Loyalty ROI – measuring and maximising loyalty points

Michael Lieberman, Multivariate Solutions, shows how statistical techniques can guide key marketing decisions

DO LOYALTY PROGRAMMES work? This question is as murky as, say, what is the meaning of life? Some work. Some don't work. Still, nearly 75% of shoppers in the US now belong to at least one loyalty programme. 'Common knowledge' in business circles states that loyalty programmes are useful, sometimes very useful.

A strong marketing programme for business loyalty has three known goals. First, to acquire new customers. Second, to keep existing customers. Third, to grow these customers into larger, and more lucrative, customer categories.

Loyalty programmes span a globe of stratagems. There are plastic cards, smart cards, thermal cards and magnetic strips. Virgin Mobile Australia, a wireless arm of UK-based Virgin Group, is using 802.11: flash your mobile device at a reader as you whiz through checkout at a Virgin music store, and you will get a couple of dollars knocked off your bill. There are frequent-buyer programmes, frequent-flier programmes, frequent-player cards and frequent-dining coupons. There are points-at-the-pump schemes, turkey giveaways at Christmas (along with $100 copper roasters in which to cook them) and, as of November 2006, donations to charities for people who use their Starbucks loyalty card to buy coffee.

As loyalty programmes and their related tactics have matured, increased attention has been placed on maximising the bang for the buck. Industry knowledge, guess work, or instinct are no longer suitable substitutes for strategic risk analysis. When marketing managers are asked how they are optimising their budgets, specific promotions are often questioned. Marketers today are under increasing pressure from their bosses to show a greater return on investment (ROI). What gets measured gets done, as the saying goes. Budgets are limited, so, how does the manager know he is doing his best?

The answer is optimisation. Simulating a loyalty promotion, restricting the budget, and searching for the peak.

This article is an abbreviated rundown of how a segmentation/ROI study functions. First, there is the market research/data-mining component of segmentation, followed by uses of Monte Carlo forecasting and optimisation of a company's promotion.

Nearly 75% of shoppers in the US now belong to at least one loyalty programme

To do this we will use a fictional example – let us call it Colossal Supermarkets and its flagship store, Food City – as an example of how to design a specific promotion, and to maximise its return within the company's promotional budget.

Food City's programme

The Colossal sales manager has a loyalty programme that he wants to deploy for the US holiday season. Food City is a regional chain, with stores concentrated in the Pacific Northwest. They want a bigger slice of the turkey and pie the locals are rushing to supermarkets to fill up on. So Colossal Supermarkets initiated a programme in which all Food City preferred customers were enrolled in a new and enhanced loyalty programme called Food City's Holiday Gift Bag. In all, nearly a million Food City customers were enrolled. They received colour-coded cards in the mail. Each colour represented the amount of money spent on a monthly basis at Food City by each customer 'unit' (a unit can be an individual, couple, family, household, and so on).

Food City's Holiday Gift Bag actually operates on a per-month baseline. Smaller spenders receive an invitation to join Preferred Membership. Those who spend more at Food City receive a Gold Gift Bag Card. There is a Platinum card, then the highest, the Food City Mayor's Club.

Each tier is tied to a level of benefits. The higher the tier, the more extras Food City will dish out. To make the programme more attractive, Colossal Supermarkets has tied its benefits to other loyalty programmes, such as car-rental discounts, frequent-flier points on partner airlines, or even discounts on menswear. Basically, the more you spend, the more you get 'free.'

Michael Lieberman is founder and president of Multivariate Solutions, a statistical and marketing research consulting firm that works with major advertising, public relations, and political strategy firms.
michael@mvsolution.com

The analysis

There are two stages to the analysis. First, to determine the confines of customer segmentation. That is, where would be the best place to draw lines among the different colours of the Gift Bag cards in order to divide up return? These customer boundaries are commonly referred to as the 'efficient frontier'. This stage uses a mixture of cluster analysis (multivariate segmentation) and Monte Carlo simulation.

The next step, how much should each point be worth? Assuming each 'point' with the programme had a cost (for example, if one point returned a 1% discount, the 'cost' of a point might be $.01). What we were assigned to do is to set the 'optimal' ratio of points to spending, so that the return on each point would be maximised.

Segmentation

Market segmentation is a behaviourally-based, statistical approach to putting respondents into baskets. Each basket is mutually exclusive, and the final 'basket' is tied to the amount of money each unit spends at Food City each month. Food City has been asking customers to fill out a small customer-satisfaction card – demographics and a short section regarding their food purchase behaviour – which is then tied to their customer identification number. This provides invaluable information about marketplace complexities facing Food City consumers.

The final groups are formed combining the results of the cluster analysis, which is tied to spending amounts. These are shown in Figure 1.

The Monte Carlo Process

Monte Carlo simulation is a method for iteratively evaluating a deterministic model using sets of statistical distributions as inputs. This is often used when a model is complex or involves more than just a couple of uncertain parameters. A simulation can typically involve over 10,000 evaluations of the model.

In the simulation, a model in spreadsheet format is set up and the cells whose values come from survey results, customer databases, financial reports, and so on, are identified. For each of these cells, a distribution of possible values using appropriate means and errors is specified. In other words, the shape (referred to in statistics as a distribution) of spending per month for Food City customers could be different from, say, the number of trips to Food City a given customer makes per month. Monte Carlo allows for these different distributions.

A series of trials is then generated, each of which represents a possible outcome of the process. Instead of a simple spreadsheet that yields one answer, Monte Carlo allows the spreadsheet to run 10,000 times, each different parameter moving within its shape, giving 10,000 different outcomes. When these are shown in a cumulative chart, the chances of a given outcome can be determined. For example, what is the chance that Food City customers will spend more than $500 a month?

FIGURE 1

Cluster analysis and spending amounts

Give bag segment levels	Monthly spending range
Preferred members	**$0–$200** Single households, lower income families, students
Gold membership	**$201–$400** Families with 2 or fewer children; lower middle-class income; blue collar workers; union members
Platinum membership	**$401–$750** Larger families. Professionals, two-income household
Food City Mayor's Club	**More than $750** Premium item purchasers; have a private school in the area. Working mothers have professional careers

Source: DWBB

In our case, parameters of spending and input are set up, using the customer satisfaction survey and customer database. To determine the output of the optimisation, these spreadsheets are run, say, 10,000 times. This is called the forecast.

Optimisation

The goal of any optimisation is to find the input values (decision variables) that make the output (forecast) as large – or as small – as possible. Figure 2 summarises the process.

There are many applications for optimisation:

- ▶ Utilisation of employees for workforce planning.
- ▶ Configuration of machines for production scheduling.
- ▶ Location of facilities for distribution.
- ▶ Tolerances in manufacturing design.
- ▶ Management of portfolios.
- ▶ Calculation of optimal price/promotional points.

In our case, as is most often the case in ROI projects, the Decision Variable is the value of the points that will be awarded for each tier in the Food City programme.

FIGURE 2

How optimisation works

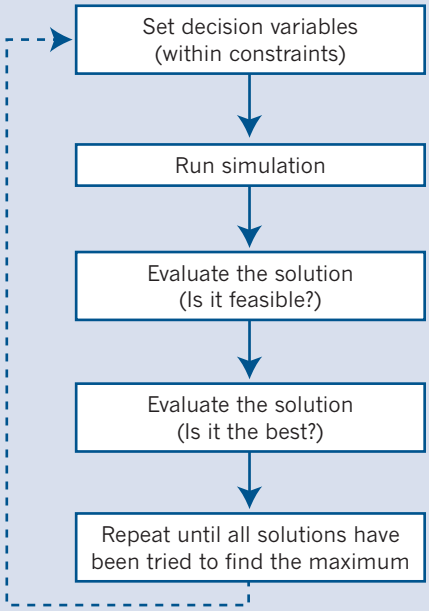

FIGURE 3

Food City promotional design

	Monthly spending range	% Customers	Lower hurdle		Higher hurdle	
Preferred members	$0–$200	45%	$140	← →	$200	
			No. of points	200	No. of points	400
	Monthly spending range	**% Customers**	**Lower hurdle**		**Higher hurdle**	
Gold membership	$201–$400	36%	$250	← →	$400	
			No. of points	500	No. of points	750
	Monthly spending range	**% Customers**	**Lower hurdle**		**Higher hurdle**	
Platinum membership	$401–$750	14%	$600	← →	$750	
			No. of points	800	No. of points	1200
	Monthly spending range	**% Customers**	**Lower hurdle**		**Higher hurdle**	
Food City Mayor's Club	More than $750	5%	$1,000	← →	$1,500	
			No. of points	1500	No. of points	2000

Final point value = $0.15

The Forecast will be the incremental increase in spending for each value of the points.

Food City's optimisation

Once all the variables are entered, the spreadsheet complete, the final step is to let the optimisation software run, and run. It is common for the forecast (which runs, say, 10,000 outputs), to run 10,000 times to find the optimal level.

Below are the constraints that are built into the optimisation process:

- ▶ Maximise ROI.
- ▶ Stay within the promotional budget.
- ▶ Try to stimulate growth of all Food City shoppers.

Given the nature of customer behaviour, it is natural to expect customers in the Platinum or Mayor's Club categories to have a higher increase in spending due to the promotion. That is all well and good. However, in market reality, the percentage of Preferred Shoppers is far greater. Food City wants them to spend more as well. In addition, it is commonly understood that not everyone will redeem every point they receive. The 'expected redemption' is another variable that is built into the model.

Now that the optimisation is run, two outputs need to be analysed. The first is called the spending 'hurdle'. That is, at what point do members move from one category to the next. These are not necessarily the same as the spending segments discussed above, for the simple reason that spending is not static, but moves up and down depending on holidays, family, and life's events (for example, the birth of a child, moving, promotion at work).

The second output we look at is the 'value' of each point when redeemed. Remember, the name of the game is return on investment: if Food City gives away too much, that return drops.

After the optimisation is complete, the top results are analysed. A few scenarios are rerun to validate the results. A few things needed to be tweaked so that they make market sense. For example, if the optimisation suggested that the spending hurdle was $563.35, it makes more sense to set it as $575. If the suggested point value was $0.0986, it makes marketing sense to set it at 10¢. Each of these things are tested.

Once these things are analysed, final decisions about where to set the spending hurdles and the value of a point are set. Figure 3 summaries the findings of the Food City study.

Conclusion

As computing power increases and marketing becomes more savvy (for example, Amazon knows your favourites), it becomes easier to facilitate an optimisation project like Food City's. Experience shows that implementing the new technique, and other risk-analysis measures, can have a high learning curve, but ultimately yield effective, cost-efficient results.

The marriage of survey research, data-mining techniques, the Monte Carlo method, and optimisation is taking more and more risk out of developing these promotional programmes, and improving return on investment for marketing managers market-wide. ■

More on loyalty marketing at WARC.com