



U] cã ã æã } Á -Á@Áæ } ~ æÁ |æ } ã *
[-Áæ* ^c^aÁ ~^!•Áæ Áã^&Á æ\^cã *
Ùc^ @æ ã Á^|æ [cÊÜ[^|Š^•Áæ aÁOæ!æ^Áæ|æP [àæ[]



DEPARTMENT OF DECISION SCIENCES AND INFORMATION MANAGEMENT (KBI)

Optimization of the annual planning of targeted offers in direct marketing

Stephanie Delanote*, Roel Leus†, Fabrice Talla Nobibon‡

Abstract. This paper develops an optimization model that captures the annual planning of product-targeting campaigns occurring in financial institutions. The formulation described in this paper extends current product-targeting optimization models found in the literature on two main aspects, namely the possibility of promoting bundled products and the use of multi-channel structures. The complexity of the problem is studied and extensive computational experiments are performed on representative datasets. The effect of client segmentation is empirically examined and some strategic applications of the model are presented. Our work is based on a practical case, but all features identifying the company have been removed.

Keywords: promotion campaigns; product bundling; multi-channel strategies; direct marketing; integer programming.

1. Introduction

The main goal of this paper is to model and analyze the problem of planning annual product-targeting campaigns encountered in financial institutions. Nowadays, the main focus of marketing departments of well-established firms is to maintain and improve relationships with existing customers by offering “the right product to the right customer at the right time” (Cohen, 2004). This process is known as *customer relationship management* (Sheth and Parvatiyar, 1995). The rationale behind this strategy is that customer retention is more profitable than customer acquisition (Reichheld, 1996). The challenge of customer relationship management, however, is to know which product(s) to offer to which customer using which channel(s) while satisfying predefined business constraints. This challenge is particularly pronounced in the retail banking industry, where the market is characterized by high switching costs and psychological factors, which cause loyalty rates to be generally higher than in most other markets (Garland, 2002). To maintain and increase the demand of the existing customers, financial institutions develop profitable long-term relationships with their customers and use customer relationship management to intensify the customer lock-in (Garland, 2002).

In this paper, we develop a mathematical model that aims at deciding which offer should be targeted to which customer segment through which sequence of channel(s) in order to maximize the bank’s expected profit while respecting predefined business constraints. This model combines the view of the product, the segment, and the channel managers in developing a campaign planning. Our model extends existing product-targeting models on two important aspects. The first extension is *product bundling*, which considers the possibility of offering a bundle consisting of several products to a client. This practice exploits, among others, complementarity relationships amongst products, economies of scope (Chae, 1992), price

*ORSTAT, Faculty of Business and Economics, KULeuven, Naamsestraat 69, B-3000 Leuven, Belgium. E-mail: stephanie_delanote@hotmail.com

†ORSTAT, Faculty of Business and Economics, KULeuven, Naamsestraat 69, B-3000 Leuven, Belgium. E-mail: Roel.Leus@econ.kuleuven.be

‡QuantOM, HEC-Management School, University of Liège, Rue Louvrex 14, Building N1, B-4000 Liège, Belgium. E-mail: Fabrice.TallaNobibon@ulg.ac.be

differentiation (Pigou, 1920), and strategic motives, such as intensifying customer lock-in and offering product differentiation (Mankila, 1999). The increasing number of complementary financial products has made bundling an important and widely used cross-selling tactic in retail banking (Hayes and Eduardo, 2003). The second extension is the use of multi-channel strategies. *Multi-channeling* involves using different channels to make offers to a client with the possibility of using them consecutively for the same offer to the same client (Frazier, 1999). Due to increased competition in the financial sector (Ennew, 1992) and the fact that the banks' traditional channels (the branches and call centers) are no longer sufficient (Faust, 1990), most of their campaigns are multi-channel, in order to reach a sufficiently large public and to enhance competitive strength. Our work is based on a practical case, but all features identifying the company have been removed.

The scientific contributions of this paper are fourfold:

1. The development of a mathematical model for the problem of planning product-targeting campaigns, including product bundling and multi-channeling.
2. The study of the complexity of the problem.
3. Extensive computational experiments on representative datasets.
4. Some strategic applications of our model.

The rest of this paper is structured as follows. Section 2 is devoted to the formal description of the problem. In Section 3, we develop a mathematical formulation of the problem. In Section 4, we review the existing literature and we study the complexity of the problem in Section 5. In Section 6 we comment the computational results and discuss some strategic uses of our model. We conclude in Section 7.

2. Problem description

This paper examines the optimization of the annual planning of targeted offers in direct marketing in a financial institution. The marketing plan is an intermediary step between the firm's marketing strategies and operational marketing. Within a bank, marketing strategies are decided by the marketing director as part of the bank's long-term strategic decisions. Operational marketing, on the other hand, refers to a short-term implementation of marketing campaign decisions. Our study focuses on the mid-term marketing planning problem, which develops an annual framework in which every operational campaign has to fit. The planning tells operational managers exactly for which offers promotions have to be sent to which segment of customers and which channels have to be used in the different promotion rounds. As a consequence, this problem has four main decision levels, each with its own set of constraints.

First, the multitude of bank products is organized into categories, with products of the same category being regarded as substitutes. As examples of such categories we have *daily banking*, containing basic financial products such as a credit card or a debit card; the *save & invest* category, containing mainly saving accounts, pension planning and investment products; the *lending* category, which consists of all types of credits; and finally *protecting*,

which deals with insurances. Banking products of the same category are rather homogeneous whereas important differences exist between products that belong to different categories. The mid-term marketing plan is developed at the level of the product categories, assuming substitutivity within and complementarity among the different categories. A campaign in a bank may involve the promotion of a product or a bundle (a set of products) resulting from combining products from different categories. Once the offers are selected, every chosen category is represented by one of its individual products; the choice of that product is made at the operational level of the bank and thus lies outside the scope of this paper. During the planning activities, product managers ensure that: (1) only offers for which a campaign outputs a positive value creation can be promoted, (2) the total cost of the promotion of all selected offers fits within the budget of the marketing department, (3) a client receives at most one offer; multiple products should be promoted in a bundle.

Second, the clients of a bank are organized into *segments*, based for example on their financial background. Mid-term promotion planning is based on these segments, with customers from the same segment being regarded as identical. A more individualized approach to the customers is only applied at the operational level. For each segment, segment managers verify that enough campaigns are scheduled in the planning to meet the expected minimum sales target.

Third, different *channels* are available within a bank for sending offers to clients, such as E-mail, banners on PC-banking accounts, direct mail and the call center. Some of these channels have restrictive capacities, and channel managers are responsible for respecting upper bounds on the number of promotions through the different channels.

Fourth, promotion campaigns in a bank are organized sequentially over time in different *rounds*. A maximum number of rounds is allowed for trying to convince a particular client. This number is set in order to avoid a saturation effect with the client and a “product push” image of the bank. Each round is characterized by a particular channel choice, which gives rise to multi-channeling. Channel managers enforce the sequential use of the available channels, based on an ordering of channels that is pre-specified (and so not open to change by the channel managers). This ordering may depend on the segment of clients and on the considered offer, and is based on the following principle: the more complex and profitable an offer is, and the better the financial background of a client, the more the bank is willing to use its most expensive channels such as its call center and direct mailing.

3. Problem formulation

In this section we propose a mathematical model for planning targeted offers that maximizes the expected profit while satisfying all the restrictions from product, segment, and channel managers. The constraints and the parameters of our model reflect the situation of a Belgian bank. Below we first present the parameters of our model, followed by the definition of the decision variables. We subsequently describe the objective function and the constraints.

3.1. Parameters

Let m be the number of segments of clients, n the number of offers, ℓ the number of channels and R the maximum number of rounds. For every segment i ($i = 1, \dots, m$) let C_i , N_i and O_i respectively be the fixed cost associated with the segment, the number of clients in the segment and the minimum sales target for the segment. We define F_j ($j = 1, \dots, n$) to be the fixed cost associated with the promotion of offer j and G_k ($k = 1, \dots, \ell$) the fixed cost associated with the use of channel k . Furthermore, UB_k is the capacity of channel k .

The marketing department budget B for the next year is fixed before starting the planning. Let CLV_{ij} ($i = 1, \dots, m; j = 1, \dots, n$) be the net present value for the bank when a client of segment i reacts positively to an offer j ('CLV' stands for *customer lifetime value*). Every offer promises a certain degree of profitability, which additionally depends on the profile of the client and on his or her expected investment capacity. The customer lifetime value is combined with the expected hit ratio P_{ijk}^r in order to obtain the expected net present value from a particular promotion. Notice that this hit probability may depend upon the round, the segment, the product, and the channel used; and is such that the more personal and direct the channel is, the more responsive the customer is expected to be (Frazier, 1999). Finally, by CSP_{ijk}^r we denote the *expected cross-selling potential*, obtained by multiplying CLV_{ij} and P_{ijk}^r . The parameters described above are based on historical data and are assumed to be available within the bank.

3.2. Decision variables

We use the following four sets of decision variables in our model.

- y_j is a binary decision variable that indicates whether a promotion of offer j is part of the planning ($y_j = 1$) or not ($y_j = 0$), for $j = 1, \dots, n$.
- u_k is a binary decision variable that indicates whether channel k is used for the planning ($u_k = 1$) or not ($u_k = 0$), for $k = 1, \dots, \ell$.
- z_{ijk} is a binary decision variable that indicates whether a promotion campaign is developed to promote offer j through channel k to customer segment i ($z_{ijk} = 1$) or not ($z_{ijk} = 0$), for $i = 1, \dots, m; j = 1, \dots, n; k = 1, \dots, \ell$.
- x_{ijk}^r is an integer decision variable indicating the number of customers of segment i that will be contacted through channel k for the promotion of offer j during promotion round r , for $i = 1, \dots, m; j = 1, \dots, n; k = 1, \dots, \ell; r = 1, \dots, R$.

3.3. Objective function

The objective of maximizing the expected profit can be modeled as follows:

$$\text{Maximize } \sum_{i=1}^m \sum_{j=1}^n \sum_{k=1}^{\ell} \sum_{r=1}^R CSP_{ijk}^r x_{ijk}^r - \sum_{j=1}^n F_j y_j - \sum_{k=1}^{\ell} G_k u_k. \quad (1)$$

The objective function (1) consists of three parts. The first part expresses the total expected cross-selling potential received from clients, the second part refers to the total fixed costs that are incurred every time a new campaign is developed and the third part represents extra investments needed to use the different channels. Note that C_i does not appear in (1) because it does not affect the optimization problem.

3.4. Business constraints

Below, we group the constraints according to the type of manager responsible.

Product managers

The first constraint states that only those offers are considered that have expected net returns coming from cross-selling potentials that offset the investments. This translates into:

$$\sum_{i=1}^m \sum_{k=1}^{\ell} \sum_{r=1}^R CSP_{ijk}^r x_{ijk}^r - F_j y_j \geq 0, \quad j = 1, \dots, n. \quad (2)$$

Second, the budget of the marketing department must not be exceeded; this is modeled as:

$$\sum_{j=1}^n F_j y_j + \sum_{k=1}^{\ell} G_k u_k \leq B - \sum_{i=1}^m C_i. \quad (3)$$

Third, each client receives the promotion of at most one offer. This is guaranteed by (4).

$$\sum_{j=1}^n z_{ij1} \leq 1, \quad i = 1, \dots, m. \quad (4)$$

Fourth, we cannot send more offers than the number of clients in a given segment.

$$x_{ijk}^r \leq N_i z_{ijk}, \quad i = 1, \dots, m; j = 1, \dots, n; k = 1, \dots, \ell; r = 1, \dots, R. \quad (5)$$

Fifth, if an offer is not part of the campaign, it is not sent to any client segment i .

$$z_{ijk} \leq y_j, \quad i = 1, \dots, m; j = 1, \dots, n; k = 1, \dots, \ell. \quad (6)$$

Finally, when a channel is not operational, no offer should be sent through that channel.

$$z_{ijk} \leq u_k, \quad i = 1, \dots, m; j = 1, \dots, n; k = 1, \dots, \ell. \quad (7)$$

Segment managers

Segment managers ensure that for every segment of clients the expected minimum sales target is reached:

$$\sum_{j=1}^n \sum_{k=1}^{\ell} \sum_{r=1}^R P_{ijk}^r x_{ijk}^r \geq O_i, \quad i = 1, \dots, m. \quad (8)$$

Channel managers

Within the bank, for each offer and each segment, internal “priority rules” have been established regarding the channels allowable for promotion of the offer to the segment, which take the form of an ordering of the channels. Concretely, the first channel in the list is the only one that can be used during the first round. The second channel can be considered together with the first during the second round, to target those customers who received the offer during the first round but did not react positively, and so on. For ease of exposition, we use k to refer to the k^{th} channel in the list for a given combination of offer and segment.

The set of constraints (9) ensures that only the first r channels in the list can be used during the first r rounds of the campaign.

$$x_{ijk}^r = 0, \quad i = 1, \dots, m; j = 1, \dots, n; r = 1, \dots, R; k = r + 1, \dots, \ell. \quad (9)$$

The set of constraints (10) enforces that in round $r + 1$ only those clients who received an offer in round r and did not react positively are targeted.

$$\sum_{k=1}^{\min\{r,\ell\}} (1 - P_{ijk}^r) x_{ijk}^r \geq \sum_{k=1}^{\min\{r+1,\ell\}} x_{ijk}^{r+1}, \quad i = 1, \dots, m; j = 1, \dots, n; r = 1, \dots, R - 1. \quad (10)$$

The capacity constraints of the channels are given by:

$$\sum_{i=1}^m \sum_{j=1}^n x_{ijk}^r \leq UB_k, \quad k = 1, \dots, \ell; r = 1, \dots, R. \quad (11)$$

In the remainder of this paper, we adopt the name *PlanCamp* for the optimization problem described by the objective function (1) and the constraints (2)–(11).

4. Literature review

This section reviews the existing literature dealing with product-targeting models in direct marketing and elaborates on two important extensions included in PlanCamp.

4.1. Direct marketing models

There are few articles dealing with the product-targeting problem in the literature. As observed by Talla Nobibon et al. (2011) and Cohen (2004), this problem is frequently approached by applying a class of existing statistical methods known as “response models.” These methods provide support for the selection process for a promotion campaign by delivering input data based on the prediction of response behavior (Piersma and Jonker, 2004). A full review of these response models is beyond the scope of this paper; below we focus only on articles dealing explicitly with the product-targeting problem.

Bhaskar et al. (2009) consider the process of selecting customers for a cross-sell campaign in a retail bank. Their objective is to maximize the total profit subject to two constraints: First, the campaign budget translates into an upper bound on the number of customers that

can be targeted. Second, the expected volume bought by the targeted customers should meet an asset growth target. The key feature of their paper is that the authors take into account uncertainty in the parameters of the model and solve a fuzzy mathematical programming problem on a real cross-selling instance from a retail bank. Cohen (2004) focuses on a similar problem with three constraints. His model maximizes the bank’s net income collected from its cross-selling campaigns under the minimum sales constraint, the budget constraint and the return-on-investment hurdle-rate constraint. Cohen proposes a mathematical programming formulation of the problem, which is applied for solving a real case at Scotiabank. Recently, Talla Nobibon et al. (2011) have proposed a model that maximizes the total profit obtained from cross-selling campaigns while satisfying the following constraints: the budget constraint, the return-on-investment constraint, the minimum quantity commitment constraint and an upper bound on the number of offers to a customer. They propose exact algorithms and several heuristics for solving the problem.

4.2. Extensions

PlanCamp extends the current literature by adding more constraints and modeling new characteristics of the problem. Below, we highlight two important extensions: the promotion of bundles of products and the use of multi-channel follow-up campaigns.

Our model includes the simultaneous cross-selling of several products under the form of bundles. This extends the marketing literature on bundling on two different aspects: first, the bundling strategy considered is subject to rigid constraints imposed by Belgian legislation, which forbids tie-in sales. Bundling becomes a tie-in when at least one component of the bundle is not available separately and, as consequence, customers are forced to buy the bundle (Burstein, 1960). Second, we complement the large existing body of marketing literature on product bundling, where prior attention is paid to the selection of an optimal bundling strategy. Adams and Yellen (1976) identify three bundling strategies: the “pure component” strategy, which is the offer of products only as separate items; the “mixed bundling” strategy that combines both individual products and their bundles; and the “pure bundling” strategy that offers only packages of products. The marketing and economic literature has mainly focused on the problem of choosing an optimal bundling strategy (Adams and Yellen, 1976; Bakos and Brynjolfsson, 1999; Banciu and Prakash Mirchandani, 2010; Ibragimov and Walden, 2010; Prasad et al., 2010). Our model does not take the bundling strategy as a variable but rather adopts the mixed bundling strategy to comply with Belgian legislation.

Our model also integrates the possibility of “multi-channeling,” where offers are promoted through many different channels that are also used for sequential follow-up promotions. The problem of finding the optimal channel(s) for contacting a particular client is neither considered in Bhaskar et al. (2009) nor in Talla Nobibon et al. (2011), who recognize it as an important limitation of their model. Cohen (2004) does consider different promotion channels with their costs and capacities, but follow-up promotions are not included. Nair and Tarasewich (2003) consider multi-period sales promotions and their optimization over time, but they ignore the optimal sequential use of channels and consider the follow-ups as uniform. Our model integrates multi-channeling and follow-up promotions, thus generalizing the listed references.

5. Theoretical analysis

This section is devoted to a theoretical study of our problem. We first observe that the integrality properties of variables y_j ($j = 1, \dots, n$) and u_k ($k = 1, \dots, \ell$) can be relaxed without impacting the convex hull of the set of feasible solutions to PlanCamp. Next, we argue that our problem contains the generalized assignment problem (GAP) as a special case and subsequently we prove the nonapproximability of the problem.

Proposition 1. *The convex hull of feasible solutions to PlanCamp is identical to the convex hull of the solutions to the problem formulation with the relaxation of the integrality constraints $y_j \in \{0, 1\}$ to $0 \leq y_j \leq 1$ for all $j = 1, \dots, n$.*

Proof: Let CH be the convex hull of the feasible solutions to PlanCamp and let CH^* be the convex hull of solutions to the problem with relaxation $0 \leq y_j \leq 1$ for all $j = 1, \dots, n$. Because any feasible solution to PlanCamp satisfies $0 \leq y_j \leq 1$, the set of feasible solutions to the original problem is included in CH^* and the convexity of CH^* implies that $CH \subseteq CH^*$. Next, let $(z^*, x^*, y^*, u^*) \in CH^*$. We prove, by contradiction, that $y_j^* \in \{0, 1\}$ for all j . If there was a j' such that $0 < y_{j'}^* < 1$, then from (6) we would have $z_{ij'k}^* = 0 \forall i, k$, and from (5) $x_{ij'k}^{*r} = 0 \forall i, k, r$. However, this will violate constraint (2) for j' ; therefore $y_j^* \in \{0, 1\} \forall j$ and hence $(z^*, x^*, y^*, u^*) \in CH$, implying that $CH^* \subseteq CH$. Since $CH \subseteq CH^*$ and $CH^* \subseteq CH$ we conclude that $CH = CH^*$. \square

Proposition 2. *The relaxation of the integrality constraint $u_k \in \{0, 1\}$ to $0 \leq u_k \leq 1$ for all k does not change the optimal extreme points in the convex hull.*

Proof: Consider an optimal solution (z, x, y, u) that is such that there exists a $k' \in \{1, \dots, \ell\}$ with $0 < u_{k'} < 1$. This would make $z_{ijk'} = 0$ for all i and j due to (7) and $x_{ijk'}^r = 0$ for all i, j and r from (5). Consider the solution (z^*, x^*, y^*, u^*) defined by $z^* = z$, $x^* = x$, $y^* = y$ and $u_k^* = 1$ if $u_k = 1$ else $u_k^* = 0$. Since $G_k > 0$ for all k the objective value of the solution (z^*, x^*, y^*, u^*) is larger than that of (z, x, y, u) . This implies that (z, x, y, u) is not optimal, a contradiction. \square

The next result proves that PlanCamp is strongly NP-hard by showing that it contains GAP as a special case.

Proposition 3. *The problem PlanCamp contains GAP as a special case.*

Proof: Consider the following special case of PlanCamp involving a single offer ($n = 1$), a single round ($R = 1$) and each channel can be used at most once ($UB_k = 1$ for all $k = 1, \dots, \ell$). Furthermore, we set $C_i = 0, \forall i = 1, \dots, m$, $F_1 = 0$ for the unique offer and $G_k = 0, \forall k = 1, \dots, \ell$. Under these considerations, we drop the indices j and r and create new parameters CSP_{ik} , P_{ik} and a new decision variable x_{ik} . Note that constraints (2)–(7)

and (9)–(10) are redundant. PlanCamp now reduces to:

$$\text{Maximize} \quad \sum_{i=1}^m \sum_{k=1}^{\ell} CSP_{ik} x_{ik} \quad (12)$$

$$\text{subject to} \quad \sum_{k=1}^{\ell} P_{ik} x_{ik} \geq O_i, \quad i = 1, \dots, m, \quad (13)$$

$$\sum_{i=1}^m x_{ik} \leq 1, \quad k = 1, \dots, \ell, \quad (14)$$

$$x_{ik} \in \{0, 1\}, \quad i = 1, \dots, m; k = 1, \dots, \ell, \quad (15)$$

which is equivalent to GAP (Kellerer et al., 2004). \square

In what follows, we prove that any polynomial-time constant-factor approximation algorithm for solving PlanCamp can be used to solve the following variant of Partition:

Instance: A finite set $A = \{1, 2, \dots, 2q\}$ (where q is an integer greater than 0) with size $s(i) \in \mathbb{Z}^+$ for each $i \in A$, and $K = \frac{1}{2} \sum_{i \in A} s(i)$.

Question: Does there exist a subset $A' \subset A$ with $|A'| = q$ and $\sum_{i \in A'} s(i) = K$?

Proposition 4. *Unless $P = NP$, there is no polynomial-time algorithm that guarantees a constant-factor approximation for solving PlanCamp.*

Proof: Given an arbitrary instance of Partition, consider the following reduction to an instance of PlanCamp with $n = 1$, $R = 1$, $m = 1$, $\ell = 2q$, $B = (2q + 1)K$, and $O_1 = q$. Furthermore, $UB_k = 1$, $G_k = 2K + s(k)$, $CSP_k = \frac{(2q+1)K}{q}$, and $P_k = 1$ for all $k = 1, \dots, \ell$.

$$\text{Maximize} \quad \sum_{k=1}^{\ell} \frac{(2q+1)K}{q} x_k - \sum_{k=1}^{\ell} G_k u_k \quad (16)$$

$$\text{subject to} \quad \sum_{k=1}^{\ell} G_k u_k \leq B, \quad (17)$$

$$\sum_{k=1}^{\ell} x_k \geq q, \quad (18)$$

$$x_k \leq u_k, \quad k = 1, \dots, \ell, \quad (19)$$

$$x_k, u_k \in \{0, 1\}, \quad k = 1, \dots, \ell. \quad (20)$$

We prove that by solving this instance of PlanCamp with a constant-factor approximation algorithm, we can infer the answer to the Partition instance.

If the instance of Partition is a yes instance then an appropriate set A' exists. We consider the solution of our problem that fixes $x_k = 1$ for all $k \in A'$ and $u_k = 0$ otherwise. The objective value of this solution is exactly 0. Therefore, a constant-factor approximation algorithm will provide an optimal solution, from which we can easily infer the elements of A' . Conversely, if a constant-factor approximation algorithm leads to an objective value different from 0 or declares the instance to be infeasible, then we conclude that the instance of Partition is a no instance. \square

6. Computational experiments

Our experiments were run on a Packard Bell laptop iPower GX Series with AMD Athlon (TM) X2 processor with 1.90 Ghz clock speed and 4 GB RAM, equipped with Windows Vista. CPLEX 12.2 is used for solving mixed integer programming (MIP) instances. Below, we first provide some details on the generation of the instances and, subsequently, we discuss computational results.

6.1. Generating test instances

For confidentiality reasons, we cannot report computational results on real-life bank data. Instead, we generate random input data that resemble real-life data reasonably well; the generation settings are based on qualitative considerations provided by managers. We generate instances with 15 offers ($n = 15$) from 4 product categories, 4 customer segments ($m = 4$), 4 channels ($\ell = 4$), and at most 12 rounds ($R = 12$). We set $N_1 = 1\,800\,000$, $N_2 = 900\,000$, $N_3 = 210\,000$, and $N_4 = 90\,000$. For every segment i ($i = 1, \dots, 4$), the minimum sales target O_i is randomly chosen between $\lfloor \frac{N_i}{100} \rfloor$ and $\lfloor \frac{5N_i}{100} \rfloor$.

The fixed costs C_i ($i = 1, \dots, 4$) and G_k ($k = 1, \dots, 4$) are proportional to the number of managers involved. In practice, up to five managers can be assigned to the same segment or channel and each manager earns about € 24 000 yearly. Therefore, C_i and G_k are multiples of 24 000, with an upper bound of 120 000. The values of F_j are determined by the magnitude of the campaign. Between € 10 000 and € 500 000 is a *small* campaign, between € 500 000 and € 1 000 000 we have an *average* campaign and a *large* campaign costs between € 1 000 000 and € 2 000 000. Furthermore, reports from financial institutions suggest that banks usually run small campaigns and can run at most 2 large campaigns during a year. We randomly associate 8 offers with a small campaign, 5 offers with an average campaign and 2 offers with a large campaign. The value of F_j is then randomly generated between the corresponding bounds. Finally, the budget B is estimated by taking into account the three types of fixed costs as follows: for both C_i and G_k , we sum the lowest value and the highest value whereas for F_j we add the highest value for small campaigns, the highest value for average campaigns and the highest value for large campaigns. The total sum over these three types of costs is taken as the budget.

For the capacities of the channels, we estimate three non-restrictive values and one restrictive channel. This represents the common setting in financial institutions, where the most frequently used channels are mail, e-mail, PC banking, and the call center; the first three are non-restrictive. The non-restrictive channels can handle promotion to all the clients whereas the restrictive one has a limited capacity of 50 000 per month. The priority order of the four channels for every combination of offer and clients segment is randomly generated.

For generating hit ratios we rely on bounds proposed in the services-marketing literature (Lovelock and Wirtz, 2006). A suitable hit ratio for pure direct marketing is estimated to be between 0 and 1% whereas for direct marketing supported by a “sales effort,” for example via a call center, hit ratios are estimated between 5% and 15%. Therefore, P_{ijk}^r is chosen between 0.00001 and 0.01 for non-restrictive channels and between 0.05 and 0.15 for the restrictive channel. To capture the saturation effect (Johansson, 1979), the hit ratios start at the lower bound (0.00001 or 0.05 depending on the channel) at round 1 and increase

to reach the upper bound (of 0.01 or 0.15) at rounds 6 and 7. From round 7 up to round 12 the saturation effects lead to decreasing hit ratios, and the ratios reach their lower bound at round 12.

Customer lifetime values within a bank range from € 10 to € 50 for the most simple products such as debit cards, and from € 1 000 to € 5 000 for the more complex offers, for instance for expensive mortgage loans. We first randomly determine for every offer j whether it is a simple offer or a complex offer. Next, we randomly generate for every client segment i a value for CLV_{ij} between the above-mentioned lower and upper bound. In order to obtain a representative set of test instances, we generate a sample of 10 different instances following the described procedure.

6.2. Computational results

Table 1 reports the objective values and running times for the ten test instances (times are expressed in seconds). All the instances are solved to optimality within an average computation time of 18.58 seconds, with the hardest instance requiring less than 45 seconds.

The size of a PlanCamp instance is dependent on the values of n , m , ℓ and R . The parameter n is not expected to drastically increase over time, because PlanCamp deals with product categories and not with individual products. Similarly, the number ℓ of actively used channels within a bank can be expected to remain limited – banks have difficulties managing multi-channel structures with more than five channels (Coelho and Easingwood, 2003). Also, the number R of promotion rounds is unlikely to increase, this because of the saturation effects (Johansson, 1979). Consequently, only the number m of customer segments may reasonably be expected to change and affect instance sizes and thus running times. Garland (2005) provides insights in the most recent segmentation approaches in retail banking, which utilize personal information and psychographic characteristics of the customers. From our test instances, we have generated new instances with m increasing from 4 to 6, 8, 10 and 12. We create these additional segments by splitting up our 4 segments until each original segment contains three subsegments ($m = 12$). We generate parameters for these new instances in two different ways: first, we keep the parameter values for the subsegments equal to the values of their original (parent) segment. This reflects a situation where customers within existing segments are already extremely homogeneous, for example in the case of specialized niche banks. In a second step, we do allow for differences in parameter values for subsegments, but we ensure that the averages are still the same as in the original instances. This represents the setting where customer segments are less homogeneous. We generate one case with a limited variety in values (up to 20% deviation from the average) and

Test	Objective value	CPU time	Test	Objective value	CPU time
1	1 062 706 158.34	10.45	6	763 074 401.53	44.39
2	1 060 020 854.34	10.06	7	990 074 321.91	09.62
3	1 082 125 558.29	43.35	8	964 934 680.02	10.81
4	1 087 667 767.52	33.38	9	1 001 388 310.74	06.86
5	1 055 155 842.87	06.58	10	1 054 603 345.92	10.34

Table 1: Objective values and CPU times of the test instances.

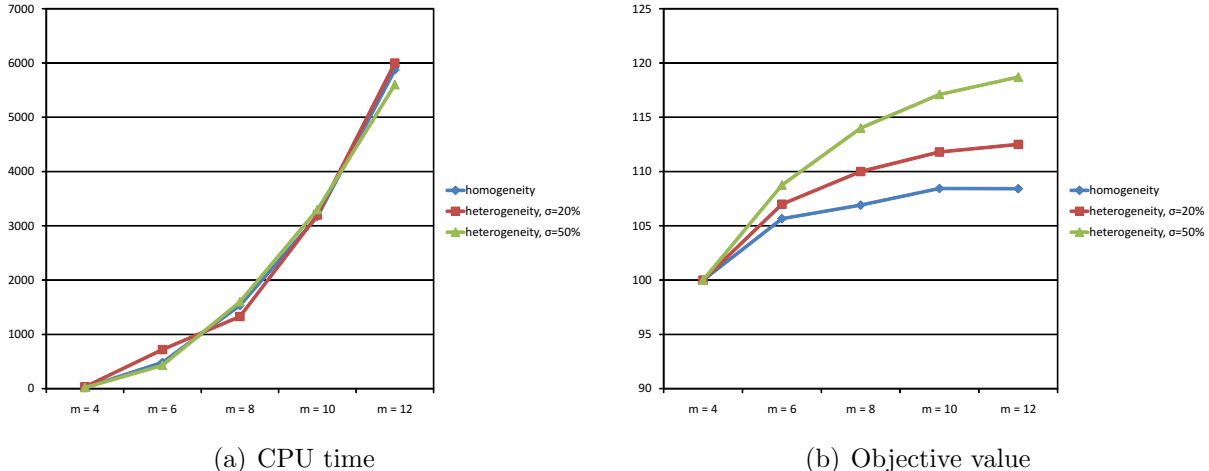


Figure 1: Evolution of the CPU time and the objective value as function of m .

another case where parameter values may deviate up to 50% from the average values. These settings will enable us to analyze the impact of the composition of the customer segments. We solve the new instances and investigate the evolution of the objective values and the computation times as a function of m . A graphical representation of the results is shown in Figure 1.

Figure 1(a) shows that when m increases the CPU time grows in a strongly non-linear fashion, and this independently of the assumptions for the parameter values of the different segments. The three plots, representing the three cases we have considered, almost coincide, meaning that there is virtually no difference in run times between the cases. Figure 1(b) allows to examine the dependence of the objective values on the degree of heterogeneity of the characteristics of the customer base. Contrary to the CPU time, significant differences are observed between the three plots. For homogeneous segments and when $m \leq 10$, the average objective value increases significantly with m , indicating that refinement of the customer segmentation leads to better results; this is in line with the observation made in Garland (2005). When m becomes greater than 10, however, the average objective value tends to stabilize and even slightly decrease. A possible explanation is that when m is greater than 10 the extra segment costs exceed the additional expected profit associated with the new segments; 10 is here a threshold value for m .

For the heterogeneous cases we observe that the level of homogeneity within a customer segment affects the threshold value of m . Indeed, when customers segments differ more from each other, more benefits can be achieved from an increased segmentation compared to the situation with homogeneous clients. Also, the threshold value of m is higher, which is in line with the marketing literature (Wedel and Kamakura, 2003), where variety among customer preferences and characteristics is considered as the fundamental driver behind customer segmentation. Regardless of the level of homogeneity of customer segments, the average expected profits increase with the number of customer segments.

We have also studied the effect of generating the hit ratios as function of the round and the channel both on the objective value and on the CPU times. The rationale behind this

Test	Objective value	CPU time	Test	Objective value	CPU time
1	1 032 641 523.93	08.76	6	812 297 415.66	10.16
2	1 078 823 994.27	09.24	7	979 169 488.17	08.26
3	993 520 716.85	11.04	8	999 511 888.25	07.11
4	1 012 562 134.03	07.96	9	995 773 234.10	04.93
5	1 076 247 963.11	04.96	10	1 003 512 379.26	07.13

Table 2: Objective values and CPU times of the test instances with adjusted hit ratios.

analysis is the fact that many banks use hit ratios that depend only on the client segment and the offer. We consider our initial test instances and aggregate the hit ratios P_{ijk}^r over the indices r and k to obtain an average value that will serve as the new hit ratio P_{ij} , which depends only on i and j . These adapted instances are solved and the results are reported in Table 2.

We observe a reduction in CPU time for all ten instances, with an average CPU time of 7.95 seconds (compared to 18.58 in Table 1); the highest CPU time is now 11.04 seconds. In summary, simplifying the assumptions regarding the hit ratios clearly reduces the running time. A possible explanation is that consumer behavior is now constant over the different promotion rounds and the responsiveness is independent of the channel used, making the capacities the only distinctive characteristic of the channels. The objective values go down by only 1.49% on average compared to the results in Table 1, which makes this simplifying approach a good alternative to more complex estimations.

Finally, we investigate the effect of channels on the outcomes of the problem. We first perform a sensitivity analysis with respect to the capacity of the restrictive channel. Figure 2 plots the ratio of the average objective values of the new instances over the average initial values as a function of the capacity of the restrictive channel. We find that when the capacity increases up to approximately 90 000 the average objective value also increases by about 20%. When we further increase the capacity until it becomes non-restrictive (for 3 000 000) we obtain an additional increase of the above ratio of approximately 10%. This extra increase stems from follow-ups that are now organized using the most responsive chan-




Figure 2: Objective value (%) as a function of the capacity of the call center.

nel. We conclude that the positive effect of increasing the capacity of the restrictive channel is mainly achieved by enlarging the reach of the bank’s campaigns, rather than having more effective follow-ups.

We have also studied the impact of fixing a priori the sequence in which the channels are used for a given offer and segment. Note that this is the current practice in several banks. In a new optimization run, we consider the promotion channels as variable and optimize the multi-channeling choices. In the new problem formulation, constraints (9) are removed and constraints (10) include only a sum over $k = 1, \dots, \ell$. We observe an increase in the average expected objective value of 18.54% compared to the initial setting. This observation suggests that abandonment of this practice may offer considerable benefits.

6.3. Strategic applications of our model

In this section we elaborate on some strategic applications of PlanCamp. We first illustrate how PlanCamp can be used to support strategic decisions concerning a bank’s product bundling strategy and subsequently we analyze the multi-channel structure of the bank.

Strategic product bundling

There are two main motivations behind product bundling in the financial industry. First, there is the high complementarity of financial products that causes reservation prices and response probabilities for their bundled offer to be significantly higher than their separate offer (Venkatesh and Kamakura, 2003). The second motivation is the high degree of fixed costs relative to variable costs that creates possible economies of scale when bundling offers (Mankila, 1999). This section studies the following questions: (1) “Is there an advantage of product bundling?” and if yes, (2) “which of the above two motivations is mostly influencing it?”.

We consider the optimization of PlanCamp on the initial dataset as the base case, where bundling is applied without assuming economies of scale in the campaign costs or complementary products with advantageous customer lifetime values or hit ratios. In order to benchmark with a situation where no bundling is applied, we only retain the single offers of the four product categories and allow every customer segment to receive multiple separate promotions, instead of imposing to promote one single bundle of these promotions to every segment. This results in a decrease of the average objective value by 27.8%, caused by a decrease in the objective value for every single instance. As no explicit financial motivation in terms of complementarity and economies of scale was imposed, the bank finds an extra, non-financial, motivation for bundling, in that it allows for a more efficient capacity usage. By promoting bundles of offers through a unit channel capacity, we generate solutions that would even be infeasible in the non-bundling case, since these offers would then all require separate promotion capacity.

Next, we further adapt our instances such that we can evaluate the effect of the two theoretical motivations of economies of scale and complementarities. In the former case, we make campaign costs of bundles subadditive and equal to 70% of the sum of the separate campaign costs. For these new instances, we obtain an average objective value that is 5% higher than the initial setting. In the latter case, we assume that customer lifetime

values increase on average with 20%, together with hit ratios that now approximate the simple sum of the hit ratios for the individual offers. For these instances, we obtain an average objective value that is about 245.8% higher. These results confirm the cost economies and the advantageous effects of complementary products as two financial drivers behind product bundling. Complementary products, however, turn out to dominate cost economies, indicating that banks should give primary attention to deciding a smart and well-considered composition of its bundles.

Strategic multi-channeling

Easingwood and Storey (1996) examine the use of promotion channels in the marketing of financial products in the UK and find that in 85% of the cases, the channel structure belongs to one of the following three categories: an “arm’s length” strategy, which employs those channels that do not involve direct contact between company staff and customers, a “network” strategy, which only utilizes channels with direct, personal contact, or a “balanced” strategy that combines the high volume, impersonal channels with direct channels. The latter is the one considered for our initial instances. We now briefly investigate the possibility of using a network or an arm’s length strategy.

For the analysis of a network strategy we adapt our instances by assuming that the four channels are restrictive direct channels with the same capacity of 50 000 and the hit ratios are chosen to be comparable to those of the only restrictive channel in our initial instances. Solving these instances yields an average objective value that is 28.3% lower than that of the initial instances. Furthermore, the reduction is observed in the objective value of all individual instances. For the instances generated, the small reach of the campaigns turns out to dominate the advantage of higher response rates.

For the arm’s length strategy, we consider the initial instances and transform the only direct and restrictive channel into an impersonal channel. We then have instances with four channels that can all handle non-restrictive volumes, but that suffer low response rates. These instances produce solutions with an average objective value that is 12.07% higher than that of the initial instances. The results for individual instances, however, are not uniformly positive.

In conclusion, we state that a network strategy may be a bad choice for banks of considerable size, due to the limited reach. An arm’s length strategy, by contrast, can be advantageous when the use of direct channels limits the reached volume. A balanced strategy that tries to improve its trade-off between responsiveness and high volumes seems to constitute a safe choice for large commercial banks.

7. Conclusions

This paper studies the development of an efficient plan for annual targeting offers in direct marketing in financial institutions. Special attention is paid to the practice of product bundling and multi-channeling, which are two important extensions of the current literature. We prove that the proposed problem is strongly NP-hard and perform extensive computational experiments. We observe empirically that intensifying a bank’s segments has

a positive impact on the expected profit. Segmentation efforts are not useful anymore, however, when the costs of managing new segments dominate the additional profit stemming from segmentation. Based on our model, we have performed a strategic analysis that leads to the recommendation of extending the capacity of the restrictive channel and to the adaptation of channel priority rules. Furthermore, we find that the choice of a multi-channel structure should strike an optimal trade-off between high responsiveness based on personal but time-consuming contacts and reaching high volumes.

References

- Adams, W.J. and J.L. Yellen (1976), ‘Commodity bundling and the burden of monopoly’, *The Quarterly Journal of Economics* **90**(3), 475–498.
- Bakos, Y. and E. Brynjolfsson (1999), ‘Bundling information goods: Pricing, profits and efficiency’, *Management Science* **45**(12), 1613–1630.
- Banciu, M. and E.G. Prakash Mirchandani (2010), ‘Bundling strategies when products are vertically differentiated and capacities are limited’, *Management Science* **56**(12), 2207–2223.
- Bhaskar, T., R. Sundararajan and P.G. Krishnan (2009), ‘A fuzzy mathematical programming approach for cross-sell optimization in retail banking’, *Journal of the Operational Research Society* **60**(5), 717–727.
- Burstein, M.L. (1960), ‘The economics of tie-in sales’, *Review of Economics and Statistics* **42**(6), 8–73.
- Chae, S. (1992), ‘Bundling subscription TV channels: A case of natural bundling’, *International Journal of Industrial Organization* **10**(2), 213–230.
- Coelho, F.J. and C. Easingwood (2003), ‘Multiple channel structures in financial services: A framework’, *Journal of Financial Services Marketing* **8**(1), 22–34.
- Cohen, M.-D. (2004), ‘Exploiting response models-optimizing cross-sell and up-sell opportunities in banking’, *Information Systems* **29**(4), 327–341.
- Easingwood, C. and C. Storey (1996), ‘The value of multi-channel distribution systems in the financial services sector’, *The Service Industries Journal* **16**(2), 223–241.
- Ennew, C.T. (1992), ‘Consumer attitudes to independent financial advice’, *International Journal of Bank Marketing* **10**(5), 13–18.
- Faust, W.H. (1990), ‘The branch as a retail outlet’, *The Bankers Magazine* pp. 30–35.
- Frazier, G.L. (1999), ‘Organizing and managing channels of distribution’, *Journal of The Academy of Marketing Science* **27**(2), 226–240.
- Garland, R. (2002), ‘Estimating customer defection in personal retail banking’, *The International Journal of Bank Marketing* **20**(7), 317–324.

- Garland, R. (2005), ‘Segmenting retail banking customers’, *Journal of Financial Services Marketing* **10**(2), 179–191.
- Hayes, J.B. and M. Eduardo (2003), ‘An economic decision-making framework for price bundling of retail banking services’, *Academy of Marketing Studies Journal* **7**(2), 91–96.
- Ibragimov, R. and J. Walden (2010), ‘Optimal bundling strategies under heavy-tailed valuations’, *Management Science* **56**(11), 1963–1976.
- Johansson, J.K. (1979), ‘Advertising and the s-curve: A new approach’, *Journal of Marketing Research* **16**(3), 346–354.
- Kellerer, H., U. Pferschy and D. Pisinger (2004), *Knapsack Problems*, Springer.
- Lovelock, C. and J. Wirtz (2006), *Dienstenmarketing*, Pearson Education.
- Mankila, M. (1999), Price bundling theory applied to retail banking, Technical report, Goteborgs Universitet.
- Nair, S.K. and P. Tarasewich (2003), ‘A model and solution method for multi-period sales promotion design’, *European Journal of Operational Research* **150**(3), 672–687.
- Piersma, N. and J.-J. Jonker (2004), ‘Determining the optimal direct mailing frequency’, *European Journal of Operational Research* **158**(1), 173–182.
- Pigou, A.C. (1920), *The Economics of Welfare*, Macmillan and co.
- Prasad, A., R. Venkatesh and V. Mahajan (2010), ‘Optimal bundling of technological products with network externality’, *Management Science* **56**(12), 2224–2236.
- Reichheld, F.F. (1996), *The Loyalty Effect: The Hidden Force Behind Growth, Profits and Lasting Value*, Harvard Business School Press.
- Sheth, J. and A. Parvatiyar (1995), ‘Relationship marketing in consumer markets: antecedents and consequences’, *Journal of the Academy of Marketing Science* **23**(4), 255–271.
- Talla Nobibon, F., R. Leus and F.C.R. Spijksma (2011), ‘Optimization models for targeted offers in direct marketing: Exact and heuristic algorithms’, *European Journal of Operational Research* **210**(3), 670–683.
- Venkatesh, R. and W. Kamakura (2003), ‘Optimal bundling and pricing under a monopoly: Contrasting complements and substitutes from independently valued products’, *Journal of Business* **76**(2), 211–231.
- Wedel, M. and W. Kamakura (2003), *Market Segmentation: Conceptual and Methodological Foundations*, Kluwer Academic Publishers.