Targeted Direct Mail Marketing

Jeffry N. Savtiz

Targeted direct mail marketing leads to lower mailing costs, higher response rates and most importantly, more profit.

With the advent of the Internet for selling products and services, it has become increasingly more difficult to run profitable direct mail campaigns.

At Savitz Research Solutions, we have an approach to direct mail list-screening that enables the marketing manager to better target a campaign by eliminating many of the addresses of people who will simply not respond to a particular direct mail solicitation. This leads to lower mailing costs, higher response rates and most importantly, more profit.

Here is an example of the model used by a firm that markets books on home improvement.

Case History[1]

The client mails out 1,000,000 pieces annually, selling a continuity series of books on home improvement. Historically, the response rate has been about 2% and, with a $25.50 net profit per order, return on investment is a mere 4%.

The only thing the firm knows about each person on the prospect list is their address. Therefore, management wants to know if zip code demographics can somehow be used to identify the addresses that will not respond, and thus avoid mailing to them.

TABLE A

ZIP CODE DEMOGRAPHICS

Last Campaign's Response Status	Median Household Income($000)	% of Houses Built 1979 of Earlier	Years of Education	% White Collar	Median Age	Average Family Size
0	44.8	52.0%	12.1	49.6%	42.6	3.2
0	41.6	65.2	12.1	53.6	39.3	3.1
0	40.5	61.9	11.5	63.6	42.3	3.0
0	45.3	58.9	12.3	46.9	47.8	3.0
0	47.1	41.8	12.3	35.5	44.1	2.9
1	43.9	77.6	10.2	42.6	41.6	3.2
1	44.8	58.0	9.6	22.3	40.8	3.2
1	48.2	44.6	12.6	55.3	42.3	3.3
1	39.9	63.3	11.5	48.2	41.4	3.2
1	35.1	64.8	12.9	64.2	40.8	3.3

Average of Zip Code Demographics

Non-Responders (0)	43.86	56.0	12.1	49.8	43.2	3.0
Responders (1)	42.38	61.7	11.4	46.5	41.4	3.2

A First Look at Demos

We studied the zip code demographics of responders and non-responders to the last campaign. A sample of 10 cases of people who were mailed the ad is shown in **Table A**. Non-responders are coded 0 and responders are coded 1.

While it is difficult to identify differences between responders and non-responders from the raw data, a comparison of the averages for each group indicates there may be some discriminating demographics. Responders tend to be younger with less education and income; they live in older houses, and are more likely to be blue collar.

1 *The data has been disguised for purposes of client confidentiality.*

A Deeper Look at the Demos

To determine demographic differences more precisely, we ran a regression analysis[2] predicting a score "Y" which indicates whether or not each individual person responded (1) or not (0) based on these demographics. The equation is shown below:

Y = - 0.14 (y - intercept)
+ 0.02 (% of houses built before 1980)
- 0.04 (years of education)
- 0.01 (% of white collar)
+ 0.10 (average family size)

The other variables, income and age, did not come into play.

After thinking about the equation, we concluded that it makes sense. People interested in information about home improvement tend to live in older homes with larger families. They are more likely to be blue collar and have less education. Not surprisingly, the average score of responders (0.49) is higher than that of non-responders (0.30).

TABLE B

Last Campaign's

Status	Score
0	0.24
0	0.45
0	0.30
0	0.38
0	0.14
1	0.90
1	0.73
1	0.02
1	0.50
1	0.33

Scoring Responders and Non-Responders

In general, we would expect the regression equation to predict values of Y closer to one for the responders and values closer to zero for the non-responders. Using this regression equation, we then proceeded to score each of the individuals used to derive the equation as shown in **Table B**. Indeed, responders tend to have higher scores than non-responders.

If we were to do this for a large sample of responders and non-responders and graph the frequency distribution, the result would look something like **Figure 1**, with two overlapping bell curves.

FIGURE 1

DISTRIBUTION OF RESPONDER AND NON-RESPONDER SCORES

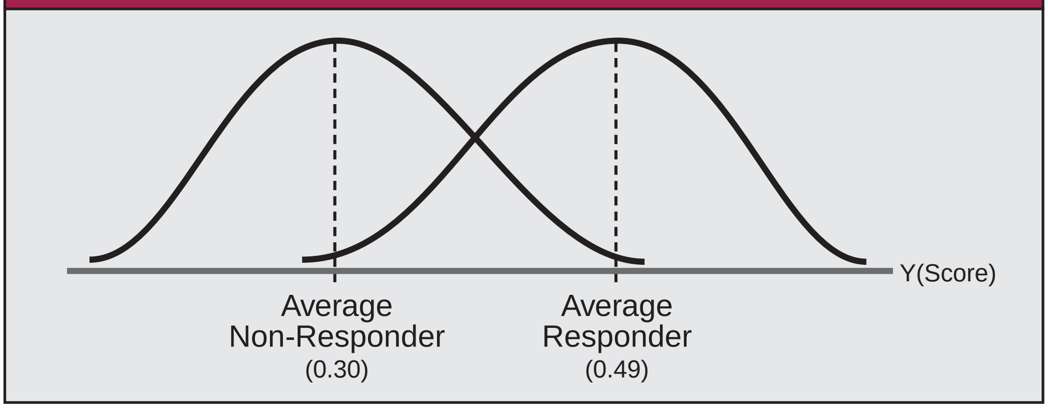

Developing a Cut-Off Score

In **Figure 2**, we have drawn a vertical line through the curves and call the value on the Y (score) line the "cut-off score." We will use the cut-off score to classify each case in the following way. If we score an individual using the regression equation and the score is high, that is, at or above the cut-off score, we will predict the person to be a responder (1). If, on the other hand, their score is below the cut-off score, we will predict the person to be a non-responder (0).

2 *Dummy variable regression or discriminant analysis can be used.*

DISTRIBUTION OF SCORES IN RELATION TO CUT-OFF SCORE

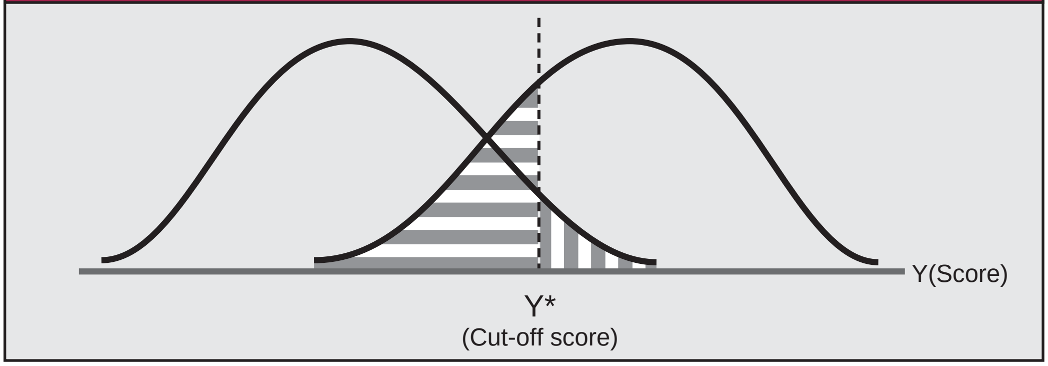

FIGURE 2

In **Figure 2**, the vertical hatching is associated with scores of non-responders who we would misclassify as responders. The horizontal hatching is associated with responders who we would misclassify as non-responders. Because the bell curves always overlap, no matter what cut-off score we choose, there will always be some degree of misclassification. Our objective will be to choose the cut-off score that maximizes the percentage of correct classifications of responders and non-responders weighted by their associated gains and losses, respectively.

Classifying Individuals

Now, let's use a cut-off score of 0.26 and classify the ten individuals as shown in **Table C**. In the small sample, we correctly predicted four out of five responders, but only two out of five non-responders. Of course, we could have picked any cut-off score between zero and one. However, as it turns out, given the percentage of responders to be mailed (2%), the profit from a responder ($25.50) and the loss from a non-responder ($0.50), the cut-off score of 0.26 is the one that maximizes the profit we expect from the campaign. The derivation of the cut-off score is shown in the appendix. But, for now, let's use the 0.26 to show how correctly we could classify all cases used to develop the regression model.

TABLE C

Last Campaign's Status	Score	Model's Prediction
0	0.24*	0
0	0.45	1
0	0.30	1
0	0.38	1
0	0.14*	0
1	0.90*	1
1	0.73*	1
1	0.02	0
1	0.50*	1
1	0.33*	1

** Correct Classification*

The Classification Matrix

Table D shows the model correctly predicts 80% of all the people who responded to the past mailing to be responders, but it only correctly predicts 30% of the non-responders to be same. This makes sense in terms of maximizing profit. At $25.50 profit per responder and $0.50 loss per non-responder, we can certainly afford to incorrectly classify a lot more non-responders than responders.

TABLE D

CLASSIFICATION MATRIX

	Model's Prediction	
Actual Status	Responders	Non-Responders
Responders	**80%**	20%
Non-Responders	70%	**30%**

The Mailing Matrix

Now, suppose instead of mailing to everyone, we mail only to those who the model predicts will respond. Let's translate the Classification Matrix into the Mailing Matrix (**Table E**). We would mail to any address the model classifies as a responder. Of the 2% of 1,000,000 or 20,000 responders, we would mail to 80% or to 16,000. Of the 98% of 1,000,000 or 980,000 non-responders, we would mail to 70% or 686,000.

TABLE E

MAILING MATRIX

	Model's Prediction	
Actual Status	Mail	Do Not Mail
Responders	16,000	4,000
Non-Responders	686,000	294,000

TABLE F

WITHOUT MODEL

Mailing to Responder	2% x 1,000,000 =	20,000
Mailing to Non-Responders	98% x 1,000,000 =	980,000
Total Mailing		1,000,000
Cost		.50
		$500,000
Gain from Responders	20,000 x $26 =	520,000
Profit		$ 20,000
Percent Return		**4.0%**

WITH MODEL

Mailing to Responder	2% x 1,000,000 x 80% =	16,000
Mailing to Non-Responders	98% x 1,000,000 x 70% =	686,000
Total Mailing		702,000
Cost		.50
		$351,000
Gain from Responders	16,000 x $26 =	416,000
Profit		$ 65,000
Percent Return		**18.5%**

Profitability Analysis

Now, let's see how the profitability of the campaign compares when we mail to the addresses the model indicates as responders versus mailing to everyone (**Table F**). Clearly, use of the model has a significant financial advantage. The mailing cost falls from $500,000 to $351,000. The return on investment increases from 4.0% to 18.5%, and most importantly, profit is up from $20,000 to $65,000.

Summary

Making profit by selling goods and services through the mail can be difficult in today's world of retail and e-commerce. We have presented a model that enables the marketer to better target a mailing to people who are more likely to buy. This has a number of applications for catalogs in general and a large range of goods such as books, specialty foods, computers, toys as well as services such as insurance, health care, home repair and maintenance. If you would like to do research to determine how to better target your direct mail marketing efforts, call us at Savitz Research Solutions. Our experience shows the results will pay for the research many, many times over.

Call
Savitz Research Solutions

972-386-4050 **Dallas**

312-377-1890 **Chicago**

TECHNICAL APPENDIX - DERIVING THE OPTIMAL CUT-OFF SCORE

Our objective is to find the cut-off score (y^*) that maximizes the expected profit EP(y) from a campaign where we mail to anyone with a score of y^* or higher. We assume the scores of the responders have a normal distribution with a mean of μ_r and a standard deviation of σ, and the scores of the non-responders have a normal distribution with a mean of μ_n and the same standard deviation.

Let's define p_r and p_n the proportion of responders and non-responders

g_r and g_n the average gain from a responder and a non-responder (a negative)

$Pr\{Y_r \geq y^*\}$ and $Pr\{Y_n \geq y^*\}$ the proportion of addresses of responders and non-responders with scores at or above the cut-off score.

Given the definitions, we want to find the y^* that will maximize the expected profit from the mailing.

$$EP(y^*) = p_r g_r Pr\{Y_r \geq y^*\} + p_n g_n Pr\{Y_n \geq y^*\} =$$

$$p_r g_r \int_{y^*}^{\infty} \frac{1}{\sqrt{2\pi}\,\sigma} \exp\{-(y_r-\mu_r)^2/2\sigma^2\}dy_r + p_n g_n \int_{y^*}^{\infty} \frac{1}{\sqrt{2\pi}\,\sigma} \exp\{-(y_n-\mu_n)^2/2\sigma^2\}dy_n$$

We find this maximum by taking the derivative of the profit function and setting it equal to 0. In general,

$$y^* = \frac{\mu_r+\mu_n}{2} - \frac{\sigma^2}{\mu_r-\mu_n} \log_e \left\{\frac{-P_r g_r}{P_n g_n}\right\}$$

In our example, $y^* = (.49 + .30)/2 - ((.63)/(.49 - .30))(\log_e(-.02 \times 25.50/(.98 \times -.50))) = 0.26$