Marketing Science Institute Special Report 07-205

Will the Frog Change into a Prince? Predicting Future Customer Profitability

Roland Rust, Raj Venkatesan, and V. Kumar

WILL THE FROG CHANGE INTO A PRINCE?:

PREDICTING FUTURE CUSTOMER PROFITABILITY

Roland T. Rust*

V. Kumar

Rajkumar Venkatesan

July 19, 2007

*Roland T. Rust holds the David Bruce Smith Chair in Marketing, is Executive Director of the Center for Excellence in Service, and is Chair of the Department of Marketing at the Robert H. Smith School of Business, University of Maryland. Address: 3451 Van Munching Hall, University of Maryland, College Park, MD 20742. Phone: 301-405-4300. Fax: 301-405-0146. Email: rrust@rhsmith.umd.edu.

V. Kumar is the ING Chair Professor, and Executive Director, ING Center for Financial Services at the University of Connecticut. Address: Department of Marketing, School of Business, Unit 1041, University of Connecticut, Storrs, CT 06269-1041. Phone: (860) 486-1086. Fax: (860) 486-8396. Email: vk@business.uconn.edu

Rajkumar Venkatesan is Associate Professor of Business Administration, Darden School of Business, University of Virginia. Address: Darden School of Business, P.O. Box 6550, Charlottesville, VA 22906-6500. Phone: 434-924-6916. Fax: 434-243-7676. Email: Venkatesanr@darden.virginia.edu

The authors thank the audience at various conferences including the EMAC 2007, QUIS 2007, and the 2007 Marketing Science Conference as well as various universities. The authors also thank the firm that provided us the access to the data.

WILL THE FROG CHANGE INTO A PRINCE?:

PREDICTING FUTURE CUSTOMER PROFITABILITY

Abstract

More and more companies have customer databases that enable them to analyze customer profitability over time. These companies often seek to determine the most important customers as indicated by their current or historical profitability, and focus attention on them. Focusing on profitable customers can result in more efficient use of marketing resources, but such an approach neglects the fact that customers can evolve over time. Some customers begin as low-profit customers, but eventually develop into high-profit customers. Others may start out as high-profit customers, but become unprofitable over time. Previous efforts to predict future profitability have been relatively unsuccessful, with relatively simple, naïve models often performing just as well as, or better than, more sophisticated ones. Our paper presents a new approach to predicting customer profitability in future periods, in the case where the company possesses a longitudinal customer database. Our model simulates future purchasing scenarios, one customer at a time. By averaging over many possible futures for each customer, we obtain a more accurate prediction of future profitability. We estimate the model on data from a high-tech company in a business-to-business context, and validate the model's prediction ability on a holdout sample. The results show that our model outperforms previous models in predicting which customers will increase in profitability and which customers will decrease in profitability over time.

WILL THE FROG CHANGE INTO A PRINCE?: PREDICTING FUTURE CUSTOMER PROFITABILITY

INTRODUCTION

In a popular Grimm Brothers fairy tale, a frog is vastly underestimated by a young princess, who does not realize that he will eventually be transformed into a prince. Like the frog, some customers may be badly underestimated (and some may be overestimated). A business would certainly like to know how many of its “frog” customers are really “princes” in disguise.

The Management Scenario

Let us consider a typical management scenario. XYZ company, a B2B marketer, directs marketing actions (e.g., sales calls, promotional mailings, telephone contacts, relationship building visits, etc.) to individual customers. The firm also keeps a database that includes a record of those actions, along with a record of sales and profitability for each customer. The firm realizes that some customers may grow in importance over time, and some customers may decline. The firm also knows, based on previous experience and academic research, that focusing more marketing effort on customers that are likely to grow is likely to yield better results (much as watering a plant that is in rich soil is likely to be more effective than watering a plant that is in depleted soil). Therefore the firm would like to identify the customers whose business is likely to grow, so that it can allocate more marketing effort toward them. Likewise the firm would also like to know the customers whose business is likely to decline, so that it can pay less attention to them.

Many businesses have become aware that targeting profitable customers can make marketing spending more efficient (Mulhern 1999; Zeithaml, Rust and Lemon 2001, Kumar et al. 2007), an approach that is increasingly possible due to the proliferation of historical customer databases. Even better is to target the customers who *will be* profitable. Accurate estimates of future profitability would allow firms to make better marketing resource allocation decisions for individual customers (Bolton et al. 2004). Unfortunately, attempts to predict future customer profitability have been relatively unsuccessful, with very simple models often performing just as well as more sophisticated ones. The purpose of this paper is to provide a new approach to predicting future customer profitability, to estimate, test, and validate it on a large-scale customer database, and to demonstrate that it performs better than existing methods.

Managing the Customer Pyramid

Managers have long been aware that some customers are more profitable than others, and have known that paying attention to more profitable customers can produce better results. One approach involves building a "customer pyramid" to classify customers into different profitability tiers (Rust, Zeithaml and Lemon 2000, pp. 187-231; Zeithaml, Rust and Lemon 2001). In the Rust-Zeithaml-Lemon customer pyramid, customers are classified into Platinum (the most profitable), Gold, Iron, and Lead (customers on whom the company loses money). Research in the banking industry (using a two-tier model) shows that increases in customer satisfaction in the top profitability tier was more effective than increases in the lower tier in increasing future profits (Rust, Zeithaml and Lemon 2000, pp. 199-200; Zeithaml, Rust and Lemon 2001).

The business world routinely discriminates among its customers based on profitability. For example, FedEx instituted an approach that categorized its customers into three groups—the good, the bad, and the ugly (Brooks 1999). The company would work hard to cultivate the "good" customers, and would actively discourage the "ugly" ones. This strategy should work well if profitability is stable, but what if it is not?

The Evolving Customer

Figure 1 shows the profitability path of two customers, culled from the database of an actual business-to-business company. We see that for the first 12 quarters, a period of three years, one customer is much more profitable than the other. The pattern reverses, however, in the next 12 quarters. If the company had based its marketing efforts on current or historical profitability, it would have wrongly identified the customer who was most profitable in the first 12 quarters as being a better target than the other customer. Because the only profits that the company can affect are the ones in the future, what is important is *future profits* rather than current or historical ones. With this in mind, the company would like to classify its customers according to their future profitability, rather than how profitable they are now, or were in the past.

(Insert Figure 1 here)

Customer Lifetime Value and Customer Equity

The attention to future profitability is also the basis for models of customer lifetime value (Berger and Nasr 1998; Jain and Singh 2002,Venkatesan and Kumar 2004, Kumar et al. 2007) and customer equity (Blattberg and Deighton 1996; Rust, Zeithaml and Lemon 2000; Rust, Lemon and Zeithaml 2004). For example, research shows that directing marketing attention to customers with a high estimated customer lifetime value can lead to higher profitability (Venkatesan and Kumar 2004). While Venkatesan and Kumar (2004) find that forward looking metrics, such as customer lifetime value, are better at identifying profitable customers, they do not evaluate the accuracy of the profitability predictions of their model.

Models for predicting customer profitability and customer lifetime value have adopted two approaches—a brand switching approach that uses customer surveys and historical data to project future customer profitability (Rust, Zeithaml and Lemon 2004) and a customer database approach that projects future customer profitability only based on the customer's interactions with the focal firm, with customer spending at competitors unknown (e.g., Fader, Hardie and Lee 2005a and 2005b; Reinartz and Kumar 2003; Rust and Verhoef 2005; Venkatesan, Kumar and Bohling 2007). In this paper, we address the latter case, in which a company possesses a historical customer database, but does not have any information about customer purchases from competitors.

Predicting Future Customer Profitability

Surprisingly, attempts to predict future customer profitability based on information obtained in customer databases have not been entirely successful. Campbell and Frei (2004) show that it is easier to predict future profitability for some customers than for others, even for customers within the same profit tier. In general they find that although using current profitability to predict future profitability explains considerable variance, the practice is problematic, because that method mis-classifies many customers. Malthouse and Blattberg (2005) build a variety of models, based on regression and neural networks, and test their ability to predict future customer profitability. They find that their best models mis-classify most customers predicted to have high profitability.

Donkers, Verhoef and de Jong (2007) build a model of customer retention and cross-buying that they use to project future profitability in a multi-service insurance industry setting. Their results show that the simplest model they tested—maintaining the same profitability over time—performed the best, and better than more sophisticated models, including regression models and models that explicitly modeled customer retention and purchase probability.

Wuebben and Wangenheim (2008) derive similar conclusions. Testing stochastic models against simple management heuristic methods, the researchers find that the simple models most often outperform the more complicated ones, when it comes to predicting repurchase, and whether a customer will be among the company's best customers in the future.

In general the existing literature tends to show that predicting future profitability from customer databases is deceptively difficult, and that simple methods tend to perform as well as more complicated ones. Building a realistic model that accurately predicts future customer profitability is, thus, an open research question, that we attempt to address.

MODEL DEVELOPMENT

The Dynamics of Customer Profitability

Before we describe the model in detail, let us quickly overview the logic of the model. Figure 2 illustrates the dynamics of customer profitability, from period t to period $t+1$. From the standpoint of the firm, what matters is whether or not the customer purchases in period t, and if so, how much profit comes from that. Both purchase and profitability are seen as being driven by marketing actions, customer characteristics (including past purchase history), and control variables such as the state of the economy. Purchase and profit at time t then alter the customer characteristics at time $t+1$, and marketing actions in time $t+1$ are affected by previous marketing actions and customer characteristics. We view these dynamics as occurring with error, necessitating simulation over multiple replications to fully evaluate the average future profitability.

(Insert Figure 2 here)

Future Profitability

In our model we adopt the *always a share* approach for predicting future profitability (Jackson 1985; Rust, Lemon and Zeithaml 2004) because it is more appropriate for the non-contractual setting of our study (Venkatesan and Kumar 2004, Venkatesan et al. 2007, Kumar et al. 2007). The "always a share" approach assumes that there is only dormancy in a customer-firm relationship and that customers never completely terminate their relationship with a firm. This assumption allows for a customer to return to purchasing from a firm after a temporary dormancy and when the customer returns to the relationship they retain the memory about their prior relationship with the firm. Hence, in this approach, we measure future profitability of a customer by predicting their purchase pattern over the prediction period, and do not predict when a customer would terminate their relationship with the firm. We measure future profitability of customer i in terms of 1) total profit (TP_i), and 2) net present value of profit (NPV_i):

$$TP_i = \sum_{t=1}^{K} (GM_t - MC_t)$$

$$NPV_i = \sum_{t=1}^{K} \frac{(GM_t - MC_t)}{(1+r)^t} \tag{1a}$$

where GM_t is gross margin (in \$s) at time t, MC_t is variable marketing cost, and r is the discount rate. The predicted future profits can then be given as:

$$T\hat{P}_i = \sum_{t=1}^{K} (\hat{\Phi}_{it} * \hat{\Pi}_{it} - X_{it} * C)$$

$$N\hat{P}V_i = \sum_{t=1}^{K} \frac{(\hat{\Phi}_{it} * \hat{\Pi}_{it} - X_{it} * C)}{(1+r)^t} \tag{1b}$$

where,

$\hat{\Phi}_{it}$	= predicted probability that customer *i* will purchase in time period *t*,
$\hat{\Pi}_{it}$	= predicted contribution to profit provided by customer *i*, in time period *t*, given purchase,
X_{it}	= marketing contacts directed towards customer *i* in time period *t*,
C	= unit marketing cost of contacting a customer,
K	= number of periods for the profitability calculation

Given information about customer behavior through time period t=0, the profitability measures defined in equation (1b) are comprised of predictions of customer as well as firm behavior for a K-period time horizon. The predictions regarding customer behavior include, (a) the propensity for customer i, to purchase in each future time period t ($\hat{\Phi}_{it}$ in equation 1b), and (b) the profit provided by customer i, given purchase in future time period t, ($\hat{\Pi}_{it}$ in equation 1b). While the predictions of customer behavior capture the revenue aspect, marketing actions originated by the firm ($\hat{X}_{it}$), such as number of sales calls or direct mail sent to a customer, capture the cost aspects, and need to be predicted for accurate CLV measurement (Venkatesan, Kumar and Bohling 2007). We propose a single model framework that provides predictions of (a) customer purchase propensity, (b) customer profit[1], conditional on purchase, and (c) firm marketing actions, and also models the potential correlation among these factors.

[1] For our empirical applications, we deal with *contribution to profit*, rather than *total profit*, as our measure of profitability, to avoid issues of allocating fixed costs to customers, which is often arbitrary, and also assume that contribution to profit is positive (or the sale would not take place). Henceforth we use the terms "profit" and "contribution to profit" interchangeably.

Purchase Propensity

Following the literature in choice modeling, we assume that at each time period, *t*, customer *i* has a latent utility (U_{it}) for purchasing from the firm. This latent utility is assumed to be a function of observed heterogeneity across firms (D_{it}) such as firm size, and industry characteristics, customer characteristics (Y_{it}) such as past purchase behavior, control variables (Z_t), such as overall macro-economic trends (e.g., GDP), and marketing efforts (X_{it}), such as number of contacts by sales personnel.

$$U_{it} = \alpha_i + \underline{I1}_{it}\, \underline{\beta 1} + \mu_{it} + \mu'_{it} \tag{2a}$$

where $\underline{I}_{it}$ is the covariate vector that consists of Y_{it}, Z_t, and X_{it}, $\underline{\beta 1}$ is the corresponding coefficient vector, μ_{it} is the error due to unobserved characteristics, and μ'_{it} is a random normal error variable. We incorporate heterogeneity across firms through the random effect interecept term (α_i). The customer-specific intercept term is modeled as a function of observed firm heterogeneity characteristics ($D1_{it}$);

$$\alpha_i = D1_{it}\, \delta 1 + u_{0i} \tag{2b}$$

where, u_{0i} = is a firm-specific random error term.

Substituting equation (2b) in (2a) we obtain,

$$U_{it} = D1_{it}\, \delta 1 + \underline{I1}_{it}\, \underline{\beta 1} + \mu_{it} + \mu^*_{it} \tag{2c}$$

where, $\mu^*_{it} = u_{0i} + \mu'_{it}$. We accommodate for potential correlation in purchase propensity over years for the same customer through the error term, μ^*_{it}. Let the vector $\mu^*_i = [\mu^*_{i1},.. \ \mu^*_{i12}]'$ represent the error terms corresponding to the customer i over the twelve quarters used in our analysis. We assume that $\mu^*_{i,}$ follows a multivariate normal distribution with zero mean and variance covariance matrix R_μ with a block diagonal structure with blocks corresponding to each individual. The off-diagonal elements in each block of R_μ, $\alpha 1_{jk}$ $(j \neq k)$, capture the correlation in purchase propensity across years for a single customer and the diagonal elements capture the residual variance.

We observe a purchase from the customer if the latent utility (U_{it}) exceeds a threshold, which is normalized to zero. Therefore, the probability of a customer *i* making a purchase in time period *t* is given by,

$$\Pr[P_{it}=1] = \Phi\,(D1_{it}\,\delta 1 + \underline{I1}_{it}\,\underline{\beta 1} + \mu_{it}) \tag{3}$$

where P_{it} is the indicator variable of whether a customer *i* made a purchase in time *t*, and Φ is the cumulative normal distribution.

Profitability

We model the log of the profit at time *t*, from customer *i* given purchase in time *t*, as a linear regression:

$$\log(\pi_{it}|P_{it}=1) = D2_{it}\,\delta 2 + \underline{I2}_{it}\,\underline{\beta 2} + v^*_{it} \tag{4}$$

Similar to the purchase propensity model, firm heterogeneity is captured through $D2_{it}$. The correlation in profits over years from the same customer is captured through the variance-covariance matrix R_v, which has a structure similar to R_μ in equation (2c). We model profitability for only those observations where there is purchase (i.e., $P_{it}=1$), and for the rest, profit is equal to zero. This is also equivalent to modeling profit conditional on a purchase.

Marketing Actions

There is a possibility of endogeneity between marketing actions (X_{it}), and customer behavior, i.e., purchase propensity (P_{it}) and profit (π_{it}) (Manchanda et al. 2004, Shugan 2004). This can be handled using instrumental variable estimation. Specifically, we model X_{it} as a function of its lagged values, and lagged values of customer characteristics (Y_{it}). We use the predicted values of X_{it} as variables in our model. Specifically, we estimate,

$$X_{it} = D3_{it}\,\delta 3 + \underline{I3}_{it}\,\underline{\beta 3} + \varepsilon^*_{it} \tag{5}$$

where $\underline{I3}_{it}$ is a linear combination of past marketing actions (X_{it-1}, X_{it-2}..) and past customer purchase history (Y_{it-1}, Y_{it-2}..), and $D3_{it}$ represents firm heterogeneity, and then substitute $\hat{X}_{it} = D3_{it}\,\delta 3 + I3_{it}\,\beta 3$ for X_{it} in the subsequent equations. Similar to the profit model and the purchase propensity model, the correlation in marketing actions over years from the same customer is captured through the variance-covariance matrix R_ε, which has a structure similar to R_μ in equation (2c) and R_v in equation 5.

The extant scanner panel literature has found that a customer's purchase propensity, purchase quantity and the instrumental variables can be correlated due to unobserved factors that are not included as drivers in the model framework (Chintagunta 1999, Villas Boas and Winner 1999). Not accounting for the correlation among the model components can lead to biased model estimates and poor predictive performance. We therefore model the intercorrelation between the various error terms μ_{it}, υ_{it}, and ξ_{it} from equations 3, 4, and 5 respectively. Our challenge to model this intercorrelation is simplified because all the error terms are assumed to be distributed normal and we can empirically observe υ_{it} (random error in the profitability model) and ε_{tt} (random error in the marketing actions model). Since μ_{it} (error in the purchase propensity model) is unobserved, we assume that μ_{it} is related to υ_{it} and ε_{it} in the following fashion,

$$\mu_{it} = \sum_j \theta_j F_{jit} + \eta_{it}$$

$$= \sum_j \theta_j * \left(\gamma_j \varepsilon_{it} + \alpha_j \upsilon_{it}\right) + \eta_{it}$$

$$\Rightarrow \mu_{it} = \left(\sum_j \theta_j \gamma_j\right)\varepsilon_{it} + \left(\sum_j \theta_j \alpha_i\right)\upsilon_{it} + \eta_{it} \tag{6}$$

where γ_j and α_j are the factor coefficients from a principal component analysis on ε_{it} and υ_{it} respectively, j indexes the principal components, F_{jit} denotes the factor scores, and η_{it} is a random normal error term.

Model Estimation

Because simultaneous estimation of all equations is not feasible, and the structure of the model is recursive, we estimate the model sequentially. Our estimation algorithm proceeds as follows:

Step 1: Estimate the marketing actions component (equation 5) and obtain the corresponding residuals ε_{tt}. (This is possible because the predictors in equation 5 are known.)

Step 2: Estimate the profit component (equation 4) using the predicted value of marketing actions, obtained from Step 1 as an independent variable. Obtain the corresponding residuals υ_{it}. (We use the instrumental variable, $\hat{X}_{it}$, in place of X_{it}. The other predictors are already known.)

Step 3: Factor analyze ε_{it} and υ_{it} and obtain the factor coefficients (denoted as γ_j and α_j respectively) as in equation (6). Set $F_{jit} = \left(\gamma_j \varepsilon_{it} + \alpha_j \upsilon_{it}\right)$.

Step 4: Substitute F_{jit} from step 3 into equation (2c), and estimate β, δ, and θ_j by maximizing the likelihood that results from equation (3). Note that (as shown in equation 6), the product of θ_j and F_{jit} provides μ_{it}.

Our proposed model framework and the corresponding estimation algorithm allow us to obtain predictions necessary for computing expected future customer profitability in a single framework, while also accommodating for the correlation between the various factors used to obtain future profitability. Next, we describe the Monte Carlo simulation algorithm we propose for predicting future customer profitability based on the estimated model coefficients.

Predicting Future Profitability

When predicting future customer profitability, there is significant uncertainty in predicting customer behavior over a multi-period horizon. Typical sources for this uncertainty relate to the poor or non-existent information on (a) customer transactions with the competition, (b) competitor marketing actions targeted at each customer, and (c) customer attitudes, in most CRM databases. Sometimes the cost of increased errors can outweigh the benefits of long term predictions. In our model framework, the uncertainty in future customer behavior is captured by the error variables, ε_{it}, υ_{it}, and μ_{it} corresponding to marketing actions, profit and purchase propensity respectively. The simulation algorithm we propose for predicting future customer profitability accounts for this uncertainty in future customer behavior. The algorithm proceeds as follows:

Step 1: Use the estimation sample, and the estimated model coefficients to obtain the empirical distribution (mean and variance) of the error variables in the marketing actions and profitability components. Let $\mu_{\varepsilon,}$ and R_{ε} represent the mean and variance-covariance matrix of the error variable (ε) in the marketing instruments component (equation 5). Similarly, let μ_{v}, and R_{v} represent the mean and variance of the error variable (υ) in the profit component (equation 4).

Step 2: For each customer i, generate a random variable from the error distribution of the marketing instruments component corresponding to future time period t and replication r, $\varepsilon_{itr} \sim N\left(\mu_{\varepsilon}, R_{\varepsilon}\right)$, and predict the marketing actions for the future time period t, and replication r, ($\hat{X}_{ijr}$), from equation 5.

Step 3: Generate a random variable from the error distribution of the profit component for customer i in time period t and replication r, i.e., $v_{itr} \sim N\left(\mu_{v}, R_{v}\right)$.

Step 4: Use equation 6, and the random variables generated in Step 2 to generate a random error variable for purchase propensity, μ_{itr}. Use equation 2 to predict the latent utility of purchase (U_{itr}), and the corresponding probability of purchase, (Φ_{itr}) from equation 3. Generate a uniform random variable, $u \sim [0,1]$, if u is less than Φ_{itr}, then customer i is predicted to purchase in time period j for replication r, i.e., $\hat{P}_{itr} = 1$.

Step 5: If, $\hat{P}_{itr} = 1$, in Step 4, then predict profit in time t for customer i, for replication r, ($\hat{\Pi}_{itr}$), from equation 4, otherwise set $\hat{\Pi}_{itr}$ to zero. When predicting profit from equation 4, use the predicted marketing action, $\hat{X}_{itr}$, from Step 2, and the random variable from Step 3.

Step 6: Repeat Steps 2-5, until the end of the prediction horizon, i.e., for 12 quarters. Using the predicted values of marketing actions, ($\hat{X}_{itr}$), purchase propensity, (P_{itr}), and profit ($\hat{\Pi}_{itr}$), compute the predicted total profitability (TP_i) and net present value of profitability (NPV_i) for customer i corresponding to replication r. The predicted values of, $\hat{X}_{itr}$, P_{itr}, and $\hat{\Pi}_{itr}$ in time t are used as lagged variables in the predictions for time periods t+1.

Step 7: Repeat Step 6 *R* times, where *R* is the total number of replications. (For our application we set *R* to be equal to 1000.) The expected values of *TP* and *NPV* for customer *i* are then the averages of the values obtained in the R replications.

The multiple replications in the simulation algorithm allow us to accommodate the uncertainty in predicting future customer behavior. Next, we describe the benchmark models we use to compare the predictive accuracy of the proposed framework.

Competing Models

To provide a fair test of our prediction method, we compare the performance of the proposed model to the performance of several models that have been previously proposed.

Current profitability. This very naïve model simply takes the profit from the most recent period for each customer and projects it indefinitely into the future. This model was found by Donkers, Verhoef and de Jong (2007) to be better than some more complicated models in mean absolute deviation of predicted customer lifetime value, and has also been supported by the research of Campbell and Frei (2004).

Average profitability. This model projects future profit as being equal to the historical average for that customer. This is very similar to the current profitability model, but may be more stable if customer profitability fluctuates.

Trend in profitability. This model estimates a simple linear regression on profitability for each customer, with time as the predictor. A variant of this model was tested by Donkers, Verhoef and de Jong (2007), and other regression approaches were tested by Malthouse and Blattberg (2005).

The BG/NBD model. This model, based on a customer dropout and repeat purchase model proposed by Fader, Hardie and Lee (2005a), assumes that purchase is according to the BG/NBD framework, and profitability conditional on purchase remains constant and equal to the customer's historical average. This BG/NBD model is known to produce results similar to the Pareto/NBD model (Fader, Hardie and Lee 2005b) but is much easier to implement (Fader, Hardie and Lee 2005a).

Proposed Model without Simulation. In this model each customer's future profitability is projected using the mean estimates from each equation. (This is equivalent to simulating one replication in which the error terms are all zero.)

Proposed Model without Customer Heterogeneity. One of the basic notions of CRM is that customers are heterogeneous, and accounting for this heterogeneity in customer valuation and in marketing actions directed towards the customers can increase firm profitability (Venkatesan and Kumar 2004). We therefore compare the performance of the proposed model framework without accounting for customer heterogeneity. In this framework, we do not estimate a customer specific intercept parameter that is specified in equations 2c, 3, 4, and 5. Instead we estimate a single intercept parameter that is common to all customers in these equations. As a result the error structure in these equations also does not accommodate for correlation for each customer over the years. For example, μ^*_{it} in equation 2c is equal to μ'_{it} in this model and is distributed normal with zero mean and variance σ_μ^2.

MODEL VALIDATION

Data

For our empirical analysis, we use customer data from a firm that sells a number of high technology products and services to other business customers. The firm's products typically require maintenance and frequent upgrades; these provide the variance necessary for satisfactory model estimation. From a cohort of customers who started purchasing in the first quarter 1997[2], we randomly sampled a set of 191 customers from the database. We use the first 12 quarters of data (1st quarter 1997 to 4th quarter 2000) as the calibration sample to estimate the proposed model framework. We use the next 12 quarters (1st quarter 2001 to 4th quarter 2003) as the holdout sample to evaluate the accuracy of our profitability predictions, since contributions to profit beyond three years in the future do not impact the net present value of profits significantly (Gupta and Lehmann 2005).

Predictor Variables

Based on past research in customer equity and CLV (Reinartz and Kumar 2003; Rust, Lemon and Zeithaml 2004; Venkatesan and Kumar 2004) we identified several antecedents of customer profitability and purchase propensity. The antecedents are classified as customer characteristics, marketing actions and control variables. The customer characteristics variables are further classified as exchange characteristics and customer heterogeneity. The exchange characteristics define and describe the nature of the customer-firm exchange, whereas firmographic variables capture customer heterogeneity. Different exchange characteristics that we include as predictors for purchase propensity (equations 2a-2c and 3) and profitability (equation 4) include past customer spending level, past purchase incidence, cross-buying,

[2] We used a cohort of customers for our analysis to avoid potential left censoring issues.

frequency of past purchase activity, and the marketing actions by the firm. Cross-buying for customer i, in time period t, is measured as the number of different product categories from which a customer purchased in until time t. Past purchase activity measures the level of customer activity in the recent two quarters. Past purchase activity in time period t is set to 2 if a customer purchased in both time period t-2 and t-1, it is set to 1 if the customer purchased in either t-1 or t-2 and is set to 0 otherwise.

The firm contacts customers through several channels, including salesperson, direct mail and telesales. In this study, we define marketing actions as the total number of firm contacts to a customer (across all channels) in a particular quarter. (We aggregate all contacts in order to maintain model parsimony.) Further, our main interest in this study is to provide a framework for accurately predicting future customer profitability while sufficiently accommodating for the endogeneity of marketing actions. Drivers of marketing actions were based both on theory and discussions with the firm regarding how they allocate marketing touches to a customer. The drivers of marketing actions include past customer spending, past levels of marketing contacts, and customer relationship (or exchange) characteristics such as cross buying. The various firmographics included were the sales of an establishment (a measure of the size of the establishment), and an indicator for whether the establishment belonged to B2B industry category.

When assigning predictors we ensured that there is at least one unique predictor for each dependent variable, i.e., marketing touches, purchase propensity and profitability, to ensure model identification (Greene 1993). Next we present results from model estimation.

Model Estimation

The first 12 quarters (1st quarter of 1997 until 4th quarter of 2000) of data are used to estimate the proposed model framework. The results from estimating the proposed framework are provided in Table 1. The table shows the results of estimating the equations for purchase incidence and profit, as well as the instrumental variables regression on marketing actions. The R^2s are equal to 0.76, 0.73, and 0.75 for the purchase profitability, profitability and marketing actions components respectively, indicating good in-sample fit[3]. The error variables from the marketing actions component (ε_{it}), and the profit component (ν_{it}) loaded onto a single factor with factor score of 0.69 for both the variables. Histogram plots indicated that both the marketing actions component error variable, (ε_{it}), and the profit component error variable, (ν_{it}), were normally distributed[4]. The coefficient of F_{1it}, was positive and significant ($\theta_1 = 0.0009$, $p<0.01$), implying a positive correlation among the error variables in the marketing actions, profit and purchase propensity components. We discuss the drivers of the each component below.

(Insert Table 1 here)

Purchase Propensity. The results indicate that customers who (a) have purchased across different product categories (i.e., higher cross-buy), (b) bought in the previous quarter (i.e., lagged indicator of purchase is equal to 1), (c) have exhibited higher levels of past purchase activity, and finally, (d) exhibited growth in their profitability over the previous two quarters, are more likely to purchase in the current quarter. We observe that the drivers included in the marketing actions model are not able to capture the relationship between marketing actions and purchase propensity (i.e., predicted marketing efforts does not have a significant influence).

---

[3] We evaluated the R^2s by comparing the residual variance obtained from the model with the residual variance obtained from a using the mean in the calibration sample across all customers as the null model.
[4] The Anderson-Darling test (Stephens 1974), failed to reject the Null hypothesis that the data is normally distributed at $\alpha = 0.05$.

However, the residual error correlation between marketing actions and profitability has a positive and significant ($\theta l = 0.0009$, $p < 0.01$) relationship with purchase propensity. This implies that our model framework is able to capture the relationship between marketing actions and purchase propensity that is not accounted for by the instrumental variable model.

Firmographics have a significant effect on purchase propensity supporting the need to account for customer heterogeneity in the purchase propensity model. We find that bigger customers or customers who have a more employees and customers who belong to the B2B industry category are more likely to purchase. Macro-economic trends (captured by GDP in this study) did not significantly influence customer purchase propensity. The correlation parameter in the error variance-covariance matrix for purchase propensity (R_{μ}) is significant and positive ($\rho_{\mu} = 0.05$, $\alpha<0.01$). This implies that customers' purchase propensities are positively correlated over the time and provide support for the random effects model structure of the purchase propensity model.

Profitability. Similar to purchase propensity, customers with a higher level of cross-buying and a higher level of past purchase activity are expected to provide a higher profit in the current quarter. Past customer spending levels, captured by, (a) the level of customer spending in the previous two quarters, and (b) the average level of spending in all previous quarters (expect the recent two quarters), positively influence a customer's spending in the current quarter. We also find a significant and positive influence of marketing actions on a customer's current spending level. Similar to purchase propensity, observed customer heterogeneity factors have a significant influence on profitability. We find that customers who have more employees and those who belong to the B2B industry are likely to provide higher profit. We observe that customers are more likely to provide higher revenue when the macro-economic trends are good (i.e., higher GDP). Finally, the correlation parameter in the error variance-covariance matrix for the profitability model (R_v) is positive and significant ($\rho_v = 0.02$, $p<0.01$). This implies that customer profitability is positively correlated over time and provides support for modeling for the random effects model structure of the profitability model.

Marketing Actions. The level of marketing contacts for a customer depends on the recent purchase behavior of the customer, which is captured through the covariates-lagged indicators of purchase and lagged profitability. While a customer's past purchase behavior in general determines whether a customer is contacted or not, the specific level of marketing touches directed towards a customer in a particular month is influenced to a large extent by the level of contacts for the customer in the two prior quarters, and by the average level of contacts allocated to the customer in the past (except the recent two quarters). Finally, the positive influence of cross-buying reflects the firm's strategy of focusing marketing activities on customers who have been actively purchasing products across several categories. The correlation parameter in the error variance-covariance matrix for the marketing actions model (R_ε) is negative and significant ($\rho_\varepsilon = -0.024$, $p<0.01$). The drivers of marketing actions included in the model imply that the firm maintains a high level of contacts for customers who were contacted frequently in the past. The correlation parameter on the other hand shows that the firm is also slightly decreasing the level of contacts for high contact customers (and vice versa) over time based on factors not captured in the model, such as salesperson inputs.

The firm seems to focus its marketing actions on customers who have more employees but does not seem to focus on any one industry category in particular. The results do not provide strong evidence that the marketing actions of the firm are influenced by macro-economic trends. We evaluate the accuracy of predictions from our proposed framework in the next section.

Accuracy of Predictions

Customer profitability is not always stable. Table 2 shows the relationship between past profitability (the estimation sample, periods 1-12) and future profitability (the validation sample, periods 13-24). Customers were classified into three groups (lowest quartile, middle half, highest quartile) by profitability. We can see from this table that customers do evolve over time. For example we see that 7 (15%) customers switched from the lowest profitability quartile to the highest quartile. We also see that 6 (13%) customers switched from the highest profitability quartile to the lowest quartile. Consistent with previous research (Malthouse and Blattberg 2005) there is considerable movement among profitability tiers. It would appear as though accurately predicting future customer profitability would require predicting movement from low profitability to high, and vice versa. An effective model of future customer profitability would be capable of predicting at least some of this movement.

(Insert Table 2 here)

The model estimates obtained from the calibration sample provided input to the simulation algorithm used to predict future customer profitability. Future profitability was estimated based on predicting the marketing actions, purchase propensity and profitability values in the holdout time frame (1st quarter of 2001 until 4th quarter of 2003). We also compared the predictive accuracy of our model to the predictions from the four alternative models.

Table 3 provides the mean absolute deviation (MAD) between the predicted and observed values of profit for the holdout period, for the proposed model and competing models. Table 3 also provides the mean absolute deviation of the net present value of the future profits, based on an annual discount rate of 15%. We see from Table 3 that the proposed model is the most accurate for both predicting future profitability and for predicting net present value of profitability. The proposed model has an MAD of 4.09 (all figures in $100,000s) for future profit, and an MAD of 3.28 for NPV of future profit. The second best model for predicting the future profit was the proposed model without customer heterogeneity (4.20), followed by the BG/NBD model with an MAD of 4.44, the proposed model without simulation (5.19), average profitability (5.23), current profitability (5.54), and trend in profitability (15.73) respectively. The second best model for predicting the NPV of future profit was also the proposed model without heterogeneity (3.40), followed by the BG/NBD model with an MAD of 3.64, the proposed model without simulation (4.05), average profitability (4.14), current profitability (4.34), and trend in profitability (12.91) respectively.

(Insert Table 3 here)

Predictions from the proposed model without the simulation algorithm performed worse than the proposed model with simulation. This indicates that the simulation algorithm is critical for obtaining accurate predictions of future customer profitability. Our model simulates future purchasing scenarios, one customer at a time. By averaging over many possible futures for each customer, the proposed simulation algorithm allows us to obtain a more accurate prediction of future profitability. The proposed model without customer heterogeneity performs better than all the other comparison models but worse than our proposed model that includes customer heterogeneity. This indicates the importance of accounting for customer heterogeneity in customer profitability models.

We treat our forecasting exercise as a within-subjects experimental design where the different models are considered as treatments administered to the same customer. For each model, the absolute deviation between the forecasted and observed future profits, and the absolute deviation between the forecasted and observed NPV of future profit from each customer are considered the response variables. As our focal model is the proposed model (with customer heterogeneity and with simulation) the most relevant comparison is between the proposed model and the best of the competing models. If the best of the competing models were known ahead of time, then it would be appropriate to do a standard pairwise comparison. However if the best of the competing models is *not* known ahead of time, then such a test will be unnecessarily conservative, because the set of competing models has multiple opportunities to capitalize positively on chance. We conducted this (unnecessarily conservative) pairwise comparison test on the differences between the proposed model and the best competing model (the proposed model without customer heterogeneity) and found that the proposed model was better at the .05 level.

(Insert Table 4 here)

Table 4 provides some additional insight into the prediction accuracy of the proposed model. In this table we compare the predicted change in profit to the actual change in profit. For simplicity of exposition, we form three groups of customers: 1) a "decline" group, defined as profitability declining more than 20%, 2) a "stable" group, defined as profitability staying within 20% of the previous value, and 3) a "growth" group, defined as profitability increasing by more than 20%[5]. We can see from this table that among the 56 customers predicted to decline in profitability, 50 (89%) customers were accurately predicted to decline in profitability. Among customers predicted to be stable in profitability, 68 (73%) were accurately predicted, and 38 (76%) were accurately predicted to grow in profitability. Thus, we can see that most predictions were correct, for each predicted category. The off-diagonal elements show customers that were

[5] We also evaluated the model's accuracy by defining more than 10% increase in profitability as "growth", more than 10% decrease in profitability as "decline" and staying within 10% profitability as "stable". The substantive conclusions of the results do not change with this alternative definition.

mis-predicted according to this three-group scheme. For example, 13 (14%) customers were predicted to be stable in profitability, but actually grew.

DISCUSSION

Customer Evolution

Based on the ideas underlying the customer pyramid, many businesses target customers according to their profitability. Typically more profitable customers receive more attention than less profitable customers. The logical extension of this approach is to target customers according to their *future* profitability, or their customer lifetime value. The problem with this approach is that most attempts to predict future profitability have been notably unsuccessful. In fact, very simple models, such as assuming that current profitability levels will continue, have been shown to be just as good as more complicated models.

Thus, obtaining a better understanding of how customers evolve and obtaining accurate predictions of future profitability, is a topic of pressing importance. We need to have a better understanding of when customers will increase in profitability over time, when they will stay more or less the same, and when they will decrease in profitability over time. The purpose of this paper is to provide a model that increases our understanding of how customers evolve over time.

Future Customer Profitability

Our research provides a new approach to predicting future customer profitability that outperforms existing methods. We accomplish this by modeling purchase incidence, and profit conditional on purchase. Given the endogeneity of marketing expenditures, we also estimate future marketing expenditures, by modeling past expenditures using an instrumental variables approach. We also model the joint error structure between the different equations in the model, assuming that the errors are correlated.

A key element of our model is a simulation approach. For each customer we use Monte Carlo simulation to project many possible futures, making random draws on the error terms. This results in many possible profitability trajectories (and total future profits) for each customer. We then average across the replications to obtain the expected future profitability for each customer, for the time horizon desired. The results show that our simulation-based model outperforms the competing models. In particular the future profitability values predicted by our model are more accurate than predictions that result from projection of current profitability, projection of average profitability, use of a linear trend in profitability, or a stochastic model based on the BG/NBD model. Our proposed model also outperforms versions of our model that do not utilize simulation or do not accommodate customer heterogeneity in the parameters. The results imply that an appropriate model specification and prediction algorithm would enable firms to obtain better returns from the investments they make in collecting detailed customer level transaction information.

The simulation algorithm allows us to obtain a distribution of future profitability for each customer, as shown in figure 3 for four randomly selected customers. The distribution of CLV (for these four customers) is not normal and has a long right tail. The CLV distribution highlights the uncertainty in future profitability within customer. The highest density of simulated CLV values (about 47%) for customer A ranges from $100,000 to $300,000. The highest density of simulated CLV values (about 50%) for customer B however ranges from $700,000 to $900,000. The distribution of CLV allows managers to determine a corresponding distribution of marketing resource allocation levels for each customer rather than a point estimate. This would reduce the risk of either under-spending or over-spending on a customer in the future.

(Insert figure 3 here)

Managerial Implications

Our research shows that using a more sophisticated, simulation-based approach provides more accurate predictions of future customer profitability than simple managerial heuristics. The research also shows that our approach is more accurate in predicting future customer profitability than a recently published stochastic modeling approach. The simulation algorithm also provides managers with a range of profitability of each customer that can be used to develop several alternative customer level resource allocation scenarios.

The model presented in this paper is applicable when managers have available a database of customer purchase history and direct marketing actions. In such a case, we conclude that prediction of future profitability is not the futile endeavor that has recently been suggested by several recently-published papers. Rather, significant gains in accuracy can be obtained over simpler approaches, by employing a simulation-based, simultaneous equations approach. The drivers of purchase incidence and profit provide managers early warning indicators of changes in customer profitability. The drivers of also present managers with levers they can influence to improve customer profitability. The accuracy of customer classification (presented in Table 2), implies that a model framework built at the customer level is also appropriate for making marketing decisions at the segment level.

Although our research focuses on predicting total profitability over a fixed time horizon, the model could also be used to predict the net present value of future customer profitability over a given time horizon. By increasing the length of the time horizon, the model can also be used to estimate the customer lifetime value, given that there is almost no contribution to customer lifetime value after about 15-20 years, given discount rates typically used by managers (Gupta and Lehmann 2005).

Limitations

Several limitations of the current research should be noted. First, our conclusions are based on the analysis of only one large-scale, industrial customer database from one particular industry. Replication across multiple databases, multiple companies, and multiple industries would be necessary to fully establish the model's general applicability.

Second, an endogeneity issue inevitably remains. That is, if a company predicts future profitability for a customer, it is also, implicitly, assuming that its own marketing actions will be determined as in the past (that is what the instrumental variables prediction of marketing actions implies). But if the company then uses the future customer profitability information to change its marketing actions, that will itself further change the customer profitability. The result is that increased attention to future customer profitability probably means that customers projected to be profitable may become even more profitable than projected. Likewise, customers projected to be unprofitable may become even more unprofitable than projected.

Third, our model is only applicable to companies that have a historical customer database that includes past purchase history, profitability (or contribution to profitability), and a history of direct marketing actions. There are many companies, especially consumer packaged goods companies that do not have such a database.

Directions for Future Research

Although we have provided a model that outperforms simple managerial heuristic models and a prominent stochastic model in predicting future customer profitability, we are well aware that further improvement is possible. This improvement can take several forms. The individual components of our model—purchase incidence, profit conditional on purchase, marketing actions, and the simulation of future customer profitability paths—all are potential opportunities for improvement. Another possible way to improve the model is to build a unified model that solves the endogeneity of marketing actions issue—somehow simultaneously building the altered marketing actions into the customer profitability projection itself. Finally there is the opportunity to validate the model across a variety of companies, countries, and industries.

Conclusions

Predicting future customer profitability is not futile, but requires a sophisticated modeling approach. Our proposed model, which includes simultaneous equations and a Monte Carlo simulation approach, outperformed the most notable competing models in a large-scale empirical test. We conclude that predicting future customer profitability is possible, and can be used to drive customer-specific marketing actions. We may be able to predict systematically whether some of our "frogs" will become "princes" after all.

Figure 1

The Evolving Customer

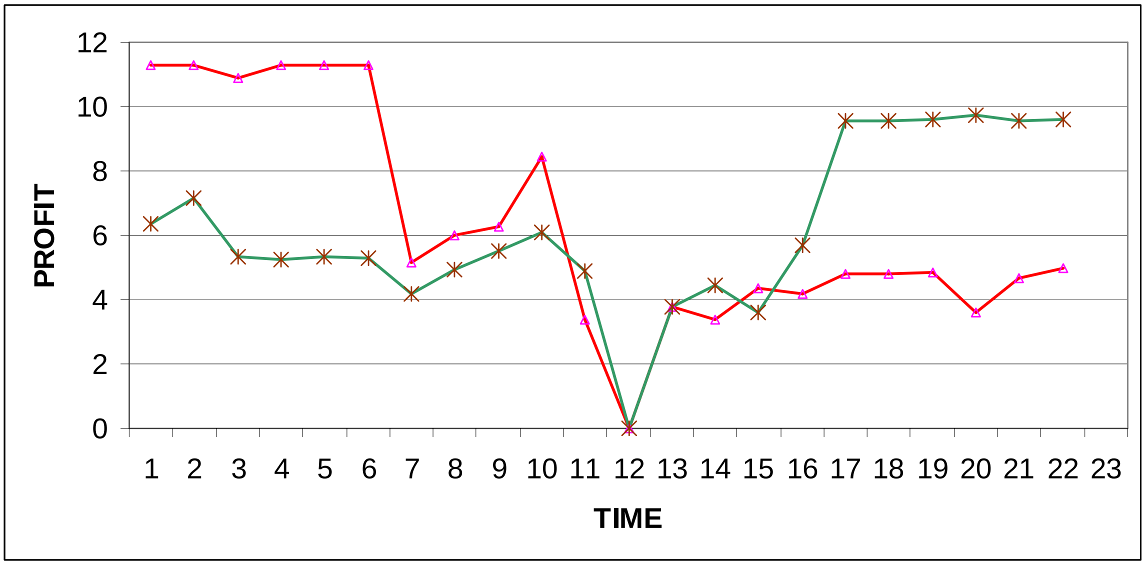

Figure 2

Dynamics of Customer Profitability

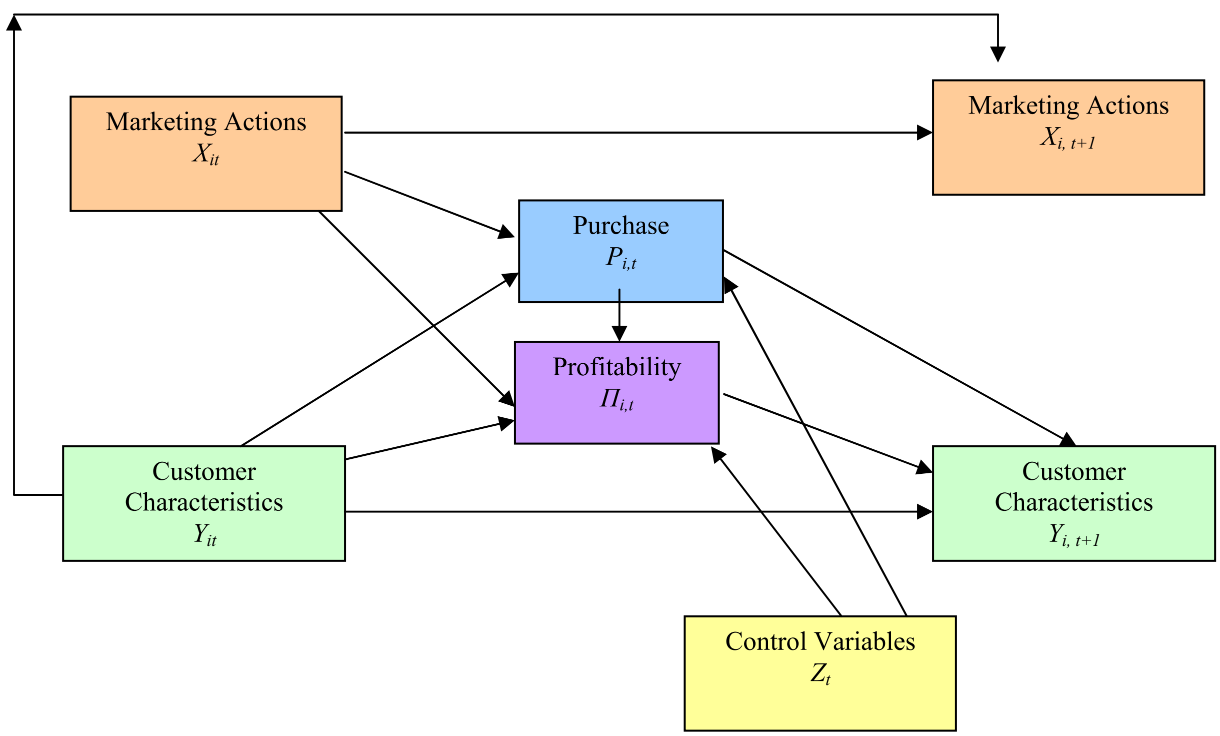

Figure 3

Predicted Distribution of Future Profitability

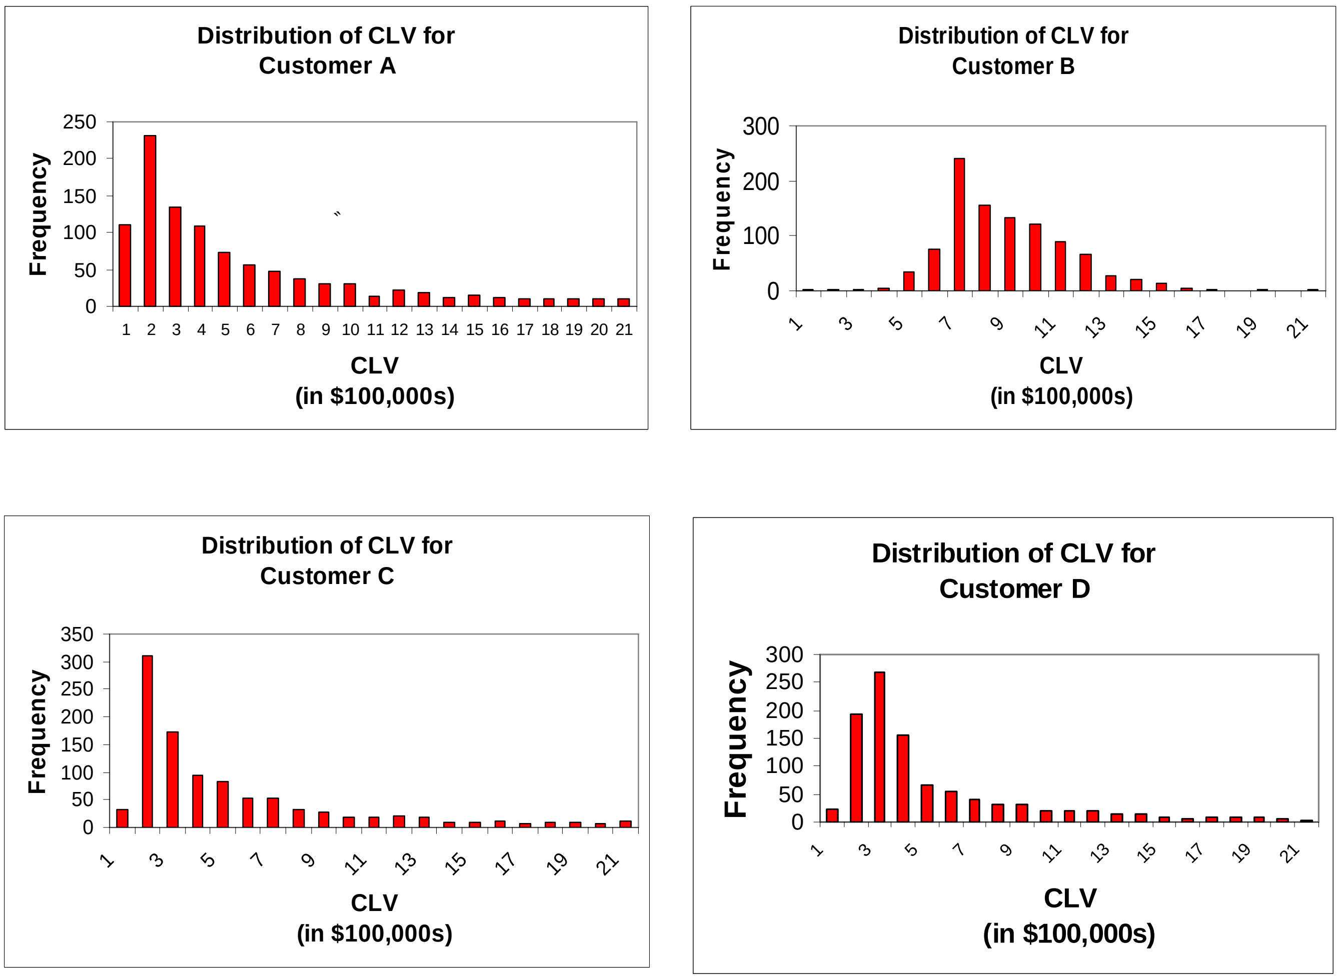

Table 1
Parameter Estimates for Proposed Model

Purchase Propensity		Profitability		Marketing Actions	
R^2= .75		R^2 = 0.72		R^2= 0.73	
Variable	**Parameter Estimate**	**Variable**	**Parameter Estimate**	**Variable**	**Parameter Estimate**
Intercept	-1.2***	Intercept	-16528	Intercept	0.99***
Cross buying	0.20***	Lagged Profit	0.387***	Log of Lagged Total Touches	0.217***
Lagged Indicator of purchase	1.55***	Predicted Total Touches	1107**	Lagged Profit	2.2e-6***
Lagged Activity	0.18**	Cross buying	2180***	Lagged Average Touches	1.91***
Lagged Growth in Profit	0.024**	Two Period Lagged Profit	0.23***	Cross buying	0.203***
Number of Employees	0.05***	Lagged Average Profit	0.26***	Two Period Lagged Total Touches	0.11***
B2B	0.18***	Lagged Growth in Profit	0.02**	Number of Employees	0.06***
GDP	0.21***	Number of Employees	1357***	GDP	0.04*
θ_l	0.0009***	GDP	0.15***		
$\rho_{R\mu}$		$\rho_{R\nu}$	3.6e-7***	$\rho_{R\varepsilon}$	-0.57**
σ^2_μ		σ^2_ν	1.067e10***	σ^2_ε	24.55***

- Significant at α = 0.10, ** Significant at α = 0.05, *** Significant at α = .01

Table 2

Changes in Profitability Over Time

Future Profitability (Periods 13 - 24)

Past Profitability (Periods 1 - 12)	Low[1]	Medium[2]	High[3]
Low[1]	23	17	7
Medium[2]	32	47	18
High[3]	6	16	25

[1]Lowest quartile
[2]Middle two quartiles
[3]Highest quartile

Table 3

Model Comparison – Predicting Future Customer Profitability

(Reported values are in $100,000s)

Mean Absolute Deviation

Model	Future Profit	NPV of Future Profit
Current Profitability	5.54	4.34
Average Profitability	5.23	4.14
Trend in Profitability	15.73	12.91
BG/NBD	4.44	3.64
Proposed Model without Simulation	5.19	4.05
Proposed Model without Customer Heterogeneity	4.20	3.40
Proposed Model	4.09	3.28

Table 4

Prediction Accuracy – Changes in Customer Profitability

Actual Change in Profitability

Predicted Change in Profitability	Decline	Stable	Growth
Decline[a]	50 (89%)	5 (9%)	1 (2%)
Stable[b]	12 (13%)	68 (73%)	13 (14%)
Growth[c]	2 (4%)	10 (20%)	38 (76%)

[a] Decline defined as >20% decline in total profitability
[b] Stable defined as <= 20% change in total profitability
[c] Growth defined as >20% growth in total profitability

REFERENCES

Berger, Paul D. and Nada I. Nasr (1998), "Customer Lifetime Value: Marketing Models and Applications," *Journal of Interactive Marketing*, 12 (Winter), 17-30.

Blattberg, Robert C. and John Deighton (1996), "Manage Marketing by the Customer Equity Test," *Harvard Business Review*, 74 (July-August), 136-144.

Bolton, Ruth N., Katherine N. Lemon, and Peter C. Verhoef (2004), "The Theoretical Underpinnings of Customer Asset Management: A Framework and Propositions for Future Research," *Journal of the Academy of Marketing Science*, Vol. 32, No. 3, 271-292

Bowman, Douglas, and Das Narayandas (2004). "Linking Customer Management Effort to Customer Profitability in Industrial Markets," *Journal of Marketing Research* 41 (4), 433-447.

Brooks, Rick (1999), "Alienating Customers Isn't Always a Bad Idea, Many Firms Discover," *Wall Street Journal*, (January 7), A1, A12.

Campbell, Dennis and Frances Frei (2004), "The Persistence of Customer Profitability: Empirical Evidence and Implications From a Financial Service Firm," *Journal of Service Research*, 7 (November), 107-123.

Chintagunta, Pradeep (1999), "Variety Seeking, Purchase Timing, and the Lightning Bolt Brand Choice Model," *Management Science*, 45(4), 486-498.

Donkers, Bas, Peter C. Verhoef and Martijn de Jong (2007), "Modeling CLV: A Test of Competing Models in the Insurance Industry," *Quantitative Marketing and Economics*, forthcoming.

Fader, Peter S., Bruce G.S. Hardie, and Ka Lok Lee (2005a), "Counting Your Customers the Easy Way: An Alternative to the Pareto/NBD Model," *Marketing Science*, 24 (2), 275-284.

Fader, Peter S., Bruce G.S. Hardie, and Ka Lok Lee (2005b), "RFM and CLV: Using Iso-Value Curves for Customer Base Analysis," *Journal of Marketing Research*, 42 (November), 415-430.

Greene, William H. (1993), *Econometric Analysis*, 3rd Edition, Prentice Hall, Inc., Upper Saddle River, NJ.

Gupta, Sunil and Donald Lehmann (2005), *Managing Customers as Investments: The Strategic Value of Customers in the Long Run***,** Wharton School Publishing, Upper Saddle River, NJ

Jackson, Barbara B. (1985), *Winning and Keeping Industrial Customers*, Lexington, MA: D.C. Heath and Company.

Jain, Dipak and Siddhartha S. Singh (2002), "Customer Lifetime Value Research in Marketing: A Review and Future Directions," *Journal of Interactive Marketing*, 16 (2), 34-46.

Kumar, V., Rajkumar Venkatesan, Timothy Bohling, and Dennis Beckmann (2007), "Managing Customer Lifetime Value at IBM," *Marketing Science*, Forthcoming.

Malthouse, Edward C. and Robert C. Blattberg (2005), "Can We Predict Customer Lifetime Value?" *Journal of Interactive Marketing*, 19 (1), 2-16.

Manchanda, Puneet, Peter E. Rossi and Pradeep K. Chintagunta (2004), "Response Modeling with Nonrandom Marketing-Mix Variables," *Journal of Marketing Research*, 41 (Nov), 467-478.

Mulhern, Francis J. (1999), "Customer Profitability Analysis: Measurement, Concentration, and Research Directions," *Journal of Interactive Marketing*, 13 (Winter), 25-40.

Reinartz, Werner J. and V. Kumar (2003), "The Impact of Customer Relationship Characteristics on Profitable Lifetime Duration," *Journal of Marketing*, 67 (January), 77-99.

Rust, Roland T., Katherine N. Lemon, and Valarie A. Zeithaml (2004), "Return on Marketing: Using Customer Equity to Focus Marketing Strategy," *Journal of Marketing*, 68 (Jan.), 109-27.

____________, and Peter C. Verhoef (2005), "Optimizing the Marketing Interventions Mix in Intermediate-Term CRM," *Marketing Science*, 24 (3), 477-489.

____________Valarie A. Zeithaml and Katherine N. Lemon (2000), *Driving Customer Equity: How Customer Lifetime Value is Reshaping Corporate Strategy*, New York: The Free Press.

Shugan, Steven M. (2004), "Endogeneity in Marketing Decision Models," *Marketing Science*, 23 (1), 1-3.

Stephens, M. A. (1974). *EDF Statistics for Goodness of Fit and Some Comparisons*, Journal of the American Statistical Association, Vol. 69, pp. 730-737.

Venkatesan, Rajkumar and V. Kumar, (2004), "A Customer Lifetime Value Based Framework for Customer Selection and Resource Allocation Strategy," *Journal of Marketing*, 68 (October), 106-25.

____________, V. Kumar, and Timothy Bohling (2007), "Optimal CRM Using Bayesian Decision Theory: An Application to Customer Selection," *Journal of Marketing Research*, forthcoming.

Villas-Boas, Miguel J. and Russell S. Winer (1999), "Endogeneity in Brand Choice Models," *Management Science*, 45 (10), 1324–38.

Wuebben, Markus and Florian Wangenheim (2008), "Instant Customer Base Analysis—Managerial Heuristics Often 'Get it Right'," *Journal of Marketing*, forthcoming.

Zeithaml, Valarie A., Roland T. Rust and Katherine N. Lemon (2001), "The Customer Pyramid: Creating and Serving Profitable Customers," *California Management Review*, 43 (4), 118-142.