
Original Article

Enhanced customer targeting with multi-stage models: Predicting customer sales and profit in the retail industry

Received (in revised form): 6th October 2009

Michael McCrary

is Modeling Manager for Limited Brands, Inc. whose brands include Victoria's Secret, Bath & Body Works, C.O. Bigelow, White Barn Candle Company, Henri Bendel and LaSenza. He is responsible for building statistical models and providing customer insights and analytics for the retail store channel. He has over 10 years experience in marketing for banking, insurance and retail services corporations, and holds a PhD in Sociology from The Ohio State University, Columbus, Ohio (USA).

ABSTRACT Drawing on the marketing, finance and retail research literature, this article uses a multi-stage targeting methodology to enhance the targeting of high-revenue customers. Such an investigation is important for both research and practice, as many retail-marketing departments use traditional response models to identify the optimal customers for marketing communications, despite the fact that maximum *revenue* is the financial objective for corporations rather than *response*. A methodology is deployed to build multi-stage models that target the most responsive, highest-spending and highest-profit customers using customer- and product-level profit data from a specialty retailer in the United States. Comparisons of the competing models show that multi-stage models identify customers with higher actual sales and profit than both traditional response models and other forms of multi-stage models in which any two combinations of customer response, spend or profit are predicted. Implications and limitations of this research are discussed in light of the results.

Journal of Targeting, Measurement and Analysis for Marketing (2009) 17, 273–295. doi:10.1057/jt.2009.22

Keywords: consumer profitability; response models; profit models; targeting techniques; multi-stage models; customer analysis

PURPOSE

This article draws insights about consumer sales and profit from the literature in the fields of marketing, finance, retail and financial services, which are used to enhance and extend the traditional targeting models deployed in marketing departments. Specifically, this article

uses two retail consumer data sets to build and validate *multi-stage* targeting models, which target more than simply customer response. The most effective of these multi-stage models uses product-specific *profit* information to predict an individual customer's future sales and profits. To preview the results, 'two-stage' models such as those that predict customer response and spend generated up to 12 per cent more sales than response models alone. 'Three-stage' models that predict customer response, spend and profit targeted customers with 17 per cent more profit

Correspondence: Michael McCrary
Limited Brands Direct Marketing, Inc., 8655 East Broad Street,
Reynoldsburg, Ohio 43068-9715, USA
E-mail: MMcCrary@limitedbrands.com

than response models alone. When comparing the three-stage model to the next-best performing two-stage model (that is, the model that predicted response and profit), the former was able to target as much as 3 per cent more profit per customer mailed. Although this appears modest, a 3 per cent profit increase for a large customer base translated into US\$264 000 more in profits in a 4-week response period, an amount that is even larger when considering the total number of marketing campaigns in a given year. Although the financial benefits of multi-stage models may be evident in the results previewed above, there are several additional reasons why consumer sales and profit research is important, and several ways that it can be used to enhance the effectiveness of traditional targeting methodology.

First, researchers are aware that marketing efforts impact the bottom line of their businesses, whether through direct mail promotions that drive store traffic and subsequent profits,¹ or through Customer Relationship Management (CRM) and its role in generating customer and therefore corporate profit.^{2,3} Because marketing drives sales and profits, professionals in this field have a special interest in optimizing their efforts, of which customer-targeting methodology is a major component, particularly for so-called 'non-contractual' retail relationships that must be constantly renewed. This article addresses the need to optimize one component of marketing tools by extending traditional response models to multi-stage models. Specifically, it enhances response models by targeting combinations of customer response, sales and profit.

Second, according to a review of 7 years of retail research literature, quantitative studies of consumer behavior are under-represented compared to qualitative research.⁴ This article on multi-stage modeling addresses the quantitative gap in the literature by building and validating statistical models of customer behavior using data from a prominent store-based specialty retailer based in the United States. Two large customer data sets are used; each data set contained different kinds of data and products. Models were built for the two data sets, and therefore the modeling results are presented separately

(hereafter referred to as the *first* and *second* analysis for the sake of brevity). The first analysis enhances traditional ('single-stage') response models by modeling customer response and customer spend. The second analysis replicates the first analysis, and also extends the logic of two-stage models to develop a three-stage model to predict customer response, spend and *profit*. This approach strengthens the findings by replicating the methodology across customers who buy two distinct products, and by building models on two types of data (campaign-response data and a 'natural' random customer sample). In sum, this article extends traditional customer response model methodology to build multi-stage models, and measures the impact of multi-stage targeting models with regard to customer sales and profit.

Third, although it is vital for predicting customer profit, marketers and statisticians are often hampered by the lack of profit data at the product (item) level.⁵ Such data unavailability forces researchers to apply the same profit rate to all customers or to assign an average profit rate to each customer in a segment.⁶ Although effective for some research questions, averages that are assigned to customers are unlikely to produce the most predictive models at the customer level because some customers buy higher-profit products while most do not. In contrast to research using averages, this article provides an opportunity to use profit data specific to each unique item that the customer has purchased. Ultimately, more accurate data should build more effective models, which in turn will drive CRM, sales and profits more generally.⁷

Finally, marketers and statisticians constantly seek to improve customer-targeting methodologies. Historically, marketers have built response models to target the optimal customers (that is, most responsive customers) for future marketing communications.^{8,9} Certainly there are several benefits of traditional response models, a topic discussed in depth during the final section of this article. However, traditional response models are insufficient to target the highest-spending or most profitable customers. In fact, response models can potentially target the *most* responsive

customers who actually spend the *least*, especially when promotional offers involve free items or when there is no purchase requirement (for the reader's information, the 'response' data in this research did not reflect such a scenario). To avoid the unnecessary marketing costs associated with targeting lower-spending and less profitable customers, statisticians in the financial service industries have enhanced response models by extending the models to predict customer spend as well as customer response. Although a more detailed discussion of this topic is found in the implications section, it is important to briefly mention that within retail businesses, this advancement has been substantially slower to materialize. Thus, this article is an important bridge between the methodologies used in the two industries, bringing successful methodology from the financial services industry into the retail sector. To build such a bridge, this article demonstrates one kind of methodology for multi-stage models with concern for maximum understanding by modelers and maximum application in marketing organizations. The need to increase profitability is not a minor issue; firms that fail to do so may find themselves irrelevant to consumers regardless of economic conditions, but especially during anemic environments.

The remaining sections of this article are organized in the following way. First, the finance and marketing research literature on the topic of customer sales and profit is explored across various industries. Second, data sources and methodology are reviewed to fit and validate individual-level customer models. These models predict combinations of customer response, sales and profit at the individual customer level. There are two distinct analyses. The first analysis uses campaign data to build a traditional (one-stage) response and a two-stage (response-spend) model. The second analysis uses a random sample of customers to replicate the methodology of the first analysis and extend the two-stage models to a three-stage model involving response, spend and profit. Third, the results section compares the models, and explains the findings. The last section discusses the findings of the research, limitations and suggestions for future analysis.

THE CONSUMER SALES AND PROFIT LITERATURE

Within the fields of marketing and finance, the study of consumer sales and profitability has spanned three lines of research: (I) descriptive and analytic research on *definitions of profit, their components and their measurement*, (II) research on the *consequences of sales and profit* and (III) research that predicts customer sales and profit (*determinants*). The next section briefly describes three areas of consumer profitability research, and their contributions and limitations. Although it focuses primarily on retail, it also incorporates banking, insurance, telecommunication and other industries.

Definitions, components, measurement of profit

Researchers in the field of finance spend much time describing and measuring customer sales and profit. This research has been conducted at the aggregate, non-customer level (for example, transactional or product level analyses),^{10,11} in an aggregated (total) customer population, or within segments of customers.^{12,13} Within this line of research, the components of sales and profit have been explored in great detail. For example, Ryals reviewed various kinds of revenue and costs for the banking and insurance industries (acquisition costs, channel costs, loan costs and other costs).¹⁴ In an analysis of the effects of store profits, it was discovered that transactional costs (for example, discounts, rebates and other costs) can reduce retailers' store profits (that is their 'pocket price') by an additional 14–39 per cent.¹⁵ Other research has reviewed various components involved in calculating profit, including product costs, capital costs, order-related marketing costs, indirect order related marketing costs and indirect customer costs.¹⁶ Finally, research from this perspective has also generated and examined concepts that are used to track customer sales and profits over time, such as Lifetime Customer Value (CLV), Net Present Value, Internal Rate of Return, Past Customer Value, Previous Period Customer Revenue and Activity Based Costing.^{6,17–21} Together, research that defines and explores detailed concepts related to profit components has

allowed us to better measure consumer profit. Research on profit components has also been vital, as it is a necessary first step before an optimal customer prioritization method can be deployed.²² However, for those responsible for executing customer-based marketing efforts, there is a need for more predictive targeting tools at the customer level,^{17,23} (see Venkatesan and Kumar),²¹ especially as marketers typically have purchase data unique to each customer. Simply put, when making customer-level decisions, assuming the same profit rate for all customers^{6,21} is likely to result in less accurate models than those with actual profit data for each customer.²⁴⁻²⁶

Consequences of sales and profit

Another focus of marketing research has dealt with the consequences of customer sales and profit, such as how customer sales and profits are related to brand loyalty and customer satisfaction^{16,27} (for a review, see Gupta and Zeithaml).²⁸ In the literature, researchers disagree as to whether profits and sales *cause* customer loyalty and satisfaction, or whether they are a *consequence* of revenue.³ Regardless of causal relationship, this literature has reminded us that customer revenue is not the only customer behavior to explain and predict. Another line of research explores the dynamics of customer retention, which follows logically from customer sales and profit at a firm. Specifically, customer 'churn' has been studied in retail sector firms,^{13,29} the telecommunications sector,^{30,31} the banking industry³²⁻³⁴ and airlines.¹³ Finally, on a topic related to customer loyalty and satisfaction, researchers who advocate for CRM argue that CRM drives customer sales and profits precisely by enhancing customer loyalty, duration and satisfaction. CRM is said to provide more relevant communication between the organization and the customer, which is optimally timed to each customer's unique purchasing needs and lifestyle. Because it delivers the 'right message to the right customer at the right time', CRM drives sales and profits.^{7,14,35} Although this article on multi-stage models does not directly involve the topic of CRM, it does involve the

enhancement of predictive models as a major CRM targeting tool (for discussions on the appropriateness, benefits and evolution of CRM, see Homburg *et al.*,²² Venkatesan *et al.*³⁶ and Berry³⁷). Altogether, research on the consequences of sales and profits has contributed to an emphasis on the total customer relationship, rather than focusing narrowly only on customers' financial outcomes. Moreover, this line of research reflects a more complex view of the way in which customer targeting is driven by prior targeting decisions.²² However, like research on the definitions of profit, studies on the consequences of sales or profit tell us less about *which* customers to target for the highest sales and profits. Rather, their purpose is to examine the reasons or causal mechanisms behind the relationship.

Determinants of customer sales and profit

Finally, finance and marketing researchers have also explored various causes or determinants of sales and profit (for a comprehensive review, see Gupta and Zeithaml).²⁸ In brief, selected studies have explored the costs of capital in determining profitability,¹¹ how credit costs determine profit in the fishing industry,²⁰ and the role of internal organizational processes and procedures within banks.³⁸ Other research studies on the determinants of customer sales and profit have explored attitudinal measures such as customer satisfaction because more satisfied customers are argued to spend more than unhappy customers;^{16,28,31,39} the impact of customer tenure for sales and profits in the retail industry,⁴⁰ the effects of past customer transactions in predicting future transactions for the airline, retail and apparel industries;¹³ attitudes regarding product perception on revenue,⁴¹ retail store traffic and direct mail promotions, both of which are positively correlated with profit and sales;¹ and the number of product categories from which customers have purchased.⁴² In the latter study, the authors found that the number of product categories from which a customer purchases was related to both revenue and CLV.

One limitation of these studies is that they typically have assumed relatively simple causality in which past customer behavior has determined future sales and profits. Several studies challenge this view, and explore the more complex relationship between past and future customer behavior.^{3,22} For example, Homburg *et al* analyzed data on 310 business-to-business and business-to-consumer firms to explore the complex determinants of customer profit. Ostensibly, marketers assume that a customer's past profit determines customer satisfaction and loyalty, which then affects his or her share of wallet with the firm. Share of wallet, in turn, affects sales and profits for the top tier of customers. However, rather than ending here, they argued that a customer's current profit affected future profits because marketers often target customers with the highest profits.²² In short, customer targeting does not occur in a vacuum or independent of past behavior. We often neglect the way in which future targeting decisions are related to past targeting decisions. Likewise, using survey data across industries, one study found that the relationships among increased purchases, tenure and satisfaction are not necessarily related to increased customer value.³ Despite a simplified causal mechanism, research that assumes a link between purchase behavior and future revenue is nevertheless predictive, and contributes a wealth of candidate variables useful when building sales- and profit-targeting models at the customer level. Using a data-mining approach, as many of these concepts as possible are offered to the multi-stage models as candidate variables.

METHODOLOGY

Data sources

Data come from a multi-channel, Fortune 500 retailer in the United States that has historically specialized in three product lines: intimate apparel (including sleepwear, swimwear and accessories), personal care items (for example, lotions and home fragrance products) and fashion clothing (men's and women's). The analyses presented here used data from customers who shopped

the store channel during a 4-week prediction (that is outcome) period. Predictor variables spanned the maximum 3 years of customer transactional history before the outcome (that is response) period. More than 500 candidate predictors included a customer's transactional, individual and household data in an attempt to draw from the literature on the drivers of customer sales and profit. A partial list includes past shopping behavior, payment information, products purchased, seasonality or time period shopped, and household demographics such as marital status, age and home value. Although customer satisfaction has been hypothesized to be a significant predictor of customer sales and profit^{28,39} (but see Singh *et al*),⁴³ along with membership in a loyalty program,³² customer satisfaction data were not directly available, and the corporation does not have a traditional 'loyalty' program. However, because of exclusive discounts and offers for credit card customers, this firm often considers the use of its proprietary credit card an indirect indicator of customer satisfaction and loyalty. The logic is that customers have many choices of credit cards to use, but continue to choose the firm's own credit card. This suggests some level of satisfaction and brand loyalty. Therefore, the multi-stage models offered up predictors related to credit card usage as an indirect indicator of loyalty and satisfaction when customers had a proprietary credit card. Individual-level predictive research has used various types of appended data (census-block, characteristics of an individual's occupation or industry, store-distance, survey data on customer satisfaction, share of wallet data from competitors and other exographic data).⁴⁴⁻⁴⁸ For purposes of illustration of the concepts in this article and for maximum generalization of the conclusions, only demographics available on the Enterprise Data Mart were used. In the following subsections, details regarding the data are specified separately because there were two data sets analyzed.

Data for the customer response-spend models

For the first analysis in which customer response and response-spend models were built, random

samples of mailed customers were aggregated from several direct mail campaigns. The total customer sample utilized for model-building and final model selection was 67 385. There were an additional 67 645 customers randomly selected from a campaign that was not used to build the model (that is, a 'hold-out' or validation campaign). One benefit of using campaign data is that we also have available a random sample of the hold-out campaign that did not receive any treatment (a control group), which allowed us to assess the impact of marketing on customers more generally (that is, to measure incremental marketing impact). All customers were 'active' purchasers of a particular specialty product line distinct from the product line purchased by customers in the second data set described below. (Owing to confidentiality, the product specialties cannot be identified.) For each customer in the random modeling sample, predictor variables at the transactional and individual customer level were extracted, aggregated and appended to each customer's record from the Enterprise Data Mart. The first dependent variable was response to the campaign (yes, no), for which responders with spend greater than \$0 were defined as 'responders'. The second dependent variable was the amount of each customer's sales during the 4-week campaign response periods for those customers who responded. Thirty thousand and three hundred customers had sales greater than \$0 in the modeling sample. Returned merchandise was not subtracted from gross customer sales because total product returns were a nominal amount of gross sales for these campaigns.

Data for the extension-to-profit models

The corporation did not have customer-level *profit* data for the specialty customers described in the first data set. However, they were available for a set of customers considered 'active' buyers of another distinct specialty product line. Although these customer data had the advantage of containing profit information at the transactional level for each customer, there were not sufficient random samples to utilize response data from direct mail campaigns. Therefore, a separate data

set was constructed using a random sample of 70 591 customers extracted from the Enterprise Data Mart. Although they were completely random, these are sometimes referred to as 'natural' data, as they are not associated with a marketing campaign *per se*. When building spend and profit models, the second analysis used the same definitions as the previous analysis, including the definition of eligible customers, the same outcome measurement period and the definition of the dependent variables. This second data set contained a random sample of 33 491 customers with spend and profit greater than \$0. In addition to the same candidate predictors used for the first analysis, the second analysis for the profit extension included independent variables related to a customer's profit margin in the 3 years before the outcome period. The definition of customer profit and the rationale for including it as a candidate predictor are discussed below in more detail.

Definition of customer profit

This research used a simple definition of customer profit, which admittedly did not address the complexities of customer profit that were cited as a strength of the finance literature regarding profit, CLV and other related concepts. However, of utmost concern for this article was the ability to communicate the multi-stage methodology and facilitate its widest possible deployment across marketing organizations. Therefore, a simple definition of customer-level profit was used, in which profit was defined as the sum of each customer's sales minus the sum of the 'cost of goods' for all items purchased by a customer. The corporate finance department determined the cost of goods for each item sold in the store, and this information was loaded to the Enterprise Data Mart for each product at the lowest possible level (that is item). An item-level analysis of the data was needed because products and customers vary in terms of profitability.¹ Indeed, in one study, 20 per cent of products generated 50 per cent of the sales and 150 per cent of the profit.¹⁷ Neither are all customers equal when it comes to what they buy and their revenue. For example, certain customers traditionally

purchase low-margin products, others buy high-margin products, and still others purchase a mix of both high- and low-margin items. To illustrate how two customers can have the same sales but very different profit, consider the following example. One customer purchased four lower-margin, lower-priced body lotions, whereas a second customer purchased a single higher-margin, higher-priced anti-aging facial care product. The two customers each spent \$80. However, the first customer purchased eight bottles of lower-margin lotion for \$10 each with a cost of goods at \$4 for each bottle, yielding a total spend of \$80 ($\$10 \text{ spend per item} \times 8 \text{ items} = \80 spend) and \$32 in the total cost of goods ($\$4 \text{ cost of goods per item} \times 8 \text{ items} = \32). Thus, this customer had a margin of \$48 ($\$80 \text{ spent} - \$32 \text{ cost of goods} = \48). In contrast, the second customer, who also spent \$80, purchased one anti-aging product with a \$20 cost of goods, and had a total profit margin of \$60 ($\$80 \text{ spend} - \$20 \text{ cost of goods} = \60 profit margin). These two customers had the same sales but significantly different profit margins based on the products they purchased: the body lotion buyer had \$48 in profit margin, whereas the anti-aging buyer had \$60 in profit. For this reason, customer profit in retail sales is best modeled according to the items that are purchased, rather than multiplying the same profit rate by each customer's sales.

The same definition of profit (customer sales minus the cost of goods) was used to construct independent variables for which prior customer profit was a candidate predictor for future profit. In addition, the profit models also included candidate variables that measured the *change* in profit for each customer in the 3-year prediction period, because customers who reduce their profit over the 3-year prediction period may generate less profit in a future outcome period (as well as the reverse). Finally, for the second analysis, both profit *dollars* and profit margin *percentages* were candidate independent variables for the predictive models. For example, a customer with \$100 spend and \$75 cost of goods had a profit margin in *dollars* of \$25 and a profit margin *percentage* of 0.25 ($\$25 \text{ margin} / \$100 \text{ spend} = 25 \text{ per cent margin percentage}$).⁴⁹

Modeling methodology

Like the logic used to define profit, a multi-stage methodology was adopted that could be well understood by the widest possible marketing audience (both statisticians and other marketers). Thus, the methodology used in this article followed common industry standards and statistical procedures available to nearly all analysts. For example, this article used the commonly practiced data-mining procedure of *sample partitioning*,⁵⁰ although more advanced data partitioning and sampling techniques are available using bootstrapping and Monte Carlo simulations, particularly for smaller data sets (for examples of these techniques in the profitability literature, see Venkatesan and Kumar,²¹ Sismeiro and Bucklin²⁵ and Lemmens and Croux³⁰). The development data sets for all models were randomly partitioned, and models were built and selected on one portion of the data and evaluated on other portions. Finally, oversamples of responders were taken because the data contained a low response rate, which would yield less robust models if not addressed.^{51,52}

Data-mining professionals have used a variety of advanced multivariate statistical algorithms when analyzing sales and profit data. These include artificial neural networks,^{29,53} decision trees,²⁹ structural equations/path analysis,^{1,22} traditional forecasting,¹³ Markov models^{26,54} and regression models.^{21,29,40,41,54,55} Another technique is latent class analysis,⁵⁶⁻⁶⁰ which is often deployed when unobserved attitudinal measures are prominent in the analysis. A subtype of latest-class models extends the technique to a regression model, which are especially useful when analysts have data sets with repeated measures and/or want an advanced alternative to a k-means clustering algorithm.⁵⁶⁻⁶⁰ Profit-modeling problems have also been considered from a Bayesian perspective, in which the timing of purchase is considered a key outcome to be modeled,^{36,61} as well as from a risk-hazards perspective in terms of profit and attrition.¹⁰ Finally, a modeling methodology that is related to the technique used in this article is worthy of mention. A study of online shopping used a series of binary probit models to predict three online

customer behaviors that followed a natural sequence. Each dichotomous dependent variable was modeled as conditional on the previous dependent variable. For example, the authors modeled the probability of completing personal information among those who first successfully configured the product. By modeling these predicted behaviors conditionally rather than using a single model to identify the ultimate outcome (placing a deposit on a new vehicle), the authors were able to realize a 15–22 percentage point improvement in the cumulative proportion of potential customers.²⁵ Such a ‘staged’ or conditional methodology is a powerful modeling approach that is related to the multi-stage models in this article.

Following the broad communication objective discussed regarding the definition of the profit and the data sampling technique, the models built in this article used modeling techniques accessible via standard data-mining software. Specifically, competing response, spend and profit models were built using decision trees (CART/CHAID variants), neural networks and multivariate regression (logistic or ordinary least squares (OLS) as appropriate).⁶² The final models presented here were logistic and OLS regression models, which were both the best-performing and most stable algorithms in the validation data.

Response models

First, to measure the performance of a traditional approach to customer targeting, response models were built to predict the probability that a customer would respond to a direct mail campaign or purchase the product of interest. Response was defined as customer sales greater than \$0 during the promotional period. The response model can be expressed by the following equation, which is the standard logistic regression expression in statistics literature:^{63–65}

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_i X_i + \varepsilon \quad (1)$$

Where Y = response to campaign (1 = yes, 0 = no), β_0 = intercept, β_1 = slope for independent variable X_1 , X_1 = independent variable₁, β_2 = slope for

independent variable X_2 , X_2 = independent variable₂, β_i = slope for independent variable X_i , X_i = independent variable_i and ε = error term. The response model was fit using a stepwise procedure, and a ‘best’ model was selected using Schwarz’s Bayesian Criterion (SBC) statistic.

For the first analysis involving the *response to campaign data*, the final *response* model included the following predictors: customer’s tenure, number of departments the customer shopped, customer’s use of house (proprietary) charge card, customer’s number of transactions, and two demographic variables: numbers of days since customers had shopped with other retailers (recency) and the vehicle composition of the customer’s neighborhood. For the second analysis involving the *random customer sample*, the final predictors in the *response* model included the customer’s number of trips, customer’s number of products purchased, customer’s average spend per trip, number of transactions customer made in past 3 months, customer’s total spend and customer’s recency, as well as indicator variables for one-trip shoppers and customers who previously shopped during sale periods.

For each analysis, customers in the validation campaign/data set were scored using the response models (equation (1)), and then ranked by the value of \hat{Y} , the predicted response rate. Performance metrics for actual spend are shown by decile, which are 10 per cent breaks in the hold-out data set. Recall that the first analysis shows actual sales statistics using a hold-out *campaign* from another time period not used to build the models and a control (no treatment) data set that was also available. The second analysis used a random customer *sample*, as adequate random samples were unavailable from direct mail campaigns, and because profit was not available for the first set of customers.

Response-spend models

Independently from the response models, statistical models were built that predicted how much a customer would spend in a 4-week future period if they responded to the campaign (equation (2)). These are also referred to as spend

Table 1: Results of swap set analyses, comparing unique customer deciles from two methodologies (Sales>\$0 and Sales>=\$0) where sales and profit per circulation reflect actual values for outcome period

Two methodologies applied to response-spend models (two-stage)	Deciles 1–4		
	Response campaign	Customer sample	
	Sales/circ	Sales/circ	
Sales>\$0	\$5.14	\$4.92	
Sales>=\$0 (all customers)	\$4.63	\$4.07	
Difference	\$0.51	\$0.85	
Confidence level	95%	99%	

Two methodologies applied to response-spend-profit models (3 stage)	Deciles 1–4		Deciles 1–2
	Response campaign	Customer sample	Customer sample
	Profit/circ	Profit/circ	Profit/circ
Sales>\$0	Not applic.	\$2.01	\$3.09
Sales >=\$0 (all customers)	Not applic.	\$1.87	\$2.72
Difference	Not applic.	\$0.14	\$0.37
Confidence level	Not applic.	< 90% (NS)	95%

NS: Not significant.

Source: Campaign validation data and random customer sample.

models that are built conditional on customer response. Two different response-spend models were estimated. The first response-spend model used the *campaign* data and the ‘no treatment’ *sample* of customers from the campaign. The second response-spend model used the *random sample of customers* from the Enterprise Data Mart. When building spend models (and later, the profit models), customers without sales during the campaign (that is non-responders) were excluded. A quantitative analysis supported this decision. It evaluated models for customers with and without sales on the dependent variable (Table 1).⁶⁶

Like the response model, several predictive algorithms were deployed to finalize an optimal spend model. This included decision trees, neural networks and OLS regression. Because OLS regression was the best-performing, the results of the OLS regression models were shown here. The spend models are expressed as the following standard OLS regression equation.^{67,68}

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_i X_i + \varepsilon \quad (2)$$

Where Y = sales per customer, β_0 = intercept, β_1 = slope for independent variable X_1 ,

X_1 = independent variable₁, β_2 = slope for independent variable X_2 , X_2 = independent variable₂, β_i = slope for independent variable X_i , X_i = independent variable_i and ε = error term. Various regression models were estimated using a stepwise technique and the SBC statistic.

The best spend model from the first analysis included the following predictors: customer transaction counts, number of items the customer purchased, the customer’s total spend, the customer’s average spend per transaction, product purchases by the customer, the customer’s payment information and a demographic-household variable (purchase amount with other catalogue retailers). For the second analysis (that is the random sample), the spend model contained the following final predictors: customer’s average spend per trip, the amount the customer spent during the last quarter and the amount the customer spent during key seasonal shopping periods.

In each data set, the predicted spend from the OLS regression model (that is, the predicted linear model score from equation (2)) was combined with the predicted score from the single-stage response model (that is, the predicted score from equation (1)) to form a ‘response-spend’ predicted

score for each customer (see equation (3) below). This was calculated in the following manner:

$$\begin{aligned} &\text{Predicted Score for 'Response-Spend'} \\ &= \hat{Y}_{\text{response model}} \times \hat{Y}_{\text{spend model}} \end{aligned} \quad (3)$$

For example, using the response model, a hypothetical customer had a probability of responding to the campaign of 10 per cent. From the response-spend model, she had an expected spend of \$100. This customer would have a score of 10 (0.10×\$100=10). Customers in the validation campaign were scored two times: once for the response model that used equation (1), and once for the spend model that used equation (2). A combined score was calculated and this was used to rank customers for the response-spend model (equation (3)). To determine the effect of the response model on customer sales, the same customers were also ranked using only their predicted response model score from equation (1).

Model performance is shown using metrics familiar to marketing professionals. The first is the customer response rate. The second metric is a measure of sales per customer mailed, and is typically referred to as 'sales per circulation' in the retail industry. This is defined as Sales per Circulation = $Sls_i / Circ_i$, where Sls_i = total sales, $Circ_i$ = quantity mailed in decile, and i = model decile. The metrics shown are cumulative by model decile because we wish to compare total model performance to a similar depth of the validation data. The sales and profit metrics used to evaluate model impact are *actual* sales and profit per customer using a validation campaign (the first analysis) or 'natural' data from the Enterprise Data Mart (the second analysis). The time period for actual customer response, sales and profit was 4 weeks for all models.

RESULTS

Response rates for the response and response-spend models

Regarding the first analysis that used campaign data, the response model performance is initially

Table 2: Cumulative response rates for response and response-spend models by decile

Decile	Cumulative response rate		
	Response model (%)	Response-spend model (%)	Difference (%)
1	12.12	11.81	0.31
2	11.31	11.01	0.30
3	10.77	10.52	0.25
4	10.32	10.08	0.23
5	9.94	9.70	0.24
6	9.52	9.35	0.17
7	9.09	8.97	0.12
8	8.72	8.63	0.09
9	8.28	8.25	0.03
10	7.85	7.84	0.00

Source: Validation campaign.

compared to the response-spend model. Table 2 shows that the actual cumulative response rates are higher for the response model, which is anticipated, as the response model ranks customers only on their predicted response.

The top decile of the response model had a cumulative response rate of 12.12 per cent, compared to 11.81 per cent for the response-spend model. This difference translates into a 2.6 per cent improvement in response for the response model compared to the response-spend model. At first glance, a higher response rate for the response model would seem to support the conclusion that a response model better targets customers. Indeed, with a large mailing of 20 million customers in total (two million per decile), this would net 6200 more responders than the response-spend model in the top decile. However, as the results will show, explicitly modeling customer spend using a response-spend model yields higher sales than the response-only model in terms of both gross and incremental sales. In the next section, the results of the response-spend model are reviewed in detail.

Sales per customer for the response and response-spend models

Using the campaign data, model performance is now compared in terms of cumulative gross and incremental sales per circulation. Incremental refers to the difference between a mail and control (that is no mail) population. Table 3

Table 3: Cumulative sales per circulation for response and response-spend models by decile

Decile	Response model		Response-spend model		% improvement of response-spend vs response model
	Cum gross sales/circ (in \$)	Cum incremental sales/circ (in \$)	Cum gross sales/circ (in \$)	Cum incremental sales/circ (in \$)	
1	8.07	4.46	9.00	5.19	16.4
2	7.22	4.10	7.77	4.50	9.6
3	6.75	3.83	7.07	4.04	5.3
4	6.35	3.58	6.57	3.76	5.0
5	6.03	3.42	6.16	3.53	3.5
6	5.72	3.27	5.78	3.35	2.5
7	5.42	3.14	5.45	3.19	1.8
8	5.17	3.04	5.18	3.09	1.6
9	4.88	2.91	4.90	2.96	1.7
10	4.62	2.78	4.61	2.82	1.3

Source: Validation campaign.

shows that the response-spend model outperformed the response-only model in terms of both gross and incremental sales.

The top decile of the response-spend model generated \$9.00 per circulation (that is per customer name mailed), compared to \$8.07 per circulation for the response model. Likewise, the response-spend model outperformed the response model in terms of incremental sales. The top decile of the response-spend model generated \$5.19 in incremental sales per circulation, compared to \$4.46 in the top decile of the response model. This is a 16 per cent improvement in incremental sales for the response-spend model over the response model. Similarly, customers in decile 2 of the response-spend model generated \$4.50 in cumulative incremental sales per circulation, whereas customers in decile 2 of the response model generated \$4.10 in cumulative incremental sales per circulation. This translates into a 10 per cent improvement in cumulative incremental sales per circulation for the response-spend model over the response model. However, the benefits of the response-spend erode as we mail deeper into a customer list. By decile 4, the response-spend model outperformed the traditional response model by only 5 per cent. Nevertheless, with large campaigns, a 5 per cent improvement can still translate into substantially higher revenue using response-spend models. To a model depth of four deciles, a \$0.18 difference in cumulative incremental sales per circulation between response

and response-spend models translates into a \$1.44 million increase in incremental revenue if eight million customers are mailed in total (two million customers per decile). This \$1.44 million is above and beyond the impact of not mailing customers whatsoever.

Swap set: Response model versus response-spend model

In addition to the cumulative sales results for the campaign data, a swap set analysis was also conducted to a hypothetical mail depth. A swap set analysis provides a view into the unique customers targeted by one model (the response model) compared to another model (the response-spend model). The swap set methodology involved the following steps. First, the random campaign sample of 67 645 customers was scored for both models. Second, a mailing depth of four deciles was selected and applied to both models. The first four deciles were selected because these were the points at which model lift dropped below random performance. Third, customers who would have been mailed using either model were removed from the analysis. That is, customers who were in the top four deciles in both the response and response-spend model were excluded from the analysis because they would have been mailed using either model. Also excluded from analysis were customers who were in the lower deciles of both the response and response-spend model

(deciles 5–10), because these customers would not be mailed using either set of model scores. Based on the cutoff described above, the swap set analysis retained customers who were in the top four deciles of the response–spend model *or* who were in the top four deciles of the response model, but *not* customers who were in the top deciles of both models. Thus, with overlapping customer deciles excluded, the swap set analysis provides insight into the unique customers targeted by each model. From the initial 67 645 scored customers in the validation campaign, the final swap set analysis contained 5040 customers. This drop in quantity is typical when one model is a component of the other. Specifically, the combined score for the response–spend model (equation (3)) contains the response score from the response model. Thus, we would expect many customers to have the same deciles for both the response and response–spend model.

The results of the swap set analysis confirm the cumulative sales results. After removing ‘overlapping’ customers from the data set that contained both sets of scores (from the response model and response–spend model), Table 4 shows that unique customers targeted by the first four deciles of the response–spend model generated substantially higher cumulative gross sales per circulation of \$5.67 compared to \$4.16 for the unique customers targeted by the response model (a statistically significant difference at the 99 per cent confidence interval).

Likewise, cumulative incremental sales were significantly higher for unique customers targeted by the response–spend model compared to the unique customers targeted by the response model. Cumulative incremental sales were \$3.56 for the unique customers targeted in the first four deciles of the *response–spend* model and \$2.34 for the unique customers targeted in the first four deciles of the *response model* (a statistically significant difference at the 99 per cent confidence interval). This translates into 37 per cent higher gross and 52 per cent higher incremental sales for the response–spend model over the traditional response model.

The cumulative sales results and the swap set results confirm that the response–spend model

Table 4: Results of swap set analysis, comparing unique customers identified by response–spend model response model

<i>Swap set analysis (Deciles 1–4)</i>		
<i>Model</i>	<i>Sales/circ (Mail)</i>	<i>Sales/circ (Incremental)</i>
Response–spend	\$5.67	\$3.56
Response	\$4.16	\$2.34
Difference	\$1.52	\$1.22
Lift over response model	36.50%	52.13%
Difference statistically significant?	Yes	Yes
Confidence level	99%	99%

Source: Validation campaign.

increased our ability to target customers based on sales. When customers’ spend model scores were combined with their response model scores (equation (3)), higher–spending customers were migrated (that is sorted) to the top deciles of the customer file using the two–stage model, yielding substantially higher gross and incremental sales per customer than a response model alone.

EXTENSION TO MULTI-STAGE PROFIT MODELS

Beyond concerns with customer sales, the literature also makes evident that customer profit is a financial objective. In seeking higher profit among customers, two–stage models (response–spend) can be extended to three–stage models to target the combination of the most responsive, highest–spending and most *profitable* customers for the firm. The methodology involves the following. First, using the second data set of random customers sampled from the Enterprise Data Mart, a new response model was built. Second, a new spend model was also built separately from the response model. Both models were built replicating the methodology used in the first (campaign) analysis. Third, competing profit models were built using various statistical techniques, as was done for the response and response–spend models. The dependent variable was a customer’s profit for their purchases during the response (outcome) period. Like the spend model in the first analysis, the profit models were

built separately from the response and spend models, and the final predictors were determined after screening several hundred potential candidate variables. Because the predictors in the other models have been described above, only the details of the final profit model are discussed here.

The best profit model contained the following final predictors: amount of customer spend, number of items customer purchased, each customer's average spend per trip, the number of products purchased per trip, and finally the customer's sum of margin dollars for his or her purchasing history leading up to the response period. Scores for each customer from this separate profit model were multiplied by scores from the response and spend models to generate a three-stage model score. The three-stage response-spend-profit model score was calculated in the following manner:

$$\begin{aligned} &\text{Predicted Score for 'Response, Spend, Profit'} \\ &= \hat{Y}_{\text{response model}} \times \hat{Y}_{\text{spend model}} \times \hat{Y}_{\text{profit model}} \end{aligned} \quad (4)$$

All customers were ranked on their predicted value of this combined score, and actual sales and profit from the validation data set were examined. In addition to the three-stage model (response, spend, profit), another two-stage profit score was calculated using only profit and sales scores as the components. That is,

$$\begin{aligned} &\text{Predicted Score for 'Response-Profit'} \\ &= \hat{Y}_{\text{response model}} \times \hat{Y}_{\text{profit model}} \end{aligned} \quad (5)$$

The purpose of this two-stage, response-profit model was to measure the effectiveness of a simpler profit model, as we might hypothesize that the most responsive and highest-profit customers can be targeted without explicitly considering their predicted sales. Such might be the case if marketers seek to maximize 'bottom line' revenue without regard to 'top line' revenue (sales).

Tables 5 and 6 contain the results of fitting several multi-stage models to the second data set. Before discussing the results of the three-stage model, note that, like the first analysis, the second analysis also found higher response rates for the traditional response model compared to other models. In the top decile, the traditional response model generated a response rate of 23.28 per cent, compared to 22.50 per cent for the three-stage model (response, spend, profit). Assuming a customer universe of two million per decile, this translated into 15 600 more customers targeted by the response model.

Once again, however, response rates alone are not the only indicators of model performance. Marketers typically focus on revenue. In this regard, the results of the second analysis, which used the random sample, found that the three-stage model generated the highest actual sales and profit in the 4-week response period. The first decile of the combined, three-stage model had \$11.22 in actual sales per circulation and \$5.10 in actual profit per circulation in the validation data. The baseline, standalone customer response model saw only \$10.05 in actual sales and \$4.37 in profit per

Table 5: Cumulative response rates by model type and decile for random sample

Decile	Response model	Response-profit model	Response-spend model	Response-spend-profit model
	Cum response rate (%)	Cum response rate (%)	Cum response rate (%)	Cum response rate (%)
1	23.28	23.00	23.09	22.50
2	21.75	21.34	21.52	20.73
3	20.34	20.06	20.15	19.48
4	19.15	18.90	18.94	18.42
5	18.03	17.86	17.92	17.43
6	16.95	16.78	16.85	16.47
7	15.90	15.77	15.80	15.54
8	14.92	14.81	14.86	14.70
9	14.03	13.94	13.97	13.88
10	13.10	13.10	13.10	13.10

Source: Validation data.

Table 6: Cumulative sales per circulation and cumulative profit per circulation by model type and decile for random sample

Decile	Response model		Response-profit model		Response-spend model		Response-spend-profit model	
	Cum sales/circ (in \$)	Cum profit/circ (in \$)	Cum sales/circ (in \$)	Cum profit/circ (in \$)	Cum sales/circ (in \$)	Cum profit/circ (in \$)	Cum sales/circ (in \$)	Cum profit/circ (in \$)
1	10.05	4.37	10.95	4.97	10.98	4.86	11.22	5.10
2	8.48	3.90	8.96	4.21	8.99	4.17	9.08	4.26
3	7.52	3.56	7.85	3.76	7.87	3.74	7.86	3.77
4	6.85	3.28	7.01	3.40	7.02	3.38	7.02	3.42
5	6.29	3.05	6.41	3.14	6.42	3.13	6.37	3.13
6	5.83	2.85	5.86	2.89	5.88	2.89	5.83	2.88
7	5.38	2.65	5.40	2.67	5.41	2.67	5.38	2.66
8	4.98	2.46	5.00	2.48	5.01	2.48	4.99	2.47
9	4.64	2.30	4.63	2.30	4.64	2.30	4.63	2.30
10	4.31	2.14	4.31	2.14	4.31	2.14	4.31	2.14

Source: Validation data.

circulation. This difference translated into 12 per cent higher sales and 16 per cent higher profit per customer for the three-stage model compared to the response model alone. The next best model performance was obtained by the various two-stage models. The *response-profit* model generated \$10.95 in sales and \$4.97 in profit per circulation, and the *response-profit* model had a similar \$10.98 in sales and \$4.86 in profit per circulation. Across these various models, the three-stage model (response, spend, profit) had the highest sales and profit compared to the standalone response model.

Although the three-stage model clearly outperformed the traditional response model in sales and profits, the performance difference was closer between the three-stage model and the next-best performing two-stage models. Specifically, decile 1 of the three-stage model had a more modest sales and profits advantage of 2.5–2.7 per cent relative to the *response-profit* model (a difference of \$0.27 in sales and \$0.13 in profit per circulation). Decile 1 of the three-stage model had a similar advantage of 2.2–5.0 per cent higher cumulative sales and profit compared to the *response-spend* model (a difference of \$0.24 in sales and \$0.24 in profit per circulation). Although sales and profit improvements of less than \$1.00 per customer (or increases of 2–5 per cent in model performance) might appear modest, they are large when estimating the model impact on the total customer base. Table 7 shows the results of a

model impact estimate assuming a total customer base of 20 million customers (two million per decile) consistent with the first analysis involving campaign data.

Under the assumption of two million customers per decile, the first decile of the three-stage model generated approximately \$2 300 000 more in sales and \$1 469 000 more in profit during a 4-week response period compared to a traditional response model. When comparing the three-stage model to two-stage models, the top decile of customers would generate a narrower range, but a nonetheless substantial amount of total sales (\$485 000–\$539 000) and total profit (\$264 000–\$489 000), depending on which two-stage model is compared. Because these model-impact estimates are for a 4-week period only, a total estimated impact of multi-stage models for a year of marketing should be substantially higher depending on the volume of direct marketing campaigns. Relative to the response-only model, the three-stage model has positive revenue that continues through all deciles in terms of cumulative profit and through decile 6 in terms of cumulative sales. However, comparing a three-stage model to a more sophisticated two-stage model (whether response-profit or response-spend), the cumulative sales and profit advantages of the three-stage model wane by approximately decile 4. Here, we observe only \$54 000–\$95 000 in higher sales and \$114 000–\$264 000 in higher

Table 7: Estimated 4-week impact of multi-stage models by model type and decile, assuming 20 million customers total, two million per decile

Decile	Estimated 4-week impact on customer file (Hypothetical 2MM per decile)						
	Cum circ (Hypothetical)	Response-spend-profit vs response model		Response-spend-profit vs response-profit model		Response-spend-profit vs response-spend model	
		Cum sales/circ (in \$)	Cum profit/circ (in \$)	Cum sales/circ (in \$)	Cum profit/circ (in \$)	Cum sales/circ (in \$)	Cum profit/circ (in \$)
1	2 000 000	2340296	1469293	539162	264220	484617	488819
2	4 000 000	2394244	1449745	476908	199392	345158	352834
3	6 000 000	2060212	1300016	78569	86033	-52318	201296
4	8 000 000	1418460	1058754	94527	113857	53697	263873
5	10 000 000	749381	763213	-430315	-130243	-546838	-31881
6	12 000 000	33714	328448	-429413	-179166	-664385	-112827
7	14 000 000	-4217	195845	-407031	-181304	-554797	-125622
8	16 000 000	55025	155271	-149403	-59594	-385505	-84339
9	18 000 000	-218700	24594	-43949	-8773	-252432	-15896
10	20 000 000	0	0	0	0	0	0

Positive dollar impacts are given in bold.

Source: Validation data.

profit. At decile 5, the three-stage model generates \$430 000–\$547 000 less in sales, and \$32 000–\$130 000 less in profits over the 4-week period. In brief, this finding reflects a trade-off that management must often make between sales and profits. As a consequence, marketing executives and statisticians should consider their financial objectives ('top line' versus 'bottom line' revenue) and use the appropriate models to best meet their needs. The trade-off in sales and profit also reflects the components of a three-stage model score, which contains a combination of customer scores for expected response, sales and profit. In contrast, the other two-stage model scores contain combinations of two scores (response and spend, or response and profit), which are simpler calculations used to rank expected sales and profits. In brief, no single multi-stage targeting tool is optimal for all marketing objectives, a topic that directly affects executive decision making with regard to analytic tools and CRM programs. In the last section of this article, both topics are discussed in greater detail.

Swap set: Response, spend and profit models

A swap set analysis was conducted on the second data set (the random sample) using the same methodology as the swap set for the

campaign data. The swap set for the random customer sample compared performance across the three-stage (response, sales and profit) model relative to the closest performing alternatives, which were the two-stage models (response-spend and response-profit). The three-stage model (response-profit-profit) served as the baseline modeling tool with the best performance.

The second swap set began with response, sales and profit model scores for all 33 491 customers in the validation data set, and was conducted to two different model depths. For consistency with the first analysis, one swap set was conducted to a depth of four deciles. Another swap set was conducted to a depth of two deciles because the lift of the three-stage model dropped below a value of 1.00 at decile 3, indicating that the profit model provided maximum lift over a random treatment at decile 2. A swap set to decile 2 also estimated the impact of a smaller, more targeted campaign. After eliminating customers with similar three-stage and two-stage model deciles to a depth of two deciles, 1997 unique customers with unique model ranks were retained from the second sample of 33 491 customers. At a depth of four deciles, 1397 unique customers were retained. Table 8 contains the results of the swap set for multi-stage profit models.

Table 8: Results of swap set analyses, comparing unique customers targeted by three-stage profit models versus two-stage models

Model	Swap set analyses			
	Deciles 1–2		Deciles 1–4	
	Sales/circ	Profit/circ	Sales/circ	Profit/circ
<i>Response-spend-profit model vs response-profit model</i>				
Response-spend-profit	\$ 24.42	\$ 11.83	\$ 14.45	\$ 7.72
Response-profit	\$ 20.06	\$ 10.01	\$ 13.85	\$ 6.99
Difference	\$ 4.36	\$ 1.82	\$ 0.61	\$ 0.73
Lift over response-spend model	21.76%	18.23%	4.38%	10.45%
Difference statistically significant?	Yes	Yes	No	Yes
Confidence level	99%	99%	—	90%
<i>Response-spend-profit model vs response-spend model</i>				
Response-spend-profit	\$ 24.29	\$ 11.92	\$ 15.02	\$ 7.85
Response-spend	\$ 21.90	\$ 9.49	\$ 14.75	\$ 6.52
Difference	\$ 2.38	\$ 2.44	\$ 0.27	\$ 1.33
Lift over response-spend model	10.88%	25.68%	1.83%	20.35%
Difference statistically significant?	Yes	Yes	No	Yes
Confidence level	95%	99%	—	90%

Source: Validation data.

In terms of profit per circulation, a three-stage profit model outperformed a two-stage model regardless of whether the model depth was limited to the top 20 per cent or the top 40 per cent of customers (unless otherwise noted, all differences were statistically significant at the 90 per cent confidence level or greater). For example, at a depth of two deciles, unique customers targeted by the three-stage model (response-spend-profit) had \$4.36 more in sales per circulation and \$1.82 more in profit per circulation than customers targeted using a two-stage *response-profit* model. The three-stage model also generated \$2.38 more in sales and \$2.44 more in profit than the two-stage, *response-spend* model. The swap set on the unique customers in the three-stage model translated into 18–26 per cent higher profit per circulation compared to the two-stage models, and 11–22 per cent higher sales per circulation depending on which two-stage model was compared.

The results of a swap set to four deciles were similar to those of the swap set at two deciles with regard to *profit*. The three-stage model generated \$0.73–\$1.33 more in profit than the two-stage models through decile 4, depending

on which two-stage model was used, which translated into 10–20 per cent higher profit per circulation than either of the two-stage models. However, at four deciles, the unique customers in the multi-stage model generated statistically similar *sales* compared to unique customers targeted using simpler, two-stage models. This indicates that the three-stage model converges to some extent with the two-stage models depending on type of model, model depth and model performance metric. Caution must be exercised, however, as the lack of a statistical difference could also be a result of the smaller sample of customers who remained after eliminating overlapping scores across models. The smaller sample could affect the magnitude of the sales/profit differences because of larger standard errors simply owing to increased variability in the smaller data set. Future work could more definitively replicate these findings using larger customer samples.

Results summary

Using retail store data, this article described a technique for building multi-stage customer-level models that can target significantly higher sales

and profits compared to traditional response models alone. An enhanced, two-stage (response-spend) model could generate up to 12 per cent more in gross sales per circulation (16 per cent more in incremental sales) compared to a response model by itself. For a typical direct mail campaign circulation with this brand, the impact of a two-stage model translated into \$1.44 million more in incremental sales (that is, sales above and beyond a 'no-mail' customer treatment). This article also presented a methodology for extending two-stage models to three-stage models, which predict a combination of response, spend and profit for each customer (see Sismeiro and Bucklin for a similar methodology using binary models).²⁵ The three-stage (response-spend-profit) model also drove higher sales and profits than the response model alone, by an estimated 12 per cent for sales and 17 per cent for profit in the top decile of the model, which was an estimated 4-week impact of up to \$2.3 million in sales and \$1.5 million in profit relative to a simpler, one-stage (response) model in the top model decile. Although a comparison of the three-stage and two-stage models yielded a much more modest 2–5 per cent increase in revenue per circulation for the top decile, this could potentially translate into \$485 000–\$537 000 in sales and \$264 000–\$489 000 more in total profit for this 4-week period, depending on which two-stage model is used as a comparison.

Why do multi-stage models generate higher sales and profits? Simply put, response models do not necessarily target the highest-spending or highest-profit customers because they are not designed to do so. Statistically, the standalone response models and two-stage models sort customers based on some combination of their probability of responding, expected sales or expected profit, but not all three predictions. By combining the three scores and ranking all customers on this single new score, we can identify customers with the highest combination of expected response, sales and profits, particularly within the top 30–40 per cent of the customer file. Beyond this proportion, some other combination of two of the scores (response and

profit, or response and spend) was a better choice than ranking customers on the three-stage model. Nevertheless, as long as a decision is made to mail a relatively low proportion of the customer base, the results suggest that marketers use multi-stage models to target higher-revenue customers and achieve revenue goals. Such profit goals take on greater urgency during economic downturns when nearly all retailers struggle to show a profit.

IMPLICATIONS

Three implications emerge from the results of the multi-stage models built to predict customer sales and profits. In light of the sales and profit advantages for two- and three-stage models, the first implication relates to the continued deployment of response models in marketing. The second implication involves observations regarding the trade-off between sales and profits in these results, and encourages marketers to align the type of model used with their campaign objectives. The third implication deals with one aspect of multi-stage model methodology. The last section reviews limitations of this research, and future directions are suggested for the topic of predictive models that target customers for their sales and profitability.

Are response models still useful?

Given the advantages of two- or three-stage models, readers may wonder whether traditional (standalone) response models are still useful marketing tools, apart from their inclusion in multi-stage models. Certainly, traditional response models can yield substantially better performance compared to single-variable targeting techniques or the first generation of recency, frequency and monetary (RFM) segmentation techniques. (Note that RFM models have since been extended to incorporate more complex data and approaches.)^{12,26,45} Even when revenue is not a primary business objective for a marketing program, there is a role for traditional response models. For example, response models are well suited for programs that seek to obtain the highest number of customer purchases. Once

consumers are regular purchasers, this can be followed with additional contacts (direct mail, telemarketing and other channels) in an effort to expand the customer relationship. Several retail brand divisions often use traditional response models in their non-customer or 'prospect' campaigns, with the objective of obtaining the highest non-customer store traffic, arguing that store sales personnel can convert them to customers. Response models also have other benefits. They can reduce the mail volume to an otherwise large customer base, increase a customer's shopping frequency, and eliminate costs associated with mailing all customers, which in turn increases the efficiency of mail campaigns. Finally, traditional response models can also help deliver more relevant customer messages. Daily, customers are bombarded with direct mail, television advertisements, telemarketing, text and email messages, and web-based advertisements. In such an environment, traditional response models can partially assist in reducing the customer experience often referred to as 'information overload'.

The appeal of traditional, standalone response models is also their relative simplicity for project managers, product managers, segment managers and other marketers with less statistical training. In this sense, response models are intuitively attractive and conceptually straightforward: a response model targets those customers who will respond to a mailing. From the experiences of colleagues in other analytic marketing roles both within and across organizations, as well as the author's own experience, it is apparent that the appeal of simpler targeting solutions may be more common in certain sectors and sub-specialties than in others. For example, marketers and statisticians in the risk-based areas of bank lending (for example, credit cards, auto loans/leases and personal credit) often build and deploy two-stage models to predict a customer's expected balance.³³ This is driven by the desire to generate customer revenue, as customers who fail to use and/or revolve their credit are typically less profitable than those who carry balances across payment cycles. Likewise, in the

property and casualty insurance industry, statisticians in marketing organizations typically build two-stage predictive models to estimate a customer's expected home or automobile insurance premium. As in the risk-banking specialty, the amount of insurance purchased (that is, premium 'spend') is linked with sales and profitability. Drivers with larger, full-coverage automobile insurance premiums are typically more profitable than those who simply carry low-priced collision-only coverage at a fraction of the cost.

Although two-stage modeling practices would therefore seem standard for statisticians working in marketing organizations, this is not the case. Even within the financial services industry, a bank in the United States ranked within the top 25 institutions with regard to deposits continued to model customer cross-sell and acquisition campaigns using traditional response models through the mid-2000s. This was based on the view that most campaign responders were sufficiently profitable to make specific revenue models unnecessary. In addition to differences across industries and subspecialties, there can also be differences in model deployment within an organization. For example, within the retailer represented in this article, different product divisions have unique historical experiences with multi-stage targeting models. Some have used multi-stage models, particularly two-stage models, whereas others have only recently made a transition from RFM segmentation/propensity modeling to traditional response modeling. In other cases, traditional response models are deployed but the campaign is measured against sales or profit metrics. Although a performance evaluation in terms of revenue requires the use of a multi-stage model to best meet the financial objectives of a campaign, the development and deployment of multi-stage models is not yet being considered. For all these reasons, this article is important because it introduces an effective method to enhance customer targeting and therefore boost campaign performance that is either overlooked or considered too complicated by marketers in certain retail organizations.

Three-stage versus two-stage models: Trade-off and objectives

The results presented in Tables 6 and 7 show how the three-stage model increased the ability to target customers with higher sales and profits through the top 30–40 per cent of the customer file compared to other two-stage models (that is, those based on customers' combined response-spend scores or their combined response-profit scores). However, once we target the top 50 per cent of customers (decile 5), the influence of the three-stage model compared to a two-stage model begins to wane. This finding is discussed in light of the way in which these scores are constructed and the targeting goal for marketers. Specifically, customer scores for the three-stage model are a combination of expected response, spend and profit. The 'best' customers in a three-stage model are sorted to the top of the file based on their combination of all three predictions. Some customers may have slightly lower expected sales, but because their expected profit is higher than average, these higher-profit customers are pushed to the top of the customer file when the three-stage model score is calculated and used to rank the customer file. This allows three-stage models to better target higher-profit customers at the top of the customer file relative to the response model, the response-spend model or the response-profit model. However, three-stage models have the potential to perform worse based on another outcome. The converse situation is also true for the traditional response model when examining response rates. The response model simply ranks customers based on their expected *response*, and therefore outperforms the two- or three-stage models when we examine cumulative *response* performance alone.

These observations explain why, at some depth of the customer file, a multi-stage or three-stage model may sacrifice any one outcome (whether response, sales or profit) compared to a simpler model that contains any of the other predicted outcomes. Because there is no one optimal model for all purposes, financial objectives should always be paramount when executives consider which models to use for their marketing programs.

Such an observation has been borne out in Homburg *et al's* analysis, which used structural equation models to understand how strategy, planning, control and management involvement were related to customer 'prioritization'.²² From the results presented here, the best marketing program performance, whether determined by response, sales or profits, is best achieved by using the most appropriate targeting model for that objective and for the targeted circulation. Although it would be ideal for this article to influence *executives* directly by demonstrating the importance of aligning marketing program objectives with the optimal targeting tool, such ideals may not always be realized in the fast-paced retail sector. *Statisticians*, who are more familiar with concerns regarding a potential disconnect between the model objective and the model type are likely to be more sensitive to such a problem in practice. Regardless of who identifies such a disconnection, a course correction taken early in the marketing cycle should improve the results of marketing programs through higher, more consistent financial outcomes.

Methodology used to build multi-stage models

This article demonstrated one method of building two- and three-stage models that predicted customer sales and profit. Specifically, three separate models were built: customer response, customer spend (for the responders) and customer profit (for those who responded and had sales). For the two- and three-stage models, combinations of two or three of these scores were multiplied together for each customer, and this resulting new score was used to rank all customers. Such an approach assumed that response, spend and profit were independent of one another. This assumption was confirmed by allowing the score from one model (for example response score) as a candidate predictor into another model (for example, the customer spend model). The predicted score did not pass the initial variable screening and thus did not enter the stepwise algorithm for the next stage of the model. Therefore, the assumption involving

the independence of response, spend and profit models and the technique of building three independent models were supported. The multi-stage model methodology used in this article was also selected to be understood and implemented by the widest possible analytic audience, particularly as a more limited group of statisticians have the necessary specialty software to build pre-packaged, multi-stage models. For analysts with such software and an interest in the topic, tests can be conducted to determine whether commercially available multi-stage model methodologies are better suited to the statistician's data.

Research limitations

This research has several limitations. First, both the campaign data and the random customer sample were used to model the *store-based* behavior for retail customers. The extent to which these conclusions also apply to other marketing channels is uncertain. Online, e-commerce and telemarketing sales data, for example, would be additional channels in which multi-stage customer models could be deployed and measured (for an example of a conditional model using online data, see Sismeiro and Bucklin).²⁵ If these conclusions are replicated across marketing channels, future multi-stage modeling efforts could be expanded even further to assist firms in achieving their maximum marketing and financial objectives across the organization.

Second, this research used a simple definition of profit (that is, sales minus the cost of goods). It did not fully measure the range of costs involved with profit, particularly from a finance perspective as reviewed in the literature. For direct mail, one of the largest costs is the discount offered to the customer. Fortunately, the data used to build the multi-stage profit models offered an opportunity to select a time period in which there were no additional discounts mailed to customers. Nevertheless, this article did not incorporate any other cost information, such as costs related to customers' payment type, past promotional costs, and other transactional and marketing costs cited in the

literature.⁶ The use of a simplified profit definition in this article means that there is substantial room for future research to evaluate the consequences of individual customer behavior on sales and profit. Future research along these lines might uncover important complexities in the profit data that, once understood, could build better targeting models. Therefore, the results from this retailer should be considered a beginning from which additional research on the components of predicted profit can be studied.

Third, this article on the extension of multi-stage *profit* models cannot draw conclusions regarding their impact on *campaign* data. This limitation was a result of data unavailability whereby a random customer sample was used rather than a direct marketing response campaign. Without campaign data and an associated control sample, our conclusions about profit models apply to gross estimates only, and cannot estimate the impact of incremental profits. With the addition of campaign data and a no-treatment sample, future research can measure the financial benefits of multi-stage profit models relative to withholding direct mail altogether.

Although there were limitations of the data used, there were two benefits to building profit models on a random customer sample. First, it allowed us to evaluate the reliability and generalizability of the initial response-spend models. If we had only used one kind of data set, we would not have known whether the benefits of two-stage models were specific to the particular campaign dynamics observed in our response data, or whether they could also be applied to marketing environments with only random customer samples. Second, although profit data availability required the use of a random customer sample, it was beneficial because it provided an opportunity to investigate trends for customers who purchased different types of products. Had we built two-stage models for customers who bought only one specific product, we might be less convinced of the reliability of the trends showing that response-spend models enhanced our customer-targeting capabilities. We would question

whether the two-stage model performance ‘interacts’ with products. Using both types of data – campaign compared to random customer samples, and data for customers who buy two different products – we concluded that two-stage (response-spend) models generated higher sales compared to response models alone. Extending the two-stage models to include expected customer profit, we also saw that the use of two- and three-stage customer-level models substantially increased sales and profits per customer, depending on the depth of model deployment. Because of the data unavailability we encountered here, the *three-stage profit model* methodology should be replicated in other contexts to determine whether the results and conclusions generalize well. If the results have a similar impact across various contexts, marketers will have taken another important step toward enhancing their traditional targeting tools.

ACKNOWLEDGEMENTS

I thank the editor and anonymous reviewers for comments and feedback on previous drafts of this article. Jack Palmer provided much assistance in data availability for the profit model extension, as well as feedback on manuscript drafts and reviewers’ comments. Conversations with Phil Krauskopf and Jack Palmer were instrumental in shaping ideas regarding the topic of multi-stage models.

REFERENCES AND NOTES

- Walters, R.G. (1988) Retail promotions and retail store performance: A test of some key hypotheses. *Journal of Retailing* 64(2): 53–180.
- While the subject is not without controversy,³ marketers usually argue that CRM grows customer sales and/or profit.
- East, R., Hammond, K. and Gendall, P. (2006) Fact and fallacy in retention marketing. *Journal of Marketing Management* 22: 5–23.
- Brown, J.R. and Dant, R.P. (2008) Scientific method and retailing research: A retrospective. *Journal of Retailing* 84(1): 1–13.
- Hart, A. and Smith, M. (1998) Customer profitability audit in the Australian banking sector. *Managerial Auditing Journal* 13(7): 411–418.
- Haenlein, M., Kaplan, A.M. and Schoder, D. (2006) Valuing the real option of abandoning unprofitable customers when calculating customer lifetime value. *Journal of Marketing* 70(July): 5–20.
- Cuthbertson, R. and Laine, A. (2004) The role of CRM within retail loyalty marketing. *Journal of Targeting, Measurement, and Analysis for Marketing* 12(3): 290–304.
- D’Souza, R., Krasnodebski, M. and Abrahams, A. (2007) Implementation study: Using decision tree induction to discover profitable locations to sell pet insurance for a startup company. *Journal of Database Marketing & Customer Strategy Management* 14(4): 281–288.
- Gunnarsson, C.L., Walker, M.M., Walatka, V. and Swann, K. (2007) Lessons learned: A case study using data mining in the newspaper industry. *Journal of Database Marketing & Customer Strategy Management* 14(4): 271–280.
- Seetharaman, P.B. (2004) The additive risk model for purchase timing. *Marketing Science* 23(2): 234–242.
- Kee, R. (2007) Cost-volume-profit analysis incorporating the cost of capital. *Journal of Managerial Issues* 19(4): 478–493.
- Ryals, L. (2002) Measuring risk and returns in the customer portfolio. *Journal of Database Marketing* 9(3): 219–227.
- Wubben, M. and Wangenheim, F.V. (2008) Instant customer base analysis: Managerial heuristics often ‘get it right’. *Journal of Marketing* 72(May): 82–93.
- Ryals, L. (2005) Making customer relationship management work: The measurement and profitable management of customer relationships. *Journal of Marketing* 69(October): 252–261.
- Marn, M.V. and Rosiello, R.L. (1992) Managing price, gaining profit. *McKinsey Quarterly* (4): 18–37.
- Helgesen, O. (2006) Are loyal customers profitable? Customer satisfaction, customer (action) loyalty and customer profitability at the individual level. *Journal of Marketing Management* 22: 245–266.
- Sievanen, M., Suomala, P. and Paranko, J. (2004) Product profitability: Causes and effects. *Industrial Marketing Management* 33(5): 393–401.
- Fader, P.S., Hardie, B.G.S. and Lee, K.L. (2006) More than meets the eye. *Marketing Research* 18 (Summer): 9–14.
- Bosch, M.-T., Montllor-Serrats, J. and Tarrazon, M.-A. (2007) NPV as a function of the IRR: The value drivers of investment projects. *Journal of Applied Finance* 17(Fall/Winter): 41–45.
- Helgesen, O. (2008) Targeting customers: A financial approach based on creditworthiness. *Journal of Targeting, Measurement, and Analysis for Marketing* 16(4): 261–273.
- Venkatesan, R. and Kumar, V. (2004) A customer lifetime value framework for customer selection and resource allocation strategy. *Journal of Marketing* 68(October): 106–125.
- Homburg, C., Droll, M. and Totzek, D. (2008) Customer prioritization: Does it pay off, and should it be implemented? *Journal of Marketing* 72(September): 110–130.
- Fader, P.S., Hardie, B.G.S. and Lee, K.L. (2005) Counting your customers’ the easy way: An alternative to the Pareto/NBD model. *Marketing Science* 24(2): 275–284.
- An additional benefit of individual-level predictive revenue models is that they can help explain customer behavior, unlike financial models built at an aggregate level. See arguments in Sismeiro and Bucklin²⁵ and Libai *et al.*²⁶
- Sismeiro, C. and Bucklin, R.E. (2004) Modeling purchase behavior at an e-commerce web site: A task completion approach. *Journal of Marketing Research* 41(August): 306–323.

- 26 Libai, B., Narayandas, D. and Humby, C. (2002) Toward an individual customer profitability model: A segment-based approach. *Journal of Service Research* 5(1): 69–76.
- 27 Reinartz, W.J. and Kumar, V. (2003) The impact of customer relationship characteristics on profitable lifetime duration. *Journal of Marketing* 67(January): 77–99.
- 28 Gupta, S. and Zeithaml, V. (2006) Customer metrics and their impact of financial performance. *Marketing Science* 25(6): 718–739.
- 29 Buckinx, W. and Van den Poel, D. (2005) Customer base analysis: Partial defection of the behaviorally loyal clients in a non-contractual FMCG retail setting. *European Journal of Operations Research* 164(1): 252–268.
- 30 Lemmens, A. and Croux, C. (2006) Bagging and boosting classification trees to predict churn. *Journal of Marketing Research* 43(May): 276–286.
- 31 Gustafsson, A., Johnson, M.D. and Roos, I. (2005) The effects of customer satisfaction, relationship commitment dimensions, and triggers on customer retention. *Journal of Marketing* 69(October): 210–218.
- 32 Verhoef, P.C. (2003) Understanding the effect of customer relationship management efforts on customer retention and customer share development. *Journal of Marketing* 67(October): 30–45.
- 33 Till, R.J. and Hand, D.J. (2003) Behavioural models of credit card usage. *Journal of Applied Statistics* 30(10): 1201–1220.
- 34 Stewart, K. (1998) An exploration of customer exit in retail banking. *International Journal of Bank Marketing* 16(1): 6–14.
- 35 Verhoef, P.C. and Leeftang, P.S.H. (2009) Understanding the marketing department's influence within the firm. *American Marketing Association* 73(March): 14–37.
- 36 Venkatesan, R., Kumar, V. and Bohling, T. (2007) Optimal customer relationship management using Bayesian decision theory: An application for customer selection. *Journal of Marketing Research* 44(November): 579–594.
- 37 Berry, J. (2006) The benefits of using a decision engine to optimize campaign planning for direct marketing. *Database Marketing and Customer Strategy Management* 13(4): 319–323.
- 38 Giltner, R. and Ciolli, R. (2000) Re-think customer segmentation for CRM results. *The Journal of Bank Cost and Management Accounting* 13(2): 3–5.
- 39 Vogel, V., Evanschitzky, H. and Ramaseshan, B. (2008) Customer equity drivers and future sales. *Journal of Marketing* 72(November): 98–108.
- 40 Reinartz, W. and Kumar, V. (2000) On the profitability of long-life customers in a non-contractual setting: An empirical investigation and implications for marketing. *Journal of Marketing* 64(October): 17–35.
- 41 Wang, Y.-M. (2008) Determinants affecting consumer adoption of contractless credit card: An empirical study. *CyberPsychology & Behavior* 11(6): 687–689.
- 42 Kumar, V., George, M. and Pancras, J. (2008) Cross-buying in retailing: Drivers and consequences. *Journal of Retailing* 84(1): 15–27.
- 43 Singh, J., Ehrenberg, A. and Goodhardt, G. (2008) Measuring customer loyalty to product variants. *International Journal of Market Research* 50: 513–532.
- 44 Greene, H. and Milne, G.R. (2005) Alternative data sources in targeted marketing: The value of exographics. *Journal of Targeting, Measurement and Analysis for Marketing* 14(1): 33–46.
- 45 Greene, H. and Greene, S. (2008) Enhancing segmentation systems. *Journal of Targeting, Measurement, and Analysis for Marketing* 16(4): 298–311.
- 46 Baumann, C., Burton, S. and Elliott, G. (2007) Predicting consumer behavior in retail banking. *Journal of Business and Management* 13(1): 79–96.
- 47 Du, R.Y., Kamakura, W.A. and Mela, C.F. (2007) Size and share of customer wallet. *Journal of Marketing* 71(April): 94–113.
- 48 McCrary, M. (1998) 'Same song, different verse:' Processes of race-sex stratification and self-employment success. *Research in Social Stratification and Mobility* 16: 319–350.
- 49 Regarding the dependent variable used in this article (profit dollars), it is theoretically possible to use multi-stage models to predict a customer's profit margin *percentage* rather than the *dollar amount*. However, a profit model that predicted profit margin percentage would potentially eliminate important variability in the data. For example, a customer with \$400 in sales and a \$300 cost of goods would have a profit margin of 0.25 in percentage terms ($\$100/\$400=0.25$), but so too would a customer with \$8 in sales and a profit margin of \$6 ($\$2/\$8=0.25$). By collapsing such different kinds of behavior into the same value on the dependent variable, we are likely to overlook the natural customer heterogeneity in the data and build poorly predicting models as a result. To avoid such a problem, the profit models built in this article predicted customers' margin *dollars*.
- 50 Berry, M.J.A. and Linoff, G. (1997) *Data Mining Techniques*. New York: Wiley.
- 51 There was an 8 per cent response rate in the first data set, and a 13 per cent response rate in the second data set. Both data sets were over-sampled to obtain a richer data set of responders and therefore build a more predictive model.^{50,52}
- 52 SAS Institute. (1998) *Data Mining and the Case for Sampling*, SAS Institute best practices paper. Cary, NC: SAS.
- 53 Liu, C and Cai, S (2008) Customer cross-selling model based on counter propagation network. *Direct Marketing* 2(1): 36–47.
- 54 Li, S., Sun, B. and Wilcox, R.T. (2005) Cross-selling sequentially ordered products: An application to consumer banking services. *Journal of Marketing Research* 42(May): 233–239.
- 55 Wang, H., Yu, Y. and Zhang, K. (2008) Ensemble PROBIT models to predict cross selling of home loans for credit card customers. *International Journal of Data Warehousing and Mining* 4(2): 15–21.
- 56 Hand, D.J. and Crowder, M.J. (2005) Measuring quality in retail banking. *Statistical Modelling* 5: 145–158.
- 57 Magidson, J. and Vermunt, J.K. (2004) Latent class models. In: D. Kaplan (ed.) *The Sage Handbook of Quantitative Methodology for the Social Sciences*. Thousand Oaks, CA: Sage, pp. 175–198.
- 58 Magidson, J. and Vermunt, J.K. (2005) An extension of the CHAID tree-based segmentation algorithm to multiple dependent variables. In: C. Weihs and W. Gaul (eds.) *Classification: The Ubiquitous Challenge*. Heidelberg, Germany: Springer, pp. 176–183.
- 59 Vermunt, J.K., Tran, B. and Magidson, J. (2008) Latent class models in longitudinal research. In: S. Menard (ed.) *Handbook of Longitudinal Research: Design, Measurement, and Analysis*. Burlington, MA: Elsevier, pp. 373–385.
- 60 Vermunt, J.K. and Magidson, J. (2005) Hierarchical mixture models for nested data structures. In: C. Weihs and W. Gaul

- (eds.) *Classification: The Ubiquitous Challenge*. Heidelberg, Germany: Springer, pp. 240–247.
- 61 Jen, L., Chou, C.-H. and Allenby, G.M. (2003) A Bayesian approach to modeling purchase frequency. *Marketing Letters* 14(1): 5–20.
- 62 Various advanced decision tree modeling techniques such as boosting and bagging are not available in many versions of data mining software. Therefore, the most common decision tree modeling algorithms and techniques were used here.
- 63 Aldrich, J.H. and Nelson, F.D. (1984) *Linear, Probability, Logit, and Probit Models*. London: Sage.
- 64 Hosmer, D.W. and Lemeshow, S. (1989) *Applied Logistic Regression*. New York: Wiley.
- 65 Menard, S.W. (1995) *Applied Logistic Regression Analysis*. Thousand Oaks, CA: Sage.
- 66 The following analysis supported the decision to exclude customers without sales when building the predictive models. Because this was a supportive and not a primary analysis, limited metrics were evaluated for select models that performed well and were the focus of this article: the two-stage model (response-spend) and the three-stage model (response-spend-profit). Two sets of sales and profit models were built: the first approach involved models for those for customers with sales greater than \$0, and the second approach built models that included customers with and without sales. All customers were scored for both model approaches. Scores for the two approaches were correlated at 0.95–0.97 as part of the two-stage (response-spend) model, and 0.94 for profit using the three-stage model. Most customers' scores (96–98 per cent) were in deciles 1–4 for both approaches (or in deciles 5–10 for both approaches). For the remaining 2–4 per cent of customers for whom this was not the case, a swap set analysis was conducted. The swap set analysis determined the impact of including or excluding customers without sales from the models. Table 1 summarizes the results. The top panel shows the impact of the two approaches as a function of the two-stage models (response-spend). The model-building approach that *excluded customers without sales* (that is, only retained customers with 'sales > \$0') yielded \$0.51–\$0.85 higher sales per circulation during the 4-week outcome period, which was statistically significant at 95 per cent confidence or higher. The second panel shows results for the three-stage models (response-spend-profit). As above, *excluding customers without sales resulted in \$0.37 higher profit per circulation*, with 95 per cent confidence, but only at a mail depth of two deciles. (Results in the top four deciles just missed statistical significance by \$0.03 profit per circulation at the 90 per cent level.) Together, the swap set results that compared the two approaches supported the decision to exclude customers without sales from the model-building process.
- 67 Mendenhall, W., Reinmuth, J.E., Beaver, R. and Duhan, D. (1986) *Statistics for Management and Economics*, 5th edn. Boston, MA: Duxby.
- 68 Lewis-Beck, M.S. (1980) *Applied Regression*. London: Sage.