08-069

Allocating Marketing Resources

Sunil Gupta
Thomas J. Steenburgh

Allocating Marketing Resources

Sunil Gupta
Thomas Steenburgh[1]

January 28, 2008

[1] Sunil Gupta (sgupta@hbs.edu) is Edward W. Carter Professor of Business Administration and Thomas Steenburgh (tsteenburgh@hbs.edu) is Associate Professor of Business Administration at the Harvard Business School, Soldiers Field, Boston, MA 02163.

Allocating Marketing Resources

Abstract

Marketing is essential for the organic growth of a company. Not surprisingly, firms spend billions of dollars on marketing. Given these large investments, marketing managers have the responsibility to optimally allocate these resources and demonstrate that these investments generate appropriate returns for the firm.

In this chapter we highlight a two-stage process for marketing resource allocation. In stage one, a model of demand is estimated. This model empirically assesses the impact of marketing actions on consumer demand of a company's product. In stage two, estimates from the demand model are used as input in an optimization model that attempts to maximize profits. This stage takes into account costs as well as firm's objectives and constraints (e.g., minimum market share requirement).

Over the last several decades, marketing researchers and practitioners have adopted various methods and approaches that explicitly or implicitly follow these two stages. We have categorized these approaches into a 3x3 matrix, which suggests three different approaches for stage-one demand estimation (decision calculus, experiments and econometric methods), and three different methods for stage-two economic impact analysis (descriptive, what-if and formal optimization approach). We discuss pros and cons of these approaches and illustrate them through applications and case studies.

1. **Introduction**

Marketing is essential for the organic growth of a company. Not surprisingly, firms spend billions of dollars on marketing. For example, in 2006, Proctor and Gamble spent over $4.9 billion in advertising alone. The total advertising budget of U.S. companies in 2006 exceeded $285 billion (Advertising Age 2007). This is more than the GDP of Malaysia, Hong Kong or New Zealand. Given these large investments, marketing managers have the responsibility to optimally allocate these resources and demonstrate that these investments generate appropriate returns for the firm.

Allocating marketing resources is a complex decision in a constantly evolving environment. The emergence of new media such as online search and display advertising, video games, virtual worlds, social networking, online user-generated content, and word of mouth marketing is creating both new opportunities and challenges for companies. It is not easy to isolate the effect of a marketing instrument in this dynamic business environment where multiple factors influence sales and profits. Consequently, many managers continue to rely on simple heuristics and decision rules for resource allocation. For example, it is common practice for managers to use "percentage-of-sales" rule for allocating their advertising budget (Lilien, Kotler and Moorthy 1992). Industry sources commonly publish such advertising to sales (A/S) ratios and managers routinely monitor them. In the sales force arena, Sinha and Zoltner (2001) report that companies typically constrain the ratio of their sales-force cost as a percentage of total sales.

An alternative approach is to arrive at the marketing budget based on a "bottom-up" method. A manager may arrive at the advertising budget based on the desired level of brand awareness and the cost of various media vehicles to achieve this awareness. Similarly, in the pharmaceutical industry a firm may decide how many physicians it wants to reach and how frequently they should be contacted. This combination of reach and frequency determines the required size of the sales force (Mantrala 2006). While such allocation methods are reasonable, they are generally sub optimal. Based on sales force size and resource allocation studies at 50 companies, Sinha and Zoltner (2001) report that, on average, optimal allocation has the potential to improve firm's contribution by 4.5% over current practices.

The approaches mentioned above have some merit. They explicitly or implicitly consider a firm's objectives (how many physicians do we wish to reach), its costs (A/S ratio) as well as its competitive environment (firm's A/S ratio compared to competitor's A/S ratio or industry benchmark). However, these methods have limitations. For example, competitive parity (e.g., A/S ratios) is useful only if competitors are equal in strength, have similar objectives and are acting optimally. Further, the methods mentioned above are incomplete since they do not account for how markets respond to marketing actions. The purpose of this chapter is to highlight practical approaches that account for costs, competitors as well as customers' reactions to marketing actions.

Marketing resource allocation decisions need to be made at several levels – across countries, across products, across marketing mix elements, across different vehicles within a marketing mix element (e.g., TV versus internet for advertising). Each decision requires some specific considerations. For example, when allocating resources across countries, managers need to account for country-specific factors (e.g., growth, local environment etc.) as well as spill-over effects of marketing actions across countries. Similarly, allocation of resources across products requires a careful consideration of substitution and complementary nature of the products (Manchanda, Ansari and Gupta 1999, Sri Devi, Ansari and Gupta 2007). In spite of these differences, there are many fundamental elements that are common across all these decisions – for example, how do customers respond to changes in a marketing action. In this chapter, we focus on these common themes. Majority of our discussion will be around marketing resource allocation for a single product, although the basic approaches can be extended to other scenarios.

Finally, this chapter will deal with rigorous, yet practical approaches to marketing resource allocation. As such we will draw on academic research and practical examples that deal with real-world situations rather than small scale lab studies or theoretical models. While the latter play a strong role in developing theories as well as improving our understanding of a certain phenomenon, we are primarily focused on how these theories can be applied in practice. Given this focus we do not intend this chapter to be a literature review of academic work, nor a road map for future research. Our purpose is simply to lay out a framework for managers who are responsible for allocating marketing resources for their products and services.

2. Approaches for Resource Allocation

The process of marketing resource allocation consists of two stages. In stage one, a model of demand is estimated. This model empirically assesses the impact of marketing actions on consumer demand of a company's product. Ideally, the model also includes competitive activities. While in some cases data on competitors' actions are available (e.g., scanner data studies for consumer packaged goods), in many other scenarios these data are not known (e.g., in database marketing).

In stage two, estimates from the demand model are used as input in an optimization model that attempts to assess the economic impact of marketing actions. This stage takes into account costs as well as firm's objectives and constraints (e.g., minimum market share requirement). While most optimization models do not account for competitive reactions to changes in target firm's marketing budget, more sophisticated models can take these reactions into consideration either through simulation or game theoretic equilibrium models.

Over the last several decades, marketing researchers and practitioners have adopted various methods and approaches that explicitly or implicitly follow these two stages. In Table-1, we have categorized these approaches into a 3x3 matrix, which suggests three different approaches for stage-one demand estimation, and three different methods for stage-two economic impact analysis. We begin by describing the pros and cons of each option at a high level in the remainder of this section. We go into greater depth in the next section by discussing specific examples of how researchers have used the techniques to address issues commonly encountered in practice.

Insert Table-1

2.1 Demand Estimation (Stage-1)

There are three broad approaches for demand estimation as shown in Table-1. Each approach has its pros and cons and each is more suitable in some situations than others.

2.1.1 Decision Calculus

In a classic article, Little (1970) lamented that "the big problem with management science models is that managers practically never use them," (p. 1841). He argued that models should be simple, robust, easy to control by managers, adaptive to changing environment, complete on important issues and easy to communicate. However, most models fail to meet these requirements. It is hard to find good models that are simple and yet include all the information relevant for a complex business environment. It is even harder to obtain appropriate data to empirically estimate these models. This prompted Little to coin the term "decision calculus" to describe models in which managerial judgment is used as input.

In many situations, the decision calculus approach is perhaps the only way to build a demand model. Consider a firm that wants to decide on the optimal number of times its sales force should call on physicians. If this firm always used a certain call frequency in the past, it has no practical way of finding how changes in call pattern may affect demand. Lack of historical variation in call patterns and practical difficulties in conducting experiments leave few options for the firm to build such a model. Decision calculus uses managerial input to estimate the demand function that can be subsequently used in stage-2 for optimization (Lodish 1971).

Since Little's 1970 article, a series of studies have used decision calculus to calibrate demand models and allocate resources successfully (Wierenga et al. 1999, Divakar, Ratchford and Shankar 2005, Natter et al. 2007). In two forecasting situations where managers made real-time forecasts, Blattberg and Hoch (1986) show that statistical models and managerial judgment achieved about the same level of predictive accuracy, while a combination of model + manager outperformed either decision input. They suggest that while models are better at combining complex data in a consistent an unbiased fashion, managers are better at incorporating intangible insights about the market and the competitive environment.

In general, decision calculus provides a useful approach for demand estimation using managerial judgment when a firm does not have historical data and can not afford, either due to lack of money or time, to do experiments. This approach is also appropriate if there are dramatic changes in the industry, a firm's environment, or a firm's strategy.

For example, managers face uncertainty and challenges in allocating resources to new media such as keyword searches, social networks or buzz marketing. However, these managers have significant experience in traditional advertising and its effectiveness. Their experience and expertise in advertising can provide them a strong benchmark for the potential effectiveness of new media channels (e.g., knowing that the traditional advertising elasticity is 0.1, a manger can judge if new media is likely to be twice as effective). These starting benchmarks can be updated as managers gain more experience with the new media channels.

Decision calculus approach might also be appropriate if managers would only be willing to use a model that considers their personal knowledge and expertise. A key strength and at the same time a key limitation of this approach is its reliance on managerial input which can be biased. We refer the interested reader to Eisenstein and Lodish (2002), who review the marketing literature on this approach and provide guidance to researchers and practitioners on how to improve them.

2.1.2 Experiments

Experiments provide a useful way to assess consumers' response to stimuli. By allowing a manager to control for factors that otherwise may influence the outcome; they enable him to isolate the impact of the marketing instrument under study. Experiments are also useful to gauge consumer response to new activities that the firm has not tried historically.

Catalog and credit card companies with millions of customers find it very useful to set up test and control samples to assess the effectiveness of various direct marketing programs. Consumer packaged goods firms have frequently conducted advertising experiments in matching cities. The advent of technology has now made it possible to conduct split-cable TV experiments with test and control households in the same city to assess the effectiveness of various advertising creatives and budgets (Lodish et al 1995). Experiments, such as conjoint analysis, are routinely used for new product design as well as to find consumers' price sensitivity. Harrah's Entertainment Inc. has used experiments very effectively to offer the right reward to the right customers at the right time (Loveman 2003).

In general, experiments provide a useful way to gauge consumers' response to a marketing action when a firm can afford to subject test and control samples to different treatments. In some situations this is not feasible. For example, if a firm wishes to test a new compensation system or organization structure for its salesforce, it may not be practically possible to have two different systems or structures for the test and control groups. Experiments are generally good at obtaining the short run impact of an action. While it is possible to find the long run effects of marketing actions through experiments, it becomes practically difficult to control environmental and competitive changes for a very long period of time. Finally, experiments can become very complex with an increasing number of factors to test. This requires a manager to carefully consider only a few critical factors that he wishes to test.

These critical factors can be determined in three ways. First, the choice of factors is governed by the decision objectives of a manager. For example, a manager in charge of allocating resources for a catalog company needs to know who to send catalogs and how often, since catalogs form a large part of his budget. Second, prior experience and knowledge of the business gives a manager a good sense of the key drivers of his business. A knowledgeable manager should know if pricing, advertising, or distribution is critical for the growth of his business. Third, similar in spirit to the multi-phase trials in the pharmaceutical industry, managers can conduct small scale experiments to determine which factors have the most impact on sales and profit. These sub set of factors can then be tested in greater detail in a large scale experiment.

2.1.3 Econometric Approaches

With the increasing availability of data, improved computer power and advances in econometrics, it is now easier for firms to harness their historical data to estimate the impact of various marketing instruments on consumer demand. In the consumer packaged goods industry, the advent of scanner data has revolutionized marketing resource allocation through this approach.[2] A large number of academic studies have built models to understand the effectiveness of sales promotions and advertising (Gudagni and Little

[2] Scanner data collect information about consumer purchases at the stores. The data also include information about consumer demographics as well as complete marketing mix information about all competitive brands.

1983, Gupta 1988, Tellis 1988). Many studies have also teased out the short and long run impact of these actions (Mela, Gupta and Lehmann 1997, Jedidi, Mela and Gupta 1999, Koen, Siddarth and Hanssens 2002). Companies such as Information Resources Inc. and Nielsen routinely offer marketing mix models based on these data as a service to their clients. The client firms, such as Campbell, actively monitor their marketing resource allocation based on the results of these models.

Econometric studies have also found significant use in database marketing. A large number of studies have used companies' historical data on RFM (recency, frequency and monetary value) to build models that estimate consumer response to marketing campaigns. These models significantly improve marketing resource allocation by providing powerful insights about who should be contacted, when and how frequently (Gupta et al. 2006, Venkatesan and Kumar 2004).

Econometric approach uses historical data of a firm and allows a manager to build models that capture the complexity of his business. These methods provide accurate and unbiased assessment of marketing effectiveness. They allow a firm to constantly learn and adapt from its previous efforts. The models are also transportable across products and geographies and thus provide a common language across the organization. When a firm has limited historical data (e.g., new product introduction), it is still possible to use this approach by using analogies or meta-analysis priors, which can be updated in a Bayesian fashion using current data on the new product.

This approach is most useful when markets are relatively stable such that historical estimates provide a good indicator of the future market conditions. A method based on historical data is unable to capture situations where the industry dynamics or a firm's strategy has undergone major changes. Therefore, model recommendations are relevant only within the range of historical data.

2.2 Economic Impact Analysis (Stage-2)

Stage-1 provides estimates of how market demand is influenced by marketing actions. These estimates become the input for stage-2 where a firm decides on optimal resource allocation that maximizes its profits. As indicated in Table-1, there are three broad approaches for stage-2.

2.2.1 Descriptive Approach

This approach uses parameter estimates of the demand equation to make directional recommendations. For example, high consumer price sensitivity for a brand may suggest allocating more promotional dollars to this brand. Parameter estimates can be converted into demand elasticities, which can be compared across various marketing instruments to guide resource allocation (Steenburgh 2007). For example, in a large scale study for the pharmaceutical industry, Wittink (2002) found very low elasticity for direct-to-consumer advertising (DTC). He further converted these elasticities into ROI to show that investment in DTC does not pay off. At the minimum this provides directional guidance to pharmaceutical managers to cut down on their DTC budget.

Descriptive approach is simple and easy to use. It is the natural outcome of demand analysis in stage-1 and requires little additional analytical work. However, as models become complex with interactions among marketing elements, descriptive approach is less suitable to isolate the effects of each marketing action. For example, it may be straightforward to see the short run effect of advertising on sales using this method. However, it is more difficult to use this approach to assess the net effect of advertising on sales that takes into account not only the short run effects but also the long run impact on brand health as well as consumers' price sensitivity. This approach also does not take competitive reactions into account, so it is better suited to understanding how the world works today rather than how it will work if major changes are undertaken.

2.2.2 Simulation or "What If" Analysis

To handle complex interactions mentioned above, optimal resource allocation can be achieved using simulations or "what if" analyses. Effectively, a manager can try various marketing plans as inputs into the demand model and simulate the effects on sales and profits. Increasing computer power makes it easy to conduct hundreds of such simulations in a short period of time. The model complexities are preserved and the user does not have to make subjective interpretations about the interactions between various marketing elements. A wide array of simulations can also build confidence in the robustness of the results. For example, if profits do not change significantly with large changes in advertising, a manager can safely conclude that other elements of marketing mix deserve more attention.

Simulations have two key limitations. First, as the number of options (marketing actions and their budget levels) increase, the potential combinations for simulation can increase exponentially. Second, simulation is effectively a coarse grid search over the profit function. In other words, it provides an approximate rather than an exact solution to the optimization problem.

2.2.3 Optimization

The most sophisticated and complex approach is the build a formal optimization model that uses demand parameters from stage-1 as inputs and sets up a profit function that is maximized using operations research algorithms. These algorithms may include linear, integer or dynamic programming methods. This approach also allows managers to put in business constraints as part of the optimization algorithm. For example, based on his understanding of the business and the market, a manager may decide that it is absolutely essential to have a certain minimum level of advertising.

Optimization methods generally take two approaches for estimation. An "elegant" approach is to find a closed-form mathematical solution to the optimization problem. However, a more practical approach is to conduct numerical grid search over the parameter space to find optimal or near optimal solutions.

On the positive side this approach provides a comprehensive solution to the resource allocation problem by searching over the entire space of options. At the same time, as the complexity of the problem and the number of options increase, it is generally harder to use this approach. Consequently, many studies simplify the problem by dealing with one or two marketing elements at a time.

3. Applications

In this section, we discuss a few applications of the previously discussed methodologies. Our goal is to provide concrete examples of each methodology to promote understanding of how it can be used rather than to provide an exhaustive list of research on each topic. For each application, we briefly describe the managerial problem, the research approach used to solve this problem, and the results obtained from this approach.

3.1 Experiments (Stage-1) and Descriptive Analysis (Stage-2)

We describe two studies that used experiments. The first study examined the effectiveness of word-of-mouth communication, and the second study allocated promotional dollars between new and existing customers of a catalog company.

Effectiveness of Word-of-Mouth Communication

Problem: The impact of "New Media" on marketing and advertising is continually being revised upward. Online search and display advertising, video games, virtual worlds, social networking, online user-generated content, and word of mouth marketing are growing by leaps and bounds and are being used by companies to address the fragmentation of consumer markets and the swing from mass to niche marketing. According to a recent survey by McKinsey & Company, a third of the companies that advertise online are already spending more than 10 percent of their advertising budgets there (Bughin, Erbenich and Shenkan 2007).

Traditional marketing giants such as Unilever and Procter and Gamble are experimenting with the new media. Unilever has gained significant attention with its campaigns for Dove and Axe, majority of which was driven by web sites, online blogs and Youtube. Since 2001, Proctor and Gamble has been building Tremor, a word-of-mouth network, which includes two consumer panels: VocalPoint, which consists of 450,000 moms of children under 19 years old[3], and Tremor Teen a network of 230,000 teenagers age 13 - 19 years old[4]. P&G claims that the members or "connectors" as P&G refers to them, are a select group of consumers that talk to their friends more than the average person. For example, P&G boasts that Vocalpoint moms talk to 20-25 people everyday, versus the average mom who only talks to 5 people.

[3] Tremor website, "VocalPoint Panel", *http://business.tremor.com//tremor_vocalpoint_panel.html*, accessed October 2007.
[4] Temor website, "Tremor Teen Panel", *http://business.tremor.com//tremor_teen_panel.html*, accessed October 2007.

Godes and Mayzlin (2007) report several other examples where companies promoted their products and services through word-of-mouth (WOM) marketing. For example, in 2001 Hasbro promoted its new handheld video game called POX through 1,600 "cool" kids of Chicago elementary schools. In September 2005, NBC promoted its reality show about weight loss, "The Biggest Loser" by encouraging its 1,000 biggest fans to throw parties during an advanced screening of the show's premiere. In 2006, WD-40 used P&G's Vocal Point to promote its new product extension, the "No-Mess" pen.

Despite companies' foray into new media and buzz marketing, several questions remain. How effective is WOM in generating sales? Can firms create WOM or is it a naturally occurring phenomenon that is not under a firm's control? Who are better disseminators of WOM?

Approach: Godes and Mayzlin (2007) address these questions through a field experiment conducted by BzzAgent, an agency, for Rock Bottom Brewery, a restaurant chain. At the time of the study, Rock Bottom Brewery did business in 15 markets across the U.S. with annual gross sales of about $100 million. The company maintained a loyalty program with several thousand customers as members.

For the field test, Godes and Mayzlin recruited 381 of the firm's loyal customers. In addition, they also recruited 692 "agents" of BzzAgent. This agency maintains a panel of agents who are encouraged to create WOM for client's products or services. The purpose of recruiting loyal customers of the firm and non-customers (agents of BzzAgent) was to compare the WOM effectiveness of the two groups. The WOM creation process ran for a product category in 15 markets for about 13 weeks (April to June 2003). Participants were asked to report their WOM creation activity by filling in a report on a web site each time they engaged in a WOM episode. Each report was graded by BzzAgent on its potential to create meaningful WOM. Participants had an incentive to create meaningful WOM since higher scores entitled them to win prizes. The average prize was valued at around $15. Godes and Mayzlin also obtained weekly sales data for each of the 15 markets for the campaign period as well as for the same time period from last year.

There was natural variation in the WOM creation across the 15 markets. To assess the impact of WOM on sales, Godes and Mayzlin ran fixed effects regression model which controlled for week and market effects.

Results: Godes and Mayzlin found that all WOM is not created equal. Specifically, the impact of WOM created by customers with no relationship to the firm had a much greater impact on sales than the WOM created by firm's loyal customers. This seemingly counterintuitive result is actually consistent with the theory of weak ties (Granovetter 1973) that suggests that WOM through acquaintances has significantly more impact than WOM to those with stronger ties in a social network. In a way this suggests that customers unconnected with the firm are likely to be less biased and more believable and therefore should be weighed more heavily in terms of their impact on WOM. Godes and Mayzlin also found that each WOM from non-customers yielded average incremental sales of $192. These results provide useful guidelines to the firm on how much resources it should spend on firm created WOM.

Kumar, Petersen and Leone (2007) used a survey method to assess the value of WOM and referral. They polled 9,900 customers of a telecom firm and 6,700 customers of a financial services firm on their referral intention. Then they tracked their behavior and the behavior of the prospective customers that the referring customers brought in over time. They also adjusted for the possibility that some of the prospecting customers would have joined the service anyway. Several interesting results emerged from their study. First, less than half of the customers who indicated their intention to refer their friend to company's services actually did so (for financial services 68% intended to refer friends but only 33% actually did; for telecom out of 81% intenders only 30% followed through). Second, very few of the referrals actually generated customers (14% at financial services and 12% at the telecom company). Further, of those prospects that did become customers very few were profitable (11% for financial service and 8% for telecom firm).

Allocating Promotional Dollars

Problem: Anderson and Simester (2004) described a situation where a catalog company wanted to understand how to allocate promotional dollars between new and established customers. The firm was also concerned that promotion may have negative long run impact on consumers' repeat purchase behavior. The company, a medium-sized firm, sold approximately 450 products targeted at well-educated older customers. The products were generally experience goods, similar to books and software. Historically, the company had offered a variety of discounts ranging from under 20% off regular price to as much as 70%. Majority of its sales were with promotions ranging between 20-60% discounts, with 20-30% discount sales accounting for almost one-third of overall sales.

Approach: To understand the long-run effects of promotion depth on new and established customers, Anderson and Simester (2004) conducted three large scale field experiments. Study A was conducted with established customers of the firm. A control version of the catalog was sent to 37,758 randomly selected customers. This catalog presented 86 products on 72 pages. The average promotional discount in the control condition was 30% off regular price. A randomly selected 18,708 customers received a promotional or test catalog where 36 of the 86 products were offered at an average discount of 60% (instead of 30% for the control condition). The price on the remaining 50 products remained the same in both the test and the control conditions.

Results: From the control group, 761 or 2.02% customers bought from the catalog, while the promotional catalog generated a response from 597 or 3.19% of the customers. Customers from the test group ordered an average of 2.14 units at an average price of $78.51, while control group customers ordered an average of 1.59 units at an average price of $124.03. Thee customers were then tracked for the next 28 months where both groups received same catalogs. Future purchase behavior revealed that compared to control group, customers from the test group purchased fewer products (6.89 versus 7.67) and less expensive items (average price $84.86 versus $95.51). As a result the overall future revenue per customer was sharply lower among customers in the test condition ($584.68 versus $733.50). In other words, increasing the depth of promotion had a negative long run effect among established customers (Table-2).

Insert Table-2

Anderson and Simester (2004) conducted two additional studies with new or prospective customers. These studies were similar to study A and their details are given in Table-2. In contrast to the results of study A with established customers, studies B and C with new customers showed exactly the opposite results. Specifically, while increasing the depth of discounts had a negative long-run effect among established customers (study A); it had a significant positive long-run effect among new customers. Although these results do not provide an optimal promotional discount value, they provide good direction to the company on how to allocate its promotional dollars among new and established customers.

It is common practice among catalog companies to conduct such experiments. In 1999, over 31% of catalog firms reported conducting split-sample experiments of pricing strategies (Direct Marketing Association 2000). Experiments are also commonly employed for testing advertising budget and creative. Eastlack and Rao (1986) report an advertising experiment for V-8 vegetable juice. In this experiment, they varied advertising budget, advertising creative as well as media mix between radio and TV across various markets to examine the effect of these elements. Using this experimental data, they also estimated the S-shaped advertising response function to determine the threshold or the minimum level of advertising as well as saturation or the maximum level of advertising.

Lodish et al. (1995) examined 389 real world split-cable T.V. advertising experiments.[5] They found that the average advertising elasticity for new products is much higher (0.26) than for established products (0.05). They also found that T.V. advertising is more likely to work when there are accompanying changes in ad creative and media strategy.

3.2 Econometric Estimation (Stage-1) and Descriptive Analysis (Stage-2)

Problem: Allocating the budget across marketing instruments is challenging for every organization. For example, for pharmaceutical firms the key issue is how to

[5] Split-cable experiments allow advertisers to stream different advertisements to different households in the same city. Purchases of these households are then tracked through store scanners to link the impact of advertising budget or creative on their purchase behavior.

allocate resources among various marketing instruments such as detailing to physicians, journal advertising, physician meetings and direct-to-consumer advertising.

Turf battles can make the process as much about building the biggest fiefdom as it is about making the right allocation decisions for the company. The sales organization is bound to ask whether marketing really needs to create more ads, and the marketing is bound to ask if more salespeople are needed in the field. When a history of data exists, estimating the historical return-on-investment (ROI) for each marketing instrument is one way to take the emotion out of the budgeting process and help the firm allocate its marketing dollars more effectively.

Approach: In a broad study of the pharmaceutical drug market, Wittink (2002) assessed the return on investment (ROI) of several marketing instruments during the period of 1995 – 2000. His study was based on drugs that produced at least $25M in revenue in 2000, which resulted in 392 branded and 127 generic drugs being included. In stage-one demand estimation, Wittink used standard regression analysis to determine the relationship between unit sales and each of the marketing instruments. Instruments used to market pharmaceutical drugs include detailing (DET), physician meetings (PME), journal advertising (JAD), and direct-to-consumer advertising (DTC). The first three of these instruments are directed only at physicians, whereas DTC advertising may be seen by physicians but is primarily directed at consumers.

The effectiveness of the marketing instruments was thought to depend on the drugs' size and the length of time they had been on the market. To control for these factors, each drug was classified into one of nine categories based on its sales revenue and launch year. The revenue categories were $25 to $100M, $100 - $500M, and greater than $500M, and the launch year categories were earlier than 1994, 1994 to 1997, and 1998 to 2000. Table-3 characterizes how the marketing expenditures varied across the various categories of drugs. The first value in a cell is the percent of products that had expenditures in at least one month during the study period for a given marketing instrument, and the second value is the average monthly expenditure. To illustrate, 94% of the drugs that produced between $25M and $100M in revenue and that were launched prior to 1994 devoted at least some money to detailing. For those brands with at least some detailing, the average monthly expenditure was $155K.

Insert Table-3

Detailing was the most well used marketing instrument, being employed by over 90% of the drugs across all classifications, and it always had the highest average spending. The other physician-directed marketing instruments, physician meetings and journal advertising, were also widely used, being employed by at least 73% of the drugs across all classifications and often by more than 90% of the drugs. DTC advertising was the least used marketing instrument, as it was used by between 10% and 78% of the drugs across the different classifications. The average monthly expenditures devoted to this instrument, on the other hand, were sometimes quite high.

Results: Let's now turn to the marketing ROIs. These calculations are interpreted as the estimated increase in revenue for a $1 increase in spending on a marketing instrument. Therefore, an ROI with a value less than one suggests that the incremental marketing spending would not pay for itself through increased sales. These numbers should not be interpreted as one marketing instrument being more effective than another, per se, as the ROI values are also influenced by the amount of money being spent on the instrument.

The ROIs for brands with revenues greater than $500 M are listed in the Table-4, Panel A. The values of the physician-directed instruments (DET, PME and JAD) are fairly consistent and are all greater than $1. Furthermore, the marketing spending is more effective for newer drugs than it is for older ones. A marketing dollar spent on a drug launched prior to 1994 returns about $3 in revenue, whereas a dollar spent on a drug launched between 1998 and 2000 returns about $12 in revenue. By comparison, the ROI on DTC advertising is much lower. Spending on DTC advertising only pays for itself when it is used on recently launched drugs, and even in this case it is a relatively small ROI of $1.3.

Insert Table-4

A similar pattern of ROIs arise for the 192 drugs that produce revenues between $100 and $500M: the effectiveness of the marketing spending is greater for the physician-directed marketing instruments than it is for DTC advertising, and the effectiveness is greater for more recently launched drugs than it is for older drugs. (See Table-4, Panel B) Nevertheless, the overall return on investment is smaller for intermediate-revenue drugs. For the physician-directed instruments, the ROIs range between $1.2 and $2.3 for the oldest drugs (as compared to $3 for the largest drugs) and between $2.1 and $4.6 for the newest drugs (as compared to $12 for the largest drugs). Furthermore, DTC advertising never pays for itself for the intermediate drugs, as the ROIs increase for more recent launch dates, but range between $0.1 and $0.2.

For the 137 brands that produce the lowest revenues, we observe a different pattern of results. (See Table-4, Panel C.) Not all physician-directed spending yields positive ROIs, as PME is $0.1 regardless of the launch date and DET ranges between $0.9 and $1.0. JAD is the only instrument that pays for itself, with ROIs ranging between $6.2 and $7.2. Furthermore, while the ROIs do increase for more recently launched drugs, they do not change much across launch dates. All else equal, firms might consider directing more money toward journal advertising and away from detailing and physician meetings for these drugs. Similar to the intermediate- and large-revenue drugs, DTC advertising does not pay for itself, as it yields ROIs of zero.

While the ROI numbers do not tell us what the optimal spending levels are for each marketing instrument, they do give us a better sense of how firms might reallocate their marketing dollars. We would expect firms to shift their spending away from instruments that produce low ROI and toward investments that produce high ROI, keeping other organizational goals in mind. For example, a firm might keep investing in detailing even if the ROI is not justified if it has strategic interest in maintaining a sales force to sell upcoming products. Calculating the ROIs, however, can help the firm make better tradeoffs when making these types of decisions too.

3.3 Econometric Estimation (Stage-1) and What-if Analysis (Stage-2)

Problem: Allocating resources between advertising and trade or consumer promotions is a topic of constant debate and discussion in most organizations. Proponents of advertising claim that advertising builds brand equity and insulates a brand from price changes in the market place. Supporters of promotions highlight dramatic market response to short term promotions as evidence of their effectiveness. While it is easier to assess the short-term effects of promotions (e.g., Gudagni and Little 1983, Gupta 1988), it is much harder to determine the long-term effects of promotions and advertising. Do promotions have a long-run negative impact on a brand? Do these long-run negative effects outweigh the short-run positive effects of promotions? Taking into account both the short and long-run effects, what is the optimal allocation of resources between advertising and promotions? Jedidi, Mela and Gupta (1999) addressed these questions for a consumer packaged goods product.

Approach: Jedidi, Mela and Gupta (1999) used eight years of disaggregate data (1984-1992) on 4 brands in a consumer non-food category for 691 households. Descriptive statistics of the data are given in Table-5. Jedidi et al. used discrete choice models to capture consumers' decisions of which brand to buy and how much quantity to buy as a function of consumer characteristics and marketing activity (regular price, temporary price reduction due to promotion and advertising). They further postulated that promotion and advertising can have long-run effects on consumer purchases in two ways – by influencing the brand equity and by affecting consumers' price sensitivity. The demand model (stage-1) was estimated using a maximum-likelihood procedure. In the second stage, they conducted simulations to assess the managerial implications of these results for resource allocation. These simulations also included competitive reaction functions. Jedidi et al. argued that simulating the effect of a change in marketing activity of a brand, say, an increase in discounts, in the absence of competitive reaction could lead to an optimistic assessment of these effects.

Insert Table-5

Results: The results of this study showed that, as expected, promotions had a positive and significant impact on consumer choice in the short-run. In the long-run, advertising improved brand equity while promotions had a negative impact on brand equity. Further, frequent promotions made consumers less promotion sensitive in their brand choice and more promotion sensitive in their quantity decision. In other words, frequent promotion of brands made it unnecessary for consumers to switch brands and made them more likely to stockpile when their favorite brand was on promotion.

These results are intuitively appealing. However, these descriptive results do not provide any specific directions for resource allocation. They still do not tell us if the short-run positive effects of promotion are outweighed by promotions' long-run negative effects. To address this question, Jedidi et al. conducted simulations. These analyses first estimated baseline sales and profits in the absence of any changes in marketing policy. Next, price, promotion or advertising of a target brand was changed by 5% and its impact on competitive response as well as consumer response was simulated based on the models of stage-1. Results of this simulation are presented in Table-6.

Insert Table-6

Results showed that increasing promotion depth or frequency decreased profits of all four brands. However, increasing advertising had mixed effects on brand profitability. It marginally improved the profits of only one brand while profits for three other brands went down.

Two broader conclusions emerge from this study. First, it is perhaps too simplistic to suggest that firms should increase advertising or cut promotions. This decision needs to be made on a case by case basis depending on each brand's current advertising and promotion budget as well as its position in the market place. Second, it is remarkable to see that 5% increase in advertising or promotions has less than 1% effect on profits. This seems to suggest that the market is operating efficiently and managers in this product category are making decisions that are close to optimal.

There are many studies that employ this approach of estimating a demand model using econometric method in stage-1 and then conducting simulations to derive optimal resource allocation in stage-2. Duvvuri, Ansari and Gupta (2007) build a model for retailers where they account for cross-category complementarities. Using data from six product categories they show that discounts in one category (e.g., spaghetti) can affect the purchase in the target category as well as its complementary category (e.g., sauce). Their simulations further show that the average profitability gain from targeted customer discounts over non targeted discounts is only 1.29% if these complementarities are ignored. However, profit gain is almost 8.26% when these complementarities are included.

3.4 Decision Calculus (Stage-1) and Optimization (Stage-2)

Problem: While every salesperson faces the problem of how to allocate effort across customers, there is little consistency in how the issue is addressed. Some salespeople essentially ignore the problem by spending most of their time with customers that like them best. Others simply base their current calling schedules on historical visitation patterns or develop rules of thumb to help manage the allocation task – say by making one call per month for every $100K that an account bills. These heuristics may be systematic, but they do not necessarily meet the overall goals of the firm. If there are diminishing returns to the number of visits, a salesperson would be better off spending less time with their biggest accounts and more time with prospects and smaller accounts.

The firm would like the salesperson to choose a calling pattern that maximizes some objective (say profit, but many others are possible), but several issues stand in the way. First, it often is difficult to build a statistical model with historical data that adequately predicts how an individual salesperson will fare with an individual customer. Personal selling is a unique endeavor, so one salesperson may thrive in a given account whereas another may struggle. Furthermore, while much progress has been made over the years, statistical models can have some difficulty in capturing account-by-account nuances and data limitations may require models to be developed on a more aggregate basis. Furthermore, salespeople tend to favor their own judgment over statistical models. Thus, they are more likely to follow a recommendation if it takes their knowledge and experience into account and if they understand why it is being made.

Approach: Lodish (1971) developed an interactive computer system, named CALLPLAN, to address this problem. CALLPLAN divided the salesperson's underlying time allocation problem into two stages. In the first stage, the expected contribution of all possible calling policies was independently determined for every prospect and account using the decision calculus approach. In the second stage, a mathematical program was used to determine the best possible calling schedule. CALLPLAN maximized the salesperson's total contribution across all accounts by considering the returns from all possible calling schedules in light of the limited amount of time that a salesperson could work. The number of possible calling schedules that an individual salesperson would be able to calculate without the aid of a computer is limited, but CALLPLAN was able to process this information very efficiently.

CALLPLAN was designed to be used by a salesperson in conjunction with his or her manager. This required the system to be easy to use and understand and the outputs to be quickly recalculated as the inputs or assumptions changed. Lodish reported that salespeople were quite comfortable using the system after a single day of training. The data required by the system were straightforward. To assess costs, salespeople were asked to input the time it took to make calls in different geographic regions. To estimate the response function in a given account, salespeople were asked to input the minimum and the maximum number of calls that could be made in a given pre-set period (typically one to three months) and to estimate various returns from different calling levels, which captured the salesperson's expert knowledge.

To make the system easier to use, salespeople could estimate the response functions in various ways. The expected returns could be directly given for all possible calling levels in each account; e.g. if the minimum number of calls was three and the maximum number was ten, then the salesperson could directly estimate the returns from each of the eight different calling levels {3, 4, 5, 6, 7, 8, 9 or 10 calls}. Alternatively, the salesperson could ask the computer to generate a best fitting response curve based on their answers to a handful of questions for each account; e.g. what would the response be if you made zero calls in this account? if you made the maximum number of calls? etc. Figure-1 illustrates some fitted response curves.

Insert Figure-1

The computer then developed a calling policy that maximized returns subject to constraints on the required number of calls in each account and on the available time than an individual could work. The computations took less than a minute in 1971, and the program would easily provide instantaneous feedback today.

Results: In his original study, Lodish (1971) reported results for eight Pennwalt salespeople who used CALLPLAN for five months. Based on questionnaires of the salespeople and their managers and his own observations, Lodish concluded the system fostered clearer and more consistent thinking about the calling patterns. Salespeople thought about tradeoffs that they had not previously considered, and the system fostered better communication between salespeople and their managers. Areas of disagreement on assumptions become explicit after using CALLPLAN. Furthermore, salespeople bought into the results because the system used their own estimates as inputs for its calculations. In some cases, CALLPLAN helped salespeople maintain a commitment to keep calling on prospects who (in the sense of expected returns) were the best place to spend their limited time. The system became a motivational tool.

Lodish also found that CALLPLAN was better suited to situations in which the selling was repetitive, as he found in plastics, dental equipment and refrigerants. The amount of time selling in an account was an important factor in predicting sales in these cases. Most participants anticipated increases of between 5% and 30%, and two salespeople reported actual increases of 15% and 30% from more efficient time allocation. Four vacuum cleaner salespeople tried CALLPLAN for four months without much success. Sales in this situation were one-time occurrences, and success was thought to be due to factors other than effort in this case; thus, CALLPLAN was not helpful in helping them allocate their time.

In a subsequent study, Fudge and Lodish (1975) designed an experiment to test the effectiveness of CALLPLAN using twenty United Airlines salespeople in New York and San Francisco. The ten salespeople who used the system were initially skeptical of its worth, but viewed it as a productive planning tool afterwards. Furthermore, CALLPLAN produced behavioral changes in these salespeople that led to significantly higher results. After six months of use, the sales results for individuals using CALLPLAN were 8.1% higher on average than they were for individuals who did not use the system. The actual dollar improvement for those ten salespeople was well into the seven figures, and the probability that such an increase occurred by chance alone was only 2.5%.

The decision calculus approach has been used in contexts other than sales force planning. For example, Little and Lodish (1969) developed an early, interactive computer program, called MEDIAC, to help managers select and schedule advertising media. In this system, the user supplied subjective and objective data about media options and the target audience as well as the firm's advertising budget. Using these data, the system scheduled a set of media options that maximized the total market response. Little (1975a, 1975b) developed BRANDAID to help managers make better decisions with regard to their total marketing plan. Analysis of each marketing element (price, promotion, advertising, distribution, etc.) was contained in its own sub-module, and each sub-module could be expanded upon or dropped as the situation required. An advantage of these early systems (and one of the reasons that we chose to highlight Lodish's paper) was that managers could understand their basic logic even if they did not understand their mathematical algorithms. Thus, working with the system fostered a constructive dialog among users, and managers were more willing to trust the results. This approach to decision making might prove to be especially useful in new media planning because it is unknown how well the established econometric results in the old media will transfer to new.

Although we have not highlighted Little's BRANDAID system in depth, we should note that it more commonly uses what-if analysis instead of optimization to complete stage-two analysis. This, however, seems to be a minor distinction between it and Lodish's CALLPLAN because both systems are used interactively to make decisions. Managers continually revisit and revise their assumptions while using the system and this process ultimately leads them to consider what would happen under new scenarios. Models based on decision calculus might be best thought of as providing a direction for improvement in an ever-changing environment.

More recent work has focused on developing computerized systems, known as Decision Support Systems (DSS), which are able to integrate a wide variety of information, including managerial judgment, to estimate demand. These systems contain data that are collected in a number of ways; for example, they may include sales and costs data from company records, subjective judgments about what would occur from increased marketing spending, as well as a database of competitors' products and sales. A variety of stage-one techniques, including decision calculus and econometric models, are used to bring this information together. In a recent DSS application, Divakar, Ratchford, and Shankar (2005) developed CHAN4CAST to forecast sales of consumer packaged goods at PepsiCo. Their system allowed managers to forecast sales across several channels and to simulate what would happen given a variety of spending and competitive response scenarios. Reinforcing the idea that the system was supposed to be used interactively, CHAN4CAST included a scorecard to track how well past forecasts had predicted the future and managers could use this scorecard to improve the accuracy of their predictions. PepsiCo estimated that the system would return benefits over 1,000% of its costs.

3.5 Econometric Estimation (Stage-1) and Optimization (Stage-2)

In this section, we highlight two econometric studies that have been coupled with optimization techniques to manage marketing investments made in different phases of the customer lifecycle. In the first study, Steenburgh, Ainslie and Engebretson (2003) develop a method that helps firms decide which customers to acquire. In the second study, Tirenni et al. (2007) develop a method that helps firms decide which of their existing customers should be targeted in loyalty program campaigns.

Which Customers to Acquire

Problem: Acquiring the right set of customers is difficult for any company to accomplish, but it can be especially challenging in the context of direct marketing. Companies typically possess limited information that can help predict how prospects will respond to offers, and acquiring third-party data, which often seems of dubious quality, is costly. Furthermore, even when companies can and do choose to buy additional information, the data can be aggregated (say to the zip-code level) to protect consumer privacy, a practice that creates additional econometric headaches for analysts to worry about. Combined, these factors lead to low rates of consumer response, with success rates commonly being under 1%. Thus, companies would like to: (1) develop methods that make the best use of whatever data they have, (2) find optimal methods of choosing which individual prospects to target, and (3) develop methods that can help them value the data they decide to purchase in terms of dollars.

Approach: Steenburgh, Ainslie and Engebretson (2003) develop a hierarchical Bayes variance components (HBVC) demand model to solve these problems. Their technique integrates data collected from multiple sources and econometrically models each set of data at the appropriate level of aggregation. They show that traditional techniques, which do not account for different levels of aggregation (say one source being tracked at the individual-level and another at the zip-code level), result in parameter estimates that are overly confident and lead to inferior predictions about which prospects should be targeted. Another advantage of their technique is that it allows the zip codes themselves to be used as explanatory variables in the demand model. The old maxim "birds of a feather flock together" holds true, as they show that the zip codes contain useful information about how the people in them will behave.

One of the reasons that Bayesian models are gaining popularity in the managerial sciences is that they can be easily combined with decision theory analysis to improve managerial decisions. In our parlance, Bayesian methods allow a seamless integration of the econometric demand estimation in stage one and the economic optimization in stage two. Steenburgh, Ainslie and Engebretson develop several decision rules based on different relationships between the marginal costs of contacting more prospects and the marginal benefits from having more positive responses, and they show how to implement these decision rules using the stage-one results. They show that the decision rules can be relatively easy to implement even when the relationship between costs and benefits are complex.

Results: In an application of their method, Steenburgh, Ainslie and Engebretson studied how a private, southern U. S. university should assess its prospects at the inquiry stage of the admissions process. Prospects at this point of the process have requested information about the university, but have not yet decided to apply. The authors used one set of 38K prospects to build the demand model and a different set of 34K prospects to test the model's predictions. The prospective students resided in approximately 7K zip codes and declared an interest in 128 different majors. These "massively categorical" variables[6] were directly included as explanatory variables in the HBVC models. Campus

[6] Steenburgh, Ainslie and Engebretson use the term "massively categorical" to describe categorical variables such as zip codes and majors that take on many possible values.

visitation data collected at the individual level and supplementary demographic data collected at the zip code level were also used to estimate the models.

The authors show that their HBVC demand model better predicts the enrollment decisions of individual prospects than the standard model does. No matter what data were used to estimate the models, the HBVC model outperformed the null model when using the same set of information. In fact, the HBVC models estimated without the supplementary demographic data were able to outperform the corresponding null models with these data. This suggests that using superior modeling techniques can be more important than purchasing more information. More important than the overall fit of the models, the authors showed that the HBVC model provided a better ordering of the prospects than the null model did. The ability of the models to order the prospects was crucial because the ordering helped determine which prospects to target.

Receiver operator curves (ROC) were used to assess how well the models ordered the prospects. (See Figure-2). These charts were constructed by repeatedly dividing the prospects into two different groups based on their estimated probability of enrollment. Prospects with probabilities below a given cutoff point were placed in one group and prospects with probabilities above the cutoff point were placed in the other. After this is done for all cutoff points between zero and one, the number of enrollees not selected is graphed against the number of enrollees selected for each division. Visually, this implies the better the model, the more the curve will move toward the bottom left-hand corner of the chart. From the figure, it is immediately clear that the HBVC model (represented by the solid line) provides a better ordering of the prospects than the null model (dotted line) does, no matter what information was used to estimate the models.

Insert Figure-2

In addition to being statistically superior to the null model, the HBVC model helped make a practical difference in the university's ability to target individual prospects. The authors derived a willingness-to-pay measure to estimate the economic impact both of buying additional data and of using different models. Averaged over an array of financial assumptions, the expected loss from using the null model instead of the

HBVC model was 43.6%. This loss was greatest when the university had the least amount of data on which to base its predictions (such as when no campus visitation data were present), which suggests, as we might expect, that finding the right model becomes increasingly important when less information is available for analysis.

Which Existing Customers to Target

Problem: Frequent-flyer programs have become ubiquitous among all airlines. Most airlines offer elite status to their customers based on how frequently customers fly with them. Typically airlines consider the upper tier of their frequent flier program to be their most valuable customer segment. Customers in the same elite level (e.g., platinum) receive the same marketing campaign as well as service. However, customers who accumulate the most miles may not pay the highest fare and may be very costly to serve. How should an airline assess the long-run profitability of its customers and how should it allocate its marketing resources across these customer groups?

Approach: Tirenni et al. (2007) address this question for Finnair. This leading Europena airline conducts numerous marketing campaigns targeting more than 700,000 customers. A typical customer receives dozens of campaigns each year. These campaigns have different goals such as cross and up-selling, minimizing attrition and tier upgrade. Campaigns are delivered through various channels such as mailings, in-cabin brochures, magazines and the Internet.

In the first stage, Tirenni et al. build a Markov decision process (MDP) which consists of a set of states, actions, transition probabilities and value functions. For example, a new customer may represent the first state S_1. A marketing action such as a special offer may move this customer to the next state S_2 (e.g., repeat purchase) with a transition probability of 0.7. A club membership may further transition this customer to state S_3 (e.g., loyal customer) with a probability 0.6. Various states and actions are obtained from historical data. Transition probabilities are estimated using a Bayesian procedure. Customer values are also obtained from the observed data. Given these estimates, future customer dynamics are simulated for a given time horizon (e.g., 12 months) to get a distribution of future values. This provides mean or expected value as well as variance of customer value.

In the second stage, Tirenni et al. set up an optimization problem where the objective is to maximize the cumulative expected value while minimizing the variance. They further add user-defined budget and other constraints. This optimization problem is solved using dynamic programming. The solution provides the optimal number of customers to be targeted in each state.

Results: Tirenni et al. apply their model to a sample of 10,000 customers of Finnair using 2 years of their historical data. Figure-3 shows how the optimal policy differs from the historical policy for customers in state S_3 (states are defined based on recency, frequency, monetary value and statistical procedure). The optimal policy suggests sending no campaigns to about 60% of the customers in state S_3, compared to only 25% under historical policy. Figure-4 shows the expected long-term value from these customers based on historical and optimal policy. The optimal policy outperforms the short-sighted historical policy. Implementation of this value-based management at Finnair resulted in more than 20% reduction in marketing costs as well as improved response rates by up to 10%.

Conclusions

Marketing has been, and continues to be, a combination of art and science. With the increasing availability of data and sophistication in methods, it is now possible to more judiciously allocate marketing resources. In this chapter we discussed a two-stage process where a model of demand is estimated in stage-1 and its estimates are used as inputs in an optimization model in stage-2. We proposed a 3x3 matrix, with three different approaches for each of these two stages and discussed pros and cons of these methods. We also highlighted these methods with various applications.

What has been the impact of these advances? Scores of studies in this area now allow us to have empirical generalizations about the impact of marketing actions on sales and profits. For example, many studies have concluded that the average advertising elasticity is 0.1, and it is is almost twice as much for new products (Assmus, Farley and Lehmann 1984, Lodish et al. 1995). Similarly, based on a series of studies, Gupta and Zeithaml (2006) conclude that 1 point improvement in customer satisfaction can potentially lead to $240-275 million gain in firm value. These are important and powerful conclusions that are not based on a single study or a single product category. Instead these are generalizable results based on several studies, products and industries. This level of generalization builds confidence in our understanding of the impact of marketing actions on firm performance.

The impact of these studies goes beyond a theoretical understanding of the phenomena. In practical terms, we have witnessed significant impact at all levels of organization. Studies such as Steenburgh et al. (2003) and Jedidi et al. (1999) can help marketing managers in better allocation of their budget for a brand. Knott et al. (2002) use a field test to show that decisions based on their model of cross-selling produce an ROI of 530% for a bank, compared to -17% based on current practices of the bank. Thomas et al. (2004) show that when budgets are allocated as per their model of customer lifetime value, a pharmaceutical company should spend 30% more on marketing to improve its profits by over 35%, while a catalog retailer should cut its marketing spending by about 30% to gain profit improvements of 29%.

Harrah's Entertainment, Inc. provides perhaps the best example of the impact of this thinking on firm performance. Harrah's drove its entire business strategy based on marketing analytics by understanding and predicting customer behavior through database analysis and experimentation. Harrah's stock price has skyrocketed from under $16 in 1999 to over $88 in January 2008. Harrah's CEO, Gary Loveman, credits Harrah's enormous success to this relentless pursuit of perfection where decisions are based on models of consumer behavior rather than hunch or judgment.

Table-1: Demand Estimation and Economic Impact Analysis

		Demand Estimation		
		Decision Calculus	**Experiments**	**Econometric**
Economic Impact Analysis	**Descriptive**		Godes and Mayzlin (2007) Anderson and Simester (2004)	Wittink (2002)
	What-if			Jedidi, Mela and Gupta (1999)
	Optimization	Lodish (1971)		Steenburgh, Ainslie and Engebretson (2003) Tirenni et al. (2007)

Table-2: **Long-run Effects of Promotion Depth on New and Established Customers**

	Study A	Study B	Study C
Customers	Established	New	New
Sample Size			
Test	18,708	148,702	146,774
Control	35,758	148,703	97,847
Average % discount in promotion version	42	47	42
# of pages in catalog	72	8	16
# of products	86	16	36
# of prices varied	36	14	32
# of months of future data	28	24	22
Purchases from the Test catalog*			
% that purchased	158	185	174
Units ordered per customer	135	116	130
Average unit price ($)	63	65	71
Repeat purchases from future catalogs*			
Units ordered per customer	90	114	136
Average unit price ($)	89	96	90

*These measures are all indexed to 100 in the respective Control condition

Adapted from: Anderson, Eric and Duncan Simester (2004), "Long Run Effects of Promotion Depth on New Versus Established Customers: Three Field Studies," *Marketing Science*, 23(1), 4-20.

Table-3: Descriptive Statistics of the Pharmaceutical Data

Panel A: Brands with over $500MM in Revenue

	Launch Year		
	<1994	1994-1997	1998-2000
DET[a]	100%[b] / $2,758[c]	100% / $3,206	100% / $6,607
PME	100% / $427	100% / $698	100% / $1,917
JAD	92% / $129	100% / $245	100% / $532
DTC	67% / $605	67% / $1,224	78% / $2,482

Panel B: Brands with $100-$500MM in Revenue

	<1994	1994-1997	1998-2000
DET	100% / $732	100% / $1,098	95% / $1,711
PME	93% / $73	95% / $210	95% / $439
JAD	79% / $41	98% / $113	100% / $161
DTC	26% / $38	53% / $272	38% / $949

Panel C: Brands with $25-$100MM in Revenue

	<1994	1994-1997	1998-2000
DET	94% / $155	92% / $460	100% / $1,144
PME	85% / $13	89% / $56	86% / $180
JAD	73% / $8	92% / $36	100% / $148
DTC	10% / $13	24% / $45	21% / $5

a) DET=Physician detailing; PME=Phyisician meetings; JAD=Journal advertising; DTC=Direct-to-consumer advertising

b) Percent of products that spent on this marketing instrument in at least one month during the study period.

c) Average monthly expenditure in thousands of dollars.

Source: Wittink, Dick R. (2002) "Analysis of ROI for Pharmaceutical Promotion (ARPP)," white paper presentation to the Association of Medical Publications, 18 September 2002, available from http://www.vioworks.com/clients/amp.

Table-4: Return on Investment for Pharmaceutical Marketing

Panel A: Brands with over $500MM in Revenue

	Launch Year		
	<1994	1994-1997	1998-2000
DET	$3.1[a]	$5.9	$11.6
PME	$3.1	$6.0	$11.7
JAD	$3.1	$6.2	$12.2
DTC	$0.4	$0.7	$1.3

Panel B: Brands with $100-$500MM in Revenue

	<1994	1994-1997	1998-2000
DET	$1.2	$1.6	$2.1
PME	$2.0	$2.7	$3.6
JAD	$2.3	$3.1	$4.2
DTC	$0.1	$0.2	$0.2

Panel C: Brands with $25-$100MM in Revenue

	<1994	1994-1997	1998-2000
DET	$0.9	$1.0	$1.0
PME	$0.1	$0.1	$0.1
JAD	$6.2	$6.7	$7.2
DTC	$0.0	$0.0	$0.0

a) To be read as: $1 increase in detailing (DET) would generate incremental revenue of $3.1.

Source: Wittink, Dick R. (2002) "Analysis of ROI for Pharmaceutical Promotion (ARPP)," white paper presentation to the Association of Medical Publications, 18 September 2002, available from http://www.vioworks.com/clients/amp.

Table-5: Descriptive Statistics of the Data for Jedidi et al.

Variable	Brand	Mean for 1984-1987	Mean for 1988-1991
Market Share	1	0.35	0.36
	2	0.11	0.13
	3	0.16	0.13
	4	0.10	0.12
Purchase Quantity per Occasion (oz)	1	27.72	28.41
	2	26.20	28.58
	3	28.04	30.27
	4	28.60	29.42
Regular Price per Ounce ($)	1	0.051	0.054
	2	0.050	0.055
	3	0.052	0.056
	4	0.048	0.053
Promotion Frequency (% of occasions)	1	15.4	33.4
	2	8.7	32.6
	3	10.2	25.3
	4	6.4	29.8
Promotion Depth (% off)	1	11.3	17.9
	2	12.1	17.3
	3	13.8	16.8
	4	29.8	20.2
Advertising[1]	1	66.61	29.78
	2	25.52	17.55
	3	45.26	26.70
	4	28.98	12.25

[1]Advertising represents average inflation-adjusted advertising dollars in thousands spent in a quarter.

Source: Jedidi, Kamel, Carl F. Mela and Sunil Gupta (1999), "Managing Advertising and Promotion for Long-Run Profitability," *Marketing Science*, vol. 18, no. 1, 1-22.

Table-6: Long-term Impact of Changes in Promotion and Advertising on Profits

% Change in Profit with 5% Increase in	*Brand 1*	*Brand 2*	*Brand 3*	*Brand 4*
Advertising	-0.02	0.13	-0.69*	-0.32*
Promotion Frequency	-0.33*	-0.33*	-0.33*	-0.31*
Promotion Depth	-0.38*	-0.49*	-0.31*	-0.31*

* Significant at 0.05 level

Adapted from: Jedidi, Kamel, Carl F. Mela and Sunil Gupta (1999), "Managing Advertising and Promotion for Long-Run Profitability," *Marketing Science*, vol. 18, no. 1, 1-22.

Figure-1: Estimating Demand from the Number of Calls using Decision Calculus

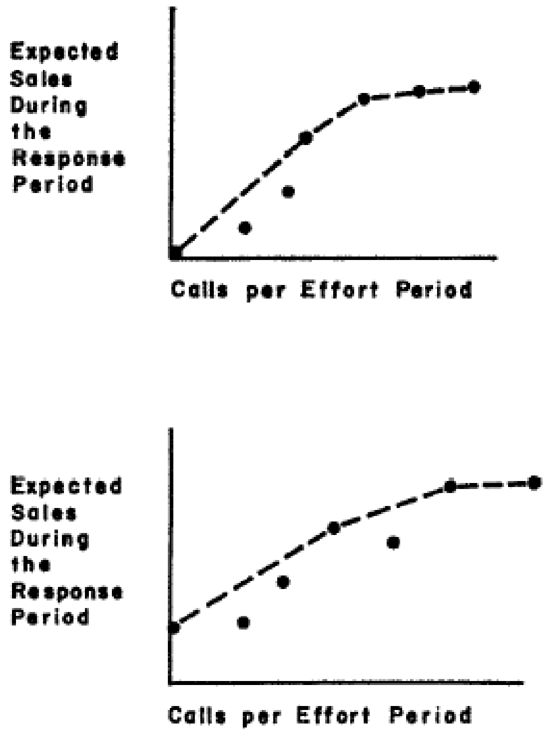

Source: Lodish, Leonard M. (1971), "CALLPLAN: An Interactive Salesman's Call Planning System," *Management Science*, vol. 18, no. 4, part II (December), pp. 25-40.

Figure-2: ROC Chart Comparing the HBVC and Null Models

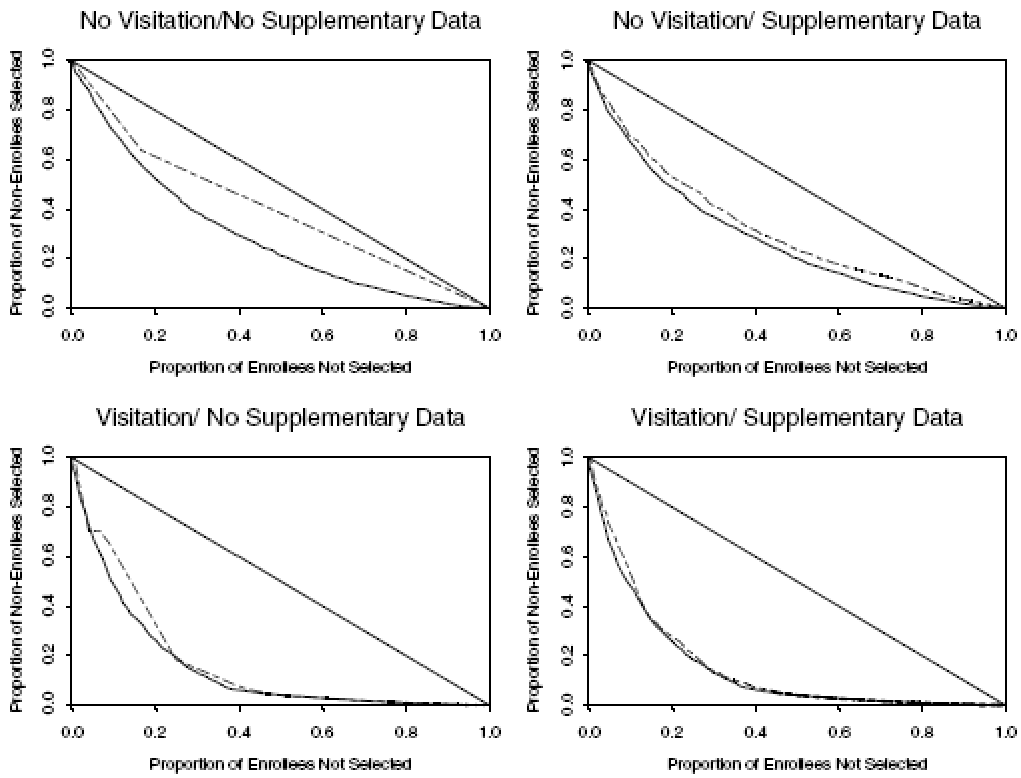

Note: These four ROC charts compare the HBVC model to the null model assuming four different sets of information. In each chart, the HBVC model is graphed with the solid line, and the null model is graphed with the dotted line. Models that better order prospects move toward the lower, left-hand corner of the chart.

Source: Steenburgh, Thomas J., Andrew Ainslie, and Peder Hans Engebretson (2003). "Massively Categorical Variables: Revealing the Information in Zip Codes," *Marketing Science*, vol. 22, no. 1 (Winter), pp. 40–57.

Figure-3: Historical and Optimal Marketing Resource Allocation for Customers in State S_3

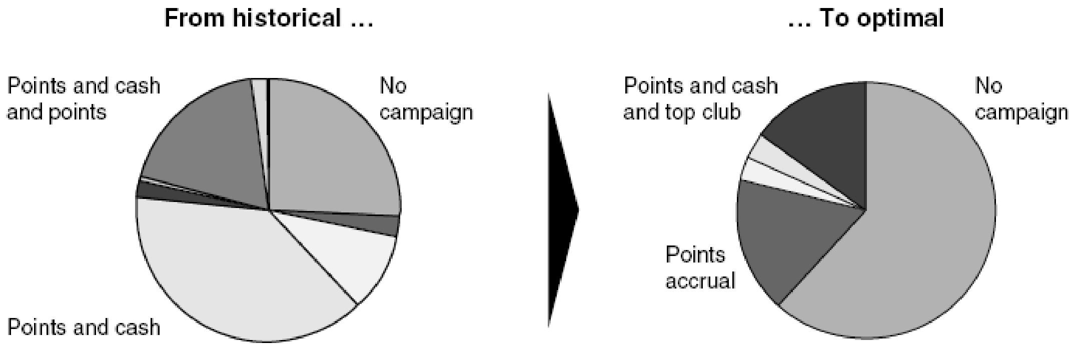

Source: Tirenni, Giuliano, Abderrahim Labbi, Cesar Berrospi, Andre Elisseeff, Timir Bhose, Kari Pauro, Seppo Poyhonen (2007), "Customer Equity and Lifetime Management (CELM) Finnair Case Study," *Marketing Science*, vol. 26, no. 4, July-August, 553-565.

Figure-4: Expected Long-Term Value Using the Historical and Optimal Policy

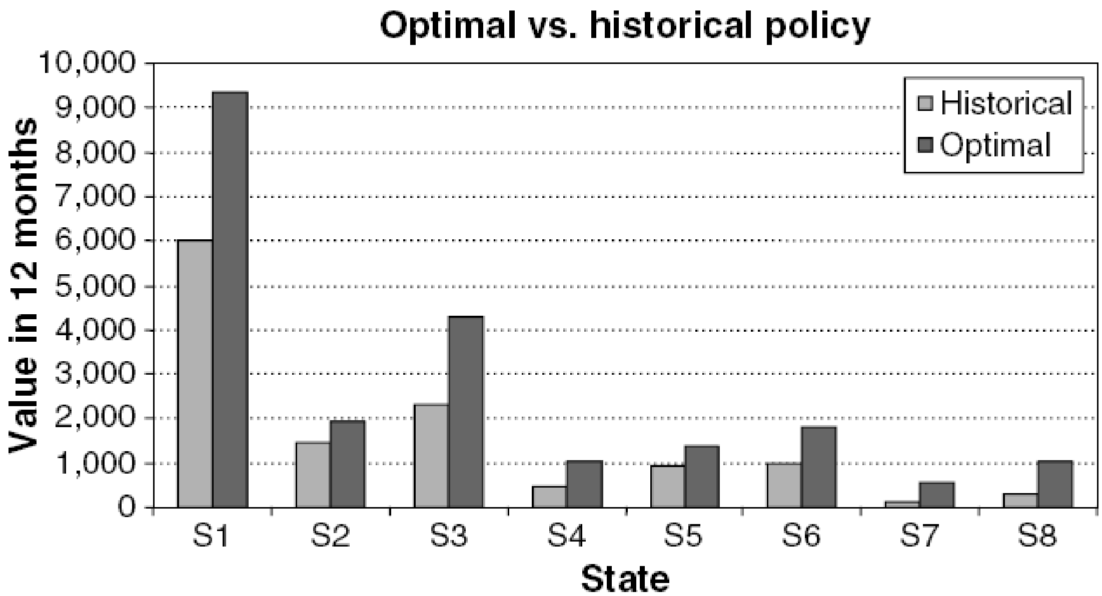

Source: Tirenni, Giuliano, Abderrahim Labbi, Cesar Berrospi, Andre Elisseeff, Timir Bhose, Kari Pauro, Seppo Poyhonen (2007), "Customer Equity and Lifetime Management (CELM) Finnair Case Study," *Marketing Science*, vol. 26, no. 4, July-August, 553-565.

References

Advertising Age (2007), 100 Leading National Advertisers, (June 25).

Anderson, Eric and Duncan Simester (2004), "Long Run Effects of Promotion Depth on New Versus Established Customers: Three Field Studies," *Marketing Science*, 23 (1), 4-20.

Assmus, Gert, John W. Farley, and Donald R. Lehmann (1984), "How Advertising Affects Sales: Meta-Analysis of Econometric Results," *Journal of Marketing Research*, 21 (February), 65-74.

Blattberg, Robert C. and Stephen J. Hoch (1990), "Database Models and Managerial Intuition, 50% Model + 50% Manager," *Management Science*, 36 (8), 887-899.

Bughin, Jacques, Christoph Erbenich, and Amy Shenkan (2007), "How Companies are Marketing Online: A McKinsey Global Survey," *The McKinsey Quarterly*, September.

Divakar, Suresh, Brian T. Ratchford, and Venkatesh Shankar (2005), "CHAN4CAST: A Multichannel Multiregion Forecasting Model for Consumer Packaged Goods," *Marketing Science*, 24 (3), 333-350.

Duvvuri, Sri Devi, Asim Ansari and Sunil Gupta (2007), "Consumers' Price Sensitivity Across Complementary Categories," *Management Science*, 53 (12), 1933-1945.

Eisenstein, Eric M. and Leonard M. Lodish (2002), "Precisely Worthwhile or Vaguely Worthless: Are Marketing Decision Support and Intelligent Systems 'Worth It'?" in *Handbook of Marketing*, Barton Weitz and Robin Wensley, eds. Sage Publications, London.

Fudge, William K. and Leonard M. Lodish (1977), "Evaluation of the Effectiveness of a Model Based Salesman's Planning System by Field Experimentation," *Interfaces*, 8 (1), part 2, 97-106.

Godes, David and Dina Mayzlin (2007), "Firm-Created Word-of-Mouth Communication: Evidence from a Field Test," *Working Paper*, Harvard Business School.

Guadagni, Peter M. and John D.C. Little (1983), "A Logit Model of Brand Choice Calibrated on Scanner Data," *Marketing Science*, 2 (3), 203-208.

Gupta, Sunil (1988), "Impact of Sales Promotions on When, What, and How Much to Buy," *Journal of Marketing Research*, 25 (4), 342-356.

-----------, Dominique Hanssens, Bruce Hardie, Wiliam Kahn, V. Kumar, Nathaniel Lin, Nalini Ravishanker, and S. Sriram (2006), "Modeling Customer Lifetime Value." *Journal of Service Research*, 9 (2), 139-155.

-----------, and Valarie Zeithaml (2006), Customer Metrics and Their Impact on Financial Performance," *Marketing Science*, 25 (6), 718-739.

Jedidi, Kamel, Carl F. Mela and Sunil Gupta (1999), "Managing Advertising and Promotion for Long-Run Profitability," *Marketing Science*, 18 (1), 1-22.

Knott, Aaron, Andrew Hayes and Scott Neslin (2002), "Next-Product-to-Buy Model for Cross-Selling Applications," *Journal of Interactive Marketing*, 16 (3), 59-75.

Koen, P., S. Siddarth, and D.M. Hanssens (2002), "The Long-Term Effects of Pricing and Promotions on Category Traffic, Brand Choice and Purchase Quantity," *Journal of Marketing Research*, 39 (November), 421-439.

Kumar, V., J. Andrew Petersen, and Robert P. Leone (2007), "How Valuable is Word of Mouth?" *Harvard Business Review*, October, 139-146.

Little, John D.C. (1970), "Models and Managers: The Concept of Decision Calculus," *Management Science*, 16 (8), B466-B485.

---------- and Leonard M. Lodish (1969), "A Media Planning Calculus," *Operations Research*, 17 (1), 1-35.

-------- (1975a), "BRANDAID: A Marketing-Mix Model, Part 1: Structure," *Operations Research*, 23 (4), 628-655.

-------- (1975b) **"**BRANDAID: A Marketing-Mix Model, Part 2: Implementation, Calibration, and Case Study," *Operations Research*, 23 (4), 656-673.

-------- (2004), "Comments on 'Models and Managers: The Concept of Decision Calculus," *Management Science*, 50 (12), 1854-1860.

Lilien, Gary L., Philip Kotler, and K. Sridhar Moorthy (1992), *Marketing Models*. Englewood Cliffs, NJ: Prentice Hall.

Lodish, Leonard M. (1971), "CALLPLAN: An Interactive Salesman's Call Planning System," *Management Science*, 18 (4), part II, 25-40.

--------, Magid Abraham, Stuart Kalmenson, Jeanne Livelsberger, Beth Lubetken, Bruce Richardson, and Mary Ellen Stevens (1995), "How T.V. Advertising Works: A Meta-Analysis of 389 Real World Split Cable T.V. Advertising Experiments," *Journal of Marketing Research*, 32 (2), 125-139.

Loveman, Gary (2003), “Diamonds in the Data Mine,” *Harvard Business Review,* May, 2-6.

Manchanda, Puneet, Asim Ansari, and Sunil Gupta (1999), “The ‘Shopping Basket’: A Model for Multi-Category Purchase Incidence Decisions,” *Marketing Science*, 18 (2), 95-114.

Mantrala, Murali (2006), “Allocating Marketing Resources,” in *The Handbook of Marketing,* Barton A. Weitz and Robin Wensley, eds. Sage Publications.

Mela, Carl F., Sunil Gupta, and Donald R. Lehmann (1997), “The Long-Term Impact of Promotions and Advertising on Consumer Brand Choice,” *Journal of Marketing Research*, 34 (2), 248-261.

Natter, Martin, Thomas Reutterer, Andreas Mild, and Alfred Taudes (2007), “An Assortmentwide Decision-Support System for Dynamic Pricing and Promotion Planning in DIY Retailing,” *Marketing Science*, 26 (4), 576-583.

Sinha, Prabhakant and Andris A. Zoltner (2001), “Sales-Force Decision Models: Insights from 25 years of Implementation,” *Interfaces*, 31 (3), Part 2 of 2, S8-S44.

Steenburgh, Thomas J., Andrew Ainslie, and Peder Hans Engebretson (2003). “Massively Categorical Variables: Revealing the Information in Zip Codes,” *Marketing Science*, 22 (1), 40–57.

-------- (2007). “Measuring Consumer and Competitive Impact with Elasticity Decompositions,” *Journal of Marketing Research*, 44 (4), 636-646.

Thomas, Jacquelyn, Werner Reinartz and V. Kumar (2004), “Getting the Most Out of Your Customers,” *Harvard Business Review*, 82 (July-August), 116-123.

Thomke, Stefan (2003) "R&D Comes to Services: Bank of America's Pathbreaking Experiments," *Harvard Business Review*, 81 (4), 70-79.

Tirenni, Giuliano, Abderrahim Labbi, Cesar Berrospi, Andre Elisseeff, Timir Bhose, Kari Pauro, and Seppo Poyhonen (2007), “Customer Equity and Lifetime Management (CELM) Finnair Case Study,” *Marketing Science*, 26 (4), 553-565.

Venkatesan, R. and V.Kumar (2004), “A Customer Lifetime Value Framework for Customer Selections and Resource Allocation Strategy”, *Journal of Marketing*, 68 (4), 106-125.

Wierenga, Berend, Gerrit H. Van Bruggen, and Richard Staelin (1999), “The Success of Marketing Management Support Systems,” *Marketing Science,* 18 (3), 196-207.

Wittink, Dick R. (2002) "Analysis of ROI for Pharmaceutical Promotion (ARPP)," white paper presentation to the Association of Medical Publications, 18 September 2002, available from http://www.vioworks.com/clients/amp .