Chapter 27
Designing Database Marketing Communications

Abstract When all the LTV calculations, predictive modeling, and acquisition and retention planning have been done, the firm ultimately must *communicate* with the customer. In this chapter, we discuss how to design database marketing communications. We discuss the planning process for designing communications, and devote most of the discussion to copy development and media selection. We pay special attention to "personalization," or individualizing the message that is communicated to each customer.

27.1 The Planning Process

Figure 27.1 presents a planning process for designing database marketing communication campaigns. The focus might be on a single campaign or a program consisting of multiple campaigns. The four major steps are: Set the Overall Plan, Develop Copy, Select and Schedule Media, and Evaluate Results. We discuss the overall plan, copy, vehicle selection, and evaluation in this chapter, and scheduling, particularly for multiple campaigns, in Chapter 28.

The process in Fig. 27.1 could apply to mass marketing as well as database marketing. However, database marketing campaigns differ from mass marketing in their emphasis on field testing and personalization. Each of the copy and media steps can be tested and personalized to individual customers. Even the evaluation can be conducted using a test. We will see illustrations of testing and personalization in this chapter.

The personalization aspect of database marketing communications is part of a broader theme: 1-to-1 marketing. That term was coined by Peppers and Rogers (1993; 1997), who applied it to creating specific products for individual customers. Hess et al. (2007) define 1-to-1 marketing more broadly, as "tailoring one or more elements of the marketing mix to each customer." Hess et al. distinguish two types of 1-to-1 marketing: "personalization" and "customization." Personalization is when firms use customer data to individualize

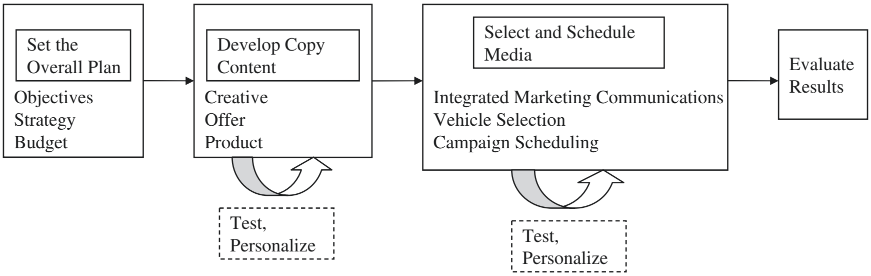

Fig. 27.1 Communication planning process for database marketing.

one or elements of the marketing mix – products, channels, prices, communications – for each customer. Customization is when the customer takes the lead, typically in the product arena, designing the product he or she most prefers. A prime example of customization is Dell's website that allows the customer to specify the product. Examples of personalization include the Amazon website's recommendations or Tesco's personalized promotion system (Chapter 22). Throughout this book, we have emphasized personalization, since it is based directly on customer databases. In this chapter, we emphasize the personalized design of communications.[1]

27.2 Setting the Overall Plan

27.2.1 Objectives

The overall plan consists of objectives, a strategy for achieving the objectives, and a budget. Objectives usually include specific goals such as 'increase sales by 10%," "acquire 20,000 customers," or "reduce churn by 5%." Campaign objectives are usually action oriented, similar to sales promotion objectives (Blattberg and Neslin 1990).

It is useful to make "qualitative" as well as quantitative considerations part of the objective. For example, the objective of a direct marketing campaign for an electronics device manufacturer might be, "Achieve a 1% response rate while *reinforcing product positioning on quality and reliability.*" The objectives of an offer for an upscale retailer's customers to join a reward program might be "induce 20% of current customers to sign up for the reward campaign by *emphasizing our quality and service.*"

Qualitative goals are important for three reasons: First, database marketing campaigns often play a dual role as promotions and advertising

[1] See Murthi and Sarkar (2003) for an excellent review of product personalization.

(e.g., Smith and Berger 1998). Promotions are most directly concerned with changing behavior, while advertising is often concerned with changing attitudes. Overly focusing on the promotion aspect of the communication can preoccupy the customer with the "deal" rather than the merits of the product.[2] Second, since database marketing campaigns are amenable to evaluation in terms of quantitative criteria such as sales, profits, ROI, etc., it is easy to lose sight of overall marketing strategy. Third, while qualitative objectives are difficult to evaluate, they guide the creative, offer, product, and copy elements of the campaign.

27.2.2 Strategy

Following Nash (2000, p. 36), we consider communications strategy to encompass: (1) Creative, (2) Offer, (3) Product, and (4) Media. Creative refers to overall tone and approach of the communication, as well execution in terms of artwork, type font, layout, etc. The offer refers to price, promotions, financial terms, shipping charges, etc. The product refers to the specific product (or service) contained in the offer, and how it is described (positioned). Media refers to the choice of media and how they will be scheduled. These aspects of the campaign can be stated in general terms when setting the overall plan, but the details are worked out as the planning process progresses. We discuss these details in subsequent sections of this chapter.

27.2.3 Budget

Budgets can be set in three ways: (1) based on last year's budget, (2) derived from objectives and strategy, or (3) derived from an optimization. From a purely scientific point of view, budgets should be derived from an optimization. However, while optimization models provide a good starting point and can help managers think "out-of-the-box," they omit strategic and organization issues that play an important role in determining the budget. Roberts and Berger (1999) discuss the role of last year's budget in determining this year's budget. If last year's results were acceptable, a small adjustment from last year might be appropriate. Also, if the objectives are for marginal growth, or emphasize ROI rather than absolute sales or profits, using last year's budget as a benchmark might be useful.

Deriving the budget from objectives and strategy involves a direct determination of what it will cost to achieve the desired objectives using the

[2] This is known as transaction versus acquisition utility. See Neslin (2002) for discussion.

prescribed strategy. For example, the plan might be to acquire 200,000 customers using direct mail. We might assume that the cost of acquiring these 200,000 additional customers via direct mail will be $20 per acquisition. This implies a $4,000,000 budget. Note we use the *marginal*, not average cost per acquisition (Chapter 26). The distinction is important. While last year the firm might have averaged $20 per acquisition, it may cost much more than that to acquire 200,000 *additional* customers this year. While deriving budget from objectives is attractive, one must be very careful to distinguish marginal from average costs.

In summary, optimization, last year's budget, or a calculation of what is required to meet objectives, all have their plusses and minuses as budgeting techniques. It is advisable to use at least two of the three techniques and check consistency.

27.2.4 Summary

The first step in developing a communications program is to develop objectives, strategy, and a budget. Objectives are almost always quantitative in a database marketing context although can contain qualitative considerations as well, even if they are difficult to measure. Strategy consists of creative, offer, product, and media plans. Budget may be derived from last year's budget, from the objectives and strategy, or from an optimization. Consider an acquisition campaign for a consumer electronics company:

Objectives:

Acquire 200,000 customers while not compromising on price.

Strategy:

Creative: Emphasize informational approach, professional tone.
Offer: Maintain standard price; perhaps offer free delivery or easy payment terms.
Product: Include three products from our product line in the offer, personalized for each customer using predictive model derived from tests.
Media: Direct mail

Budget:

Direct mail historically acquires customers at a cost of $20 per customer, so the budget is $200{,}000 \times \$20 = \$4,000,000$. However, money is very tight this year and our hope is through personalization we can lower that cost. Cost savings will go toward R&D for various product lines.

27.3 Developing Copy

27.3.1 Creative Strategy

As discussed earlier, the creative strategy represents the overall tone or approach to the communications piece, as well as the particular executions in terms of artwork, layout, etc. Various creative approaches include:

- *Informational*: Emphasizes the provision of information. This might be appropriate for a B2B product or a "serious" product such as a pharmaceutical.
- *Humorous*: Humor might be appropriate for a gift-oriented catalog.
- *Testimonial*: Emphasizes testimonials by satisfied customers. This would be appropriate for a high-risk product such as consumer electronics.
- *Deal-oriented*: Emphasizes price and getting a good deal. This might be appropriate when the target group is price sensitive or for a sale catalog.
- *Authoritarian*: Relies on an authoritarian figure such as an expert. For example, a campaign designed to enhance drug compliance might use a physician as a spokesperson.
- *Glossy*: Emphasizes an expensive look. This would be appropriate for a high-end cataloger.
- *Personal*: Oriented toward the individual's needs. This would be appropriate for a cross-sell appeal.
- *Comparative*: Compares the company's product to a competitor's.

There are several other possibilities. Nash (2000, chapter 9, p. 217) thinks of the creative strategy as consisting of three major elements: the product (what does it deliver), the customer (how does it make the customer feel emotionally), and company credibility (does the company have the credibility to deliver what it claims). Nash suggests that a relatively small company selling risky products might "allocate" most of its creative effort toward establishing product credibility. This might be accomplished through a testimonial, authoritarian, or comparative approach.

Lewis (1999) distinguishes among "romantic," "sedate," and "flat" approaches to devising copy for direct marketing pieces. Romantic approaches emphasize an idealistic, dreamy tone. An example for cookware would be (Lewis 1999, p. 57):

> C'est Merveilleuse! Bourgeat Copper Cookware from France! One of the world's greatest chefs cooks with the world's best copperware. Now you can too.... This master chef's solid copper gourmet cookware ensemble cooks as beautifully as it looks.

A sedate approach emphasizes the benefits of the product but in an informative, expository style. For example (Lewis 1999, p. 57):

> All-Clad Copper Cookware. No other material spreads heat faster or more evenly than copper, which is why you'll find a gleaming battery of copper pots in every famous restaurant.

All-Clad lines their copper cookware in stainless steel, making them easy to clean, impervious to scratches, and non-reactive with food.

A flat approach lays out the bare facts. Following is an example (Lewis 1999, p. 58):

Five-piece place setting in 18/10 stainless steel with pearlized resin handles. Choose lime, yellow, blue, or clear. Silver-plated connectors. Imported.

The flat approach emphasizes just the facts. It might be appropriate for a business-to-business communication or in any case where the customer is an expert and just needs to know the facts to make a purchase decision.

Fiore and Yu (2001) compare "imagery" and "descriptive" approaches (analogous to romantic versus sedate) in designing a clothing catalog. They conduct an experiment in which one catalog used just descriptive ads (communicating the benefits of the product but in a factual way) while another included both the descriptive and imagery approaches (also adding a romantic story about wearing the clothing on an adventurous vacation, etc.). The authors found there was no difference in respondent attitudes or intentions to buy between the two catalogs. This may be due to (1) heterogeneity in the target group, i.e., some resonate with the romantic approach while some are turned off, or (2) the combination of descriptive and imagery diluted their individual contributions.

Creative approaches should harmonize with the positioning of the product. An upscale cataloger should use a glossy approach. A price-oriented catalog should adopt a deal-oriented approach. A cross-sell by a local bank might utilize the personal approach (e.g., "As part of our continued efforts to attend to your personal banking needs, we thought you might be interested in our IRA transfer service"). Casual experience suggests this may not always be the case. Credit card solicitations often are deal-oriented ("Get 5% APR and bonus miles when you sign up for the ABC credit card") rather than informational ("The ABC credit card is accepted nearly everywhere and rewards you with airline miles for every dollar spent – as a special introductory offer, you get a 5% APR and bonus miles when you sign up"). See Schlosser et al. (1999) for a discussion of Internet shopper preferences for various creative approaches.

Comparative advertising has received much attention in the advertising literature. Database marketing provides a particularly attractive venue for comparative advertising, because appropriate comparisons can be targeted to the right customers. The general benefits of comparative advertising are increased attention and more effective communication of information, while its pitfalls include competitive response, customer confusion, and even customer sympathy for the "attacked" brand (Assael 1995, pp. 439–441; 726–728). While the evidence is somewhat mixed, it appears that comparative advertising can produce higher sales when the source of the information is credible, or when the target product is new or lower share (Assael 1995, pp. 439–441; 726–728; Engel et al. 1995, pp. 569–570).

Smith and Burger (1998) investigated comparative advertising in the context of a direct mail offer for a new video camera. The authors found using a laboratory experiment that the effectiveness of comparative advertising depended on what was being compared (price, product attributes, or product experiences) and the target customer (highly knowledgeable about the category, medium knowledge, or low knowledge). The results also depended on whether the dependent variable was purchase intentions for a single product or choice between two side-by-side products. In terms of purchase intentions (most relevant for direct marketers), price comparisons increased purchase intentions for low and high knowledge customers. Attribute comparisons were effective for low knowledge customers, while experience comparisons were effective for high knowledge customers. When choosing between products (as in retail shopping), experience comparisons were effective for low-knowledge customers.

This study highlights that comparative advertising may be an appropriate creative approach for database marketing provided the appropriate comparisons can be targeted based on product knowledge. Database marketers may therefore find it advantageous to test various comparative advertising executions. However, the potential pitfalls of comparative advertising (sympathy and hence sales for the competitive brand, and competitive retaliation) are difficult to measure in a test. More research is needed to gauge these effects in a database marketing context.

While the above discussion pertains to the overall creative appeal, there are several execution details that need to be worked out. These include everything from the type font to specific colors to the wording of sentences (length, use of active verbs, etc.) to the placement of larger pictures on the bottom versus the top of a catalog page. Usually these decisions on the basis of judgment; sometimes specific elements might be tested if deemed important enough.

Gatarski (2002) presents a method for developing over time the optimal creative for an Internet banner ad. A crucial measure of success for an online ad is its click-through rate (CTR). CTR is defined as the number of click-throughs divided by the number of exposures. Gatarski describes the use of a genetic algorithm to design online advertising copy as factorial combinations of ad attributes. The genetic algorithm systematically weeds out under-performing ads over time when it creates a new "generation" of ads. A given ad may generate a high CTR and therefore replicate itself in succeeding generations. However, as its CTR decreases perhaps due to wear out, it is less likely to replicate itself and it dies out, replaced by another ad whose CTR is higher.

Gatarski applied the method to a Compact Disk retail website. The author selected a particular artist and tested the ability of the algorithm to create banner ads that would maximize CTR for the artist's CD. Ads were created as factorial combinations of ad attributes called "genes." Each gene consisted of a number of levels (as in a conjoint analysis) called "alleles." For example, one gene (attribute) was the artist's picture, of which there were four alleles (versions). Other genes included the product picture and use of action phrases such as "Click to Play." A collection of genes with their specific alleles would constitute a particular ad, called a "chromosome." The algorithm to create and evolve chromosomes was as follows:

1. Start with 20 randomly generated banners. This is the first "generation" of chromosomes.
2. Place those banners on the website (the algorithm does not optimize the placement of ads, just the ad copy).
3. Observe response (click-throughs) to the banners.
4. After a period of time, generate the next generation of banners:

   a. Calculate click-through rate for each banner (chromosome).
   b. Calculate the "Fitness" for each chromosome, proportional to its CTR. Fitness represents the probability each chromosome will be selected as "parents" for the next generation.
   c. Select two chromosomes randomly but proportional to fitness.
   d. Randomly select the "cross-over point" gene.
   e. Generate a new child that has the genes of one parent from the left side of the cross-over point, and of the other parent to the right side of the cross-over. (The genes for each chromosome are displayed in a sequence, so one can think of the left or right side of the cross-over gene.)
   f. "Mutate" the child by randomly changing a few of its genes.
   g. Return the parents to the "pool" of prospective parents.
   h. Repeat steps c–g 19 more times, generating a total of 20 children, the next generation of banner ads.
   i. Repeat steps 2–4 for as many generations as desirable.

The algorithm ensures that the more successful a banner is at generating click-throughs, the more likely it is to be selected as a parent for the next generation of banners. Twenty children are created each generation, and each is generated by two randomly selected parents. As a parent, the banner gets to pass on its genes (a randomly generated amount depending on the cross-over point). The mutation step represents the key step in evolution whereby potentially valuable new banner characteristics become part of the banner population. If these mutations are not effective, the chromosome is less likely to be selected as a parent and the mutation dies out. In general, the chromosomes with valuable genes replicate and their genes are perpetuated as long as they are effective.

It is interesting to note the genes that flourished. Use of a click-through sign became more and more prevalent over time. A close-up of the artist's face also became prevalent, as did "Click to Play." Figure 27.2 shows the average click-through rate achieved for the successive generations. The average CTR starts at about 1% and grows to 1.66% by the 16th generation. There are some "lapses," possibly caused by mutations that didn't work or by wear out. The average CTR for the standard banner used by the company was 0.68% over the course of the experiment, and the CTR for the first generation of

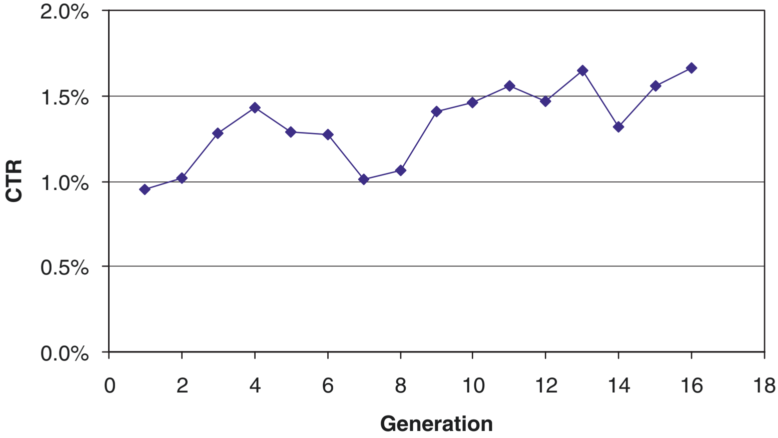

Fig. 27.2 Click through rates for 16 generations of optimally designed banner ads (From Gatarski 2002).

banners when they are not allowed to change over time averaged 1.00%. The genetically engineered banners outperform both these benchmarks.

This example presents promising evidence that the creative details of advertising copy can be optimally generated and adjusted over time. The adjustment is driven by customer data in the form of click-through rate. Ads would have to be quantifiable as factorial combinations of attributes, which may not always be possible. However, Gatarski's work suggests the practicality of the approach for banner advertising.

27.3.2 The Offer

Price is often a critical component of a database marketing offer. Research has investigated how to communicate price in a catalog or direct mail offer, particularly the use of end-nine prices and Sales signs.

Simester and Anderson (2003) investigate the use of end-nine prices by conducting a field experiment with three catalogs. The control catalog contained 50 items whose price ended in a nine. In the two treatment catalogs, prices of these items were either increased or decreased by $1. The three catalogs were each mailed to 20,000 customers. The authors found that after controlling for the popularity of specific items, price *per se* had insignificant impact on sales within the plus-or-minus $2 range in which it was manipulated. However, an end-nine price generated on average 35% higher sales. That is, if an item was presented at $58, $59, or $60, the $59 version would generate 35% higher sales than either the $60 or $58 version! They conducted a second test where prices were varied in a broader range, and the incremental sales due to end-nine was about 15%. In addition, the authors found that the end-nine effect was particularly pronounced for new items.

In another test, the authors manipulated the use of a "Sale" cue along with the price. They found that the use of the Sale cue decreased the end-nine effect. For example, an end-nine price would increase sales of a new item by 8.5%, but if a "Sale" message is already included, the incremental gain is only 3.9%.

The authors posit two explanations for the end-nine effect. The drop-off explanation is that customers round down, process left to right, or just process the left-most digit because of the cognitive cost of processing all digits. If they round down, then $59 is the same as $50. Similarly, if they process left to right, $55 will look different from $45 but the same as $59. The information explanation is that customers use an end-nine as a cue that the item is specially priced or on sale. They may have been conditioned by previous promotional pricing that used end-nine prices.

In a *post hoc* analysis, Simester and Anderson find support for the information explanation and not the drop-off explanation. They use a regression analysis of sales to show that the tens and ones digits of a price are equally important. That is, on a per dollar basis, a change in the tens digit of the price is equally important as a change in the ones digit. Moreover, the end-nine and Sale effects remain significant in this analysis.

If indeed the end-nine effect is due to the information it conveys, one would naturally be concerned about over-using end-nines. If the customer reacts positively to end-nines because he/she infers the product is specially priced or on sale, one cannot use all end-nine prices because it is doubtful that all items could be on sale.

The authors follow this line of reasoning in their investigation of the use of Sale signs in catalogs (Anderson and Simester 2001). They find in a regression analysis that the impact of a Sale sign on sales of a particular catalog item is smaller if more items on the page have Sale signs. Similarly, in a field experiment, they create a control and a treatment catalog. In the control catalog, three items on four pages are depicted as on sale and the remaining five items are not depicted as on sale. In the test catalog, all eight items are depicted as on sale. The results showed that sales decreased for the three items on sale in both catalogs when the other five items are put on sale.

There are three possible explanations for the results: substitution, attention, and credibility. Substitution means that Sale signs become less effective because items on sale cannibalize each other. Attention says that consumers pay less attention to the Sale sign when more brands have them. Credibility says that consumers do pay attention, but don't believe the Sale sign because it is not plausible that so many items could have a reduced price at the same time.

The authors discriminate between substitution and credibility/attention effects by noting the impact of Sales signs on total category demand. The substitution argument would predict category demand to be constant no matter how many Sale signs the catalog has. However, the credibility and attention arguments would suggest that category sales will increase up to a certain point and then decrease when there are too many Sale signs that are either ignored or not believed. In a study of packaged goods, the authors indeed find an inverted-U relationship between the number of Sale signs and category sales, supporting the credibility/attention arguments.

To discriminate between attention and credibility explanations, the authors run a laboratory experiment in which they vary the number of items on a catalog page that contain Sale signs. They find that when an item is on sale and additional products are also on sale, subjects think it more likely that the item will be on sale next period and that the average future price will be lower. This implies they pay attention to the Sale sign in that it influences their perception of future prices. However, they don't believe the current price is "special" in that the item is more likely to be on "sale" next period. That is, Sale signs lose their credibility as signaling a special price the more they are used.

The authors conclude that the credibility effect is the best explanation for Sale signs becoming less effective the more items use them. One avenue for future research is to investigate the substitution effect in a catalog environment rather than packaged goods. Packaged goods are purchased in a supermarket, and category effects may dominate since most of these products can be stored in household inventory (see Van Heerde et al. 2004). In buying from a catalog, shoppers may focus on which computer to buy, or which dress to buy, so substitution would dominate.

In summary, a key tactical issue in copy design for database marketing is how to communicate price. The evidence is that (1) end-nine prices increase sales, (2) the end-nine effect is particularly strong for new items, (3) Sale signs decrease the impact of end-nine prices, (4) Sale signs increase sales of particular items until too many items have Sale signs, at which point its effectiveness per item actually decreases. The reason for the end-nine effect appears to be the information conveyed that the item is at a special price. The reason for the Sale-sign-saturation effect seems to be a loss in credibility. While no research has investigated the loss of credibility for end-nine prices, this is an obvious avenue for future study. The above conclusions are based on two excellent, carefully researched studies, but only two. The results need replication and extension.

The implication is that database marketers should use end-nine and Sale signs albeit judiciously, in designing their offer. In the case of a catalog, the database marketer can make sure he or she does not overuse end-nine prices or Sale signs. However, in designing a single offer, the decision is either/or, not how many. The individual database marketer may feel he or she is not overusing this tactic, but if all database marketers behave the same way, the result is a collective saturation of the impact – customers grow weary of so many end-nine prices and Sale signs. It would be valuable to see this issue explored further.

27.3.3 The Product

27.3.3.1 Deciding Which Products to Feature in Communications

Often the product to be included in the communication is known beforehand. For example, a B2B electric generator company knows the set of products it will include in its catalog. There may still be issues such as which product(s) to promote. Lin and Hong (2008) use market basket analysis (Chapter 13) to select various bundles of products to promote in an electronic catalog. They utilize the "support" (P(AB), see Chapter 13) for various combinations of products to create the bundles and provide evidence that this method increased catalog sales.

Database marketing provides the capability to personalize the firm's product to be offered. For example, a fundamental tenet for Capital One's credit card business is that the particular credit card to be included in a direct mail offer is determined by test mailings. Financial services routinely decide which products to cross-sell to which customers (Knott et al. 2002).

Ansari and Mela (2003) develop a model and optimization for deciding which URL links to include in customer-specific e-mails. The setting is a news/information website that sends e-mails to its users to entice them to the site. The company wishes to gain "readership," thereby increasing its advertising revenues. The website has 12 content areas (products) and it must decide which k content areas to feature in the e-mail ($k \leq 12$) and in which order. The content areas are featured by providing a URL link to the page in the company's website that includes this content. The authors collect data on response to previous e-mails, and estimate the following model:

$$U_{ijk} = X'_{ijk}\mu + Z'_{jk}\lambda_i + W'_{ik}\theta_j + \gamma_k + e_{ijk} \tag{27.1}$$

where:

U_{ijk} = Utility of customer i for content link k in e-mail j.
X'_{ijk} = Vector of attributes describing e-mail j with content link k seen by customer i. These include for example the content area, the position of the content link (1st, 2nd, 3rd, etc., on the list), the number of content area links featured on the e-mail, whether the e-mail was text or HTML, the length of time since customer i last clicked through an e-mail, etc. These represent "fixed effects" in that μ is the average response to these variables.
Z'_{jk} = A subset of the X variables that vary across content areas or e-mails, for which response is assumed to be heterogeneous across customers. Customer response to content link, number of content areas in the e-mail, and position was assumed heterogeneous, captured by λ_i.
W'_{ik} = A subset of the X variables that vary across customer or content-area levels, whose effect could be heterogeneous across e-mails. This could include position as well as dummy variables for the content area. For

example, it might be that two e-mails featuring the same content area might differ in the response they elicit (due to unobserved factors such as season, etc.). The degree that the effects of these variables varies across e-mails is reflected in θ_j.

γ_k = Unobserved factors specific to content area k.

e_{ijk} = Unobserved factors specific to customer i, e-mail j, content area k.

The authors estimate Equation 27.1 using Bayesian techniques and experiment with various ways to capture heterogeneity. Some of their key results are: (1) Few content areas are universally preferred by all customers. (2) There is ample heterogeneity in content preferences across customers. (3) The lower the position of a content area in an e-mail, the less likely it is to be clicked through. (4) The number of content areas featured in an e-mail had no effect *on average* across customers, but there was some heterogeneity across customers.

The estimation generated individual-level response coefficients to content areas featured, position of a particular content area in the e-mail, and number of content areas featured in the e-mail. The authors then develop an optimization model to design each customer's e-mail. They considered two objectives – maximizing the expected number of clicks per e-mail, or maximizing the likelihood of at least one click through. We show the formulation for the latter objective, because it is unlikely the customer would come back to the e-mail after clicking through one of the content area URL's.

To make the optimization easier, the authors undertake a two-phase approach where for a fixed number of items in the e-mail (k), they find the best content areas to include in which position. They vary k from 1 to 12 (since there are 12 content areas, $k \leq 12 \equiv n$). For a given customer and a given number of items, the optimization is:

$$\underset{x_{ij}}{Maximize} \sum_{i=1}^{n}\sum_{j=1}^{k}\left[1 - \prod_{i=1}^{n}\prod_{j=1}^{k}(1-p_{ijk})^{x_{ij}}\right]$$

$$\Rightarrow \underset{x_{ij}}{Minimize} \sum_{i=1}^{n}\sum_{j=1}^{k}\left[x_{ij}\log(1-p_{ijk})\right] \qquad (27.2)$$

subject to:

x_{ij} = 1 if content area i is include in the e-mail in position j. There are $n = 12$ content areas and k possible positions, with $k \leq n$.

p_{ijk} = Probability the customer clicks through content area i if it is in position j when the total number of content areas included is k.

The optimization is performed for all 12 values of k for the 100 customers in the sample. The authors then test their predictions on a holdout sample using three approaches: (1) the optimization described above, (2) an ordering algorithm which took takes each e-mail in the holdout and re-orders the content areas according to the customer's model-inferred preferences. Whether

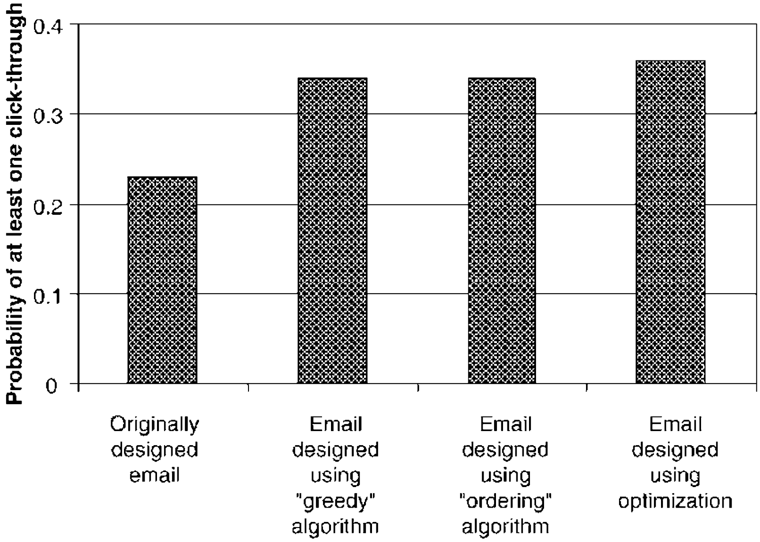

Fig. 27.3 Probability of at least one click-through in a personalized e-mail (From Ansari and Mela 2003).

the most preferred was placed on top or bottom depended on the sign of the position variable (a positive sign implies the customer has higher utility if the content link is toward the bottom of the set of links, i.e., the customer is a bottom-up processor, while a negative sign implies the customer is a top-down processor), and (3) a greedy algorithm that kept the number of content areas the same, i.e., k, but included the k highest-preferred content areas for each customer and inserted them in the e-mail in order of preference. The results are shown in Fig. 27.3.

All three algorithms improve significantly over the baseline click-through rate: the rate in the holdout data using the e-mails as they were originally designed. The strong performance of the ordering and greedy algorithms, which took the number of links in the e-mail as given, is probably due to the fact that the number-of-items variable turned out not to be crucial – its coefficient on average was not different than zero, and although there was customer heterogeneity on this coefficient, it was not enough to make a big difference. However, the success of the ordering and greedy algorithms shows that the customer-level content preference and top-down versus bottom-up processing measures (reflected by the position variable) were important.

Ansari and Mela's work is very important because it shows how to personalize a communication to include the products most relevant to a customer. The model could be extended to include creative and offer variables as well (in fact, the authors included one creative detail – text versus html – in their model). Two avenues for future research building on this work would apply it to multi-campaign scheduling – i.e., scheduling several e-mails over time (see Chapter 28), and scaling the heterogeneity up to a company's full house list. One approach would be to use latent class segmentation and calculate posterior probabilities of segment membership using Bayes theorem.

27.3.3.2 Allocating Catalog Space to Various Departments

Single Catalog

Squeezed by mailing costs that limit the size of catalogs coupled with expanded product lines, allocating catalog space to departments or product lines has become a challenging issue for catalogers. The problem is somewhat similar to the space allocation problem in bricks-and-mortar retail stores. However, as Desmet (1993) points out, the space decision is fixed for the duration of a catalog whereas it is flexible in a store. In addition, catalog space allocations draw the customer into the catalog, whereas shelf space allocations in stores influence the customer who has already decided on a shopping trip. Of additional interest to database marketers is that catalogs can be tailored to individual customers, which cannot happen in a bricks-and-mortar store.

While from a database marketing perspective, personalization is the ultimate goal, it is useful to examine first the situation of a single mass-marketed catalog. Rao and Simon (1983) pioneered this case. They derive the simple rule that under certain assumptions, the optimal allocation of space is such that the profit per square inch of allocated space is equal for all departments. Their decision problem is as follows:

$$\underset{x_i}{Max}\ Z = \sum_i m_i S_i(x_i) - \sum_i C_i(x_i) \tag{27.3}$$

$$\text{subject to:} \sum_i x_i = P \tag{27.4}$$

where:

m_i = Profit margin for category i.
S_i = Sales of department i.
x_i = Space allocated to department i.
C_i = Cost function for department i.
P = Total page limit for catalog.

This can be solved using Lagrange multipliers. The most important implication is that the space allocation should be such that the marginal profit contributions for all departments is the same. If at a given allocation, the marginal contribution of one department is higher than the others, we would add more space to that department and reduce space for the others. That is, we are not at an "equilibrium" allocation until all the marginal contributions are equal.

Rao and Simon (1983) show that if sales response is a simple power function, $S_i = \alpha_i x_i^{\beta_i}$ with all β_i's less than 1,[3] and if the β_i's are equal for all departments, then the optimal allocation is:

$$x_i^* = \frac{P m_i S_i}{\sum_j m_j S_j} \tag{27.5}$$

or,

$$\frac{m_i S_i}{x_i^*} = \frac{\sum_j m_j S_j}{P} \tag{27.6}$$

Equation 27.6 says that at the optimal allocation, all departments will generate the same profit per square inch (or other unit of area) of allocated space. This is a very usable rule of thumb since profit per square inch is easily calculated. If the space response elasticities are not all equal, Equations 27.5 and 27.6 are modified so that $\beta_i m_i$ replaces m_i, and $\beta_j m_j$ replaces m_j. Equation 27.6 then becomes:

$$\frac{m_i S_i}{x_i^*} = \frac{\sum_j \beta_j m_j S_j}{\beta_i P} \tag{27.7}$$

Equation 27.7 implies that profit per square inch for a given department is inversely proportional to its elasticity. The reason for this is that as elasticity increases, we tend to allocate more space to the department, but the increase in profits has diminishing returns (since $\beta_i < 1$). As a result, the profit per square inch decreases with higher elasticity.

An important assumption is the absence of cross-department effects.[4] There could be complementarity or substitution effects, which would make the allocation more complicated. The objective would still be the same as Equation 27.3, except there would probably not be a closed-form solution. Corstjens and Doyle (1981) or Bultez and Naert (1988) solve similar problems in a retail store environment.

A critical empirical question for the above analysis is, what is the space response elasticity (β)? Equation 27.6 for example is predicated on the assumption that elasticities are all less than one and equal to each other across departments. Equation 27.7 abandons the equality assumption but still assumes concave returns.

---

[3] This assumption means diminishing returns to space allocation, or concave returns. If space elasticities are >1, the optimal allocation is to devote all possible space to the highest elasticity department.

[4] See Hofacker and Murphy (2000) for a discussion of banner ad cannibalization in the context of Internet Web page design.

The evidence is limited in scope but it appears to be that elasticities are less than one. Desmet (1995) found that space elasticities for a book catalog were 0.43 for new books and 0.43 for old books. Sokolick and Hartung (1969) found an elasticity of 0.56. Desmet (1993) found an elasticity of 1.18 for clothing, although suspects that this is over-stated. Desmet (1993) finds managerial judgments of elasticities range from 0.55 to 0.77. This work is suggestive, but much more work is needed to measure own and cross-department elasticities. To obtain these estimates, managers must be willing to vary space allocation systematically in a field test. See Seaver and Simpson (1995) for an experiment to study catalog design.

Personalized Catalogs

Of keen interest to the database marketer is how to personalize catalog space for individual customers. We present an illustration of how this might be done. We assume there are J market segments, each with a different response function. The goal is to design a catalog tailored to each segment by allocating the number of pages from each of D departments to each of the J catalogs. Note the number of catalogs equal the number of customer segments, so J is both the total number of segments and the total number of catalogs to be specified. The problem can be expressed as:

$$\underset{\vec{x}_j}{max} \sum_j m\delta_j S_j(\vec{x}_j) \tag{27.8}$$

subject to:

$$\sum_i x_{ij} = P \qquad \forall j \tag{27.8a}$$

$$CL_i \leq x_{ij} \leq CU_i \qquad \forall i, j \tag{27.8b}$$

$$DL_i \leq \sum_j x_{ij} \leq DU_i \qquad \forall i \tag{27.8c}$$

$$x_{ij} \geq 0 \qquad \forall i, j \tag{27.8d}$$

$$x_{ij} \quad integer \qquad \forall i, j \tag{27.8e}$$

where:

$\vec{x}_j = \{x_{ij}\} = (x_{1j}, x_{2j}, \ldots, x_{Dj})$, where x_{ij} is the number of pages for department i allocated to catalog j. D is the number of departments.

S_j = Sales per customer in catalog j as a function of the allocation $\vec{x}_j$. The sales response function differs for each catalog because each catalog represents a different market segment.

δ_j = Percentage of customer base represented by segment $j (\sum_j \delta_j = 1)$.

m = Profit margin, assumed to be equal across departments.

P = Number of pages to be allocated to each catalog.

Table 27.1 Parameters and optimal solution for catalog customization

	Response functions				
	Constant	Dept A elasticity	Dept B elasticity	Dept C elasticity	Segment percentage
Segment 1	0.01	0.4	0.5	0.2	30%
Segment 2	0.015	0.3	0.1	0.7	50%
Segment 3	0.0325	0.3	0.7	0.8	20%

	Catalog page constraints			
	Within catalog Lower bound	Upper bound	Across catalogs Lower bound	Upper bound
Department A	10	40	30	120
Department B	10	40	30	120
Department C	10	40	30	120

Other parameters	
Profit margin	30%
Number of customers	100,000
Page limit per catalog	50

	Optimal number of pages allocated			
	Dept A	Dept B	Dept C	Sales/customer
Segment 1	18	22	10	\$0.236
Segment 2	12	10	28	\$0.410
Segment 3	10	19	21	\$5.818
			Profit/customer	\$0.431
			Total profit	\$43,186

Constraint 27.8a ensures that each catalog has P pages. Constraints 27.8b and 27.8c are departmental exposure requirements that arise because the firm wants to make sure that each department gets at least minimum exposure, but is not over-emphasized. This reflects long-term considerations in the positioning of the catalog and inventory management. Constraint 27.8b ensures that each catalog contains at least a minimum number of pages from each department, and no more than a maximum. Constraint 27.8c ensures that each department has minimum and maximum representation across all catalogs. Constraints 27.8d and 27.8e require that page allocations take on integer values greater than zero. Therefore, this is an integer programming problem.

Note this is a joint optimization does not maximize profits from each segment separately. Therefore, there is the possibility that one of the generated catalogs, say j', generates more profit in segment j than the catalog generated for segment j! This is because the catalogs are being generated to satisfy departmental representation constraints (27.8b and 27.8c) so segment j's catalog may contain a large number of pages from department i, not because this maximizes sales from segment j but because segment j has the best response to department i and we need a minimum representation of department i among the J catalogs.

Table 27.1 summarizes the parameters for an illustrative analysis of the model. We have three segments so there will be three catalogs ($J = 3$). There are three departments ($D = 3$). Each catalog consists of 50 pages, and must contain between 10 and 40 pages of each department (given the 50-page total, the effective upper bound is 30 for each department). Across all catalogs, we must have between 30 and 120 pages for each department. The space response function is:

$$S_j = \alpha_j \prod_{i=1}^{D} x_{ij}^{\beta_{ij}} \tag{27.9}$$

This is a multiplicative response function with constant elasticities shown in Table 27.1. The elasticities and segment percentages show that Segment 2 is the largest segment and is particularly responsive to Department C space. Segment 3 is the smallest and likes Departments B and C. Segment 1 is mid-sized and likes Departments A and B.

The bottom of Table 27.1 shows the optimal allocation of pages. The allocations are intuitive. Segments generally are allocated pages proportional to their elasticities. For example, Segment 3 gets pages mostly from Departments B and C and has the minimal amount required for Department A. Segment 2 receives 2 more pages than the minimum for Department A even though its elasticity is the same as Segment 3's. This is because Segment 2 has a very low elasticity for Department B. None of the across-catalog departmental representation constraints is binding, although the within-catalog requirements sometimes are binding on the low side. Clearly it isn't optimal for Segment 2 taken in isolation to receive 10 pages of Department B, since its elasticity is only 0.1. However, long-term considerations dictate the minimum requirement. Finally, note that Segment 3, although it is smallest in terms of percentage of customers, has a high constant term so contributes the most to overall profit.

Figure 27.4 shows how the optimal solution changes as elasticity increases. In particular, we vary Segment 2's elasticity for Department B. The results are predictable. As elasticity increases, the pages allocated to Department B for Segment 2 increases. Since we have a 50-page limit, something has to give. The optimal solution takes away pages from both departments. After the allocation to Department A hits its lower bound, more pages are taken from Department C. Profits increase exponentially after elasticity becomes greater than 1, because then we have a convex demand function. While not shown in Fig. 27.4, the allocations to the other segments remain about the same.

Figure 27.5 shows how the optimal solution changes as a lower bound constraint increases. In particular, we vary the across-catalog minimum for Department B. This minimum is currently at 30 while the current solution allocates 51 pages. So the constraint is non-binding and the optimal solution does not change until the constraint moves above 50 pages. At 60 pages, the model has to allocate an additional 9 pages, and allocates them to Segment 1's catalog. As the constraint increases, it allocates next to Segment 2's

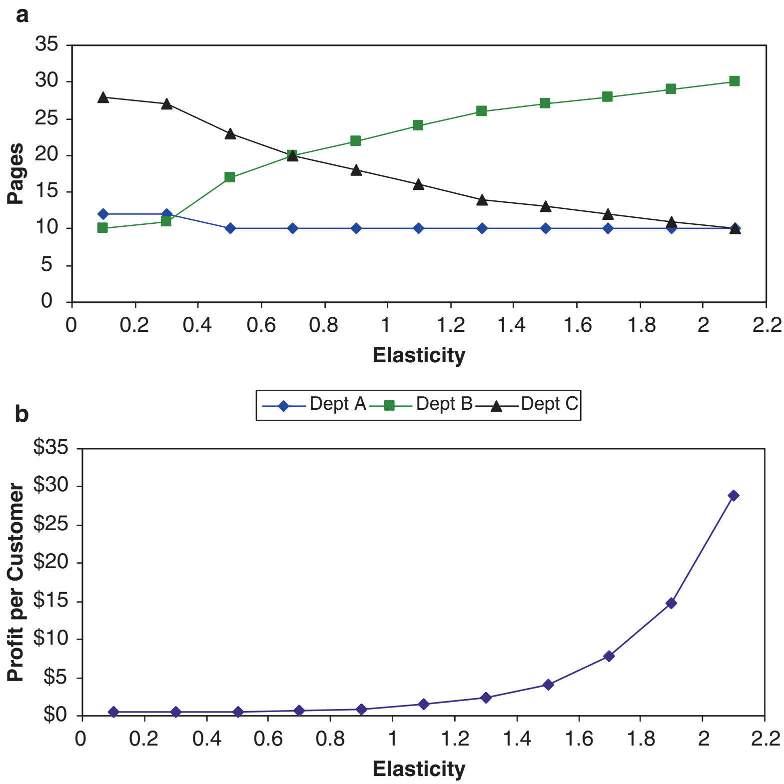

Fig. 27.4 Optimal allocation and profits as a function of elasticity. (**a**) Page allocations as a function of Department B, Segment 2 elasticity; (**b**) Catalog profit as function of Department B, Segment 2 elasticity.

catalog, and finally to Segment 3's.[5] The reason for this order of allocation is that Segment 1 has a fairly high elasticity for Department B, so we can take away pages from the other departments in Segment 1's catalog without too much sacrifice. Next the model allocates additional pages to Segment 2's catalog although this is a low elasticity department for this segment. The reason is that although we will have to take away from Department C, which is popular with this segment, the segment does not influence total profits as much. Finally, the model allocates more Department B pages to Segment 3's catalog. This severely diminishes profit because Segment 3 represents the majority of firm profits and we are using a sub-optimal catalog for that segment in order to fulfill the lower bound requirement.

[5] As noted above, while the explicit upper bound is 40 pages per catalog per department, the effective constraint is 30 because the other departments have to have at least 10 pages and the total catalog has 50 pages.

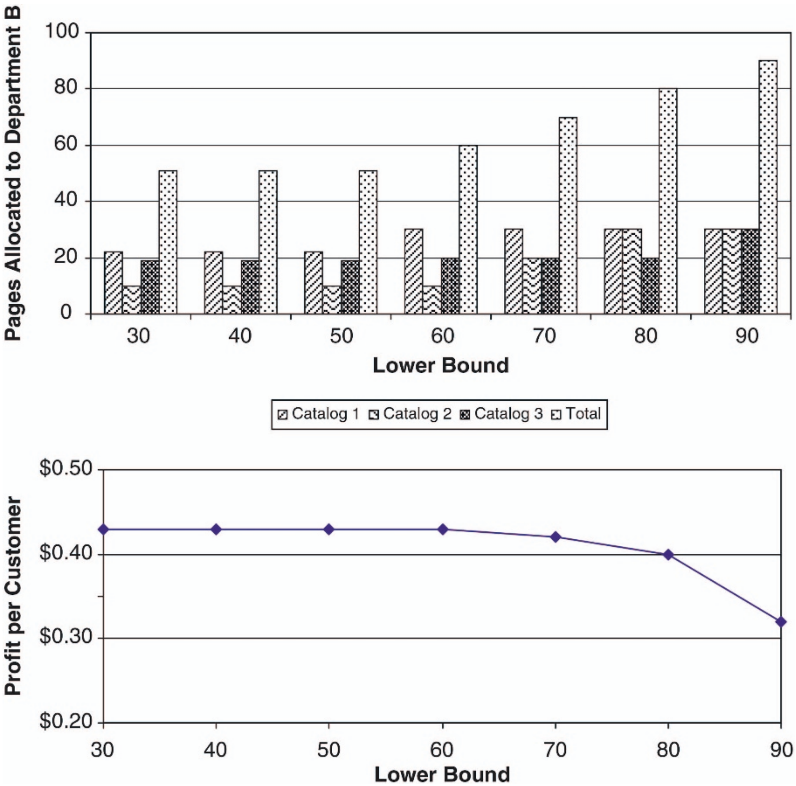

Fig. 27.5 Optimal allocation and profits as function of lower bound page requirement. (**a**) Page allocation as a function of Department B lower bound; (**b**) Profit per customer as function of Department B lower bound.

The example illustrates that pages are allocated roughly proportional to elasticities, but minimum and maximum requirements create distortions that reverberate across as well as within segments. One might therefore question why we have these constraints. The reason is that catalogers wish to maintain a presence in several departments, they may have plans to grow various departments in the future, and they need to make up for the fixed costs of running a particular department.

While the above demonstrates that the optimization in Equation 27.8 is quite rich, it does not contain explicit cross elasticities between departments. The demand function is for the catalog as a whole, not individual departments. Extensions of this model would include departmental level demand functions with cross-elasticities.

Another extension would be to apply the model to the customer level, where segments now correspond to individual customers. With 1,000,000 customers (N = 1,000,000), we could have 1,000,000 different catalogs. This would be Equation 27.8 with J = 1,000,000. Another possibility would be to specify J < N: We might have 1,000,000 customers with 1,000,000 different response functions, but only have say 10 catalogs (J = 10). This would be a more difficult problem because the objective function would have to distinguish between catalogs and customers. It would be:

$$\max_{\vec{x}_j, M_{cj}} \sum_c \sum_j mS_{cj}(\vec{x}_j)M_{cj} \tag{27.10}$$

where:

$S_{cj}(\vec{x}_j)$ = Sales for customer c if mailed catalog j and catalog j has page allocation $\vec{x}_j$.

M_{cj} = 1 if catalog j is mailed to customer c; 0 if not. We require of course that only one catalog can be mailed to each customer, and all customers must receive a catalog.

This would be a more complicated optimization because we have two decision variables, the page allocation to creating each catalog, and the decision of what catalog to mail to each customer. In the problem analyzed in this section, catalogs corresponded to segments, which made the allocation and mailing decision the same.

27.3.4 Personalizing Multiple Components of the Communication

Ansari and Mela (2003) focused on personalizing the product(s) featured to the customer, but their model also contained creative elements (their "Text" variable). It also could have included aspects of the offer such as price (see Rossi et al. (1996) for a price-personalization model, and Chapter 29). Therefore, if a customer response model contains variables representing creative, product, and offer components, it can be used to personalize communications in terms of all these elements.

Customer-level response models are attractive for personalizing communications. However, they rely on a fair amount of data per customer and the optimization can become complex when tailoring more components. Another approach is to use machine learning tools to personalize communications. One domain where there has been initial work on this is in personalizing websites.

Ardissono and Goy (2000; see also Ardissono et al. 2002) describe "SETA", a system that personalizes the products featured, the use of graphics, and the length and terminology of product descriptions on an e-tailing website. A demonstration can be found at http://www.di.unito.it/∼seta/seta.htm. SETA employs a user-based collaborative filtering system that associates each user with a "stereotype," then predicts preferences for product attributes, styles of presentation, etc., based on the probability the user belongs to each stereotype (see also Kohrs and Merialdo 2001). The probability is based on how well the available information for the user matches each stereotype. The more information available for the user the more accurate this probability distribution is, and hence the predictions are more accurate. If the user registers at the website, the information can be stored and grows over time. If the user does not register, the system starts with a generic presentation and then learns more about the customer as the customer navigates through the site.

Between customer-response models and machine learning systems, there is great potential for personalizing all aspects of a database marketing communication. More work is needed both to develop and test these methods.

27.4 Selecting Media

27.4.1 Optimization

Database marketers have a plethora of outlets through which to communicate with customers, including direct mail, e-mail, telephone, the Internet, and catalogs (Tellis 1998). Selecting which media vehicle(s) to use is amenable to the same approaches that have been used for vehicle selection in mass marketing (see Tellis, Chapter 16; Lilien et al. 1992, Chapter 6). The basic requirements are response functions for each medium, and the degree of overlap between media. For example, if a database marketer purchases 1,000,000 Internet impressions, and e-mails a list of 1,000,000 addresses, what would be the net number of individuals exposed to the campaign? To our knowledge, no research has yet investigated this issue or the general problem of vehicle selection in a database marketing context. To formulate the problem, let

E = Number of times customer is exposed to the offer.
X_i = Number of insertions of the offer in vehicle i.
$i = 1, \ldots, N$ where N is the total number of vehicles.
$P(E|X_1, X_2, \ldots, X_N)$ = Probability a randomly selected customer will be exposed to the offer E times, given the vehicle selection strategy $X_1, X_2, \ldots, X_N$.
λ = Parameter depicting decreasing returns to multiple exposures.
$(1 - e^{-\lambda E})$ = Probability a randomly selected customer will respond to an offer, given he or she is exposed to it E times.
C_i = Cost per insertion in vehicle i.
B = Budget.

Then the optimization problem is:

$$\underset{X_1 X_2 \ldots X_N}{Max} \sum_{E=1}^{N} (1 - e^{-\lambda E}) P(E|X_1 X_2 \ldots X_N) \tag{27.11}$$

such that

$$\sum_{i=1}^{N} C_i X_i = B \tag{27.12}$$

Equation 27.11 shows the goal is to select media $\{X_1, X_2, \ldots, X_N\}$ to maximize average response rate (equivalent to maximizing the total number of responses). The average response rate is the sum of the customer response rates for each number of possible exposures, weighted by the probability the customer receives that number of exposures. The optimal vehicle selection also must satisfy the budget constraint (Equation 27.12).

The key parameters are the response parameter λ and the exposure probability function P. λ could be estimated judgmentally or through testing. The exposure probability function is a frequency distribution estimated for example by Rice (1988) or Rust et al. (1986) in the context of mass media.

The approach can easily be made more complex, and potentially more effective. First is that individual customers may differ in their response rates. We would have λ_c, where c labels the customer. Second, λ may depend on vehicle, so that we have λ_i. Finally, λ could depend both on customer and vehicle, λ_{ci}.

The importance of adjusting λ by vehicle is illustrated by Sherman and Deighton (2001). They consider the problem of selecting websites on the Internet for a banner advertising campaign. The goal of the campaign was to induce web surfers to visit the drugstore.com website. The task was to identify the websites on which drugstore.com should place banner ads in order to maximize visits to its website. The approach was to estimate a predictive model predicting visits to the drugstore.com website as a function of visits to other websites. The results yielded an "affinity" between the websites and drugstore.com. Then a cluster analysis was used to group websites by various website attributes, and the affinity for each cluster was then calculated. This enabled the researchers to identify clusters, and hence websites, that had particular affinities for drugstore.com. These websites were then targeted with banner ads.

The results were quite favorable. The high-affinity sites generated 10 times the purchases per impression as were generated by the low-affinity sites. One site generated 43% of all orders using just 32% of the overall budget.

Sherman and Deighton show that response rates differ across websites and hence should be incorporated in Equations 27.11 and 27.12. The study also shows that predictive modeling can be used to increase website visits. The study focuses only on the response of individual sites and does not consider decreasing returns to repetition as in Equation 27.11. It could be that placing ads on several websites within an affinity cluster creates duplication. This means that it may have been possible to improve the allocation. However, the example shows clearly that starting with a relatively simple approach can generate important gains.

27.4.2 Integrated Marketing Communications

The analysis in Sect. 27.4.1 assumes media vehicles have to be selected to distribute a given offer. The more general case is that different messages might be distributed in different vehicles. For example, mass television advertising may be used to create awareness for the product, and the Internet may be used to extend a promotional offer. The "ads" themselves would be different. However, the concept of "integrated marketing communications" (IMC) posits that the content of these ads should be coordinated. The basic tenet of IMC is that the firm benefits by consistency, or at least complementarity, in its advertising messages.

IMC is especially important in database marketing because the modern database marketer sends out many different communications. Sheehan and Doherty (2001) discuss IMC in a database marketing context. They measure the degree of coordination in advertising messages between print and online advertising for 186 companies. They collected their data by first finding a print ad; then going to the company website to assess the consistency between the ads. The authors find a higher degree of coordination in certain aspects of the advertising but not in others. For example, in 82.8% of cases, the same logo could be found on both the website and the print ad. However, the "promise or single most important message of the print advertisement" was easily found on the website only 38.6% of the time. It was found albeit with difficulty 33.9% of the time, and not found at all 27.5% of the time.

Sheehan and Doherty show that companies do not necessarily coordinate their communication strategy across advertising vehicles. The next step is to assess empirically the importance of coordination in a database marketing context.

27.5 Evaluating Communications Programs

Four methods for evaluating a communications program are: (1) calculation of profits, (2) *post hoc* statistical analysis, (3) embedded testing, and (4) surveys.

Often a database marketing communications campaign can be evaluated directly by calculating profits. For a direct marketing campaign, the general form of the calculation is:

$$\Pi = NrM - Nc \tag{27.13}$$

where

Π = Profits
N = Number of contacts (e.g., number of pieces mailed)
r = Response rate
M = Profit margin per response
c = Contact cost

Another important measure is return on investment or ROI:

$$ROI = \frac{NrM - Nc}{Nc} = \frac{rM - c}{c} \tag{27.14}$$

ROI is the profits earned by the program divided by its cost. It has become popular to talk about "marketing ROI," and the strong advantage of the concept is that it is comparable across programs, companies, and industries. It also has high face validity: "For every dollar invested, we made [ROI] dollars in profit." However, one problem is that ROI is oblivious to scale. Is it better to make $2,000,000 on a $1,000,000 investment or $200,000 on a $10,000 investment (ROI = $2 per dollar versus $20 per dollar)? As discussed in Chapter 10, "Type I error" decision-making – where the goal is to avoid investments that would not have paid out, rather than avoid not making investments that would have paid out – emphasizes ROI. Shrinking the denominator of Equation 27.14 may inherently increase ROI, because the smaller investment is made on the customers who are sure to pay out. However, the total profits may not be impressive.

Another important issue with both Equations 27.13 and 27.14 is they typically are single channel, not multiple channel calculations. The multichannel company runs the risk that a marketing effort in Channel A takes away from Channel B. For example, a bank conducting a direct mail campaign to sell credit cards may take away from credit card sales in the its branch offices. Communications evaluation in a multichannel context requires data on *total* sales, not just those generated through the channel through which the communication was distributed.

Statistical analysis to evaluate communications falls under the rubric of "marketing mix modeling." These analyses take the form of regression models such as:

$$Sales_t = \beta_0 + \beta_1\ Marketing_t + \beta_2\ Seasonality_t + \varepsilon_t \tag{27.15}$$

The coefficient β_1 measures the increase in sales per dollar increase in marketing expenditure. To calculate total profits, one would multiply the total marketing expenditure during the period of interest times β_1 to yield total sales generated, and then apply the margin percentage to calculate total profits or ROI. Equation 27.15 is an outline of a marketing mix model. There are several other variables that might be included (e.g., other marketing efforts besides the campaign being evaluated, competitive marketing efforts, and carryover effects of marketing). The reader is referred to Leeflang et al. (2000) for a thorough discussion (also see Dean 2006).

Equation 27.15 can be estimated for each of the company's channels. For a retail bank, there would be equations for mail-order and branch sales of credit cards. It is possible that β_1 would be positive in the mail-order equation but negative in the branch equation, signifying cross-channel cannibalization. This provides important learning as well as the "bottom line" profits of the campaign.

In practice, regression modeling can become quite complex and one is never sure that all possible influencers of sales during the period of interest have been included. Therefore, a good practice is to embed a "control group" within the full launch of the campaign to provide an experimental test of the campaign. That is, the full credit card campaign might have been mailed to 1,000,000 customers, but an additional 10,000 were not mailed the campaign to provide a control. Profit is calculated by comparing sales, and hence profit, per customer in the campaign group versus the control group.

A few issues are important in using embedded tests. First, sales data need to be collected for all customers across all channels. In the banking example, the control group will be of no use if the bank cannot compile credit card sales for *all* its customers across *both* direct mail and in-branch channels. Second, the control group must be randomly selected from the group *eligible* for the campaign. For example, the customers selected to receive the direct mail credit card piece might have scored in the top four deciles as classified from a predictive model (Chapter 10). The 10,000 control customers must be selected from that group, not from the bottom six deciles. This means that the company is incurring an opportunity cost in not mailing to the control group, since they are prime prospects for the campaign (Chapter 9). That is why the sample size for the control group in an embedded test should be much smaller than the full mailing.

Testing plays a crucial role both in campaign design *and* evaluation. In the banking example, a test mailing to 30,000 customers could provide data for a predictive model. That model would be used to score the company's 1,000,000 customers, 400,000 of whom might fall in the top four deciles that are profitable. Of those 400,000 customers, the mailing might go to 390,000, with 10,000 held out as a control. The first test is crucial for targeting the campaign; the second test is crucial for evaluating it.

Customer surveys are especially important for evaluating qualitative aspects of the campaign such as reinforcing product positioning. The company may survey a subgroup (e.g., 500–1,000) of its customers after the campaign to make sure the product and company is being perceived according to plan. The time and expense of conducting surveys may preclude evaluating every communication using a survey. However, periodic use of surveys can be an important way to ensure a firm's database marketing efforts are not leading the company astray from its desired positioning and target group.

In summary, communications campaigns can be evaluated using direct calculations, *post hoc* statistical analysis, embedded tests, and surveys. Direct calculations are the easiest when there are no "externalities" to worry about, such as cross-channel cannibalization. Embedded tests are easy to "read" and the opportunity cost is worth the clean calculation of profits. However, testing may not always be practical. For example, a direct response television campaign cannot be randomly targeted to some customers within a given locale and not targeted to others. *Post hoc* statistical analyses become essential (Tellis et al. 2000). Surveys are especially useful for evaluating the qualitative objectives of a campaign, and can be an important source of insight for the database marketing firm that can easily get caught up with the immediate "bottom line" while losing sight of the long term.