

Driving Profitability by Encouraging Customer Referrals: Who, When, and How

Many firms are now using referral marketing campaigns to harness the power of word of mouth and to increase referrals to acquire new customers. Prior research has identified a method of computing the value of referrals using only a customer's actual past referral behavior to compute customer referral value (CRV). In this article, the authors develop and test a new four-step approach to compute CRV. In addition, they determine the behavioral drivers of CRV and then identify the most effective methods of targeting the most promising customers on the basis of their customer lifetime value (CLV) and CRV scores. The authors illustrate and test this approach through four separate field experiments with firms from two industries: financial services and retailing. They find that to maximize profitability, it is critical to manage customers in terms of both their CLV and CRV scores and that understanding the behavioral drivers of CRV can help managers better target the most profitable customers with their referral marketing campaigns.

Keywords: customer referral value, field studies, Bayesian Tobit model, customer lifetime value

Consumers frequently rely on word of mouth and referrals (active) and observational learning (passive) from other consumers when making a purchase decision (Dichter 1966). In some cases, the referral has a positive impact on their purchase decision, and in other cases, it does not alter their decision. In this article, we define a referral as when a new customer enters into a transaction with a firm and attributes the motivation for the transaction to a current customer. Research has shown that customers spread word of mouth and generate customer referrals in various situations, including during new product diffusion (Mahajan, Muller, and Bass 1995) or when making general purchase decisions, in which they often seek others' opinions (Feick and Price 1987).

Thus, we know that word of mouth and referral behavior play a role in consumer decision making and purchase behavior. The challenge that practitioners face is finding a way to use the information about word of mouth and refer-

ral behavior when managing customers. Although the existence of word of mouth and referral behavior has been studied in the marketing literature, only a few studies have examined the monetary value of word of mouth and referral behavior or how information about word of mouth or referral behavior can be used to manage customers effectively. This is important because, as Rust, Zeithaml, and Lemon (2000, p. 46) note, "the effect [of word of mouth] is notoriously hard to measure, but is frequently significantly large." Research has also shown that firms that acquire customers through referrals are acquiring more profitable customers, which makes it desirable to run referral marketing campaigns (Villanueva, Yoo, and Hanssens 2008).

Several studies in the marketing literature have shown that measuring the value of word of mouth or referrals can provide several key implications for managing customers and that referral marketing campaigns can be highly valuable. However, these studies do not help marketing managers with the selection of the "best" customers for these referral marketing campaigns. We summarize the key contributions and implications from these prior studies in comparison with the contributions and implications of the current study in Table 1.

During the past few years, many managers have adopted metrics such as "Net Promoter score" (Reichheld 2003) to strategically manage firm revenue growth. However, as Keiningham and colleagues (2007) empirically show, the Net Promoter score shows only a weak link to a firm's revenue growth, leading to an ongoing debate over the issues with measuring and managing metrics that predict firm and customer growth.¹ Given this finding and the interest in

V. Kumar is Richard and Susan Lenny Distinguished Chair in Marketing, Executive Director of the Center for Excellence in Brand & Customer Management, and Director of the PhD program in Marketing, J. Mack Robinson School of Business, Georgia State University (e-mail: mktvk@langate.gsu.edu). J. Andrew Petersen is Assistant Professor of Marketing and Assistant Director for the Center of Integrated Marketing and Sales (CIMS), Kenan-Flagler Business School, University of North Carolina at Chapel Hill (e-mail: Andrew_Petersen@unc.edu). Robert P. Leone is a professor and J. Vaughn and Evelyne H. Wilson Chair of Marketing, Neeley School of Business, Texas Christian University (e-mail: r.leone@tcu.edu). The authors thank Wagner Kamakura, Roland Rust, Ruth Bolton, Robin Coulter, Susan Spiggle, Jeff Inman, Bob Burnkrant, Barak Libai, Carl Mela, Linda Price, Sunil Gupta, Peter Verhoef, John Mathieu, and Kay Lemon for their comments on a previous version of this article. Special thanks are also owed to a retail firm and a financial services firm for sharing their data. Marnik Dekimpe served as guest editor for this article.

¹See the *Journal of Marketing* blog at http://www.marketingpower2.com/blog/journalofmarketing/2007/07/a_longitudinal_examination_of.html (accessed May 13, 2010).

TABLE 1
Contributions of the Current Study Relative to Prior Studies

Criteria	Helm (2003)	Hogan, Lemon, and Libai (2003)	Hogan, Lemon, and Libai (2004)	Ryu and Feick (2007)	Kumar, Petersen, and Leone (2007)	This Study
1. Computation of referral value	N.A.	Value lost from customer churn	Word-of-mouth multiplier from advertising	Referral likelihood based on rewards	Actual past referral behavior	Predicted future referral behavior
2. Key objective of the article	Reviewing three approaches to computing referral value	A method to account for social effects in the management of customers	Quantifying the advertising ripple effect	Determine the effectiveness of rewards for referral marketing	Comparison of customer segmentation based on CLV and CRV	Optimal customer targeting for referral marketing campaigns
3. Contribution to marketing	An extensive literature review	Customers cannot be valued in isolation because of word-of-mouth effects	Managers should account for customer word-of-mouth interactions when advertising	Shows how extrinsic motivations affect customer referral behavior	Computing past referral value of each customer	Determining behavioral drivers of referral value, which can help in targeting customers for referral marketing campaigns
4. Field experiment(s) and/or empirical application	N.A.	A simulation exercise using data collected about the online banking market	A survey of students based on their hair salon choice and the amount of word of mouth they engaged in following the last haircut	A series of four laboratory experiments that investigate the effects of different types of rewards on different types of customers	Simple field experiment with a telecommunications firm based on computed CLV and CRV in which all customers were targeted with an incentive to make referrals	A series of field experiments with a financial services firm that uncover the most effective way to target customers for referral marketing and cross-sell/up-sell campaigns and a driver-based field experiment with a retailing firm that shows the effectiveness of using the drivers of CRV as a method for customer selection and resource allocation
5. Key learning	The three models discussed offer insights into the determinants and effects of word of mouth	The value of a lost customer changes as a function of the time in the product life cycle	The word of mouth generated after an advertising-motivated purchase can be significant	Offering a reward increases referral likelihood, but the size of the reward is not significant	Customers with high CLV are not the same as customers with high CRV	Firms can effectively select customers within segments of high and low CLV/CRV for referral marketing campaigns using the drivers of CRV (dynamic targeting)

Notes: N.A. = not applicable.

developing a better understanding of these issues, it is important for research in marketing to address the following three research questions:

- Q₁: Is it possible to predict each customer's indirect impact on the firm's future profits (i.e., through referral behavior)?
- Q₂: Which current customers are likely to add the most value to the firm through referral activities?
- Q₃: How can managers design marketing campaigns that most effectively select customers for referral marketing campaigns?

We answer Q₁ in the following way: We measure customer lifetime value (CLV) and customer referral value (CRV) of each customer. To predict CLV, we used a method already established in the marketing literature (Kumar et al. 2008). We then develop a new four-step approach for predicting CRV with information from a customer's past referral behavior. We show that this new approach offers accurate predictions of a customer's future referral behavior. We then use each customer's predicted CLV and CRV scores in three field experiments with a financial services firm. The results of these field experiments show the following:

- Customers with a high CLV are not necessarily the same customers as those with a high CRV. Therefore, a customer must be managed using both CLV and CRV.
- It is important not to saturate each customer with multiple marketing messages but rather with messages that have the highest likelihood of increasing CLV or CRV.

Next, we empirically determine the behavioral drivers of predicted CRV using a retail firm's customer database to answer Q₂, which includes each customer's transaction and referral behavior. Then, we run a fourth field experiment with the same retailing firm to answer Q₃. The insights gained by answering Q₁–Q₃ can help firms target customers who are most likely to respond positively to a referral marketing campaign. We also demonstrate the positive financial consequences of adopting the proposed process in creating and targeting an effective marketing strategy.

Measuring CLV and CRV

Customer Data

We collected the data used in this study for measuring CLV and CRV between July 1, 2001, and June 30, 2005 (four years) for 14,160 customers on a semiannual basis from a financial services firm that ranks high among the global *Fortune* 1000 firms. This financial services firm has a wide range of products and services, including banking, insurance, and investments. We provide a list of products this financial institution provides along with the average profit per month for each product for each customer in Table 2. As Table 2 shows, the average monthly profit for each product ranges from approximately \$13 (current account) to approximately \$92 (home loan).

Additional customer data from the financial services firm include data collected from those same 14,160 customers during the three field experiments. We conducted these experiments over a one-and-a-half-year period (July 1, 2005–December 31, 2006) immediately following the

TABLE 2
Products Currently Offered by the Focal Financial Services Firm

Product Category	Average Monthly Profit (SD)
1. Current	13.20 (25.2)
2. Saving	28.09 (36.8)
3. Deposit	42.77 (77.8)
4. Invest	35.90 (40.1)
5. Insurance	28.70 (35.6)
6. Overdraft	48.70 (63.4)
7. Personal loan	40.40 (59.6)
8. Home loan	92.72 (118.1)
9. Vehicle loan	31.06 (52.5)
10. Cards	25.78 (41.4)

Notes: Average monthly profit across product categories per customer.

end of the initial data collection. The data included each customer's own purchase behavior, the estimated marketing costs associated with attracting and retaining each customer, the actual referral behavior of each customer, and the transaction data for any new customer who was referred to the financial services firm by one of the original 14,160 customers in the sample and purchased at least one product. The data set also included the findings from a survey given to each of the 14,160 financial services firm's customers in July 2004 that provided information on their "willingness to refer" new customers. We used all the initial customer data from the first four years for three purposes: (1) to calibrate and test the predictive accuracy of the computed CLV values, (2) to calibrate and test the predictive accuracy of the computed CRV values, and (3) to empirically test whether there is a link between "willingness to refer" and actual referral behavior. We then used each customer's CLV and CRV scores as a basis for customer selection in the first three field experiments.

Computing CLV

To measure the CLV of each customer, we used Kumar and colleagues' (2008) model, which calculates the expected future profitability of a customer purely on the basis of behavioral and demographic customer information. This is given in Equation 1:

$$(1) \quad CLV_i = \sum_{j=T+1}^{T+N} \frac{p(\text{Buy}_{ij} = 1) \times \hat{CM}_{ij}}{(1+r)^{j-T}} - \frac{\hat{MC}_{ij}}{(1+r)^{j-T}},$$

where

- CLV_i = lifetime value for customer i,
- p(Buy_{ij}) = predicted probability that customer i will purchase in period j,
- CM_{ij} = predicted contribution margin provided by customer i in period j,
- MC_{ij} = predicted marketing costs directed toward customer i in period j,
- j = index for time periods (semiannual in this case),
- T = the end of the calibration or observation time frame,
- N = total number of prediction periods, and

r = semiannual discount factor (.07238 in this case; this amounts to a 15% annual rate).

Prior research has empirically demonstrated that this model accurately measures each customer's CLV; the model can also be used to maximize CLV by optimizing all the marketing resource variables (e.g., e-mail, direct mail, face-to-face costs and frequencies) that managers can control in terms of marketing expenses. For complete details of the procedure for estimating the model, see Kumar and colleagues (2008).

We tested the predictive accuracy of this model on the customers from the financial services data set. We used the first three years of data to determine the drivers that help predict the CLV for each customer. We then used this information to predict the CLV of each customer for the fourth year, which represents the holdout period. We then compared this predicted value with the actual discounted profit calculated for each customer during the one-year holdout period. The absolute value of the difference between the actual CLV and the predicted CLV for each customer from the financial services firm was just under 10% (mean absolute percentage error [MAPE]), or approximately \$27 (mean absolute difference [MAD]) per customer (average CLV was approximately \$274). This provides strong evidence that the CLV model can be used to accurately predict each customer's future value to the firm.

Computing CRV

We measured CRV for each customer using the objective function from Kumar, Petersen, and Leone (2007) given in Equation 2. We advance the measurement of CRV from the previous study by introducing a new four-step approach for predicting each customer's CRV using past customer transaction and referral data. This is a key advancement of and improvement over Kumar, Petersen, and Leone's model because they simply use the previous period's actual referral behavior value, and in the current research, we predict referral behavior into the future using past customer referral data.

$$(2) \quad CRV_i = \sum_{t=1}^T \sum_{y=1}^{n1} \frac{A_{ty} - a_{ty} - M_{ty} + ACQ1_{ty}}{(1+r)^t} + \sum_{t=1}^T \sum_{y=n1+1}^{n2} \frac{ACQ2_{ty}}{(1+r)^t},$$

where

T = the number of periods that will be predicted into the future (e.g., years),

A_{ty} = the gross margin contributed by customer y who otherwise would not have bought the product,

a_{ty} = the cost of the referral for customer y ,

$n1$ = the number of customers who would not join without the referral,

$n2 - n1$ = the number of customers who would have joined anyway,

M_{ty} = the marketing costs needed to retain the referred customers,

r = semiannual discount factor (.07238 in this case; this amounts to a 15% annual rate),

$ACQ1_{ty}$ = the savings in acquisition cost from customers who would not join without the referral, and

$ACQ2_{ty}$ = the savings in acquisition cost from customers who would have joined anyway.

To predict CRV, it is important to consider four aspects of the referral process. First, the timing of the referral is important to determine the proper discounting. Second, knowing how many referrals will be made each year and how each of these referrals will contribute to the gross margin of the firm is critical. Third, if a firm rewards referral behavior (i.e., offers any incentive to the customer making the referral), the cost of that incentive must be subtracted from any value attached to the referring customer. Fourth, if a new customer comes directly from the referral process, the money the company saves in the acquisition cost must be considered. These acquisition cost savings can be applied to two groups of people: customers who would have joined anyway at some time in the future and those who joined only because of the referral. For customers acquired only as a result of the referral, it is important to count the entire transaction value of the referred customer because the company would not have realized this sale otherwise. However, in the case of a customer who would have joined anyway, only the savings in acquisition cost is counted, even if the customer was acquired through a referral. When all these components are considered, a customer's referral value to the firm can be calculated as in Equation 2. We provide a sample calculation of CRV for a single customer in Web Appendix W1 (<http://www.marketingpower.com/jmsept10>).

Rather than using a measure of willingness or intention to refer, we use the past referrals and the corresponding referral values to predict the number of future referrals and the value of those referrals. We believe that this is a better approach because Keiningham and colleagues (2007) find that stated intentions of future referrals are poorly linked to revenue growth and Kumar, Petersen, and Leone (2007) find that stated willingness to refer is poorly linked to actual referral behavior. We also found supporting evidence about the gap between stated intentions to refer and actual referral behavior for the data in this study. The focal firm in this study conducted a survey of the 14,160 customers in the sample to better understand what mechanisms were driving the growth of new customers from referrals. In July 2004, the entire sample from the financial services firm (14,160 customers) was asked whether they were willing to recommend new customers to the firm. We analyzed the data from July 2004 to June 2005 for each customer's actual referral behavior, along with the purchase behavior and the profitability of the referred customers.²

The findings from the analysis show that there is a definite gap between a customer's stated willingness to refer

²The customer's referral behavior was tracked electronically, which enabled us to observe both attempts at making referrals when e-mails were forwarded to potential customers and whether that referred customer actually purchased.

and his or her actual referral behavior. For example, 70% (9492) of customers stated that they intended to make referrals. However, only 44% ($n = 4204$) of those who stated this actually made attempts at referrals. In addition, only 42% ($n = 1763$) of these attempts at referrals actually resulted in prospects making a purchase (i.e., a successful referral). Finally, only 77% ($n = 1356$) of customers who made successful referrals brought in new customers who were profitable. These findings show that using a measure such as willingness to recommend (Reichheld 2003) exaggerates what would actually happen and thus provides misleading information for use when managing customers for profitability.

Determining an Individual Customer's CRV

To predict the value of a customer's referrals, we must understand how the flow of information moves from customers to potential prospects. Thus, we propose a four-step approach to determine each customer's CRV that will make it more straightforward to implement within different organizations. First, we determine whether the referred customer would have bought the product or service without a referral. Second, we predict the average future value of each referred customer. Third, we predict the future number of referrals each customer will make in a given period. Fourth, we predict the timing of the predicted future referrals.

Step 1: Determine whether a customer would have bought anyway. To determine whether the referred customer would have purchased without the referral, the firm asked each newly acquired customer the following question: "How likely is it that you would have subscribed to this service in the next 12 months without a referral?"³ This question used a five-point scale (1 = "definitely no," and 5 = "definitely yes"). The purpose of using a five-point scale was to determine whether there was any incremental difference between the strength of the referral as the responses moved from 1 to 5. That was not the case, because the responses were bimodal, with the modes being 1 and 5. Therefore, we mapped the results from the five-point scale to a binary (0/1) scale for the analysis.⁴ A different approach could be used if there were a situation in which the distribution of responses was not bimodal.⁵

Step 2: Predict the future value of each referred customer. We used the following approach to predict the future

³Customers are being asked to think about whether the purchase they made was either an accelerated purchase (i.e., would have joined anyway) or an unplanned purchase (i.e., would not have joined without the referral). In other words, this is a recall measure in which customers are reflecting on a recently made purchase. We believe that this recall measure is more accurate than the purchase intention measure.

⁴If the customer responded with a 4 or 5, we considered him or her a customer who would have joined without a referral. Alternatively, if the customer stated that he or she became a customer only because of the referral (1, 2, or 3), we estimated his or her future value and included this as part of the CRV for the referring customer.

⁵In certain instances, the probabilistic method would be a better choice for determining which referred customers would have bought anyway. These situations could include where the distribution of responses to the survey question is more uniform (versus

value of each referred customer: First, using the three years of data for each customer (July 2001–June 2004), we computed the average semiannual contribution margin from referrals for each customer. This was done by examining each customer's referrals and computing the average gross profits that each referral brought to the firm in each period during the three years. To check the validity of using this as a measure of future gross profits (A_{ty}), we used the simple six-period moving average values for each referral per period in the holdout year (July 2004–June 2005).⁶ Second, we used the past marketing costs (M_{ty}) to predict the average cost of retaining each customer's referral. We did this by examining the past marketing expenses for each period and projecting that average marketing cost into the future. Although we used the averages of the past values to predict the future values of contribution margin and marketing costs for referrals, we found a high correlation between the actual and predicted values of A_{ty} and M_{ty} . The correlation coefficient was .79 between the actual and predicted A_{ty} and .83 between the actual and predicted M_{ty} across all customers. In terms of predictive accuracy, the MAPE between the actual and predicted values is 14.2% for A_{ty} and 11.1% for M_{ty} . This provides further evidence that this method can accurately predict contribution margin and marketing costs for customers.

Step 3: Predict the number of referrals. We used a similar method to predict the future number of referrals (n) each customer was likely to make. This number is based on using the average number of referrals per year across the first three years for each customer (July 2001–June 2004) and projecting that number into the holdout year (July 2004–June 2005).⁷ We found that the average number of past referrals was also a good predictor of the actual number of referrals a customer made in the future. When we computed the average number of referrals on the basis of past referral behavior, we rounded it off to the nearest integer. The correlation between the actual and predicted number of referrals is .80. In terms of predictive accuracy, the MAPE between the actual and predicted number of referrals is 17%, which shows that this method is accurate in predicting the number of referrals a customer will make in the future.

Step 4: Predict the timing of customer referrals. The timing of the referrals was based on the average timing of past referrals by each customer. For example, because the

bimodal), where the prediction of number of referrals is not merely based on a past average (e.g., number of referrals is modeled as a function of a set of variables), or in an environment that is less stable than this current case (e.g., high-growth industries). In this case, a firm could predict the probability of each newly referred customer joining anyway (p) and then compute the probability of that customer joining only because of the referral ($1 - p$). Then, CRV could be computed in the following manner: For each newly referred customer, the acquisition cost savings would be added to $(1 - p) \times$ (expected value of the newly referred customer).

⁶The one-year prediction of A_{ty} starting in semiannual period 7 ($A_{7y} + A_{8y}$) would be computed in the following manner: $A_{7y} = \text{avg}(A_{1y}, \dots, A_{6y})$ and $A_{8y} = \text{avg}(A_{2y}, \dots, A_{7y})$.

⁷We did this using a three-period simple moving average because the data on number of referrals were at a yearly level.

financial services firm collected data semiannually, we were able to split the prediction of timing of referrals into two six-month periods, July–December (“the first half”) and January–June (“the second half”). If the customer averaged six total referrals per year and all the referrals came in the first six-month period, in the CRV prediction, all the referrals would be credited to that specific period. Splitting up the timing of referrals into two periods is important because it can have a dramatic impact on the value of each customer’s CRV. The earlier a customer makes a referral not only accelerates the time when the firm receives the profit from the referred customer but also adds the benefit of allowing more time for the new customer to make his or her own referrals. The average proportions of first-half and second-half referrals from the first three years were projected into the holdout year to compute CRV. Again, the average was a good predictor of future referral timing because the correlation between the actual and predicted number of referrals in each period was .82 and .79, respectively. Similarly, the MAPE between the actual and predicted number of referrals in each period is 16% and 19%, respectively, showing a strong fit between the model predictions and the actual results.

Predictive Accuracy of CRV

In addition to showing that the results of the individual steps (e.g., number of referrals, timing of referrals) produce predictions that are close to the actual values for each customer, we must also determine whether the overall prediction of CRV is able to capture the actual value of the referral behavior of each customer. The results of the analysis show that we were able to accurately predict the CRV for each individual customer for the holdout year. The correlation and the MAPE between the actual and predicted values of CRV across individuals is .77 and 18%, respectively, showing good predictive accuracy. Across all the customers from the sample of customers we received from the financial services firm, we compared the actual value of each customer’s CRV with the predicted value of that customer’s CRV. The MAD between the actual and predicted values of CRV was approximately \$26 per customer based on the CRV of each customer (the average CRV for each customer was approximately \$145). This shows that the new four-step approach for computing CRV we propose in this research not only accurately predicts the value of each part of the CRV model but also accurately predicts each customer’s CRV in the holdout year using three years of past data.

Field Experiments

We conducted four field experiments with the goal of quantifying the impact of measuring and managing both CLV and CRV simultaneously and determining the behavioral drivers of predicted CRV.⁸ We predict the value of CRV for

⁸Note that in the previous section, we used only a one-year-ahead prediction of CLV and CRV as a way to validate the model. Given that we have four years of data for the field experiments, CLV and CRV are predicted out for three years, which represents a long-term strategic approach to customer management.

each customer by projecting his or her past behavior into future periods using the four-step approach described previously. The end result of these field experiments will show whether an intervention of a marketing campaign in a given period can cause a significant increase in the predicted CLV or CRV of a given customer. In the case of CRV, for which the prediction is expected to mirror the past behavior, the results will show whether a marketing campaign can cause the current path of a customer’s referral behavior to change (increase or decrease) during a given period.

We conducted the first three field studies in succession using different customers of the financial services firm. The first study took place in the third quarter of 2005, the second study took place in the fourth quarter of 2005, and the third study lasted for all of 2006. For each of these three studies, we discuss which customers participated in the study, which campaign each customer received (or did not receive), and the effect of the campaign. We then discuss the retail firm data for the fourth field experiment in detail at the beginning of Study 4.

Study 1

Purpose and Method

We used the first study as a pretest to help determine which customers would be most responsive to specific marketing campaigns (stimuli) that focused on increasing CLV or increasing CRV. The participants of the experiment were 300 customers from the full sample of customers (14,160) from the financial services firm. The experiment was a marketing campaign directed at these specific customers. In July 2005, we randomly selected and sent the 300 customers a targeted marketing campaign that encouraged them to increase both CLV and CRV. The customers were sent direct mail pieces with incentives to purchase bundled offerings for one or more products, such as savings accounts, checking accounts, and/or liability accounts, which, if effective, would increase their CLV. The customers were also sent direct mail pieces with incentives to refer new customers by offering both them and the people they successfully referred a \$20 credit toward their accounts, which, if effective, would increase the customers’ CRV.

Results

At the beginning of this marketing campaign, we measured the CLV and the CRV for each of the 300 customers in the study. Using a median split of their initial July 2005 CLV and CRV values, we placed 75 customers in one of four segments: (1) high CLV/high CRV, (2) high CLV/low CRV, (3) low CLV/high CRV, and (4) low CLV/low CRV. At the conclusion of the marketing campaign (September 2005), we again measured each customer’s CLV and CRV. If the campaign had any significant effects on customers in any of the four segments, we should observe significant differences between the CLV and the CRV values before (July 2005) and after (September 2005) the campaign.

For the high-CLV/high-CRV customers, there was no significant difference between the post- and the precampaign CLV ($t = .92, p = .36$) or CRV ($t = .98, p = .33$). In

both cases, the average gains in CLV and CRV were less than 1%, meaning that the campaign to this segment was not effective in increasing either CLV or CRV. For the high-CLV/low-CRV customers, the campaign was successful at significantly increasing the average CRV ($t = 174.2, p < .001$), an average increase of approximately 20%. However, it was not successful at significantly increasing the average CLV ($t = 1.1, p = .27$), an average increase of less than 1%. For the low-CLV/high-CRV customers, the campaign was successful at significantly increasing the average CLV ($t = 114.3, p < .001$), an average increase of approximately 2.5%. However, it was not successful at significantly increasing the average CRV ($t = 1.26, p = .21$), an average increase of less than 1%. For the low-CLV/low-CRV customers, the campaign was successful at significantly increasing both the average CLV ($t = 212.3, p < .001$) by approximately 60% and the average CRV ($t = 186.9, p < .001$) by approximately 130%.

Discussion and Limitations

Study 1 showed that some aspects of the marketing campaign were more effective for certain segments of customers. In all cases, when the average customer's CLV (CRV) was below the median split, the campaign was effective at increasing his or her CLV (CRV). When the average customer's CLV (CRV) was above the median split, the campaign was not effective at further increasing his or her CLV (CRV). This suggests that low-CLV customers should be targeted with marketing campaigns to encourage them to cross-buy and add on services and that low-CRV customers should be targeted with marketing campaigns to encourage them to refer new customers to maximize the return on marketing investment. In addition, this suggests that the two groups, which are high on one value and low on the other (e.g., high CLV/low CRV), should not receive a campaign with both marketing messages.

This raises the following question: Although the firm might be tempted to send a direct mail piece with incentives to increase both CLV and CRV, does the effectiveness of a single campaign directed to these two segments, which are high on one value and low on the other, suffer when a second campaign aimed at increasing the high-scoring metric (e.g., increase CLV for the high-CLV/low-CRV group) is also sent to them? This is possible because, as the number of incentives increase, the effective incentive could get lost in the clutter with other less effective incentives (Brown and Rothschild 1993). As a result, we expect that both the response rate and the overall profits gained from the campaign aimed at increasing a low-scoring metric would decrease. To answer this question, we conducted a second study during the fourth quarter of 2005 (October 2005–December 2005).

Study 2

Purpose and Method

The second study with the financial services firm acted as a pretest to determine whether there was a significant difference in the effect of the marketing campaigns when the

high-CLV/low-CRV and low-CLV/high-CRV segments are targeted with two marketing campaigns—one to increase CLV and one to increase CRV—versus when the segments are targeted with only a single campaign intended to increase the low-scoring metric. There is almost no additional marketing costs to the firm to target customers with both marketing campaigns (the direct mail piece can easily include both campaigns). Thus, if the average increase in the low-scoring metric is not statistically different when one or both marketing campaigns are received, the firm might want to send an incentive with both campaigns. However, as we noted previously, if there is a negative consequence because of the clutter, both should not be sent to the same customer.

To conduct this second study, we randomly selected 600 customers from the remaining sample of 13,860 (14,160 less the 300 from the Study 1) and, as in Study 1, sorted the customers into the same four segments according to their CLV (high/low) and CRV (high/low) as of October 2005. Each of these four segments had 150 customers. We then chose the 150 customers from the high-CLV/low-CRV segment and the 150 customers from the low-CLV/high-CRV segment to investigate the research question. For these segments, we randomly split each into two groups of 75 customers. We compared the average CLV and CRV values from the two groups of 75 customers within each segment and found no significant difference between the average values of CLV and CRV. We did this to verify that we had two sets of matched pairs. At the beginning of October 2005, we sent one group of 75 customers within each segment a targeted marketing campaign with two incentives—one to increase CLV and one to increase CRV. The other group of 75 customers within each segment received only one targeted marketing campaign, an incentive to increase only the low-scoring metric (either CLV or CRV).

Results

We compared the postcampaign (December 2005) values of average CRV and average CLV within each segment across the two groups of 75 customers. For the high-CLV/low-CRV segment, there was no significant difference between the postcampaign average CLV values for the two groups of 75 customers ($t = .57, p = .57$); both groups from the high-CLV/low-CRV segment had increases in average CLV of less than 1%. However, the postcampaign average CRV values for the group targeted only with the campaign to increase CRV had a significantly higher average gain in CRV than the group targeted with both marketing campaigns ($t = 3.50, p = .001$). For the group that received only one campaign, the increase in average CRV was approximately 50%, while for the group receiving both campaigns, the increase in average CRV was approximately 20%.

The findings were similar for the low-CLV/high-CRV segment. There was no significant difference between the postcampaign average CRV values for the two groups ($t = .83, p = .41$); both groups showed increases in average CRV of less than 1%. However, the postcampaign average CLV values for the group targeted only with the campaign to increase CLV had a significantly higher average CLV than the customers targeted with both campaigns ($t = 3.40, p =$

.001). For the group targeted with only one campaign, the increase in average CLV was approximately 10%. For the customers targeted with both campaigns, the increase in average CLV was approximately 2%, similar to the 2.5% increase we observed Study 1.

Discussion

The results of Study 2 showed that targeting the high-CLV/low-CRV and low-CRV/high-CLV customers with both marketing campaigns generates less profit than targeting these two segments with only the marketing campaign aimed at increasing the low-scoring metric. This result is noteworthy because it suggests that even though the cost of sending a direct mail piece with both campaigns is about the same as sending a direct mail piece with one campaign, adding the second campaign does not increase the high-scoring metric. It actually reduces the effectiveness of the campaign aimed at increasing the low-scoring metric. This most likely happens because of information clutter (or information overload). By adding the second campaign to increase the high-scoring metric, customers do not focus on (or notice) the campaign that should be effective in increasing the low-scoring metric (Brown and Rothschild 1993). This suggests the need for a third study in which the high-CLV/low-CRV and low-CLV/high-CRV segments are targeted only with campaigns to increase the low-scoring metric.

In addition, given the design of both Studies 1 and 2, there is a potential threat to internal validity due to regression to the mean. In Study 1, we observed that for all segments with a low-scoring metric, the average value of that metric after the campaign increased. In Study 2, we again observed that the average values of the low-scoring metric for customers from both segments increase after the campaign. Given that the objective of these studies was to suggest the most effective campaigns for each customer segment, this is less of an issue. However, the way we conducted Study 3 removes the potential concern of regression to the mean and provides stronger, additional evidence for these findings.

Study 3

Purpose and Method

The purpose of the third field study is to investigate whether targeting the right customer segment with the right marketing campaigns can lead to significant increases in both CLV and CRV for each customer segment. In addition, the third field study aims to increase the internal validity of the previous two field studies by removing any effect related to regression to the mean. The results from Studies 1 and 2 indicated that the best strategy is to target only three of the four segments with targeted marketing campaigns aimed at improving the low-scoring metric: high CLV/low CRV, low CLV/high CRV, and low CLV/low CRV. This is because the high-CLV/high-CRV segment did not show significant gains in either CLV or CRV in Study 1 and because Study 2 showed that targeting only the low-scoring metric was an optimal strategy. For the case of CRV, although the four-step approach predicted that these customers would have a

high CRV, it was not incrementally higher than in the previous period. There are two possible reasons for this: (1) These customers are already referring a large number of customers (or a few customers with large value), and (2) these customers have already exhausted their social network with referrals.

However, in this particular case, we believe that application network exhaustion should not be an issue for several reasons.⁹ First, with regard to the financial services industry, there tends to be a high degree of switching behavior between customers and their banks. This switching behavior can arise from (1) changes in the customer's financial situation (e.g., change in job), (2) changes in the customer's debt structure (e.g., mortgage refinance), (3) changes in the customer's geographic location, and (4) changes to take advantage of different product terms (e.g., high interest savings accounts, free checking accounts).

Each of these situations opens up ample opportunities for customers to make referrals to people within their social network. Furthermore, at the time of the study, this particular financial services institution was expanding its offerings at a rate higher than its competitors. This increases the market potential because more prospects view this as an opportunity and may switch to this institution to capitalize on the benefit of consolidating their accounts. Second, the CRV prediction extends for only three years into the future. We expect that most customers have a large enough social network that the time to exhaust the network through referrals is likely to take at least three years—and likely more time. Finally, with the onset of many social networking Web sites (e.g., Facebook, LinkedIn), many consumers have a social network that is actually increasing over time.

The third field study was conducted using the following four steps: (1) customer sampling, (2) customer segmenting, (3) customer targeting, and (4) customer measurement. We discuss each in turn.

Customer sampling. The first step was to draw two random samples of customers from the database of the financial services firm. These samples came from the remaining 13,260 customers from the financial services firm (14,160 – 300 [Study 1] – 600 [Study 2]). The two samples were randomly assigned to a test group and a control group. After we evaluated the sample composition, the test group included 6700 customers, and the control group included 6560 customers. The test group was targeted with the customized marketing campaigns, and the control group received no marketing campaigns. The purpose of assigning the two groups randomly to the test and control cells was to

⁹When a firm observes a decrease in referral behavior by an average customer over time, it is possible to accommodate this in the current framework. For example, it would be possible to model the decay of the referral rate in a sample of customers and apply that rate of decay across the customer base when computing CRV. Consider the following simple example of this: If a customer is making approximately four referrals per period and we observe in the general customer database that, on average, the number of referrals decreases by 50% each period, we could predict two referrals in the next period, one referral in the following, and so forth.

test whether there was a problem with regression to the mean—that is, whether low CLV or CRV scores regress to the mean value. However, if the results show that the test group’s CLV and/or CRV increases and the control group’s does not, this provides evidence that regression to the mean is not producing the effects observed in this field study (Shadish, Cook, and Campbell 2002) and that the campaign caused the increase in CLV and/or CRV.

Customer segmenting. After calculating the CLV and CRV values for each customer in both the test group and the control group, we again used a median split of the values of CLV (high/low) and CRV (high/low) and placed the customers in the four segments. We conducted a test to determine whether the average CLV and CRV scores for customers in each of the matched segments were equal across the test and control groups. We found that there was not a significant difference between the CLV and the CRV scores across the test and control groups (see Table 3).

Customer targeting. Next, we targeted the customers of the test group with the customized marketing campaign. The customers in the three segments with at least one low-scoring metric received different targeted campaign in this study based on which campaign would likely be most effective. The three different campaigns were as follows:

1. *Targeting low CLV/low CRV.* We targeted these customers with bundled offerings for one or more financial services products. This was done through personalized communications sent through direct mail and e-mail. One offer was a small interest rate reduction on loans if the customer owned multiple products. This was to entice the customer to open a new account or transfer an account from another financial institution. We followed up with a second direct mail or e-mail piece after two weeks. We also returned telephone calls to customers who sought additional information on the offer. In all communications, the value of making referrals for new customers was also emphasized by telling them that a \$20 incentive would be given to them and to the referred customer. (*Objective: increase both CLV and CRV*)
2. *Targeting high CLV/low CRV.* We targeted these customers with an offering emphasizing the referral incentive for both the current customers and the referred customers. We also did this with direct mail and e-mail communications, followed by additional direct mail and e-mail communications after two weeks. The main goal of the communications was to emphasize the \$20 incentive that both the referring customer and the referred customer would receive. (*Objective: increase CRV*)
3. *Targeting low CLV/high CRV.* These customers received personal communications in the form of direct mail and e-mail pieces that included only an offer for bundling one or more products. There was no mention of any referral offer

in any of the information these customers received. We followed up with additional direct mail and e-mail pieces after two weeks. We also returned telephone calls to customers who sought additional information on the offer. (*Objective: increase CLV*)

Customer measurement. At the end of the one-year period (December 2006), we calculated the CLV and the CRV for every customer in the test and control groups. We discounted these values to the same period (January 2006) as the initial CLV and CRV measurement, which we calculated before the campaign.

Results

At the end of the one-year period, it was clear that each of the three targeted marketing campaigns with the three customer segments had a significant impact on the targeted metric (either CLV or CRV). Evidence of the impact is based on the differences between the pre- and postcampaign CLV/CRV, both discounted to January 2006. The following figures show how the customized marketing campaigns affected each of the three segments in the test. We also show how CRV and CLV across all three groups compared with those of the control sample.

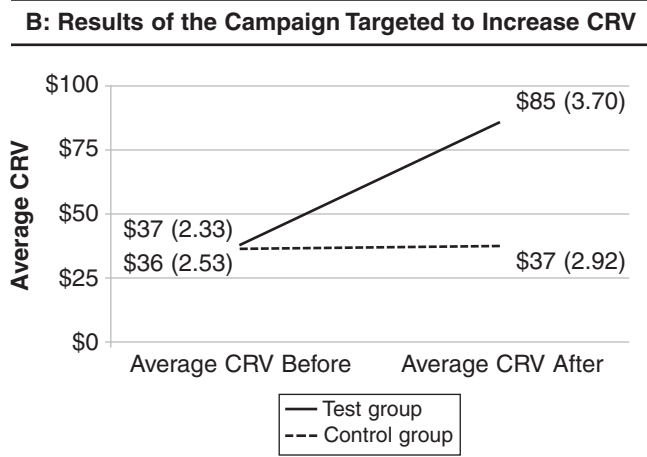

Panels A and B in Figure 1 show how both the CLV and the CRV of the low-CLV/low-CRV group changed over the period of the marketing campaign. For CLV, the average low-CLV/low-CRV customer in the test group had a CLV of \$70 before and \$120 after, while the average low-CLV/low-CRV customer in the control group had a CLV of \$68 before and \$68 after. The difference between the postcampaign CLV for the test and control groups is significant ($t = 289.74, p < .001$). For CRV, the average low-CLV/low-CRV customers in the test group had a CRV of \$37 before and \$85 after, while the average low-CLV/low-CRV customer in the control group had a CRV of \$36 before and \$37 after. The difference between the postcampaign CRV for the test and control groups is significant ($t = 392.73, p < .001$). In terms of overall gains in dollars, for the test group, the total CLV increased from \$103,000 to \$177,000 and the CRV from \$55,000 to \$125,000, while for the control group, the total CLV remained almost unchanged at \$103,000 and the CRV only went from \$54,000 to \$55,000.

These are dramatic increases, and an explanation is warranted. The percentage of people who responded to the offerings in each of the field studies was approximately 5%. If we assume that the other 95% in the sample group on average produced no increase in CLV, this would mean that the 5% who responded to the offering must have increased their profit to the firm a significant amount to explain the

TABLE 3
CLV and CRV for the Test and Control Groups

Segment	Test Group (n = 6700)		Control Group (n = 6560)	
	CLV	CRV	CLV	CRV
High CLV/high CRV	\$190	\$302	\$194	\$296
High CLV/low CRV	\$604	\$26	\$609	\$28
Low CLV/high CRV	\$101	\$350	\$103	\$347
Low CLV/low CRV	\$70	\$37	\$68	\$36

FIGURE 1
Results for the Low-CLV/Low-CRV Group




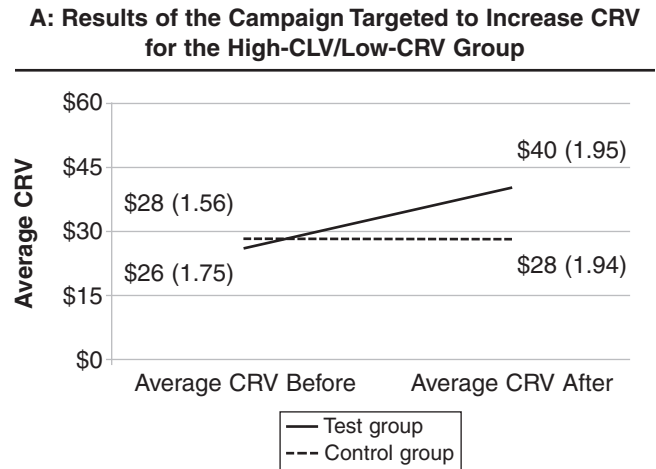
Notes: Figure shows the mean value (standard deviation).

increase in the overall average. For example, in the first case (Figure 1, Panel A), an increase of CLV from \$70 to \$120 means that the 5% who responded needed to increase their profit to approximately \$36 per month for 36 months (or \$432 per year for three years). Although this increase in the amount of profit from this group may seem large, in this industry, large increases can easily result when a customer adds a new product (or a few new products) to an existing portfolio of products (see average monthly profits from each product in Table 1).

Figure 2, Panel A, shows how the CRV for the high-CLV/low-CRV group changed from before to after the period of the marketing campaign. For CRV, the average high-CLV/low-CRV customers in the test sample had a significant increase in CRV from \$26 before to \$40 after, while the average high-CLV/low-CRV customer in the control sample had no change in CRV; it was \$28 before and \$28 after. The difference between the postcampaign CRV for the test and control groups is significant ($t = 189.60, p < .001$). For the test group, the total CRV increased significantly from \$51,000 to \$77,000, while for the control group, the total CRV remained almost unchanged at \$52,000.

Figure 2, Panel B, shows how the CLV of the low-CLV/high-CRV group changed before to after the period of

FIGURE 2
Results for the High-CLV/Low-CRV and Low-CLV/High-CRV Groups



Notes: Figure shows the mean value (standard deviation).

the marketing campaign. For CLV, the average low-CLV/high-CRV customer in the test group had a CLV of \$101 before and \$111 after, while the average low-CLV/high-CRV customer in the control group had a CLV of \$103 before and \$103 after. The difference between the postcampaign CLV for the test and control groups is significant ($t = 104.98, p < .001$). For the test group, the total CLV increased from \$196,000 to \$215,000, while the total CLV remained almost unchanged at \$196,000 for the control group.

Panels A and B in Figure 3 show how the aggregated CLV and CRV across the three targeted segments changed before to after the period of the marketing campaign. For CLV, the average across the three segments of the test group customers was \$274 before and \$291 after. The average across the three segments of the control group customers was \$270 before and \$271 after.¹⁰ The difference between the postcampaign CLV for the aggregated test and control

¹⁰The results for Figure 3, Panels A and B, include customers from the high-CLV/low-CRV, low-CLV/high-CRV, and low-CLV/low-CRV cells.

groups is significant ($t = 238.80, p < .001$). For CRV, the average for customers across the three targeted segments in the test group was \$145 before and \$164 after. The average for customers across the three segments of the control group was \$145 before and \$146 after. The difference between the postcampaign CRV for the aggregated test and control groups is significant ($t = 237.80, p < .001$). For the aggregated test groups (across the three segments), the total CLV increased from \$1.47 million to \$1.56 million and the CRV from \$785,000 to \$882,000, while for the control group, the total CLV remained almost unchanged at \$1.42 million, and the CRV only went from \$766,000 to \$767,000.

These results show that each of these campaigns was successful in significantly increasing the average customer's CLV, CRV, or both CLV and CRV in each of the test groups. However, although the results show that the campaigns worked in generating new profit, we need to determine the return on investment (ROI) to know whether such campaigns are worth the investment (these marketing dollars could have been spent elsewhere in the firm). The discounted costs of carrying out the three campaigns (including direct mail pieces, e-mails, and selected telephone calls

for the 5360 customers in the test group¹¹), which were implemented twice during the course of the year (once in January 2006 and once in July 2006), were \$21,256. This means that the campaign cost was approximately \$3.97 per targeted customer. The overall profit during the one-year campaign generated from the increase in each customer's CLV and/or CRV across the three campaigns was \$216,410. These profits include the cost of the referral incentive in the case of CRV and the discount from bundling products in the case of CLV. These incentives were not considered marketing costs in the ROI computation, because they were conditional on a purchase occurring and, in this firm, came out of a different budget. They were accounted for by reducing the profit from the item purchased.¹² Therefore, ROI can be computed as follows: $(\text{profit} - \text{marketing costs}) / \text{marketing costs}$. In this case, the ROI was just over 10. Because this campaign is above the firm's general ROI target for investments, it would be viewed as a success.

What does this mean for the financial services firm? Should the firm extend this campaign to a larger group of customers from its customer database or to its entire customer base? If we project these gains across 1 million (approximately 7%) of the firm's customer base, we would estimate incremental gains over a one-year period of approximately \$19.4 million in CLV and approximately \$20.9 million in CRV. When we project these gains across the entire customer base of 15 million people, we estimate incremental gains over one year of approximately \$291 million in CLV and \$314 million in CRV.

Although these campaigns were customized for a specific firm in the financial services industry, these findings are consistent in magnitude to those that Kumar, Petersen, and Leone (2007) report for a similar type of campaign conducted by a telecommunications firm. Thus, there is evidence across two industries that the success of campaigns targeted to increase referrals can play a key role in new customer acquisition and in generating profits. This strengthens the external validity of this study and suggests that these findings can be extended to other firms in industries in which referrals are important for new customer acquisition. Although the tactics used to generate the referrals may vary from firm to firm, these findings confirm that firms should be able to harness the power of CRV when implementing referral marketing campaigns.


Discussion

The third field study shows the importance of measuring CLV and CRV independently to understand their distinct impact on overall profits. As we expected, there was no significant change in CLV or CRV for the control group, which was not exposed to the targeted marketing campaigns; val-

¹¹Only customers from three of the four segments were contacted; the 1340 customers from the high-CLV/high-CRV cell were not targeted.

¹²This is similar to retailers that issue coupons for which the cost of printing and mailing the coupon makes up the marketing cost portion of ROI computation and the cost of the redemption of the coupon is subtracted from the gross margin received from the purchase to generate profit.

FIGURE 3
Results Across All Three Targeted Cells (Low CLV/Low CRV, Low CLV/High CRV, and High CLV/Low CRV)



Notes: Figure shows the mean value (standard deviation).

ues changed by only approximately \$1 per customer. However, we observed significant gains in CLV, CRV, or both CLV and CRV due to the targeted marketing campaigns used for the test group. This finding is consistent with that from the first two field studies, and given the design of the third study, we are confident that the possibility of regression to the mean as a potential threat to the internal validity is not an issue that biased the results.

As we reported previously, we know that all customers who were targeted to increase CRV from the low-CRV segment did not make referrals. Only a few customers (approximately 5%) accounted for the significant overall increase in the mean CRV. This means that if we can identify an extrinsic motivation (i.e., referral incentive) that is appealing to the small set of customers who are most likely to make referrals, managers can more efficiently target only those customers for referral marketing campaigns.

Therefore, to better understand the profile of the people who actually made successful referrals, we conducted an exploratory analysis of the differences in the behavioral characteristics (e.g., recency or frequency of purchase) of the people giving referrals versus people not giving the referrals. To validate and generalize these findings, we ran a fourth field study with a firm in a different industry.

Study 4

Purpose

We designed Study 4 to first determine which customers are most likely to provide future referrals before a campaign. We accomplish this by using a sample of customers from a retailing firm and relate each customer's predicted CRV (as computed by the four-step approach outlined previously) to a set of behavior variables identified after reviewing the relationship marketing literature. We then use the results to run a fourth field experiment that targets only the customers most likely to respond to referral marketing campaigns in a one-year-long field study. In this field study, we observe not only the effectiveness of the campaign but also the impact of understanding the drivers of CRV on helping managers target individual customers for referral marketing campaigns.

Data

The data used in Study 4 to determine the drivers of CRV were collected from January 1, 2005, to December 31, 2007 (three years). The retailing firm reported the data on a quarterly basis for 40,000 customers. The data were collected from those customers during the fourth field study, which occurred over a one-year period following the end of the initial data collection from January 1, 2008, to December 31, 2008. These data included each customer's own purchase behavior, the marketing costs associated with attracting and retaining each customer, the actual referral behavior of each customer, and the transaction data for the customers referred by the original 40,000 customers in the retailing firm sample.

Method

Before we conducted this field study, we used the first three years of data to (1) predict each customer's CRV on the

basis of the four-step approach described previously and (2) determine empirically the drivers of each customer's predicted CRV (calculated in Step 1) on the basis of the set of observable variables in the firm's customer database.¹³ After predicting each customer's CRV in Step 1, we needed to identify a set of appropriate variables to use in the drivers model. To do this, we reviewed empirical studies in the relationship marketing literature (e.g., Venkatesan and Kumar 2004) and identified a list of variables to use as potential observable behavioral predictors of CRV. The drivers of CLV that help explain a customer's own purchase behavior are a good place to start when trying to understand a customer's actual referral behavior. In addition, we reviewed the word-of-mouth literature for possible predictors of CRV. We restrict the list of variables to those that are regularly contained in a firm's customer database. We begin by selecting variables from each of the three main categories of relationship marketing variables as they relate to the firm-customer exchange process. These include the following:

- *Firm-specific factors*: These include any actions the firm takes to manage the relationship (e.g., marketing communications).
- *Exchange characteristics*: These include any observed customer behaviors (e.g., purchases, product returns, referral behavior) that help describe the state of the customer's relationship with the firm.
- *Customer characteristics*: These include any demographic variables (e.g., income).

In addition to selecting the linear and quadratic terms for each of the three key categories of relationship marketing variables, we introduce variables that analyze how changes in customer behavior over time affect CRV (e.g., change in spending). We describe each variable and provide an expected effect and a rationale for that effect in Table 4.

Because CRV is a censored variable, the value can be greater than or equal to 0. We model each customer's CRV as a function of predictor variables using a Bayesian Tobit model (see Equation 3). We chose a Bayesian Tobit model rather than a Heckman two-stage model for the following reason: The problem of selection bias can be corrected within a single step of data augmentation (see Tanner and Wong 1987) in a Bayesian framework, and it requires a two-step process in a maximum likelihood approach. Chib (1992) provides the details on estimation.

$$(3) \quad CRV_i^* = X_i\beta + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2),$$

where

$$CRV_i^* = \begin{cases} CRV_i, & \text{if } CRV_i > 0 \\ TN_{[-\infty, 0]}(\mu_{CRV}, \sigma_{CRV}^2), & \text{if } CRV_i = 0, \end{cases}$$

- X_i = a matrix of predictor variables,
- β = a vector of parameter estimates, and
- ε_i = an error term that is distributed normally.

¹³We use a three-quarter simple moving average to predict the components of CRV. We do this because recency of the referral behavior is the best predictor of future referral behavior in the retailing context.

TABLE 4
Variables Used as Drivers of CRV

Variable	Operationalization (M) [Variable Name]	Expected Effect	Rationale
Firm-Specific Factors			
Marketing spend	Average spent on marketing per year to customer <i>i</i> from 2005 to 2007 (\$42.60 per year) [Avg_Mktg]	∩	Marketing will have a saturation effect on customers, such that too much will cause strain in the relationship (Fournier, Dobscha, and Mick 1997).
Exchange Characteristics			
Monetary value	Average profit per year from customer <i>i</i> from 2005 to 2007 (\$182.40 per year) [Avg_MV]	∩	Customers who spend moderately are most likely to have the highest referral value (Kumar, Petersen, and Leone 2007).
Change in monetary value	Difference in previous two quarter's profit: $MV_t - MV_{t-1}$ (\$-2.78) [Change_MV]	-	As a customer reduces his or her contribution to a firm's profits through purchases, he or she will likely increase his or her referrals.
Frequency	Average interpurchase time of customer <i>i</i> from 2005 to 2007 (1.6 months) [AIT]	∩	Higher (to a threshold) purchase frequency results in more customer interaction with the firm, which will also increase a customer's referrals.
Change in frequency	Difference in previous two interpurchase times: $AIT_t - AIT_{t-1}$ (-.16 months) [Change_AIT]	-	Increases in purchase frequency (lower AIT) will increase the interaction of the firm and customer and, in turn, likely increase referrals.
Recency	Number of months since previous purchase (3.46) [Recency]	-	The more recent a customer purchase (i.e., the more recent the customer interaction), the more likely a referral will result.
Cross-buy	Cumulative number of departments customer <i>i</i> purchased in from 2005 to 2007 (5.8 categories) [Cum_CB]	+	Shopping across more categories will strengthen the relationship and, in turn, is more likely to increase referrals.
Multichannel shopping	Cumulative number of channels customer <i>i</i> purchased in from 2005 to 2007 (2.1 channels) [Cum_MC]	+	Shopping across more channels will strengthen the relationship with the firm and, in turn, is more likely to increase referrals.
Product returns	Average amount of product returns per year from customer <i>i</i> from 2005 to 2007 (\$70.40 per year) [PR]	∩	A moderate number of product returns signals a strong future relationship and a higher number of referrals.
Past referral	A binary variable that equals 1 if a customer has made a referral in the past and 0 if no referral has been made (.11) [Referral]	+	Customers who have made referrals in the past are most likely to continue to refer new customers in the future.
Customer Characteristics			
Age	Age of customer <i>i</i> in 2007 (39.8) [Age]	N.A.	We include these as covariates in the model merely to control for customer heterogeneity. Thus, we do not provide any hypothesized effects.
Income	Reported income (in thousands of dollars) of customer <i>i</i> in 2007 (56.9) [Income]	N.A.	

Notes: N.A. = not applicable.

We estimate this model using the statistical program R. Using standard protocol, we run 10,000 iterations, with the first 5000 used as burn-in and the remaining 5000 used as posterior.¹⁴

¹⁴We provide the results of two separate null models, along with the full model, in Web Appendix W2 (<http://www.marketingpower.com/jmsept10>). Model 1 has only a dummy variable for customers who have made past referrals. Model 2 has a dummy variable for customers who have made past referrals and variables that represent marketing cost and marketing cost squared.

Estimation Results

We present the parameter estimates along with model fit statistics in Table 5. The parameter estimates from the model provide several key insights into the effect of each variable on a customer's CRV.

Firm-Specific Factors

We find that Avg_Mktg has an inverted U-shaped relationship to CRV, which is consistent with what would be expected—namely, diminishing returns. Therefore, spend-

TABLE 5
Parameter Estimates and Model Fit for the CRV Model

Variable	M (SD)
Intercept	41.40 (1.28)
Firm-Specific Factors	
Marketing (Avg_Mktg)	.512 (.086)
Marketing squared (Avg_Mktg ²)	-.196 (.042)
Exchange Characteristics	
Monetary value (Avg_MV)	.426 (.109)
Monetary value squared (Avg_MV ²)	-.094 (.046)
Change in monetary value (Change_MV)	-.164 (.049)
Frequency (AIT)	.238 (.032)
Frequency squared (AIT ²)	-.067 (.021)
Change in frequency (Change_AIT)	-.201 (.056)
Recency (Recency)	-.138 (.056)
Cross-buy (Cum_CB)	1.304 (.108)
Cross-buy squared (Cum_CB ²)	-.045 (.012)
Multichannel shopping (Cum_MC)	1.287 (.136)
Multichannel shopping squared (Cum_MC ²)	n.s.
Product returns (Avg_PR)	1.018 (.204)
Product returns squared (Avg_PR ²)	-.029 (.005)
Past referral (Referral)	3.281 (.736)
Customer Characteristics	
Age (Age)	.298 (.138)
Age squared (Age ²)	n.s.
Income (Income)	.326 (.108)
Income squared (Income ²)	-.019 (.007)
Fit	
Log-likelihood	-53.08
Akaike information criterion	146.16

Notes: n.s. denotes not significant at $p < .10$.

ing on marketing has a diminishing return on CRV, which is similar to the effect of marketing spend on CLV (Venkatesan and Kumar 2004). This suggests that it is important to optimize the marketing spend with each customer because a customer can become saturated with too many marketing touches, possibly even negatively affecting the relationship (Fournier, Dobscha, and Mick 1997).

Exchange Characteristics

We find that Avg_MV has an inverted U-shaped relationship to CRV, which is consistent with Kumar, Petersen, and Leone's (2007) findings. They find that customers in Deciles 5, 6, and 7, when rank-ordered by CLV, are most likely to have the highest CRV. We find that Change_MV has a negative relationship to CRV, which matches our expectation. This suggests that as spending with the firm slows, the customer is more likely to translate that loss in purchases into an increase in referrals. We found evidence of this phenomenon in several other industries in which referral behavior is encouraged with incentives. For example, we spoke with a manager of a major U.S. airline about referral incentives that help customers attain airline miles for free flights. We were told that customers who begin to fly less frequently (i.e., lower spend over time) react positively to opportunities to make referrals for mileage that can be used for free flights.

We find that AIT has an inverted U-shaped relationship to CRV, which matches our expectation. Thus, customers who purchase with moderate frequency are the customers with the highest CRV, similar to the effect of AIT on CLV (Venkatesan and Kumar 2004). When we examine the change in AIT (Change_AIT), we find a negative effect on CRV, which matches our expectation. This suggests that while a customer with a moderate AIT has the highest CRV, all else being equal, as the AIT decreases with time, referral behavior tends to increase. This indicates that as customers increase their purchase frequency, and likely the strength of their relationship with the firm, they increase their referral behavior. An explanation for this finding is that each additional purchase occasion (because of the decrease in the interpurchase time) may act as a trigger for these customers to talk more to their friends and/or relatives.

We find that Recency has a negative relationship to CRV. This matches our expectation. This suggests that the longer the time since the last purchase, the lesser is the strength of the firm-customer relationship, and in turn, the less likely the customer is to make referrals. We find that Cum_CB has an inverted U-shaped relationship to CRV, which partially matches our expectation. We expected this relationship to be positive. However, we find that customers who purchase in several (but not too many) categories are the best candidates for making referrals. It is likely that customers who have purchased in several categories have sufficient relationship strength and product knowledge. However, it is likely that customers who already purchase in all categories do not perceive value in the referral incentives because they already have purchased a significant amount of products from the company. We find that Cum_MC has a strictly positive relationship to CRV, which matches our expectation. Thus, customers who buy across more channels are more likely to engage in referral behavior, similar to the effect of multichannel shopping on CLV (Venkatesan, Kumar, and Ravishanker 2007). We find that Avg_PR has an inverted U-shaped relationship to CRV. This matches our expectation. Thus, customers who return a moderate number of products are more likely to have a stronger future relationship with the firm (Petersen and Kumar 2009) and, in turn, are more likely to engage in referral behavior. Finally, we find that Referral has a positive relationship to CRV, which matches our expectation. This shows that customers who already have demonstrated that they are able and willing to make referrals are the most likely to continue to make future referrals.

Customer Characteristics

Taking into account demographic variables, we find that Age is positively related to CRV, and Income has an inverted U-shaped relationship to CRV. This suggests that older customers and customers with moderately high incomes are most likely to engage in referral behavior for this retailing firm. Undoubtedly, the results of these variables could vary across firms and industries. However, it is important to consider these variables for two reasons: (1) It helps explain some of the customer heterogeneity within the firm's customer database, and (2) it helps the firm target the best prospects for acquisition who not only purchase

products but also are likely to refer even more customers to the firm.

Field Study Purpose and Method


After completing this stage, we conducted a year-long field study with the retailing firm. The purpose of this field study was to determine whether we could successfully target individual customers according to their expected referral behavior. In the first three field studies, we found that segmenting customers into low- and high-CRV groups and targeting those customers with low CRV was an effective method of increasing profit to the firm through referrals. However, it was not the case that all low-CRV customers increased their referral behavior. If we can understand the drivers of CRV, we can target only customers in the low-CRV segment who offer the highest potential referral behavior. Thus, the purpose of the fourth field study is to investigate whether there is a difference in the impact on a firm's profit through CRV when individual customers are targeted with referral marketing campaigns using the behavioral drivers of CRV versus when low-CRV customers are simply randomly targeted with referral marketing campaigns.

We split the 40,000 customers into two equally sized (median split) groups of 20,000 customers according to their CRV scores (i.e., high or low). We then randomly split the 20,000 customers in the low-CRV group into two groups, each with 10,000 customers. One group served as the targeted group, and the other served as the random group. For the targeted group, we used the drivers of CRV to predict which customers would be most responsive to a referral marketing campaign according to past behavior, and we rank-ordered those customers from most to least responsive. Then, for the targeted group, we selected 5000 customers (of the 10,000) with the highest potential to make referrals according to their CRV score, and for the random group, we randomly selected 5000 (of the 10,000) customers to include in the referral marketing campaign. Over the course of the one-year field study with the retail firm, each of the customers in the two groups (targeted and random) received a direct mail piece (each month), which provided a \$20 incentive that would be given to both the customer who makes the referral and the referred customer on purchase. We then tracked the referral behavior of the customers in both samples over the course of one year.

Field Study Results

We found that knowing the drivers of CRV helped significantly in targeting the right customers for referral marketing campaigns. Figure 4 shows evidence of this. Before we conducted the field study, the average CRV for the targeted group was approximately \$191, and the average CRV for the random group was approximately \$192. After the field study, the targeted group increased to \$249 (discounted to the beginning of 2008), or approximately 30%. The CRV of the control group increased to \$212, or approximately 10%. The difference between the CRV of the targeted and random groups in the posttreatment period is statistically significant ($t = 327.43, p < 0.001$). Note that customers who increased their CRV also on average increased their CLV. In

FIGURE 4
Average CRV Before and After Campaign with Retail Store: Targeted and Random Groups



Notes: Figure shows the mean value (standard deviation).

other words, the act of providing referrals led to increases in their own purchases, so the firm in effect “hit the jackpot twice.”

Discussion and Limitations

These results show that targeting low-CRV customers randomly can be beneficial (increase in CRV by approximately 10%). However, understanding exactly which customers in the low-CRV segment who are most likely to respond to referral marketing campaigns can further increase each customer's CRV. In this case, the increase was three times greater (30%), or an additional 20%. This also does not include the savings from the efficiency gained from the targeting or the prevention of lost customers because of the “harassment factor.” In addition, we did not target any customers who were segmented into the high-CRV group because, as an aggregate, the group did not increase their CRV in Study 1 or 2. However, it is likely that there are customers—even if only a small group—in the high-CRV segment who would be responsive to referral marketing campaigns. A future field experiment could potentially address this issue.

Managerial Implications

This research has several key implications for marketing theory and practice. First, we propose a new four-step approach to measure CRV. Using data from a financial services firm, we show how to quantify the incremental gain in profits when marketing strategies are managed according to each customer's CRV and CLV. The findings show that it is important for managers to measure both CRV and CLV separately to better understand their customers and to maximize their profits. There are two key advantages of using the process proposed in this article to understand and manage both CLV and CRV:

1. Because CRV is a prediction of a referred customer's profitability, simply adding CRV to CLV to derive total customer value, as some researchers have suggested, would only lead to double-counting the contribution. Therefore,

summing CLV and CRV yields a measure of total customer equity.

2. Managers need to have separate metrics to implement an explicit strategy that targets customers according to their referral behavior or their own transaction behavior. Without using these separate metrics (CLV and CRV) to help make decisions about allocating resources, a company may actually ignore customers or provide bad service to those who would offer great returns from referred customers.

We were also able to illustrate empirically how to determine a set of behavioral drivers of CRV using data from a retailing firm to target marketing campaigns more efficiently. This was demonstrated through a field experiment that showed the following:

- Targeting customers from the low-CRV group using the behavioral drivers of CRV significantly increased profits over targeting customers randomly.
- Marketing to random customers in the low-CRV group was profitable but not nearly as profitable as targeting customers according to the behavioral drivers of CRV.
- Customers who made referrals during the field experiment actually increased their purchases, adding additional profits to the firm.

It is clear that in many situations, managers have ignored the customers who have the potential to make referrals. These customers may have become “turned off,” causing them to feel alienated from the firm and potentially spread negative word of mouth. If this happens, these firms can lose future revenue because these customers (1) might stop purchasing and referring customers (Hogan, Lemon, and Libai 2003), (2) might drive away potential prospects who would have joined even without a referral, and (3) could develop negative attitudes toward the firm, leading the acquisition costs of these prospects to increase.

There are clear opportunities available for firms to harness the power of CRV. One way is to use CRV as a metric for introducing new products by using the referral base of customers. Firms that are new entrants to a market can take leadership not just with an innovative product but also by using innovative marketing tools that enable them to harness the power of referrals (Shankar, Carpenter, and Krishnamurthi 1998). Research has shown the value of these referrals through indirect and direct network externalities (Sun, Xie, and Cao 2004), network effects (Godes and Mayzlin 2009; Hogan, Lemon, and Libai 2003, 2004), and a faster diffusion of new products through accelerated purchase behavior (Mahajan, Muller, and Bass 1995).

Another question this research addresses is how a firm might develop a referral program that enables it to maximize profitability across all the customers within the firm. One method to enhance the effectiveness of the referral program is to differentiate the value of the referral incentive.¹⁵ It is possible to use the behavioral drivers of CRV to customize rewards to the customer level and to strengthen the

¹⁵For details of the referral program by Bank of America, which differentiates referral amounts to different types of customers, see https://epreferences.bankofamerica.com/asbs/html/landing_page/Phase3_Referral_Landing_page.html?rc=CRPOE&adlink=000309029c910000a641.

overall customer relationship (Verhoef 2003). Similar to approaches that manage loyalty program rewards on the basis of each customer’s CLV (Kumar and Shah 2004), building referral reward programs based on each customer’s CRV can also be straightforward.

Building a strong social network can be a long-term competitive advantage for both the customers and the firm because it enables the firm’s relationship with each customer to evolve over time (Ganesan 1994). It becomes more difficult for competitors to lure away customers who are tightly locked in to a company’s social network. At the same time, this benefits customers because a strong social network allows for ease of information sharing about products and the use of common products across a set of customers. Thus, a firm should view its customers as skilled resources and work with them to build strong social networks in which both can benefit.

Limitations and Suggestions for Further Research

We should point out a few limitations that firms should consider when measuring CRV and implementing marketing programs to increase customer referrals. First, measuring CRV is only one method of determining whether a customer can bring a firm new customers and might not be applicable to all firms or industries. For example, for many products, their value does not invite rewards for referrals. In these instances, managers should consider exploring what drives referral behavior with their customers (as we showed in this study) and determine an appropriate method to measure this referral behavior at a customer level. In this situation, managers could target prospects who fit the profile of high-CRV customers currently in their database.

Second, not every business has the ability to track referral behavior, and even if they do, they may not have systems in place to do so or a long history of tracking customer referral behavior to estimate the models we recommend. For the firms in this study, we were able to use three years of referral behavior history to calibrate the predictions of future referral behavior. In other situations, such as the roll-out of new products, firms may not have sufficient data to build strong customer insights and therefore would not be able to target customers accurately.

Third, the method we suggest for measuring CRV only takes into account the extrinsic motivation of a customer to make referrals (i.e., through incentives) and does not try to determine which customers generate referrals through an intrinsic motivation. However, there is no reason to expect that the presence of intrinsic motivations would alter the findings pertaining to the effect of extrinsic motivations. This could be a worthwhile area for further research.

Fourth, the four-step approach for computing and validating CRV worked well in a fairly stable environment (i.e., the financial services industry). In this industry, the average customer purchases and holds onto a product (e.g., checking account) for a long time. In cases in which the market is more turbulent, it may be necessary to use a driver-based approach to compute the four parts of CRV to increase the accuracy of prediction.

Finally, this study does not directly investigate the level or type of reward that should be given to a customer. It is also important to understand what the referral incentive should be when dealing with customers who communicate with prospects in different social networks (Ryu and Feick 2007). Although Ryu and Feick (2007) find that the likelihood of a customer making a referral does not change much given a small or large incentive, it is important to understand the value of the customers when designing the incentives for the referral program. For example, the customers targeted with the referral incentives from the financial services and retailing firms in this study had vastly different

customer profiles. Some of the customers were students, some were large family households, and some were small business owners (in the case of the financial services firm). However, in this study, all the customers were targeted with the same incentive level (\$20 for referring customer and \$20 for referral). Perhaps differentiating the level of referral incentive on the basis of the type of customer is a strategy that can further maximize the gains from a referral marketing campaign. Further research might examine this possibility to help firms maximize their profits. Another potential research idea is to explore the cause behind the “hitting-the-jackpot-twice” phenomenon.

REFERENCES

- Brown, Tom J. and Michael L. Rothschild (1993), “Reassessing the Impact of Television Advertising Clutter,” *Journal of Consumer Research*, 20 (1), 138–46.
- Chib, Siddhartha (1992), “Bayes Inference in the Tobit Censored Regression Model,” *Journal of Econometrics*, 51 (1), 79–99.
- Dichter, Ernest (1966), “How Word-of-Mouth Advertising Works,” *Harvard Business Review*, 44 (6), 147–66.
- Feick, Lawrence F. and Linda L. Price (1987), “The Market Maven: A Diffuser of Marketplace Information,” *Journal of Marketing*, 51 (January), 83–97.
- Fournier, Susan, Susan Dobscha, and David Glen Mick (1997), “Preventing the Premature Death of Relationship Marketing,” *Harvard Business Review*, 75 (January–February), 2–8.
- Ganesan, Shankar (1994), “Determinants of Long-Term Orientation in Buyer–Seller Relationships,” *Journal of Marketing*, 58 (April), 1–19.
- Godes, David and Dina Mayzlin (2009), “Firm-Created Word-of-Mouth Communication: Evidence from a Field Test,” *Marketing Science*, 28 (4), 721–39.
- Helm, Sabrina (2003), “Calculating the Value of Customers’ Referrals,” *Managing Service Quality*, 13 (2), 124–33.
- Hogan, John E., Katherine N. Lemon, and Barak Libai (2003), “What Is the True Value of a Lost Customer?” *Journal of Service Research*, 5 (3), 196–208.
- , ———, and ——— (2004), “Quantifying the Ripple: Word-of-Mouth and Advertising Effectiveness,” *Journal of Advertising Research*, 44 (3), 271–80.
- Keiningham, Timothy L., Bruce Cooil, Tor Wallin Andreassen, and Lerzan Aksoy (2007), “A Longitudinal Examination of Net Promoter on Firm Revenue Growth,” *Journal of Marketing*, 71 (July), 39–51.
- Kumar, V., J. Andrew Petersen, and Robert P. Leone (2007), “How Valuable Is Word of Mouth?” *Harvard Business Review*, 85 (10), 139–46.
- and Denish Shah (2004), “Building and Sustaining Profitable Customer Loyalty for the 21st Century,” *Journal of Retailing*, 80 (4), 317–29.
- , Rajkumar Venkatesan, Timothy R. Bohling, and Denise Beckman (2008), “The Power of CLV: Managing Customer Value at IBM,” *Marketing Science*, 27 (4), 585–99.
- Mahajan, Vijay, Eitan Muller, and Frank M. Bass (1995), “Diffusion of New Products: Empirical Generalizations and Managerial Uses,” *Marketing Science*, 14 (3), 79–88.
- Petersen, J. Andrew and V. Kumar (2009), “Are Product Returns a Necessary Evil? Antecedents and Consequences,” *Journal of Marketing*, 73 (May), 35–51.
- Reichheld, Frederick F. (2003), “The One Number You Need to Grow,” *Harvard Business Review*, 81 (12), 46–54.
- Rust, Roland T., Valarie A. Zeithaml, and Katherine N. Lemon (2000), *Driving Customer Equity*. New York: The Free Press.
- Ryu, Gangseog and Lawrence Feick (2007), “A Penny for Your Thoughts: Referral Reward Programs and Referral Likelihood,” *Journal of Marketing*, 71 (January), 84–94.
- Shadish, William R., Thomas D. Cook, and Donald T. Campbell (2002), *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. New York: Houghton Mifflin.
- Shankar, Venkatesh, Gregory S. Carpenter, and Lakshman Krishnamurthi (1998), “Late Mover Advantage: How Innovative Late Entrants Outsell Pioneers,” *Journal of Marketing Research*, 35 (February), 54–70.
- Sun, Baohung, Jinhong Xie, and H. Henry Cao (2004), “Product Strategy for Innovators in Markets with Network Effects,” *Marketing Science*, 23 (2), 243–54.
- Tanner, Martin A. and Wing Hung Wong (1987), “The Calculation of Posterior Distributions by Data Augmentation,” *Journal of the American Statistical Association*, 82 (398), 528–40.
- Venkatesan, Rajkumar and V. Kumar (2004), “A Customer Lifetime Value Framework for Customer Selection and Resource Allocation Strategy,” *Journal of Marketing*, 68 (October), 106–125.
- , ———, and Nalini Ravishanker (2007), “Multichannel Shopping: Causes and Consequences,” *Journal of Marketing*, 71 (April), 114–32.
- Verhoef, Peter C. (2003), “Understanding the Effect of Customer Relationship Management Efforts on Customer Retention and Customer Share Development,” *Journal of Marketing*, 67 (October), 30–45.
- Villanueva, Julian, Shijin Yoo, and Dominique M. Hanssens (2008), “The Impact of Marketing-Induced Versus Word-of-Mouth Customer Acquisition on Customer Equity Growth,” *Journal of Marketing Research*, 45 (February), 48–59.