Available online at www.sciencedirect.com
ScienceDirect
Journal of Interactive Marketing 29 (2015) 39 – 56
JOURNAL OF INTERACTIVE MARKETING
www.elsevier.com/locate/intmar

Managing Customer Acquisition Risk Using Co-operative Databases

Hongju Liu [a,1] & Joseph Pancras [a,*] & Malcolm Houtz [b,2]

[a] Department of Marketing, University of Connecticut, 2100 Hillside Rd., U-1041, Storrs, CT 06269-1041, USA
[b] Alliant Inc., 301 Fields Lane, Brewster, NY 10509-2621, USA

Abstract

Acquisition of new customers involves both opportunity and risk, and it is important for firms to predict and manage the risks involved in customer acquisition. Despite its importance, the management of customer acquisition risk has not been the subject of much academic research. This paper develops a framework for firms to manage customer acquisition risk using co-operative databases. We illustrate this framework in the context of the optimal selection of customers for direct mail with a 'buy now, pay later' payment option when the acquisition risk manifests as bad debt risk. Using data from a large scale direct marketing campaign, we show that our empirical model that incorporates bad debt risk substantially outperforms suboptimal targeting schemes that overlook bad debt risk. We also demonstrate how alleviating bad debt risk is one beneficial outcome of a fairly recent trend in database marketing, namely the emergence of co-operative databases.

Keywords: Customer acquisition; Acquisition risk; Bad debt risk; Co-operative database; Direct mail; Direct marketing; Targeting

Introduction

Incorporating a customer risk forecast into assessment of the economic value of a customer is an important conceptual development that has both theoretical and practical implications. In developing relationship marketing strategies, prior research has focused on selective customer retention through a customer value analysis and a risk-adjustment process (Ryals and Knox 2005). The risk of an existing customer can originate from factors that affect the volatility of future revenue, e.g., the probability of filing insurance claims and the probability that the customer will not be retained.

While risk in customer retention is an important aspect of managing customer relationships, risk in customer acquisition is arguably at least as important as that in customer retention. When acquiring new customers, firms face higher uncertainty in customer responses to a promotional campaign. Without past customer relationships, the risk of undesirable behaviors such as bad debt (inability or unwillingness to pay) or product return is more prominent, and meanwhile it is challenging for firms to forecast such risk because of the limited information. Despite the importance of risk management in customer acquisition, it remains an understudied area in the literature possibly due to the lack of appropriate data.

In this paper we develop a framework to manage customer acquisition risks, and illustrate our framework in the empirical context of a direct mail campaign. A key decision for direct marketing firms is to choose the right consumers to target. Previous studies strive to identify variables that can predict consumer responses so that firms can improve response rates and reduce wasteful mailings. Although response rate is a key metric for direct marketers, the literature has hitherto largely ignored the risk management in a customer acquisition campaign, especially the identification of consumers who respond but do not pay. While the related but different behavior of product return has been the subject of several recent papers (e.g. Anderson, Hansen, and Simester 2009; Petersen and Kumar 2009), the widespread phenomenon of bad debt in direct marketing has not been the subject of much academic research.

The management of acquisition risk through prediction of bad debt is important for direct marketers since the cost of bad

* Corresponding author: Tel.: +1 860 486 8397; fax: +1 86 486 5246.
E-mail addresses: hliu@business.uconn.edu (H. Liu), jpancras@business.uconn.edu (J. Pancras), mhoutz@ALLIANDATA.com (M. Houtz).
[1] Tel.: +1 860 486 8397; fax: +1 86 486 5246.
[2] Tel.: +1 845 276 2600x238.

http://dx.doi.org/10.1016/j.intmar.2014.09.002
1094-9968/

debt can be very high. Delinquencies and bad debts cost the industry at least tens of millions of dollars annually (Acxiom Corporation 2008). This problem is exacerbated by the common practice of many firms selling small ticket items with a 'buy now, pay later' payment option. Given the relatively small amount involved per customer, direct marketing businesses typically spend limited effort to recover bad debts, unlike businesses selling big ticket items who may expend substantial recovery effort through repossession and debt collection agencies. This increases the importance of accurately forecasting bad debt risk for direct marketers. For example, Fingerhut, a direct marketing company that targets lower income households, estimated that its bad debt customers accounted for about 15% of its base (West 2006). The cost of bad debt soon became as high as 40% of Fingerhut's sales, and the parent company (Federated) was forced to take a $150 million charge for bad debts in the second quarter of 2000. Fingerhut began layoffs in 2001 and discontinued most operations in the next year.

We believe that one reason why previous academic research has not studied bad debt risk in customer acquisition is that historical information on bad debt behavior across many campaigns from multiple firms is hard to come by. The data used for a research study are typically from one direct marketing firm.[3] This is also a quandary for direct marketing firms when they buy mailing lists for the purpose of customer acquisition. Since by definition these lists do not contain their own customers, firms often do not have information on transaction and payment history of individuals on the purchased list. Limited by the information available, firms seem to have their hands tied in terms of being able to take the crucial step in managing acquisition risk by identifying potential bad debt customers.

A fairly new and interesting solution has emerged in the industry to help direct marketing firms address this issue. This solution is in the form of *co-operative databases*, which can be described as *pooling of data across direct marketing firms by a third party vendor in order to provide a broader view of customer transactions and thus enable direct marketing firms that have access to the co-operative databases to refine their promotional strategies*. Basically contributing direct marketing firms give their historical customer transaction data to an independent co-operative database firm, which provides data warehousing and data analytic services to the contributors. By keeping track of customers' transaction and payment history across direct marketing firms, such co-operative databases can offer a major advantage in customer acquisition in that they provide customer information to help direct marketing firms manage potential bad debt risk.

To analyze the optimal selection for direct mail under bad debt risk, we first develop an analytical model to reveal how the benefit from incorporating bad debt risk varies according to the characteristics of a direct mail campaign. We then apply our model to empirical data and study consumers' responses to a direct mail campaign and their bad debt behavior. Our empirical analysis is based on data from a direct mail campaign for a magazine offer in which 3.56% of the targeted consumers responded and paid while another .86% responded but never paid. Given that a significant portion of the responses eventually became bad debts, it is critical to account for bad debt risk and screen out potential bad debtors. We match the target consumers with a co-operative database and obtain their historical purchase and payment information. Such information allows us to identify variables that predict consumers' responses to the campaign and their bad debt behavior. Our model accurately predicts consumers' decisions and ranks consumers according to the expected return. By targeting those consumers with a positive expected return only, the proposed targeting scheme is effective in screening out potential bad debtors.

Instead, if the firm follows a traditional binomial response model that focuses on response vs. non-response, consumers will only be selected based on their response probabilities. In effect the model will overlook bad debt risk and miscalculate the true profitability of a customer because those who are quick to respond could well be bad debtors. We show that our proposed targeting scheme would generate 5.3% more profits than such a suboptimal targeting scheme for the specific magazine offer under study. Note that the average loss from a bad debt is limited to $10 for this magazine offer because the firm can stop sending magazine issues as soon as a bad debt consumer is identified. In other product categories a bad debt is often more costly, in which cases it will be even more beneficial to apply the proposed targeting scheme and alleviate bad debt risk. For example, at an average loss of $15 from a bad debt, the proposed targeting scheme could improve profits by 23.7% over the suboptimal targeting scheme. Therefore, our results highlight the benefit of accounting for bad debt risk in direct mail, and such benefit increases when the loss from a bad debt becomes larger.

We utilize historical purchase and payment information from a co-operative database to predict consumer choices among non-response, paid response and bad debt. To demonstrate the value of co-operative databases, we try to predict consumer choices without accessing the information from the co-operative database. The alternative to using a co-operative database would be to use other information available such as demographic variables, and we therefore collect these variables from census data based on consumer addresses. Demographics can explain some variation in bad debt behavior, but are much less effective in screening out bad debt consumers. In fact, using demographics the firm can avoid a mere 2.5% of bad debtors and increase its profit by .9%, while in contrast, using information from the co-operative database the firm can avoid 38.2% of bad debtors and improve its profit by 5.3%.

Literature Review

Although there are many papers devoted to the study of customer relationship management and customer retention, there has been a comparatively lesser degree of focus on customer acquisition. Related to customer acquisition, Blattberg and Deighton (1996) present a managerial approach to balance

[3] Only recently have there been studies of competitive direct marketing (e.g. Van Diepen, Donkers, and Franses 2009).

resource allocation between acquisition and retention, while Reinartz, Thomas, and Kumar (2005) develop an empirical modeling framework to achieve the same objective by extending the work of Thomas (2001). Lix, Berger, and Magliozzi (1995) discuss the use of commercial databases and statistical models for customer acquisition by modeling response probability. Lewis (2006a) argues that promotionally acquired customers have lower repeat buying rates and smaller lifetime values, and similarly Villanueva, Yoo, and Hanssens (2008) show that customer acquired through marketing efforts are about half as valuable as those acquired through word of mouth. Lewis (2006b) studies the effect of shipping fees on online acquisition and retention and finds that base shipping fees affect customer retention more than customer acquisition, with the latter being more sensitive to order size incentives. Schweidel, Fader, and Bradlow (2008) study the relationship between a prospect's time to acquire a cable subscription service and the subsequent duration that the same customer retains the service using a split hazard model. Recent research has started reemphasizing the importance of acquisition. For example Voss and Voss (2008) argue that customer acquisition strategies become more effective than retention strategies as the number of competitors increases. Arnold, Fang, and Palmatier (2011) show that a firm's customer acquisition orientation (as compared to retention orientation) enhances the performance of its radical innovation.

This paper also contributes to the extensive literature of targeted marketing and direct marketing (e.g., Bult and Wansbeek 1995; Gönül and Shi 1998) by demonstrating how firms can incorporate bad debt risk into their targeting algorithms. Recent developments in this literature have studied undesirable behaviors by consumers conditional on purchase. To the extent that such behaviors may be exhibited during customer acquisition, they constitute an important aspect of acquisition risk management. Several recent papers study the effect of product return risk on firm profitability (e.g. Anderson, Hansen, and Simester 2009; Ofek, Katona, and Sarvary 2011; Petersen and Kumar 2009; Shulman, Coughlan, and Savaskan 2011) and identify conditions under which returns are undesirable to firms. Zhao, Zhao, and Song (2009) develop a model to predict risk types in the credit card market based on prior transactions and payment behavior. Braun and Schweidel (2011) separate nonpayment/abuse into a unique category in their study of multiple causes of churn for a telecommunications service provider. These papers rely on internal transaction and payment data in their prediction model, which is natural in the context of customer retention. By contrast, we focus on the domain of customer acquisition, where it is crucial to incorporate the information from co-operative databases.

Conceptual Framework

Fig. 1 describes our conceptual framework for acquisition risk management using co-operative databases. We extend the framework proposed by Ryals and Knox (2005) to the customer acquisition context. Typically assessing the economic value of a customer requires forecasts of revenue and cost (Jain and Singh 2002). However, as observed by Ryals and Knox (2005), this assessment needs to be expanded to incorporate a customer-specific risk forecast. While their discussion focuses on risk adjustment in customer retention decisions, their insight regarding customer-specific risk management is also applicable to customer acquisition. Theoretically acquisition is almost entirely within the control of the firm (Hansotia and Wang 1997), but it is important to recognize that firms face considerable uncertainty in selecting the right consumers to target. Therefore, a key component of the targeting process is the development of an acquisition risk profile forecast for each prospective consumer. Acquisition risks can arise due to various undesirable behaviors on the part of consumers who are targeted and respond, chief among these risks being the risk of bad debt, and the risk of product return. While we conceptualize acquisition risks in general, in our empirical analysis we confine our model to bad debt risk and leave it to further research to calibrate the relative importance of different risk factors. As we demonstrate later in the empirical results, it is crucial for firms to incorporate an acquisition risk forecast into the calculation of the economic value of a customer and thus into their targeting decisions.

Next we address the process by which the acquisition risk forecast is to be obtained. It is often challenging to forecast a prospect's acquisition risk profile when the prospect has not transacted before with the firm in other contexts, e.g., in buying a different product from the firm. The key resource to overcome this lack of information for obtaining this forecast is the historical transaction data for each consumer that is provided by contributors to the co-operative database. The data may include past purchase and payment behaviors of targeted customers in various product categories. Because a co-operative database has data that span several product categories, it enables a broader view of a customer's behavior than the profile that a single firm can obtain from its own database. Typically the co-operative database firm will collate its information into a predictor selection process that will involve creation of traditional Recency–Frequency–Monetary (RFM) Value type predictors (e.g. Fader, Hardie, and Lee 2005; Gönül and Shi 1998). To account for acquisition risks, we need predictors that reflect relevant historical behaviors including traditional purchase behaviors as well as undesirable bad debt and return behaviors. Thus predictors could include variables such as the recency and frequency of past bad debts. These predictors will be used to construct an acquisition risk profile forecast for each consumer, and this forecast becomes an input into the evaluation of the economic value of the consumer and therefore into the firm's decision on whether to acquire a given consumer.

Optimal Selection for Direct Mail Under Acquisition Risks

In the direct mail industry, firms have two broad sources from which they can select names to mail, one consisting of their existing customers, collectively called 'house' names, and the other consisting of names rented from third parties, typically called 'outside' or 'acquisition' lists. For names on acquisition lists, firms do not observe the past purchase and

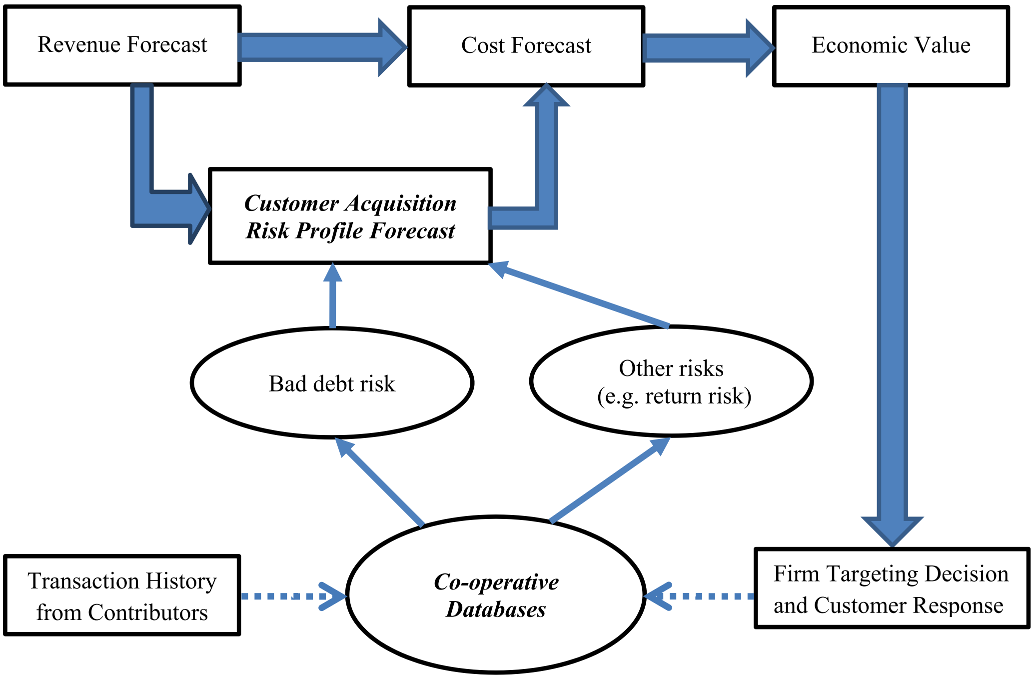

Fig. 1. Conceptual framework for managing acquisition risk using co-operative databases.

payment history, leading to significant uncertainty in predicting consumer purchase decisions and subsequent payment behaviors. Direct marketing firms selling small ticket items such as books, music CDs, DVDs, and magazines typically offer the payment option of 'bill me later', or 'buy now, pay later', which allows consumers to pay for the product after receiving it. This encourages impulse purchases and increases the response rate dramatically, as direct marketers have known for a long time. The tradeoff that marketers face that can offset the increased response rate is that this leads to an increased risk of non-payment. Because the items are small-ticket, firms tend to write off non-payment as bad debt rather than turning to credit collection and repossession agencies as would be the case for large ticket items. Therefore, the accurate forecast of bad debt risk becomes even more important when direct marketing firms select consumers for acquisition campaigns.

Previous studies have formulated the outcome of a direct mail as a binary choice of response and non-response. A positive return is assumed if a consumer responds to a direct mail campaign. Typically marketers would identify a threshold on the probability of response based on breakeven criteria (Hansotia and Wang 1997), and then select households that have a probability of response greater than the threshold. For example, let P_0 be the probability of non-response and P_1 be the probability of response ($P_1 = 1 - P_0$). If we plot P_1 on a line segment as shown in Fig. 2, all consumers to the right of the threshold Z will be targeted, while all consumers to the left of the threshold Z will not be targeted.

We can expand this framework and introduce another alternative that captures undesirable behavior of acquired consumers such as bad debt. Let the probability of such undesirable behavior be P_2. A discrete choice model can be applied to study consumer choices between different options. In our empirical setting we focus on bad debt risk and use options 0, 1 and 2 to indicate non-response, paid response and bad debt respectively. If a consumer i with characteristics x_i chooses an option $j \in \{0, 1, 2\}$ after receiving a direct mail, the consumer's latent utility function is specified as

$$u_{ij} = x_i{'}\beta_j + \varepsilon_{ij}. \tag{1}$$

Assume that the random error vector ε_{ij} follows a joint distribution with density function $f_\varepsilon(\cdot)$. The probability that consumer i chooses option j can be expressed as

$$P_{ij} = P\left[u_{ij} > u_{ik}, k \neq j | x_i\right] = P\left[\varepsilon_{ij} - \varepsilon_{ik} > -x_i\beta_j + x_i\beta_k, k \neq j | x_i\right]. \tag{2}$$

Given that P_0, P_1 and P_2 are collinear ($P_2 = 1 - P_0 - P_1$), the response probabilities will correspond to the triangular area in the (P_0, P_1) space as shown in Fig. 3. Following Equation (2), each consumer i with characteristics x_i can be mapped to a point within that area. Thus Equation (2) defines a mapping from the x-space of consumer characteristics to the two-dimensional (P_0, P_1) space. Consumers are distributed in the (P_0, P_1) space

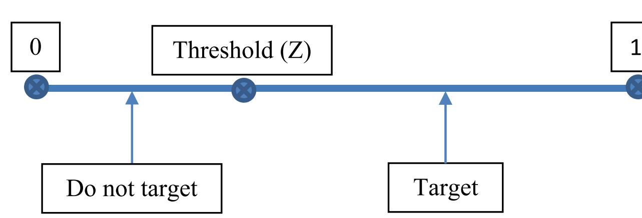

Fig. 2. Traditional targeting scheme with a threshold on response probability.

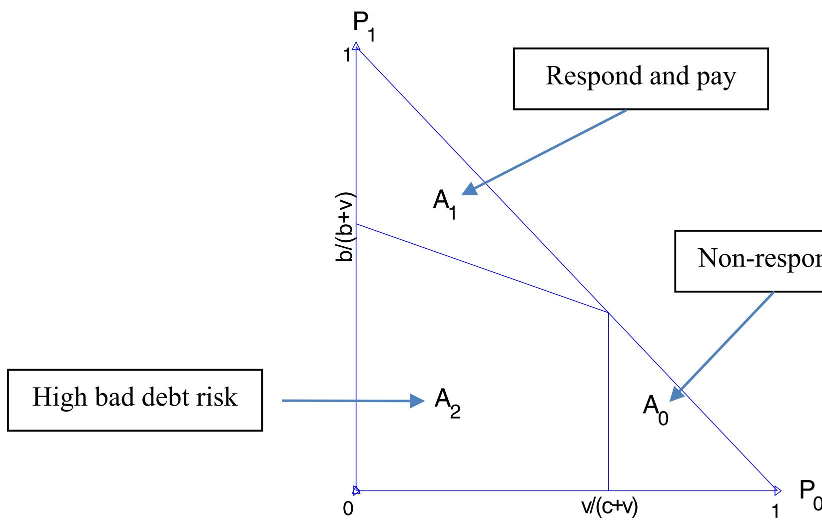

Fig. 3. Partition of the probability space spanned by (P_0, P_1).

according to a density function $f(P_0,P_1)$. If there is no bad debt risk, i.e. $P_2 \equiv 0$, then all consumers are located on the line segment $P_0 + P_1 = 1$ in Fig. 3, and the model will be reduced to the traditional binary choice between response and no response as a special case.

The set of payoffs for the different outcomes {0, 1, 2} are specified as $\{-c, v, -b\}$, where c is the mailing cost of direct mail to a consumer, v is the profit from a paid purchase, and b is the loss from a bad debt. Note that $b > c$ because the loss from a bad debt includes the mailing cost. Given the probabilities for different outcomes of a direct mail and their relationship $P_2 = 1 - P_0 - P_1$, the expected return from a consumer (suppressing subscript i) is

$$r = -cP_0 + vP_1 - bP_2 = (b-c)P_0 + (b+v)P_1 - b. \tag{3}$$

We can rank consumers according to their expected return, and then set a threshold or cutoff point to select a subset of consumers based on desired criteria. To maximize the total expected return, a firm would target all consumers that satisfy $r > 0$, or equivalently

$$P_1 > \frac{b}{b+v} - \frac{b-c}{b+v} P_0. \tag{4}$$

This corresponds to the area A_1 indicated in Fig. 1, corresponding to a high probability of paid response and low probability of bad debt.

For comparison we now consider the scenario in which the issue of bad debt is overlooked, and as a result, the option $j = 2$ is mistakenly assigned the same payoff as the option $j = 1$. In other words, the firm does not differentiate between paid response and bad debt, and assumes that the same return would accrue from paid response and bad debt customers, so that the expected return from a customer would be naively assumed to be:

$$r* = -cP_0 + vP_1 + vP_2 = -cP_0 + v(1 - P_0). \tag{5}$$

In this case the target consumers would satisfy $r^* > 0$, or equivalently

$$P_0 < \frac{v}{c+v}. \tag{6}$$

Equation (6) defines the area $A_1 + A_2$ in Fig. 3, while consumers in the complement area A_0 would not be targeted due to low probability of response.[4] However, as indicated by Equation (4) above, it is unprofitable to target the consumers in the area A_2 because of high bad debt risk. Thus a naïve firm would incur a loss from targeting bad debt customers, which we call the *cost of bad debt*.

In summary, the area A_1 in Fig. 3 represents the profitable consumers who are likely to respond and pay, whereas the area A_0 represents those consumers who are unlikely to respond, and the area A_2 represents the consumers with high bad debt risk. Given the distribution of consumers in the (P_0, P_1) space, the fraction of consumers within each area j can be obtained by integrating over the joint density function of P_0 and P_1:

$$Q_j = \iint_{A_j} f(P_0, P_1) dP_0 dP_1, \quad j = 0, 1, 2. \tag{7}$$

The expected return from each area j can be written as

$$R_j = \iint_{A_j} r(P_0, P_1) f(P_0, P_1) dP_0 dP_1, \quad j = 0, 1, 2. \tag{8}$$

[4] It can be verified that the three straight lines $P_0 + P_1 = 1$, $P_0 = \frac{v}{c+v}$ and $P_1 = \frac{b}{b+v} - \frac{b-c}{b+v} P_0$ intersect at the same point $(\frac{v}{c+v}, \frac{c}{c+v})$. Consequently the triangular area consists of A_0, A_1 and A_2 as depicted in Fig. 3. Otherwise the shape of A_0, A_1 and A_2 may be different.

By definition we would expect a positive R_1 but negative R_0 and R_2.

Next we analyze the comparative statics of these quantities. Note that

$$Q_0 = \int_{\frac{v}{c+v}}^{1} dP_0 \int_0^{1-P_0} f(P_0, P_1) dP_1. \tag{9}$$

Clearly Q_0 depends on the ratio of v to c, not individual values of v, c or b. If we fix c and let v vary, we can show that

$$\frac{\partial Q_0}{\partial v} = -\frac{c}{(c+v)^2} \int_0^{\frac{c}{c+v}} f\left(\frac{v}{c+v}, P_1\right) dP_1 < 0.$$

Intuitively, if the profit from a paid response becomes higher relative to the mailing cost, more consumers will be targeted, thus reducing the fraction of consumers in the area A_0. This can also be seen graphically from Fig. 3. If v increases relative to c, then the vertical line $P_0 = \frac{v}{c+v}$ moves to the right, reducing the fraction of consumers in A_0.

Similarly we can inspect how Q_1 and Q_2 change with b, c and v. Note that

$$Q_1 = \int_0^{\frac{v}{c+v}} dP_0 \int_{\frac{b}{b+v} - \frac{b-c}{b+v} P_0}^{1-P_0} f(P_0, P_1) dP_1;$$
$$Q_2 = \int_0^{\frac{v}{c+v}} dP_0 \int_0^{\frac{b}{b+v} - \frac{b-c}{b+v} P_0} f(P_0, P_1) dP_1.$$

We can fix c and$v$ and vary b, the loss from a bad debt. We can show that

$$\frac{\partial Q_1}{\partial b} = -\int_0^{\frac{v}{c+v}} \frac{v-(v+c)P_0}{(b+v)^2} f\left(P_0, \frac{b}{b+v} - \frac{b-c}{b+v} P_0\right) dP_0;$$
$$\frac{\partial Q_2}{\partial b} = \int_0^{\frac{v}{c+v}} \frac{v-(v+c)P_0}{(b+v)^2} f\left(P_0, \frac{b}{b+v} - \frac{b-c}{b+v} P_0\right) dP_0.$$

Within the integration interval of $\left(0, \frac{v}{c+v}\right)$, we have $v - (c+v)P_0 > 0$. Therefore $\frac{\partial Q_1}{\partial b} < 0$ and $\frac{\partial Q_2}{\partial b} > 0$. Intuitively, as the loss from a bad debt increases, a smaller fraction of consumers will be targeted, while more consumers will be screened out due to bad debt risk. Again this can be visualized in Fig. 3—As b (loss from a bad debt) increases relative to v (profit from a purchase), the line $P_1 = \frac{b}{b+v} - \frac{b-c}{b+v} P_0$ will have a larger intercept $\frac{b}{b+v}$, thus reducing the area corresponding to paid response customers (A_1) but increasing the area for bad debt customers (A_2).

In previous discussions we use $f(P_0,P_1)$ to represent the distribution of consumers in the (P_0, P_1) space, but remain agnostic on the functional form of $f(P_0,P_1)$. To get more detailed insights we need to assume a specific density function for $f(P_0, P_1)$. For simplicity we assume that consumers are uniformly distributed within the triangular area in the (P_0, P_1) space, in which case $f(P_0, P_1) \equiv 2$ and the fraction of consumers in each area has a closed form[5]:

$$Q_0 = \int_{\frac{v}{c+v}}^{1} dP_0 \int_0^{1-P_0} 2dP_1 = \frac{c^2}{(c+v)^2};$$
$$Q_1 = \int_0^{\frac{v}{c+v}} dP_0 \int_{\frac{b}{b+v} - \frac{b-c}{b+v} P_0}^{1-P_0} 2dP_1 = \frac{v^2}{(b+v)(c+v)};$$
$$Q_2 = \int_0^{\frac{v}{c+v}} dP_0 \int_0^{\frac{b}{b+v} - \frac{b-c}{b+v} P_0} 2dP_1 = \frac{cv}{(c+v)^2} + \frac{bv}{(b+v)(c+v)}.$$

Note that Q_0 is independent of b but depends on the ratio of c to v, while both Q_1 and Q_2 depend on c/v and b/v. Therefore we hold $c = 1$ and plot Q_0 against v in Fig. 4(a). As we have shown above, Q_0 decreases as v increases. If v becomes relatively large, e.g. for $v > 10$, a firm would only skip a very small fraction of consumers due to likelihood of non-response. Because in this case the profit from a potential paid response dominates the mailing cost, a consumer will be skipped only if she is extremely unlikely to respond. Similarly in Fig. 4(b), we hold $c = 1$ and $v = 10$ and plot $Q_0/Q_1/Q_2$ against b. As expected Q_0 is independent from b, while Q_1 decreases and Q_2 increases as b increases. That is, if the potential loss from a bad debt increases, marketers have to be more cautious so that more consumers will be screened out due to bad debt risk and fewer consumers will be selected for direct mail.

Now we examine the expected return from each area. Again a closed form can be obtained based on a uniform distribution of consumers in the (P_0, P_1) space:

$$R_0 = \int_{\frac{v}{c+v}}^{1} dP_0 \int_0^{1-P_0} [(b-c)P_0 + (b+v)P_1 - b] 2dP_1 = -\frac{c^3(b+c+2v)}{3(c+v)^3};$$
$$R_1 = \int_0^{\frac{v}{c+v}} dP_0 \int_{\frac{b}{b+v} - \frac{b-c}{b+v} P_0}^{1-P_0} [(b-c)P_0 + (b+v)P_1 - b] 2dP_1 = \frac{v^3}{3(b+v)(c+v)};$$
$$R_2 = \int_0^{\frac{v}{c+v}} dP_0 \int_0^{\frac{b}{b+v} - \frac{b-c}{b+v} P_0} [(b-c)P_0 + (b+v)P_1 - b] 2dP_1$$
$$= -\left[\frac{bcv}{(c+v)^2} + \frac{(b-c)^2 v^3}{3(b+v)(c+v)^3}\right]$$

As expected $R_0 < 0$, $R_1 > 0$, and $R_2 < 0$. To see how the expected returns vary according to different values of v and b, we fix $c = 1$ and plot them against b for different levels of v in Fig. 5. Fig. 5(a) shows that R_0 declines linearly with respect to b, and the rate of decline is slower for larger values of v. In general the small magnitude of R_0 indicates that the consequence of targeting those consumers who are unlikely to respond is not substantial, especially when the mailing cost is small compared with the profit from a paid response. Fig. 5(b) shows how R_1 changes with b at different levels of v. Because R_1 is the expected return based on the optimal targeting scheme, we can see that a direct mail campaign will generate a higher

[5] Admittedly the uniform distribution is a simplifying assumption used to illustrate the comparative statics. A more realistic distribution will be used in our subsequent empirical analysis.

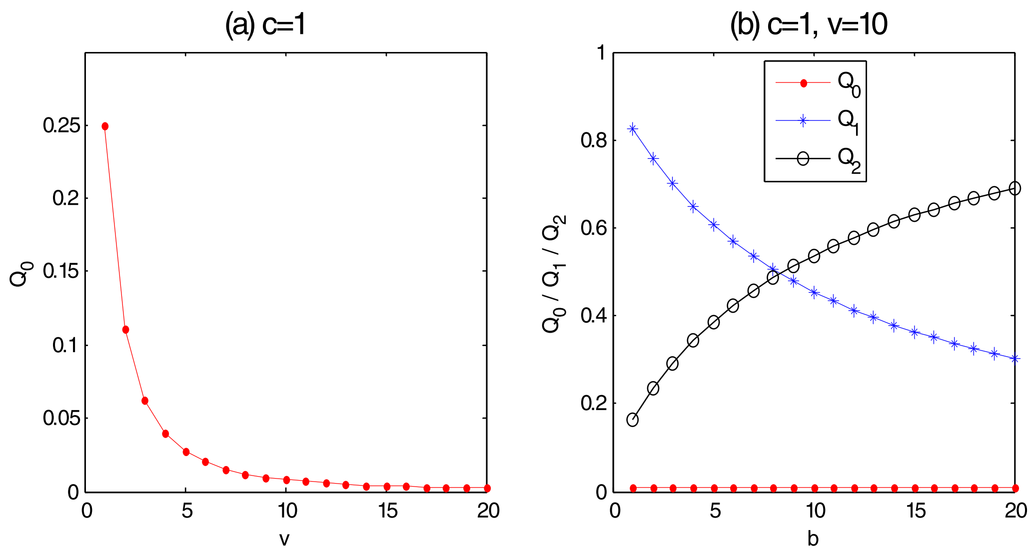

Fig. 4. Fraction of Consumers in each area (Q_0, Q_1, Q_2).

return when the loss from a bad debt is lower, and when the profit from a paid response is higher. Fig. 5(c) is a plot of R_2, the expected loss from targeting those consumers with high bad debt risk. It can also be interpreted as the expected savings by taking into account bad debt risk compared with a naïve targeting scheme that overlooks such risk. The expected loss will be deeper when a bad debt costs more and when a paid response brings less profit. Compared with Fig. 5(b), we can see that R_2 is more sensitive to the changes in b than R_1 is, especially for smaller values of v.

Empirical Analysis

An optimal targeting scheme involves identifying the profitable consumers in A_1 by following Equation (4). Given that Equation (4) requires the response probabilities P_0 and P_1 for each consumer, we need to calculate the response probabilities defined by Equation (2), which in turn depend on the distribution of the error term in the utility function (1). Therefore, for empirical applications the key is to specify $f_\varepsilon(\cdot)$, the distribution of the error term ε_{ij} in Equation (1). Different assumptions on the error term distribution lead to different model specifications such as a multinomial probit model (if ε_{ij} follows a joint normal distribution) or a multinomial logit model (if ε_{ij} follows independent type-I extreme-value distributions). As described by Anderson, de Palma, and Thisse (1992), the logit demand system has a number of theoretical and practical advantages, many of which derive from the closed-form expressions for the probability of choosing any alternative. Thus we employ a logit demand system with appropriate covariates to model consumers' choice probabilities P_0, P_1 and P_2.

Using data from a direct mail campaign, we first estimate a multinomial logit model to characterize consumers' choice probabilities, so that we can calculate the expected return from each consumer and target those profitable consumers following Equation (4). Because this targeting scheme takes into account bad debt risk, it is expected to outperform the traditional binary response model that ignores bad debt risk. We thus compare the proposed targeting scheme with the traditional one, and illustrate the benefit of incorporating bad debt risk. The implementation details and results are discussed in the following subsections.

Data and Estimation

We use the data from a co-operative database company in the U.S. about a direct mail campaign for a magazine offer. The offer was sent to a list of more than 3.6 million consumers. Approximately 3.56% of consumers responded to the campaign and paid for the product, while another .86% responded but eventually did not pay for the product, resulting in bad debts. Because almost 20% of the responses turned into bad debts, it is important for the firm to take into account bad debt risk and avoid potential bad debtors.

By matching the target consumers with a co-operative database, we obtain historical information on consumers' past purchases and payments. The co-operative database is compiled by pooling transactional data across various catalog marketing firms (contributors). Based on the information from the co-operative database we set up a number of variables to predict a consumer's purchase and payment behavior. We follow the Recency–Frequency–Monetary Value framework that has been widely used in the direct marketing literature to include three variables: days since last order (recency), number of orders in the last 2 years (frequency), and dollar amount spent in the last 2 years (monetary value). In addition to the three RFM-type variables often used to predict response, we also construct two variables that reflect the historical bad debt profile of a customer across different campaigns: number of bad debts in the past 2 years, and percent of orders that resulted in bad debts in the past 2 years. The summary statistics of these variables are presented in Table 1. We calibrate our model

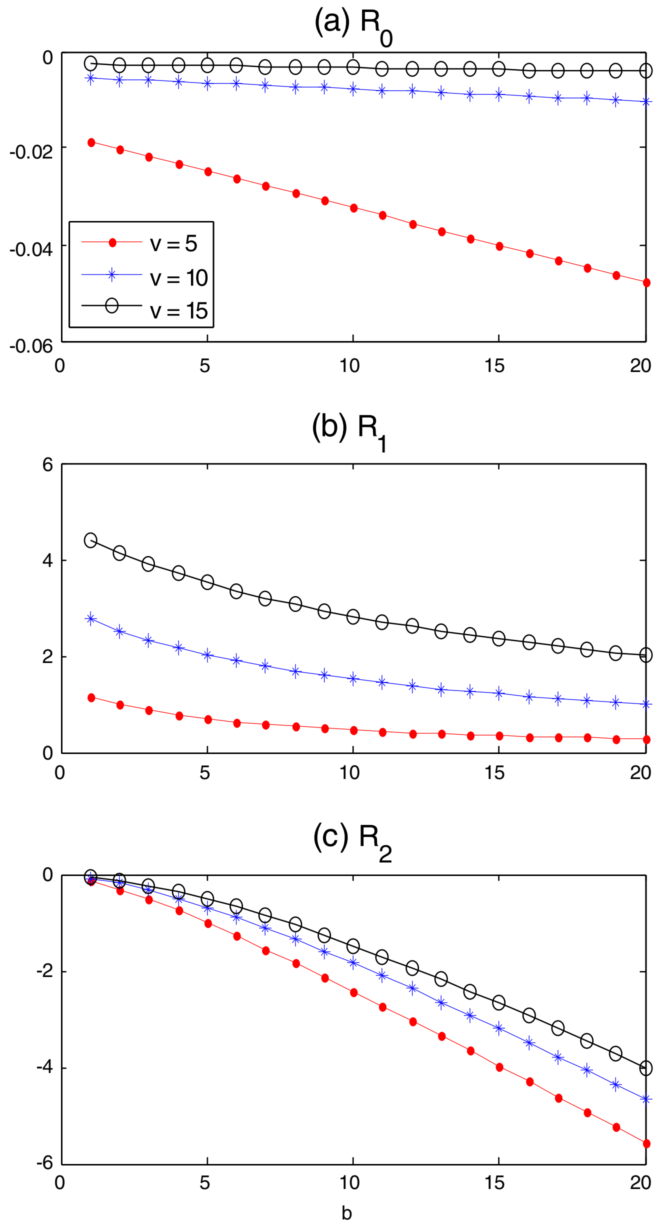

Fig. 5. Expected return (R_0, R_1, R_2) from each area.

using a 10% random sample and then use another 10% for validation.

We estimate the model parameters in the latent utility function (1) by maximizing the joint likelihood of the observed outcomes. Note that there are three alternatives in the choice set with 0, 1 and 2 correspond respectively to non-response, paid response and bad debt. Following the standard practice in discrete choice models we normalize the option 0 to have a zero mean utility, i.e., $u_{i0} = \varepsilon_{i0}$, because we can only identify the differences between the utilities of these options based on consumers' discrete choices. For the purpose of normalization we take the log transformation of the first four variables (plus one) during the maximum likelihood estimation.

In the first column of Table 2, we present the estimates for model parameters based on the calibration sample. Consistent with previous studies, we find that if a consumer has not ordered recently, she is less likely to respond to the mail campaign. On the other hand, a larger number of orders in the last two years (indicating higher frequency of purchase) increase the purchase probability. After controlling for the number of orders, the total spending has a negative effect on the purchase probability, possibly due to the fact that a consumer who has already spent a lot of money would probably have satisfied their needs in related categories and would have less incentive to make further purchases.

To our knowledge this study is the first to measure the impact of RFM variables on bad debt risk, a result that contributes to the empirical direct marketing literature. We find that days since last order (recency) has a negative impact on bad debt risk, while the number of recent orders (frequency) has a positive impact. This reflects the common wisdom in the catalog marketing industry that *those who are most likely to*

Table 1
Summary statistics of key variables.

Variable	Full sample		Calibration sample		Validation sample	
	Mean	St. dev.	Mean	St. dev.	Mean	St. dev.
Paid response (%)	3.56	48.68	3.51	48.42	3.57	48.70
Bad debt (%)	.86	24.30	.85	24.12	.87	24.37
Days since last order	301.51	932.88	299.76	925.20	300.52	929.67
Number of orders in the past 2 years	2.68	9.27	2.70	9.65	2.68	9.10
Dollar amount spent in the past 2 years	32.26	145.95	32.56	150.05	31.88	142.88
Number of bad debts in the past 2 years	.47	3.73	.48	3.84	.46	3.51
Percent of bad debts in the past 2 years	8.16	53.82	8.32	54.58	8.22	54.26

Table 2
Parameter estimates.

Variable	Proposed model	Traditional model
For option 1: paid response		
Intercept	−3.230 **	−2.913 **
Days since last order	−.059 **	−.074 **
Number of orders in the past 2 years	.355 **	.341 **
Dollar amount spent in the past 2 years	−.068 **	−.106 **
Number of bad debts in the past 2 years	−.024	.160 **
Percent of bad debts in the past 2 years	−.064	.302 **
For option 2: bad debt		
Intercept	−4.188 **	
Days since last order	−.130 **	
Number of orders in the past 2 years	.237 **	
Dollar amount spent in the past 2 years	−.263 **	
Number of bad debts in the past 2 years	.821 **	
Percent of bad debts in the past 2 years	.591 **	
Number of observations	361,210	361,210
Log likelihood	−71,583.9	−64,571.3

The log likelihood functions of the two models are constructed differently and hence cannot be directly compared.
** Significance at .01 level.

respond are often least likely to pay. If a consumer has spent more recently (monetary value), the bad debt risk becomes lower. The two bad debt variables constructed from the co-operative database, number of recent bad debts and percent of recent bad debts, turn out to be strong predictors of bad debt risk. Consumers with high values of these variables in the last two years have a higher probability of bad debt in the focal campaign.

Optimal Selection of Target Consumers

Now we consider the optimal selection of target consumers for direct mail. We ensure that our analysis reflects empirical reality by obtaining the actual profit and cost information from the firm, according to which the mailing cost c = \$.50, the loss from a bad debt b = \$10, and the profit from a purchase v = \$20 for the magazine offer under analysis. Using the parameter estimates in Table 2, we can calculate the expected return from each consumer and target those with a positive expected return. Our analysis shows that, out of the 361,210 consumers in the validation sample, a positive return is expected from 285,232 consumers. The realized return from this group of consumers is aggregated to be \$56,777, 5.3% higher than the total realized return of \$53,913 from targeting all 361,210 consumers in the validation sample.

To analyze the performance of our targeting scheme, we divide the consumers into deciles according to their expected returns. For each group, we report the actual profits and percentages of non-response, paid response and bad debt. The results in Table 3 show that our method accurately ranks consumers based on their profit potentials, with the expected rates of non-response, paid response and bad debt being very close to the observed rates. As the firm mails deeper into the groups from the top to the bottom, it will be confronted with diminishing incremental profit for each additional group.

It is also worth noting that the ranking of consumer groups according to their expected returns is not monotonic in non-response rate, paid response rate, or bad debt risk. Instead, it is the combination of these factors that determines the attractiveness of targeting a specific consumer group, although the top group does have the highest paid response rate, and the bottom group does have the highest bad debt risk.

Comparing with a Traditional Targeting Scheme

For comparison we consider a naïve firm that overlooks bad debt risk and fails to differentiate bad debt consumers from paid responses. According to Equation (6), the firm would focus on the issue of non-response, and screen out those consumers with a non-response probability that is higher than 97.56% (i.e., $P_0 > 97.56\%$), or equivalently a response probability that is lower than 2.44%. For the direct mail campaign that is under consideration, the 3.56% purchase rate combined with the .86% bad debts leads to an average response probability of 4.42%. Compared with the 2.44% threshold, we would expect very few consumers, if any, to be screened out by this targeting scheme.

In fact, we estimate a traditional binary logit model of response vs. non-response using the same calibration sample and the same set of variables. The parameter estimates are provided in the second column of Table 2. Note that past bad debts have a *positive* effect on the response probability because bad debts are mistakenly treated as a response by the naïve firm. We then calculate the expected payoff from each consumer in the validation sample following Equation (5). The results suggest that, out of the 361,210 consumers in the validation sample, 361,200 should be targeted because a positive return is expected from them. Therefore, if the issue of bad debt is ignored, almost all consumers in the validation sample should be targeted and the total realized return would be \$53,918, almost the same as the total return of \$53,913 from all consumers.

Again we divide the consumers into deciles according to their expected returns, and present the results in Table 4. Because in this case the firm focuses on response vs. non-response only, groups are strictly ranked by their predicted response probabilities. Clearly this scheme does not rank the groups correctly in terms of profitability. For example, the top

Table 3
Decile analysis of the proposed targeting scheme.

Decile	Profit	Realized (%)			Expected (%)		
		No response	Paid response	Bad debt	No response	Paid response	Bad debt
Top	17,423	94.2	5.2	.6	94.2	5.2	.7
2	9465	95.3	4.0	.7	95.1	4.3	.6
3	6975	95.7	3.6	.6	95.4	3.9	.6
4	7213	95.8	3.7	.6	95.7	3.7	.6
5	5815	95.7	3.6	.7	95.9	3.5	.7
6	5380	95.9	3.4	.6	96.0	3.3	.7
7	2713	96.2	3.1	.7	96.3	3.1	.7
8	1560	95.8	3.1	1.0	96.2	2.9	.8
9	1153	96.3	3.0	.8	96.6	2.6	.7
Bottom	−3783	94.5	3.0	2.5	95.0	2.7	2.4

Table 4
Decile analysis of the traditional targeting scheme.

Decile	Profit	Realized (%)			Expected (%)		
		No response	Paid response	Bad debt	No response	Paid response	Bad debt
Top	6733	93.1	4.6	2.3	93.3	6.7	–
2	8123	94.6	4.2	1.3	94.5	5.5	–
3	8318	94.9	4.1	1.0	95.0	5.0	–
4	8895	95.2	4.0	.8	95.4	4.6	–
5	6588	95.8	3.6	.6	95.7	4.3	–
6	5360	96.0	3.4	.6	95.9	4.1	–
7	5083	96.0	3.4	.6	96.2	3.8	–
8	1900	96.4	3.0	.6	96.4	3.6	–
9	1770	96.4	3.0	.6	96.7	3.3	–
Bottom	1145	96.7	2.8	.5	96.9	3.1	–

decile is not the most profitable due to its highest bad debt risk. Its profit of $6733 is less than that of the next three deciles, and in fact it is the fourth decile in Table 4 that has the highest profits of $8895. In contrast, results in Table 3 demonstrate that our proposed targeting scheme that incorporates bad debt risk ranks groups accurately in terms of profitability.

Recall that our proposed targeting scheme would lead to a realized return of $56,777, which is 5.3% higher than the realized return of $53,918 from the naïve targeting scheme that ignores bad debt risk. A comparison of these two schemes is provided in the first two rows of Table 5. We can see that the proposed targeting scheme that uses information from a co-operative database is able to screen out a significant portion of the bad debt consumers and improve the profitability.

Our empirical context is a direct mail campaign for a magazine offer, for which the average loss from a bad debt is limited to $10 (i.e. $b = 10$) because the firm can stop sending more issues as soon as a bad debt consumer is identified. For other single-item offers (e.g. book, clothing) the loss from a bad debt can be significantly higher. As we show previously in the analytical model, if b (loss from a bad debt) becomes larger, both Q_1 (fraction of paid responses) and R_2 (return from bad debt consumers) would decrease, i.e., less consumers would be targeted and more savings would be generated by accounting for bad debt risk. In such cases it is even more important to account for bad debt risk. For example, consider $b = 15$ instead of $b = 10$. Under the suboptimal targeting scheme that ignores bad debt risk, again we would target 361,200 consumers in the validation sample, now realizing a total return of $38,218. In contrast, if we follow the proposed targeting scheme that accounts for bad debt risk, we would only target 260,407 consumers from whom the realized return would be $47,270, $9052 or 23.7% higher than the realized return based on the suboptimal targeting scheme.

In Fig. 6, we compare the realized returns from alternative targeting schemes and plot the gain from the proposed targeting scheme over the suboptimal targeting scheme for various values of b and v. As b (loss from a bad debt) increases, it is even more beneficial to follow the proposed targeting scheme and take into account bad debt risk. Furthermore, for a smaller value of v (profit from a purchase), it also becomes more costly to ignore bad debt risk by following the suboptimal targeting scheme. This pattern is similar to the one predicted by our analytical model and depicted in Fig. 5(c), except that in Fig. 6 we plot the gain as positive numbers, while in Fig. 5(c) we plot R_2 as negative numbers.

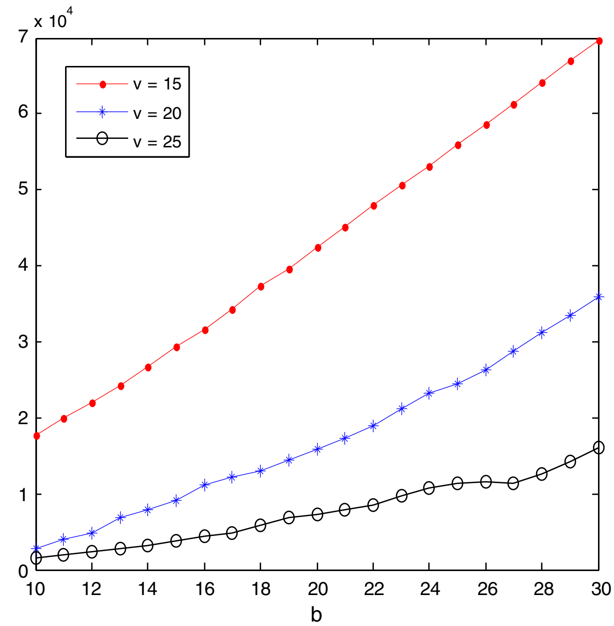

Fig. 6. Gain from the proposed targeting scheme.

Robustness to Sampling

Our results highlight the significant advantage of our proposed targeting scheme that alleviates bad debt risk. Such advantage increases with the relative importance of a bad debt, namely when the loss from a bad debt increases or the profit from a purchase decreases. Our analyses are based on a 10% random sample for calibration and another 10% random sample for validation. To verify that our results are not sensitive to the random sampling, we repeat our analyses multiple times with different random samples. The differences between these sets of results are negligible. This observation is consistent with the summary statistics in Table 1. Given the large sample size, a 10% sample actually produces very similar summary statistics to the full sample.

Table 5
A comparison of targeting schemes.

	# of targets		# of bad debts		Total profit	
Ignore bad debt risk	361,200	≈100%	3140	100%	$53,918	(Base)
Incorporate bad debt risk						
Use co-operative database	285,232	79.0%	1941	61.8%	56,777	5.3%
Use census data	359,425	99.5%	3061	97.5%	54,408	.9%
Two-stage process						
No bad debt history	284,140	78.7%	1693	53.9%	46,580	−13.6%
Low bad debt risk (top 99%)	357,595	99.0%	2865	91.2%	55,565	3.1%
Low bad debt risk (top 95%)	343,148	95.0%	2437	77.6%	55,702	3.3%
Low bad debt risk (top 90%)	325,086	90.0%	2092	66.6%	55,055	2.1%
Low bad debt risk (top 80%)	288,966	80.0%	1686	53.7%	49,727	−7.8%

Value of Co-operative Database

In the previous section, we have compared the proposed targeting scheme with a suboptimal targeting scheme that ignores bad debt risk. Both schemes are based on the same set of variables collected from a co-operative database. The profit gain resulted from the proposed targeting scheme illustrates the importance of taking into account bad debt risk in a direct mail campaign. In order to alleviate bad debt risk, the information from the co-operative database is crucial. To demonstrate the value of the co-operative database, we conduct a separate comparison. We compare our proposed targeting scheme, which uses the variables collected from a co-operative database, with an alternative targeting scheme that attempts to account for bad debt risk without the information from a co-operative database. If the information from the co-operative database enables our targeting scheme to outperform the alternative targeting scheme, we have evidence for the value of the co-operative database.

The basic information required for a direct mail campaign includes consumers' names and addresses. We match consumers' addresses with census data to construct a set of variables at the census tract level, and use these variables to predict consumer purchases and bad debts. The summary statistics of these variables are provided in Table 6. We estimate the multinomial logit model using the same calibration sample as before but with these new variables collected from census data. According to the parameter estimates in Table 7, rural population is more likely to purchase the product, while those more affluent are less interested. Bad debt risk is lower for college graduates, white collar employees, rural population and the more affluent, while unemployment and low income are associated with higher bad debt risk.

Using these parameter estimates, we can determine the target consumers in the validation sample and examine the effectiveness of this targeting scheme as an alternative to our proposed targeting scheme based on a co-operative database. The comparison is summarized in Table 5. Although this alternative scheme performs better than a naïve scheme that ignores bad debt risk entirely, clearly census data are not effective in separating out those consumers with high bad debt risk. According to this alternative scheme based on census data, 99.5% of the consumers in the validation sample should be targeted, resulting in a mere .9% improvement over the naïve targeting scheme. By contrast, using the information from a co-operative database, we are able to improve the total return by 5.3%. This gain of 4.4% relative to the census data method illustrates the value of this co-operative database.

Table 7
Parameter estimates with census data.

Variable	Estimate
For option 1: paid response	
Intercept	−3.321 **
Percent of college graduates (age 25+)	−.031
Percent of white collar employees (age 25+)	.071
Percent of unemployed (age 16+)	−.168
Percent of rural population	.066 **
Percent of households with income < 25K	.042
Percent of households with income > 100K	−.234 *
For option 2: bad debt	
Intercept	−4.580 **
Percent of college graduates (age 25+)	−.725 **
Percent of white collar employees (age 25+)	−.638 **
Percent of unemployed (age 16+)	1.637 *
Percent of rural population	−.212 **
Percent of households with income < 25K	1.400 **
Percent of households with income > 100K	−1.041 **
Number of observations	361,210
Log likelihood	−72,520.7

** Significance at .01 level.
* Significance at .05 level.

We have assumed a $10 loss from a bad debt, i.e. $b = 10$, in the above analysis. As the loss from a bad debt increases, it becomes more important to accurately predict bad debt risk from potential customers. In Fig. 7, we plot the realized returns from the validation sample for different values of b based on different targeting schemes. As b increases, all schemes lead to smaller returns, but the advantage of incorporating bad debt acquisition risks using information from a co-operative database becomes even more salient, as indicated by the increasing gap between these curves.

Alternative Specifications and Robustness Checks

In the proposed logit framework we focus on predicting consumers' choices among the three distinct options, while abstracting away from a detailed investigation in consumers' decision-making processes. In this section, we explore variations of the logit framework that may be consistent with different behavioral assumptions on how consumers choose

Table 6
Summary statistics of census variables.

Variable	Full sample		Calibration sample		Validation sample	
	Mean	St. dev.	Mean	St. dev.	Mean	St. dev.
Percent of college graduates (age 25+)	24.97	43.92	24.88	43.99	25.03	43.93
Percent of white collar employees (age 25+)	39.56	26.53	39.50	26.58	39.64	26.38
Percent of unemployed (age 16+)	2.92	6.38	2.93	6.51	2.92	6.42
Percent of rural population	30.88	111.75	30.94	111.01	31.00	110.77
Percent of households with income < 25K	24.65	38.78	24.78	39.13	24.61	38.67
Percent of households with income > 100K	13.76	35.96	13.69	36.00	13.84	36.14

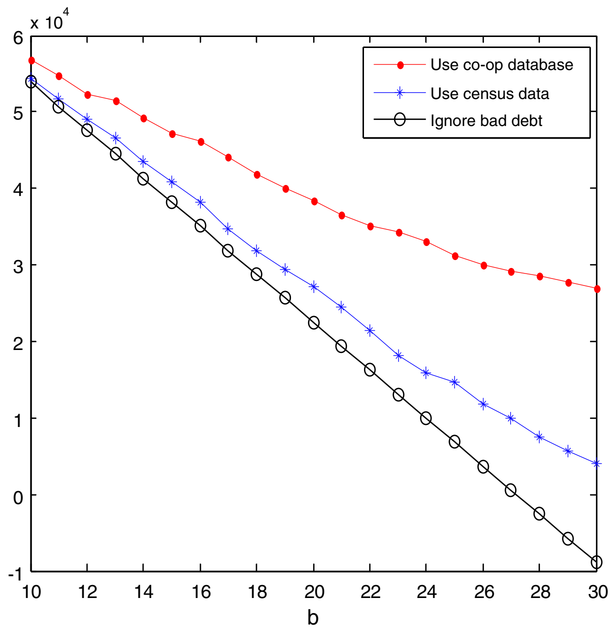

Fig. 7. Returns from alternative targeting schemes.

among the three options. We also discuss two-stage targeting schemes and potential endogeneity.[6]

Nested Logit

Rather than deciding between the three options in one shot, consumers could be making sequential decisions of purchase and payment, i.e. consumers first decide whether to purchase and then whether to pay conditional on purchase. In this sequential decision process, paid response and bad debt are more closely related to each other than they are to the non-response option. This correlation structure is not reflected in the standard logit model. Instead, it fits a nested logit model that places paid response and bad debt in one nest while the non-response option is placed in a separate nest.

Specifically, we assume that the random error vector ε_i in utility function (1) has the following cumulative distribution function:

$$\exp\left(-\exp(-\varepsilon_{i0})-\left[\exp(-\varepsilon_{i1}/\lambda)+\exp(-\varepsilon_{i2}/\lambda)\right]^{\lambda}\right).$$

This distribution is a type of generalized extreme value distribution that induces correlations within nests. That is, ε_1 and ε_2 are correlated while both are uncorrelated with ε_0 (with options 0, 1, and 2 corresponding to non-response, paid response and bad debt respectively). The parameter λ is an independence parameter that governs the degree of independence between ε_1 and ε_2. A higher value of λ indicates greater independence and smaller correlation. When $\lambda = 1$, all options become independent and the model reduces to a standard logit model. The choice probability of each option can be written as

$$P_{i0}=\frac{1}{1+\left(\exp(x_i{}'\beta_1/\lambda)+\exp(x_i{}'\beta_2/\lambda)\right)^{\lambda}};$$

$$P_{ij}=\frac{\left(\exp(x_i{}'\beta_1/\lambda)+\exp(x_i{}'\beta_2/\lambda)\right)^{\lambda-1}\exp\left(x_i{}'\beta_j/\lambda\right)}{1+\left(\exp(x_i{}'\beta_1/\lambda)+\exp(x_i{}'\beta_2/\lambda)\right)^{\lambda}},\ j=1,2.$$

Based on these choice probabilities, we can compute the likelihood of the observed consumer choices in the calibration sample and estimate the model parameters through maximum likelihood. The results are reported in the second column of Table 8. For comparison, the standard logit estimates in Table 2 are copied into the first column of Table 8.

We can see that the coefficients for the predicting variables are very close in these two specifications. The independence parameter λ is estimated to be 1.095. At a standard error of .151, we cannot reject the hypothesis that $\lambda = 1$ (p-value = .53). In other words, we do not find evidence for the correlation between the unobserved error terms for paid response and for bad debt. Despite a small improvement in the log likelihood function by the nested logit model, both AIC and BIC favor the original standard logit model. In terms of prediction using the validation sample, the nested logit model hardly produces any changes to our previous results generated by the standard logit model.

In summary, consumer choices in this direct mail campaign seem to be consistent with simultaneous decisions on purchase and payment, i.e., some consumers choose to purchase with an intent and ability to pay (paid response consumers), whereas some other consumers choose to respond with an intent *not* to pay (bad debt consumers). Sequential decisions of purchase and payment may be more appropriate in certain contexts such as big ticket purchases, where a random shock (e.g. stock market plunge) *after* the purchase of a house or a car may impair the consumer's ability to pay. A simultaneous decision process maps more closely to our proposed logit model setup, and is more reflective of the reality in the product category under study according to our data provider.

Mixed Logit

We further test the robustness of our results to an even more flexible mixed logit specification. According to McFadden and Train (2000), mixed logit can allow for flexible substitution patterns and approximate any random utility model. Using a mixed logit framework, we can control for possible correlations between consumer preferences for different options, while remain agnostic about the concrete decision process. Specifically, we estimate a flexible covariance structure on the error components:

$$u_{ij}=x_i{}'\beta_j+\xi_{ij}+\varepsilon_{ij},\quad j=1,2;\ \text{and}\begin{pmatrix}\xi_{i1}\\ \xi_{i2}\end{pmatrix}=\xi_i\sim N(0,\Sigma).$$

When estimating the elements in the covariance matrix Σ, we need to ensure that the covariance matrix is positive

[6] We thank two anonymous reviewers for pointing out these possibilities.

Table 8
A comparison of alternative model specifications.

Variable	Logit	Nested logit	Mixed logit	Add census variables
For option 1: paid response				
Intercept	−3.230 **	−3.258 **	−3.249 **	−3.241 **
Days since last order	−.059 **	−.058 **	−.059 **	−.056 **
Number of orders in the past 2 years	.355 **	.358 **	.356 **	.361 **
Dollar amount spent in the past 2 years	−.068 **	−.065 **	−.068 **	−.068 **
Number of bad debts in the past 2 years	−.024	−.040	−.025	−.024
Percent of bad debts in the past 2 years	−.064	−.098	−.062	−.068
Percent of college graduates (age 25+)				−.065
Percent of white collar employees (age 25+)				.080
Percent of unemployed (age 16+)				−.211
Percent of rural population				.067 **
Percent of households with income < 25K				.003
Percent of households with income > 100K				−.277 *
For option 2: bad debt				
Intercept	−4.188 **	−4.316 **	−4.191 **	−4.155 **
Days since last order	−.130 **	−.134 **	−.129 **	−.106 **
Number of orders in the past 2 years	.237 **	.224 **	.238 **	.279 **
Dollar amount spent in the past 2 years	−.263 **	−.275 **	−.264 **	−.264 **
Number of bad debts in the past 2 years	.821 **	.877 **	.822 **	.793 **
Percent of bad debts in the past 2 years	.591 **	.623 **	.590 **	.569 **
Percent of college graduates (age 25+)				−.575 *
Percent of white collar employees (age 25+)				−.618 *
Percent of unemployed (age 16+)				1.568 *
Percent of rural population				−.174 **
Percent of households with income < 25K				1.135 **
Percent of households with income > 100K				−1.050 **
Independence parameter (λ)		1.095 **		
Covariance matrix				
σ_{11}			.208	
σ_{12}			.076	
σ_{22}			.016	
Number of observations	361,210	361,210	361,210	361,210
Log likelihood	−71,583.9	−71,583.6	−71,582.5	−71,574.7
AIC	143,191.8	143,193.2	143,195.0	143,197.4
BIC	143,321.4	143,333.6	143,357.0	143,456.5

** Significance at .01 level.
* Significance at .05 level.

definite. Therefore we employ its Cholesky decomposition and estimate the elements in the lower triangular matrix L where

$$\Sigma = \mathrm{L} * \mathrm{L}'; \quad \mathrm{L} = \begin{pmatrix} \sigma_{11} & 0 \\ \sigma_{12} & \sigma_{22} \end{pmatrix}.$$

Note that we only know about the distribution but not the individual value of ξ_i for each consumer. We can integrate over its distribution in order to form the expected choice probabilities:

$$P_{i0} = \int \frac{1}{1 + \sum_k \exp(x_i{}'\beta_k + \xi_{ik})} dF(\xi_i);$$

$$P_{ij} = \int \frac{\exp\left(x_i{}'\beta_j + \xi_{ij}\right)}{1 + \sum_k \exp(x_i{}'\beta_k + \xi_{ik})} dF(\xi_i), \qquad j = 1, 2.$$

The model parameters can be then estimated by maximizing the joint likelihood of the observed outcomes. During the estimation procedure, we use Monte Carlo integration with 200 random draws to evaluate the integrals.

The estimation results are reported in the third column of Table 8. The mixed logit model slightly improves the log likelihood, but due to its greater number of model parameters, the original logit model is still preferred as indicated by the AIC and BIC values. The coefficients for the predicting variables are very close to those generated by the original logit model. As a result, the mixed logit model leads to very similar predictions in the validation sample. Therefore, our results are robust to controlling for possible correlations between consumer preferences for different options through a mixed logit specification.

Incorporating Demographic Variables

When we examine the value of co-operative database, we estimate the model using demographic variables at the census tract level, and compare with model estimates using historical variables from a co-operative database. We show that historical variables from the co-operative database outperform demographics in predicting consumer choices. In principle, we could estimate the model with both demographics and the historical variables. We report the parameter estimates for such a model in the last column of Table 8.

The results in Table 8 indicate that including demographics in addition to the historical variables from the co-operative database does improve the model fit but not by a substantial margin. Using either AIC or BIC as a criterion, we prefer the model with historical variables from the co-operative database only. On the one hand, we expect that incorporating both sets of variables could augment the prediction performance of the model even further. On the other hand, the additional gain from demographics is limited after taking into account consumers' past purchase and payment behavior.

Two-stage Targeting Schemes

Alternatively, direct marketers could use a two-stage process and deal with bad debt risk and response probability separately in each stage. Comparing with a two-stage approach, our model incorporates bad debt risk and response probability simultaneously, which can be advantageous depending on the strength of correlation between bad debt risk and response probability.

We first analyze a simple two-stage selection process: in the first stage we screen out all consumers with a history of bad debt; and in the second stage we use a logit model of response vs. non-response to select consumers to target. Following this scheme we would target 284,140 consumers and receive a profit of $46,580. The results are included in Table 5. It is interesting to note that this scheme targets a similar number of consumers to our proposed model, and is more effective in screening out bad debt consumers, but leads to a much lower profit. It seems that too many paid responses are screened out along with bad debts.

Instead of ignoring all consumers with a bad debt history, we explore another two-stage process. In the first stage we estimate a logit model to predict bad debt risk and screen out those consumers with high bad debt risk. The second stage is still a logit model of response vs. non-response. Because in the first stage it is unclear how to determine the threshold on bad debt risk, we try different thresholds that result in different percentages of consumers in the second stage. Again we find that, if we focus on screening out bad debt consumers in the first stage, we may lose paid responses at the same time. As we successfully screen out more and more bad debts, the profit actually declines.

Endogeneity

In our model we use bad debt history variables to proxy for the unobserved customer characteristics that predict the likelihood of a bad debt. If the bad debt history does not perfectly reflect such customer characteristics, they will enter the error term and cause correlations between the error term and the bad debt history variables. The severity of the resulting endogeneity issue may depend upon how well the bad debt history captures the underlying customer characteristics. To address the potential endogeneity bias, ideally we hope to find customer and time specific variables that are correlated with past bad debts but not current bad debt likelihood. Due to the lack of detailed information on past bad debts, we could not identify a proper instrument in our dataset. Therefore we follow recent developments (Dong 2010; Yang et al. 2012) on using control functions to address the endogeneity concern using a semiparametric approach. Specifically, in this approach we first use a kernel regression of the potential endogenous variable on the exogenous variables, and then add the residuals from the first step into the discrete choice model as a predictor. The results reported in Table 9 indicate that the coefficients of the residuals are not significantly different from zero, suggesting that we cannot reject the null hypothesis that the bad debt history variables are exogenous. Also, our parameter estimates do not change substantially after including the residuals as a predictor. Overall, we find that the bad debt history does a reasonably good job of capturing the customer characteristics that predict the likelihood of a bad debt.

Table 9
Control for endogeneity of bad debt history.

Variable	Original estimates	Control for # of bad debts	Control for % of bad debts
For option 1: paid response			
Intercept	−3.230 **	−3.230 **	−3.234 **
Days since last order	−.059 **	−.059 **	−.059 **
Number of orders in the past 2 years	.355 **	.353 **	.358 **
Dollar amount spent in the past 2 years	−.068 **	−.068 **	−.068 **
Number of bad debts in the past 2 years	−.024	−.004	−.044
Percent of bad debts in the past 2 years	−.064	−.097	−.001
(Control)		−.045	−.083
For option 2: bad debt			
Intercept	−4.188 **	−4.187 **	−4.212 **
Days since last order	−.130 **	−.129 **	−.130 **
Number of orders in the past 2 years	.237 **	.235 **	.253 **
Dollar amount spent in the past 2 years	−.263 **	−.264 **	−.259 **
Number of bad debts in the past 2 years	.821 **	.831 **	.712 **
Percent of bad debts in the past 2 years	.591 **	.577 **	.903 **
(Control)		−.021	−.378

** Significance at .01 level.

In addition, despite the common use of RFM variables in the literature to proxy for a customer's propensity to respond, there could still be relevant factors that are imperfectly represented by these RFM variables and thus cause them to be correlated with the error term. If the RFM variables are potentially endogenous, we cannot apply the above semiparametric approach due to insufficient exogenous variables in the model. Cui, Wong, and Lui (2006) apply the control function approach suggested by Blundell and Powell (2004) to deal with the potential endogeneity of RFM variables. The instruments that they use are lifetime orders and lifetime contacts, indicating previous selection and responses. Similar to lifetime orders, we observe the total number of orders and total order amount in our dataset. Although we do not have the lifetime contacts, we do observe flags indicating whether a customer was responsive to mailings in the recent past. We believe that these flags affect the lifetime contacts and hence can be valid instruments as in Cui, Wong, and Lui (2006). Following Blundell and Powell (2004) and Cui, Wong, and Lui (2006), we use the above instruments to test and control for the potential endogeneity of the RFM variables. According to the results reported in Table 10, we do not find significant endogeneity at the .05 level, and the parameter estimates do not change much from the original ones.

Discussion and Conclusions

In this paper we propose a framework to manage customer acquisition risks using co-operative databases. In the empirical context of a direct mail campaign, we analyze the bad debt behavior of consumers and demonstrate the importance of alleviating bad debt risk. Our empirical model that accounts for bad debt risk substantially outperforms the traditional modeling framework that restricts attention to the binomial outcome of non-response vs. response. Our model can be directly applied by marketers to select optimal targets for direct mail. Once they calibrate the model using observed outcomes from a smaller set of consumers or from a similar campaign, they can use the parameter estimates to examine a mailing list and predict how each consumer would respond to the focal campaign.

Table 10
Control for endogeneity of RFM variables.

Variable	Original estimates	Control for # of bad debts
For option 1: paid response		
Intercept	−3.230 **	−3.231 **
Days since last order	−.059 **	−.059 **
Number of orders in the past 2 years	.355 **	.356 **
Dollar amount spent in the past 2 years	−.068 **	−.068 **
Number of bad debts in the past 2 years	−.024	−.026
Percent of bad debts in the past 2 years	−.064	−.059
(Control for recency)		−.007
(Control for frequency)		−.039
(Control for monetary value)		.007
For option 2: bad debt		
Intercept	−4.188 **	−4.191 **
Days since last order	−.130 **	−.130 **
Number of orders in the past 2 years	.237 **	.237 **
Dollar amount spent in the past 2 years	−.263 **	−.265 **
Number of bad debts in the past 2 years	.821 **	.819 **
Percent of bad debts in the past 2 years	.591 **	.603 **
(Control for recency)		−.014
(Control for frequency)		−.046
(Control for monetary value)		.028

** Significance at .01 level.

Impact of Bad Debt Behavior on Iso-value Curves

Fader, Hardie, and Lee (2005) demonstrate the utility of iso-value curves linking RFM variables and the customer lifetime value (CLV). The insight that they provide is that while RFM variables are sufficient statistics for a customer's transaction history, customers with different values of the RFM variables could actually yield the same CLV to the firm. Thus RFM iso-value curves that trace contours of customers with differing RFM values but the same CLV provide a useful summary of the customer base to guide managerial decision-making. Such curves can provide quantitative benchmarks to gauge return on investment decisions on company initiatives to manage portfolios of customers.

We now investigate the impact of incorporating bad debt risk on the iso-value curves. We first plot a three-dimensional customer value surface to show how the expected return from a customer varies according to the number of orders (a measure of frequency) and the days since last order (a measure of recency). Consider a customer with two bad debts in the past that accounted for 50% of the previous orders. In Fig. 8(a) we ignore bad debt risk while in Fig. 8(b) we incorporate bad debt risk. By comparing the two surfaces in Fig. 8 it is clear that bad debt behavior has a substantial impact on the customer value surface, shifting the entire surface downward to lower customer values.

Tracing out points that have the same expected return helps us to obtain two iso-value curves, one ignoring bad debt risk and the other incorporating bad debt risk. Both curves are provided in Fig. 9. The differences between these iso-value curves show some interesting implications of bad debt risk. To reach the same iso-value of .1, the number of orders required in Fig. 9(b) would be about three times as many as that in Fig. 9(a). Thus companies that make investment allocation decisions based on a naïve model that overlooks bad debt risk are likely to mistakenly expect higher returns on their investment than would actually be the case.

Co-operative Databases and the Elusive 360-degree View of a Customer

Both marketing researchers and practitioners have recognized that firms operating in data-rich marketplaces labor under significant disadvantages when they have only a small subset of the transactions of their customers. Obtaining the elusive 360-degree view of consumers has been the 'holy grail' for both researchers and practitioners in targeted marketing. Typically researchers have treated the problem as a missing data problem and attempted to impute the missing data using statistical and data-mining approaches (e.g., Chen and Steckel

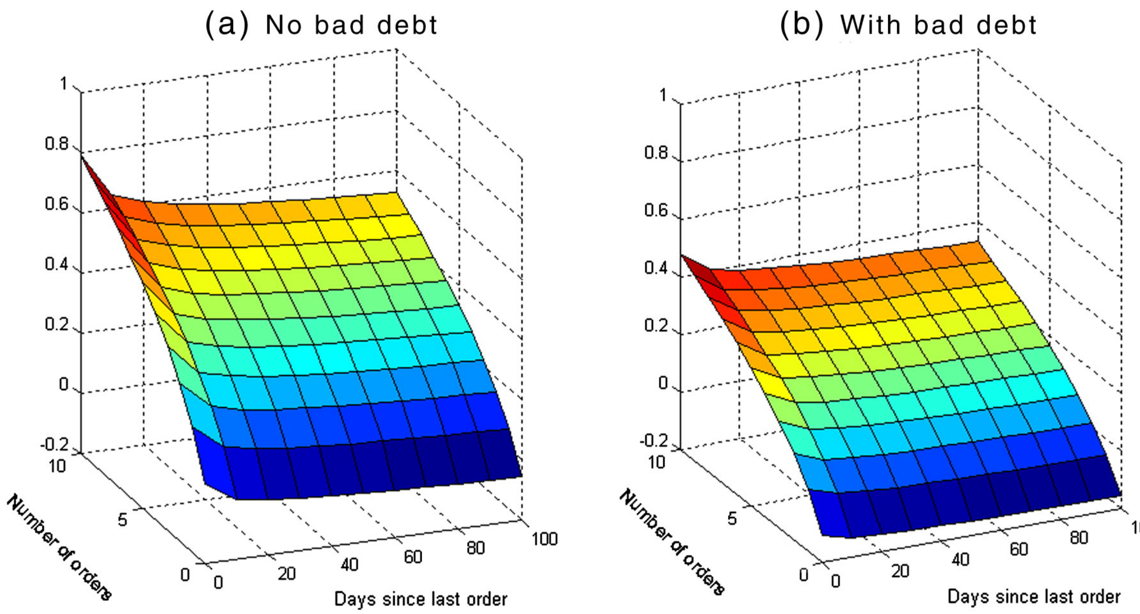

Fig. 8. Customer value surface — impact of bad debt information (number of bad debts = 2; percent of bad debts = 50%).

2012). While these approaches can be useful and interesting, the introduction of modeling error in such approaches has led many practitioners to search for better data to solve the problem. Co-operative databases have been developed in the past decade or so as a market solution to the issue of obtaining the elusive 360-degree view of relevant customer transactions, with vendors specializing in the collection, storage, retrieval and analysis of co-operative data for firms in specific industries.

Co-operative databases originated in the credit rating industry, where updated transaction and credit information across a wide range of financial transactions are used to inform decisions on whether to provide credit to consumers purchasing autos, homes, etc. Co-operative databases often require firms to contribute data in order to benefit from the pooled data. Practitioners realize that the pooling of transaction information across even competing firms provides efficiencies to targeted marketing that outweigh the possible loss of competitive advantage from sharing data about its own customers. This is especially the case in those industries where any single firm captures only a small proportion of the transactions of a consumer.

One of the reasons why the co-operative database industry has expanded rapidly in recent years is that these co-operative databases have enabled firms to be more effective in avoiding undesirable customers. In the direct marketing industry, firms have been buying response lists or buyer's lists to supplement

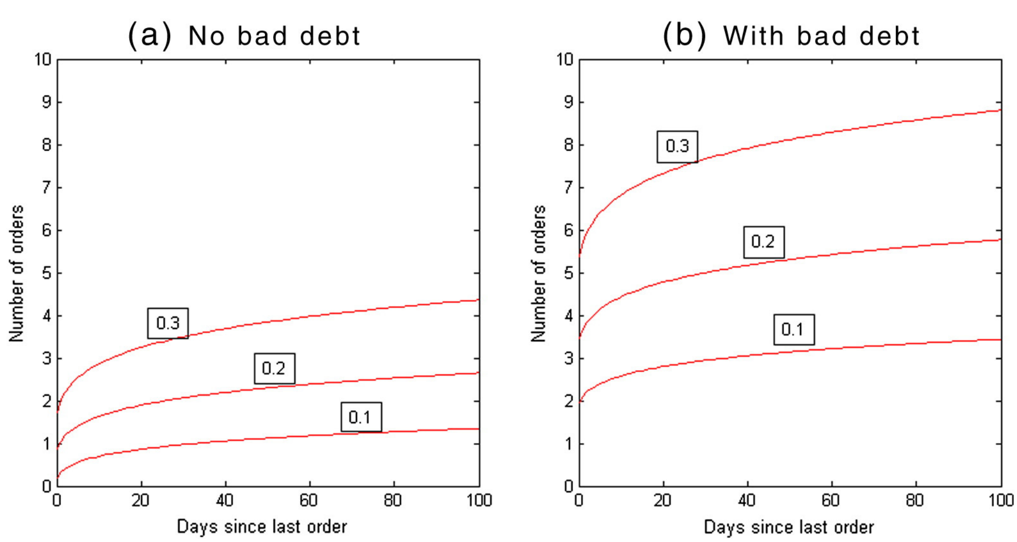

Fig. 9. Iso-value curves — impact of bad debt information (number of bad debts = 2; percent of bad debts = 50%).

their house lists for many years, especially for acquisition purpose. A significant drawback of such response lists has been that they tend not to have detailed customer purchase history information. This becomes especially costly where a large proportion of undesirable customers are expected to be attracted by an offer, and where the undesirability is of the extreme kind, i.e., consumers sign up for the offer, receive the product but refuse to pay. As we demonstrate in this study, a co-operative database can help all firms in the industry through the pooling of historical information about undesirable customers.

The framework developed in this paper will be useful to co-operative database firms, as well as to client firms that use co-operative databases to mitigate acquisition risks. We show how the cost of bad debt varies according to the characteristics of a direct mail campaign. Client firms could conduct cost–benefit analyses on alleviating bad debt risk by developing curves similar to Fig. 5 in this paper, and then determine the level of investment for bad debt mitigation. Co-operative database firms, on the other hand, could utilize this framework to determine the value of their service to clients and thereby optimize their pricing strategies.

Limitations

Our empirical setting is a one-shot magazine offer in which we model consumer responses to the acquisition campaign. For each paid response we assume a value of $20 without taking into account the probability of future customer retention. One way to incorporate retention rate into our model is to assume a uniform retention probability, which simply increases the value of a paid response. However, this approach does not capture the fact that retention probability differs across consumers. Limited by the information available, we focus on the acquisition campaign while abstracting away from the retention aspect.

The acquisition risk manifests as bad debt risk in our application. Thus we conceptualize acquisition risk broadly but restrict attention to bad debt risk in our empirical model. In the sense that product return also leads to a negative return to the firm, return risk can be handled similarly and our model bears resemblance to models of product return.

We observe bad debts in the data but remain agnostic between unwillingness and inability to pay by consumers. Using information from a co-operative base our model may pick up both types of risks if present. Credit score can potentially be a strong predictor of inability to pay but typically not available to direct marketers selling small ticket items.

Our model is applicable to non-contractual settings in which small ticket items are being sold. Because of the small amount involved in each case, typically sellers do not check credit history of buyers and expend limited effort to recover bad debts. By offering a ‘buy now, pay later’ payment option, a seller relies on consumers’ conscience for them to pay for their purchases. Many consumers may still worry that a bad debt may affect their credit ratings or have other potential consequences. Therefore incentives to pay can differ across consumers depending on their socio-economic and educational background. Due to data limitation we are not able to explore this interesting aspect of consumer heterogeneity in this study.

References

Acxiom Corporation (2008), “Finding credit-worthy consumers,” *Newsletter* winter 2008.

Anderson, Simon P., Andre De Palma, and Jacques François Thisse (1992), *Discrete choice theory of product differentiation*. MIT press.

Anderson, Eric T., Karsten Hansen, and Duncan Simester (2009), “The option value of returns: Theory and empirical evidence,” *Marketing Science*, 28, 3, 405–23.

Arnold, Todd J., Eric Er Fang, and Robert W. Palmatier (2011), “The effects of customer acquisition and retention orientations on a firm’s radical and incremental innovation performance,” *Journal of the Academy of Marketing Science*, 39, 2, 234–51.

Blattberg, Robert C. and John Deighton (1996), “Manage marketing by the customer equity test,” *Harvard Business Review*, 74, 4, 136.

Blundell, Richard W. and James L. Powell (2004), “Endogeneity in semiparametric binary response models,” *The Review of Economic Studies*, 71, 3, 655–79.

Braun, Michael and David A. Schweidel (2011), “Modeling customer lifetimes with multiple causes of churn,” *Marketing Science*, 30, 5, 881–902.

Bult, Jan Roelf, and Tom Wansbeek (1995), “”Optimal selection for direct mail,” *Marketing Science*, 14, 4, 378–94.

Chen, Yuxin and Joel H. Steckel (2012), “Modeling credit card share of wallet: Solving the incomplete information problem,” *Journal of Marketing Research*, 49, 5, 655–69.

Cui, Geng, Man Leung Wong, and Hon-Kwong Lui (2006), “Machine learning for direct marketing response models: Bayesian networks with evolutionary programming,” *Management Science*, 52, 4, 597–612.

Dong, Yingying (2010), “Endogenous regressor binary choice models without instruments, with an application to migration,” *Economics Letters*, 107, 1, 33–5.

Fader, Peter S., Bruce GS Hardie, and Ka Lok Lee (2005), “”RFM and CLV: Using iso-value curves for customer base analysis.”,” *Journal of Marketing Research*, 42, 4, 415–30.

Gönül, Füsun and Meng Ze Shi (1998), “”Optimal mailing of catalogs: A new methodology using estimable structural dynamic programming models,” *Management Science*, 44, 9, 1249–62.

Hansotia, Behram J. and Paul Wang (1997), “Analytical challenges in customer acquisition,” *Journal of Interactive Marketing*, 11, 2, 7–19.

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.

Lewis, Michael (2006a), “”Customer acquisition promotions and customer asset value.”,” *Journal of Marketing Research*, 43, 2, 195–203.

——— (2006b), “The effect of shipping fees on customer acquisition, customer retention, and purchase quantities,” *Journal of Retailing*, 82, 1, 13–23.

Lix, Thomas S., Paul D. Berger, and Thomas L. Magliozzi (1995), “New customer acquisition: Prospecting models and the use of commercially available external data,” *Journal of Direct Marketing*, 9, 4, 8–18.

McFadden, Daniel and Kenneth Train (2000), “Mixed MNL models for discrete response,” *Journal of Applied Econometrics*, 15, 5, 447–70.

Ofek, Elie, Zsolt Katona, and Miklos Sarvary (2011), “”“Bricks and clicks”: The impact of product returns on the strategies of multichannel retailers,” *Marketing Science*, 30, 1, 42–60.

Petersen, J. Andrew and V. Kumar (2009), “Are product returns a necessary evil? Antecedents and consequences,” *Journal of Marketing*, 73, 3, 35–51.

Reinartz, Werner, Jacquelyn S. Thomas, and Viswanathan Kumar (2005), “Balancing acquisition and retention resources to maximize customer profitability,” *Journal of Marketing*, 69, 1, 63–79.

Ryals, Lynette J. and Simon Knox (2005), “Measuring risk-adjusted customer lifetime value and its impact on relationship marketing strategies and shareholder value,” *European Journal of Marketing*, 39, 5/6, 456–72.

Schweidel, David A., Peter S. Fader, and Eric T. Bradlow (2008), "A bivariate timing model of customer acquisition and retention," *Marketing Science*, 27, 5, 829–43.

Shulman, Jeffrey D., Anne T. Coughlan, and R. Canan Savaskan (2011), "Managing consumer returns in a competitive environment," *Management Science*, 57, 2, 347–62.

Thomas, Jacquelyn S. (2001), "A methodology for linking customer acquisition to customer retention," *Journal of Marketing Research*, 38, 2, 262–8.

Van, Diepen, Merel, Bas Donkers, and Philip Hans Franses (2009), "Dynamic and competitive effects of direct mailings: A charitable giving application," *Journal of Marketing Research*, 46, 1, 120–33.

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, 1, 48–59.

Voss, Glenn B. and Zannie Giraud Voss (2008), "Competitive density and the customer acquisition-retention trade-off," *Journal of Marketing*, 72, 6, 3–18.

West, June (2006), *"Fingerhut's Pricing Strategy (B): Epilogue."*. Darden Case . No. UVA-E-0289. Available at SSRN: http://ssrn.com/abstract=1278374.

Yang, Sha, Mantian Hu, Russell S. Winer, Henry Assael, and Xiaohong Chen (2012), "An empirical study of word-of-mouth generation and consumption," *Marketing Science*, 31, 6, 952–63.

Zhao, Yi, Ying Zhao, and Inseong Song (2009), "Predicting new customers' risk type in the credit card market," *Journal of Marketing Research*, 46, 4, 506–17.