branch}} - \text{Prob}_{\text{DM}}) > \text{Cost}_{\text{Branch}} - \text{Cost}_{\text{DM}} \\ \Rightarrow & \text{Profit} \cdot (\text{Prob}_{\text{Branch}} - \text{Prob}_{\text{DM}}) > 3 - 1 \\ \Rightarrow & \text{Profit} \cdot (\text{Prob}_{\text{Branch}} - \text{Prob}_{\text{DM}}) > 2 \end{aligned}$$

Either the profit given that the customer accepts (*Profit*) or the incremental effect on the response probability of the Branch follow-up ($\text{Prob}_{\text{Branch}} - \text{Prob}_{\text{DM}}$) is large enough to overcome the \$2 increase in contact costs ($\text{Cost}_{\text{Branch}} - \text{Cost}_{\text{DM}}$). The greater the difference in this inequality the more beneficial the follow-up contact is. The difference can be large if a highly profitable product is being offered, or the increase in the probability of response is high. For reasonable values of the probability of response from the direct mail (Prob_{DM}), the increase in response rate due to a branch follow-up ($\text{Prob}_{\text{Branch}} - \text{Prob}_{\text{DM}}$) will rise with an increase in the probability of response from the direct mail (Prob_{DM}). From a business perspective, this means that either highly profitable offers and/or customers who are most likely to respond to the offer will tend to be given a follow-up treatment. This is a particularly important result when dealing with the business owners of the call centre and branch channels.

3.5 Summary of Tactical Example

In summary, the solution provided an approximately optimal solution to the ideal capacitated assignment problem. The output is a decision, for each customer as to which, if any, product to offer and through what channel. The campaign expected profit is \$3.58 million on \$1 million invested, for a return on investment of 258%. Using an ad hoc approach, that utilized response models and near term profit, and met all the business constraints the most profit that could be generated was \$2.65 million on \$1 million invested for a campaign return on investment of 165%. The boost in the campaign return on investment of 93% is entirely attributable to the quality of the solution produced by the optimization process as opposed to the ad hoc approach.

4. STRATEGIC USES

Although this technique was developed primarily for its tactical application, as described above, it has some significant strategic applications too. The strategic applications are in the area of capacity planning. Two insights will be discussed, one dealing with campaign budgeting and the other with channel capacity planning.

4.1 Campaign Budget Allocation

In general, campaign budgets are determined prior to the campaign design. The degree of analysis that goes into determining specific campaign budgets, or annual campaign budgets, can vary greatly from institution to institution. As a strategic tool, the optimization technique provides an opportunity to determine the effects of making different budget allocations – in the budgeting process.

For example, in determining how much money to invest in a campaign, it would be useful to know the marginal return on an additional dollar investment. This is reported at \$1.53 in **Figure**

3 as the marginal value of the \$1 million cost constraint. Given the campaign definition, all of the other constraints and the current customer base, investing one more dollar would result in an increase in profit of \$1.53. If marginal return on investment is greater than the corporate hurdle rate then that supports an argument to increase the investment.

This technique provides a compelling and empirically based process for altering, positively or negatively, the campaign budget. Of course, other considerations go into budget allocations but the technique could shed light onto the impact of such decisions.

4.2 Channel Capacity Planning

Similarly, channel capacity planning can benefit from this technique. It is understood that in the short run, channel capacity is fixed, although in the long run, or planning mode, channel capacity can be changed.

Again, from the Constraint Report, **Figure 3**, if it appears as though a specific channel is used to capacity, then we can look at the marginal value of the constraint. The marginal value of these constraints gives the increase in profit, the objective function, given a one-unit increase in channel capacity. With the cost of this increase in capacity quantified, one can determine if the additional investment in the channel is warranted. This also helps to quantify the opportunity costs of having branch staff shift away from non-campaign related work. Again, this is from the perspective of campaign execution, there are other benefits to channel capacity augmentation that would also have to be taken into account, but at least the campaign benefits would be understood.

4.3 Sensitivity Analysis

In addition to using the solution for capacity planning the approach lends itself to analyzing the sensitivity of the solution to changing assumptions. This is particularly useful for evaluating the solution to changes in the business constraints.

For example, a scenario can be optimized with several alternative budget constraints. **Figure 5** shows the results from solving a model that is very similar to that discussed above. Each row in the table corresponds to a different limit on the total cost of the campaign. All the other constraints and data in the model are held constant. The first scenario allocates at most \$1,000,000 on the campaign, the second \$1,250,000 and the third \$1,500,000. For each of these scenarios the optimal solution is calculated and the expected return from that solution, the number of offers through the branch, and the ROI, are also shown in the figure.

Cost	Return	Branch Activity	ROI
1,000,000	3,756,135	94,088	275%
1,250,000	4,132,821	148,518	230%
1,500,000	4,426,410	250,000	195%

Figure 5. Effects of increased campaign size.

It is interesting to note that for each of these scenarios, as for the example above, the branch capacity is 250,000 offers. However, until the size of the campaign, as measured by the cost of the campaign, reaches \$1,500,000, the branch channel is not used to

its capacity. This means that for campaigns with a budget less than \$1,500,000 there is excess capacity in this channel and its value is lost. The additional \$500,000 expenditure results in an additional \$670,000 return with a 34% ROI and no excess branch capacity.

5. SHORTCOMINGS

The shortcomings of this approach can be broadly classified into two categories: business and technical. From a business perspective there are three perceived shortcomings of this solution. These shortcomings are related to the required inputs to the solution, changes in the types of campaign design decisions, and post analysis, which are performed by the business user. From a technical perspective there are two general shortcomings of this approach. The first is related to the constraints, and the second is related to the acceptable problem size.

5.1 Business Perspective Shortcomings

The approach requires as an input the expected incremental profit associated with each offer and customer combination. Fundamentally this implies the creation, and maintenance, of offer specific probability of response models and detailed profitability measurements. Both of these inputs require ongoing maintenance and enhancement, as the marketplace is ever changing. This requirement does not seem to be too onerous as most organizations that would consider implementing this solution are likely already producing these inputs.

The types of decisions made by business users will be altered with the adoption of this approach. The business users will be making decisions that affect the objective function and constraints but not *directly* about how many offers to make of each type or through which distribution channel. At first the business users might resist this solution, as the decisions that are being made by the business are more abstract than those made during the business as usual campaign design process. Although the business users are making more abstract decisions about campaign design, they are actually gaining more *effective control* over the campaign. Thus the successful adoption of this process would likely involve some amount of business user education with respect to the design of the optimization problem.

The solution was designed to explicitly maximize the expected incremental profitability from running a campaign, not to maximize the overall response rate of the campaign. For instance, a customer might have a higher estimated response probability to a low rate credit card offer than a high rate card. However, the profitability given offer acceptance could be substantially higher for the high rate card. If the goal were to maximize the response rate, then the low rate card would be offered. Whereas, if the goal were to maximize profitability then the offer with the greatest expected profitability would be made. This is only a shortcoming in the sense that business users' expectations have to be managed as campaign response rates are more easily and quickly measured than is campaign profitability.

5.2 Technical Perspective Shortcomings

As this solution was developed in a linear programming framework the constraints must be expressed as linear functions of the choice variables. A couple of examples of plausible

business constraints that are non-linear and therefore do not work with the linear programming approach are:

- The number of credit card offers must be greater than 20,000 or equal to zero.
- The cost of a specific offer is a function of the number of those specific offers being made.

The scalability of this approach has not been exhaustively explored. Tests have been run on data sets with 2.5 million customers, 11 offers and 3 distribution channels and the solution is generated in an acceptable amount of time and resources. The solution was explicitly designed to scale well to the number of customers. Scalability in the offer and channel dimensions is significantly more expensive than along the customer dimension. Although, the number of distribution channels that a company can utilize in an automated campaign is not too large; the number of possible products that could be offered could grow well beyond eleven.

6. CONCLUSION

This offer optimization approach provides three significant improvements over other, more standard, approaches to the problem of campaign design.

- First and foremost, the developed solution produces significantly more incremental profit than competing solutions. As demonstrated in the tactical example, the campaign incremental profit is almost twice as high as that of the standard approach.
- Secondly, this technique is designed to implement multiple constraints and therefore affords the business more control over the direct marketing process. Attempting to satisfy several business constraints simultaneously using ad hoc techniques is a very labor-intensive task and generally produces poor results.
- Finally, the additional information that can be presented as a part of this solution can provide the business with more insight into the customer base, product offerings and the effects of the constraints.

This insight can be used to guide the company to craft better investment decisions in order to make future campaigns even more successful.

7. ACKNOWLEDGMENTS

Our thanks to the three anonymous referees whose suggestions helped clarify our discussion.

8. REFERENCES

- Campell, D., Erdahl, R., Johnson, D., Bibelnicks, E., Haydock, M., Bullock, M., and Crowder, H. (2001), "Optimizing Customer Mail Streams at Fingerhut," *Interfaces*, (31), 77-90.
- Cohen, M. (2000), "Offer Optimization, Optimizing customer value," *SAS Internal Technical Paper*. May 2001.
- Cohen, M., and Parks, J. (2000), "Optimizing the Allocation of Cross-Selling Effort," *Proceedings of the DiamondSug 2000 Conference*, San Francisco.
- Cohen, M. (2000), "Targeted Marketing, Optimizing the customer/event/channel assignment," *SAS Internal Technical Paper*. December 2000.
- Nemhauser, G., and Wolsey, L. *Integer and Combinatorial Optimization*, John Wiley & Sons, New York, 1988.