The benefits of using a decision engine to optimise campaign planning for direct marketing

Received (in revised form): 24th February, 2006

Julian Berry

is a founding director of The Customer Partnership and has been employed in customer marketing for over 25 years. For more detail on The Customer Partnership's specific approach and tools see www.campaigncalculus.co.uk

Abstract This paper is about a new technique for undertaking campaign planning for direct marketing, and the benefits that this can provide, as evidenced by case studies in financial services and utilities. It demonstrates a very considerable commercial uplift in terms of the value generated from direct marketing, but does not require any change in operational technology.

Top down spreadsheet-based planning is very sub-optimal because it cannot take account of the opportunity costs of sending customers or prospects a particular communication, even if it manages to make a small profit at the time. By building campaign plans using individual level data, the impacts of alternative planning scenarios can be demonstrated, and the one that best fits business requirements selected.

The capabilities required to do this, besides a 'decision-engine' to develop scenarios, are samples of customer or prospect data with individual propensities and eligibilities attached.

INTRODUCTION

Most direct marketing activity, all around the world, works off the concept of a campaign. These campaigns are usually planned on a weekly or monthly basis, with set volumes of customers or prospects offered up to each one, and with expected response results and costs linked to them.

Normally, within a period, campaigns are selected from a customer or prospect database in a sequence linked to their perceived value or importance. There are often some contact interval rules in play that determine the minimum time between repeating a message to an individual, or the number of messages the individual can receive in a period.

Over the last ten years, an alternative approach has been developed, whereby a decision engine looks at each individual customer and determines the optimal contact approach based on criteria provided (usually return on investment [ROI]). This partially optimised view is then usually modified to take account of business requirements such as budget, sales targets and channel capacities.

Only a very small proportion of direct marketing, the author suspects less than 1 per cent, is directed on the basis of decision engines selecting is the optimal communications approach for an individual. There are a number of reasons why this is the case, including

Julian Berry
The Customer Partnership,
212 Piccadilly,
London W1J 9HG, UK
Tel: 2(0)208 878 9350;
e-mail Julian.Berry@thecustomerpartnership.com

the fact that substantial investments have gone into buying campaign management software and there is a natural reluctance to replace it; the decision engines require that one's data-mart and propensity models are in perfect order; and moving away from 'business as usual' requires process change, as well as significant IT capability enhancement.

With major ROI improvements to be gained from the introduction of individual level decision engine-based campaigning, however. The Customer Partnership (TCP) has been examining whether there is a more incremental approach that could be adopted, using the decision engine to support the *planning* process only, whilst transferring the results of TCP has called 'optimised planning' back into the conventional campaign management process.

This paper, examines how much of the benefit that can be obtained from individual level optimisation can also be obtained from optimised planning, and what is required to achieve that.

HOW DOES THE DECISION ENGINE SUPPORT CAMPAIGN PLANNING?

The decision engine used by The Customer Partnership starts by calculating the potential economic benefit of each potential communication over the planning period in question. So, for 5 million customers, 10 key propositions to market and 12 planning periods there are 600 million potential ROIs to calculate. The formula used is

$$(\text{propensity to respond}) \times (\text{value of the response}) - (\text{cost of the communication}).$$

It will be evident that this requires response propensity models to be applied to the customer or prospect base for each of the ten propositions. An even more effective result is achieved if there

are individual level models for potential value as well, but these are not necessary and an average value for each proposition can be used.

The next step is for the decision engine to eliminate a proportion of these potential communications by applying rules that deal with areas such as:

- The eligibility of the person for the treatment
- Contact interval rules
- Minimum ROI requirements
- Segment communication protocols.

The decision engine is then set a number of 'constraints' that direct the way in which it forms scenarios. These normally include:

- Budget limits
- Which propositions are deemed appropriate for which time period
- Minimum sales targets by product
- Maximum channel capacities.

The decision engine then processes a sample of the customer or prospect data to develop scenarios that are, in effect, campaign plans for the period in question, but built up from individual data. Scenarios are reviewed and changes to rules and constraints made until they are acceptable to the marketer.

The next step in The Customer Partnership methodology is, however, atypical. It has been discovered that by extrapolating from the individual level data-based scenario to the campaign plan as executed, a very considerable optimisation uplift can be achieved. In other words, the decision engine is used to build the plan; the plan is then applied to a normal campaign management system. As a result no operational software changes are made as the scenarios are developed offline and on a sample by sample basis.

Table 1: Information on case study companies

Financial services company	Utility company
Eight main products Annual direct marketing spend >£10m Additional contribution +24% Spend constant	16 main treatments Annual direct marketing spend >£50 Additional contribution +27% Spend -6%

MAGNITUDE BENEFITS EXPECTED FROM OPTIMISED PLANNING?

There follow two case study examples; one from the financial services industry, the other a utilities company (see Table 1). In both cases, the clients had working individual-level response propensity models. The Customer Partnership was able to use the companies' data to develop a sample file to feed into its decision engine. The companies' existing, detailed company plans (known as 'business as usual') were also taken.

Using the same budget, The Customer Partnership's objective was then to see, how much more ROI optimised planning could provide compared with business as usual (the 'additional contribution'). It is important to note that these comparisons are made using the outputs from the decision engine, not from actual contact made. The basis for comparison should only change if the propensity models used are not accurate, however.

It is also interesting to compare the contribution from business as usual in a single time period with an optimised plan, because this gives more insight into the kinds of changes that are required.

The case outlined in Tables 2 and 3 takes just one month in a year of the utility company, working on a sample basis. In order to further anonymise the result the results from only a small sample are shown. In this example, optimised planning gave an improvement of 29 per cent, with a 6 per cent reduction in budget.

The difference between the business as

usual campaign and the optimised campaign is the level of investment going into each campaign and the order in which they are selected.

WHAT ARE THE UNDERLYING DRIVERS THAT TAKE OPTIMISED CAMPAIGN PLANNING AHEAD OF NORMAL SPREADSHEET BASED CAMPAIGN PLANNING?

There are a number of underlying factors that cannot normally be understood or resolved within conventional campaign planning spreadsheet processes, but which can be improved with optimised planning. The following explains how these benefits are achieved.

The distribution of an overall budget between individual campaigns

If one starts from the position of having a fixed budget and or fixed rules about levels of allowable contacts per individual prospect or customer, and if one is trying to allocate the budget optimally between campaigns, then it is found that the larger the budget given to any one campaign, the lower the response rates, assuming that selections are being made on the basis of propensity models picking from the best prospects downwards.

Additionally, each time you settle on a volume for campaign A, that decision will impact the volumes available for all the other campaigns and hence their ROI will have to be recalculated. This can lead to the need to change volumes for campaigns B, C etc to meet sales and

Table 2: Business as usual scenario

Treatment	Contrast volume	Take up	Total spend £
T1	350	–	126
T2	1,440	–	518
T4	6,150	44	2,214
T6	490	–	245
T8	9,690	49	13,082
T9	8,320	70	6,739
T11	7,440	185	10,937
T12	1,800	39	4,914
T13	1,750	82	4,778
T14	50	–	137
Total	37,480	469	43,689

Table 3: Impact of campaign optimisation

Treatment	Contrast volume	Take up	Total spend £
T1	10	–	4
T2	10	–	4
T4	8,410	36	3,028
T6	20	–	10
T8	1,740	–	2,349
T9	15,310	195	12,401
T11	2,260	149	3,322
T12	2,590	56	7,071
T13	2,490	115	6,798
T14	2,270	124	6,197
Total	35,110	675	41,183

ROI targets, and hence the iteration back to recalculating campaign A.

To get the best mix between campaign volumes requires endless iterations, as, in effect, all campaigns are competing against each other for budget, or contact opportunity. Using the decision engine, the best distribution of budget between campaigns is achieved at once and is based on individual level propensities, so that the impact of any changes can be seen immediately.

The sequence of campaign selection within a period

All campaign management systems need to prioritise campaigns within a time period, such as a month, in which selections are being made. In effect, the user sets up a selection sequence moving from campaign A to B to C

etc, and this is usually done on the basis of the anticipated ROI from each campaign.

As described above, however, there is only a finite amount of available customer contact availability and purchase propensity in any database, and therefore campaign A may in effect be robbing campaigns B and C of a large proportion of this.

It is self evident that although one campaign may be expected to have a better overall ROI than another, this is entirely dependent on which individuals are selected for each campaign. Consequently there are very considerable gains and losses to be obtained from altering the sequence of campaign selection within a time period. The optimised planning approach picks the best sequence in which to select campaigns within a period.

The impact of campaign interval rules on overall ROI through a period

Almost every large direct marketing company works with a set of contact interval rules. These may be set as:

- A minimum time period between repeating a campaign to an individual
- A maximum number of contacts for an individual in a single time period
- The minimum time period before repeating the use of a channel for a particular individual.

A recent survey undertaken by The Customer Partnership has found that there are very few norms when it comes to contact density rules, however, and that most companies intuit their own, with little or no research backing. The impact of these rules on the overall ROI to be derived from a customer or prospect base with finite response propensities, and within an overall budget ceiling for the period, are substantial.

The optimised planning approach allows different contact rules to be tested, albeit inside the decision engine and not on live customers, and their cumulative impact over the whole planning period to be examined. As an example there may be one very successful campaign, which, if given shorter repeat intervals, could considerably boost the overall ROI to be obtained. Also, allowing more than one campaign communication to an individual within a time period may extract more value from the best responding customers.

The impact of the decision engine on the campaign planning process itself

Underlying the examples given above is the considerable benefit (in terms of reduced resources) required to do planning. Optimised planning using a decision engine can iterate different scenarios in real time and can quickly look at alternative ways of distributing a budget or selecting campaigns. This will reduce the resources going into planning at the same time as delivering a much stronger and more realistic result.

CONCLUSION

There are virtually no risks associated with introducing optimised campaign planning and a high potential ROI gain. The scenarios run in this example show high level gains and we would expect similar levels of uplift in most cases. The extent of the gain compared with business as usual is dependant on the quality of the business as usual plan and the propensity models that have been developed.

Trialling optimised campaign planning is simple. Usually, it is found that the main capability requirements, other than a suitable decision engine, will be to get the propensity models and individual level data items in place.

Useful websites

www.campaigncalculus.co.uk
www.dm-attitudes.com