
Practitioner Article

Using a decision support optimisation software tool to maximise returns from an overall marketing budget: A case study from a *B-to-C* marketing company

Received (in revised form): 21st March 2012

Paul Rowson

is an experienced independent Marketing Analyst and developed the mathematical formulation for Budget Calculus. Before setting up Knowledge Based Marketing in 1995, he was Head of Database Marketing at GE Capital.

Howard Thompson

is a professional services Consultant at TCP Marketing Solutions. He has an analytical-based marketing background, and a wealth of experience in a variety of sectors including home shopping, retail and financial services.

Julian Berry

is Founder of TCP Marketing Solutions, and has previously been MD of Berry Consulting, Joint MD of Information for Marketing, and Head of Direct Marketing at TSB and Nat West Banks.

ABSTRACT Most marketers go through an annual budget-setting round in which they allocate sums of money to a number of different marketing activities. This process is usually undertaken using historic experience of results, combined with last year's budget, to arrive at a decision about how the money should be distributed. The authors have used a new approach to solving the budget-allocation problem, using a decision support system that they have developed using integer linear programme techniques. This article is a case study of one application of this technology in the *B-to-C* field; it describes the problems that were encountered in getting the right metrics to go into the system, and the ways in which these were overcome. It goes on to show how the optimised outputs differed from business as usual, and the types of scenario outcome that resulted.

Journal of Database Marketing & Customer Strategy Management (2012) **19**, 138–142.

doi:10.1057/dbm.2012.10; published online 30 April 2012

Keywords: marketing budget optimisation; software; metrics; halo

BACKGROUND TO THE *B-TO-C* MARKETING COMPANY, AND HOW IT DEPLOYS ITS MARKETING BUDGET

The company distributes a small range of financial services products using six channels (web, telephone, door-drops, direct mail,

DR TV and paid search). It has no retail presence and as such is entirely dependent on direct and direct response advertising.

It markets using direct mail through several corporate partners, who provide names and addresses, through bought lists, and to its own enquirers and policy holders.

Correspondence:

Julian Berry
TCP Marketing Solutions,
37 Sheen Road, Richmond,
Surrey TW9 1AJ, UK
E-mail: jberry@
tcpmarketingsolutions.com

It views each month in the calendar as a discreet entity, because of the seasonal variations in response.

In all, the permutations of actively used spend destinations, and calendar periods, come to 250 alternative ways in which the marketing budget can be spent.

Across all these destinations, which we call treatments, the overall budget is around £40 million. On average, it generates an ROI multiple in excess of three, in terms of the longer-term value of sales over the marketing cost.

The company does its own marketing planning, and has achieved a high level of accuracy in source coding of response, the exceptions being independent web search enquiries, and the halo effects of interactions between treatments.

As a result, the company has reasonably accurate figures for the product units sold through each treatment, as well as the average longer-term value to the company of those sales; in fact, the value per unit product sale has a range of 1–2.9 from the lowest to the highest treatment.

Clearly, many of the channels and treatments have maximum and minimum spend levels; this can, for example, be because of a shortage of names and addresses, or existing contracts with media owners.

In addition, there is demand elasticity; as more is spent in a channel the returns per unit spend is less. We describe these as saturation curves.

Lastly, there are halo effects; for instance, spend on TV was found to be influencing response in other channels.

We learnt that the company had developed good spreadsheet-based skills in budget allocation.

However, the problem they faced was getting beyond the range of normal spreadsheet computational powers, with such a large number of possible treatments, many with maximum and minimum spends, combined with saturation curves, halo effects, as well as recruitment targets.

OBJECTIVES FOR THE PROJECT

The primary objective we were set was to increase the value generated from a fixed budget by 5 per cent while living within the company's business constraints. Alternatively, for a fixed value generated we were to show that we could reduce costs by 5 per cent.

We were asked to provide these gains in a more reliable and auditable way than the current exercise of judgement by the company's planners.

We were also asked to show how the same budget could be re-allocated to increase or reduce the number of new recruits obtained, and the impact that changing the level of recruits obtained had on value.

Finally, we were to show efficiency gains in terms of planning resources used, and time to deliver new budget scenarios.

All of this had to be done within the constraints, discussed in the background, such as treatment minima and maxima.

We used our own optimisation software called Budget Calculus to run these scenarios.

KEY COMPONENTS OF THE DECISION SUPPORT SYSTEM

The problem of maximising contribution for a given marketing spend is an optimisation problem that is formulated in the decision support system as an integer linear programme.

For each treatment the user provides a demand saturation curve: the expected sales for different given values of spend (spend points). The main other user inputs consist of a series of constraints such as:

- maximum overall budget;
- maximum and minimum spends for a treatment;
- maximum and minimum spends for specific channels (combinations of treatments);

- maximum and minimum sales for specific treatments or channels.

The integer linear programme then determines how much (if anything at all) to spend on each treatment in order to maximise contribution subject to all the specified constraints. The binary decision variables a_{ij} determined in the integer linear programme take the value 1 if spend point j is selected for treatment i and 0 otherwise.

Alternative objectives such as minimising spend to achieve a given level of sales can also be set.

The user can also provide further parameters so that constraints relating to the number of recruits, the number of enquirers and even the volume of carbon units can be considered.

Finally, the integer programme formulation also supports the inclusion of a halo effect whereby spend for one group of treatments (for example, TV) can produce an uplift on the sales for specified other treatments (for example, direct mail).

OBTAINING THE MARKETING PERFORMANCE METRICS

Budget optimisation can get extremely complicated when all interactions are considered, that is, the effect of multiple treatments on an individual. Our approach was to start by regarding each treatment as independent, and then to consider layering on the halo effects.

For a company trying to build a true set of saturation curves for each treatment (ignoring halo effects for the time being) the challenge is that they cannot be tested in the real world. For a more mature business, however, such as this one, there is a great deal of very useful historic information.

The budget, when broken down, contained a total of c. 250 treatments, each requiring its own saturation curve. Fortunately, the client had a very detailed description of all campaigns and treatments

that had been used in the past 5 years; over this time each treatment had been given, at different times, different budgets, achieving various levels of sales value.

The 250 treatments were split into around 30 groups, and each group was then put through a process of normalisation. This was to take account of the changes in overall budget, and the resultant changes in overall value of the campaigns. By extrapolating the results, we could understand the overall ‘shapes’ of the saturation curves at the treatment group level.

Once we had the group-level saturation curves, the curve of each individual treatment was calibrated to match the points of the last year’s budget, and to determine the appropriate spend level.

This simulation does not, however, fully represent the problem, as it does not take account of the interdependence between the treatments. For example, in this instance, without the halo effects, budgets for TV were reduced by 40 per cent; to address this, halo effects need to be applied.

Practically speaking, the relationships between most treatments are lost in ‘noise’. If a potential prospect were to be exposed to three treatments, there would be nine separate interactions to be considered, yet each customer would have each spend allocated to the effect of just one treatment type.

In addition to this, there is the further noise of competitor campaigns, which are entirely ignored in this kind of simulation.

Viewing these relationships from a higher level reduces the noise. In this case, we set the halo effect as dependent on spend on TV as a whole. When reviewing the rules that created each of the 30 groups’ initial saturation curves, it was found that several of them were clearly ‘out of sync’. These could be matched to trends in the TV spend, and thus the halo effect was able to be derived.

RESULTS ACHIEVED AND OVERALL PERFORMANCE

The initial scenario was based on the previous year's results, that is, spend c. £40 million to return three times as much longer-term value.

Optimisation performance was measured either as a percentage of the savings from budget reductions to achieve the same value, or as the uplift in value for the same budget spend.

In all there were 10 scenarios developed by Budget Calculus: each set out to answer or react to a specific possible business objective:

- Maximise value/minimise budget without halo effects.
- Maximise value/minimise budget with halo effects.
- Maximise value/minimise budget with halo effects and maximizing recruit LTV.
- Maximise value for a change in budget (increase and decrease).
- Maximise value after removal or adjustment of specific treatments.

The initial scenarios (1 and 2) reallocated spend from TV; this was a result of the halo effect not being included. Once the halo effect was introduced, Budget Calculus could save 8 per cent from the budget or

add 5 per cent to the overall profit (value minus spend).

Scenarios 5 and 6 include the LTV of recruits, and therefore favoured those treatments with a higher recruit ratio. The scenarios after six are also not necessarily comparable to the base. These cases were undertaken to show that the contribution levels can be maintained with tactical budget cuts.

In the table below we have indexed the previous year's spend, and total value, each as 100.

AN ANALYSIS OF HOW THE BUDGETS WERE REDISTRIBUTED ACROSS 10 SCENARIOS

The results from the latter scenarios proved that the redistribution of budgets can either maintain the contribution or profit levels, even if the overall budget is cut, or maintain the budget and increase contribution. We now want to examine the frequency and size of these changes.

As mentioned earlier, each spend will settle onto a point. Figure 1 shows the distribution of change from business as usual, for all treatments in each scenario.

Typically, at least 25 per cent of all treatments in each scenario did not change.

<i>Title</i>	<i>Scenario</i>	<i>Objective</i>	<i>Total spend</i>	<i>Total value</i>
Initial scenario	0	The base spend	100	100
Minimum spend	1	Initial optimisation fixing contribution	95	100
Maximum contribution	2	Initial optimisation fixing Budget	100	102
Base and Halo (minimum spend)	3	Copy of Scenario 1, including Halo affects (4 per cent increase in certain treatments)	92	100
Base and Halo (maximum Contribution)	4	Copy of Scenario 2, including Halo affects (4 per cent increase in certain treatments)	100	103
Base Halo and LTV (minimum spend)	5	Copy of Scenario 3, including LTV	82	100
Base Halo and LTV (maximum contribution)	6	Copy of Scenario 4, including LTV	100	107
Base Halo and LTV increase in budget	7	Copy of Scenario 6 – increase of budget to £42 million	116	112
Base Halo and LTV decrease in budget	8	Copy of Scenario 6 – decrease of budget to £30 million	83	100
Base Halo and LTV removal of treatments	9	Copy of Scenario 6 – removing treatments	100	104

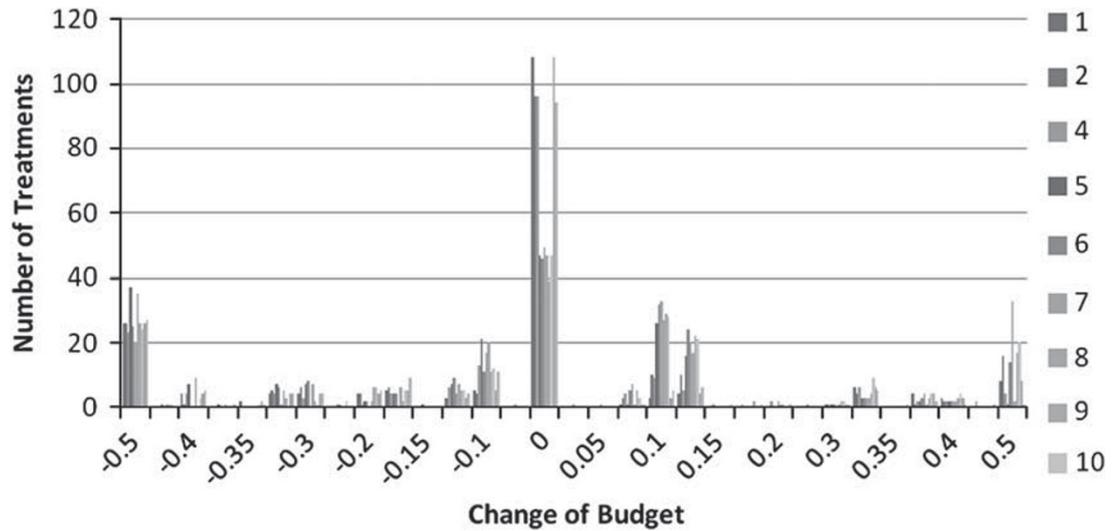


Figure 1: Budget change.

A further 25 per cent were pushed to an extreme upper and lower limit, while in the majority of scenarios 50 per cent of treatments only had a slight change.

In the case of Scenarios 5 and 6, where LTV was to be considered, up to 50 per cent of treatments do not change spend. This implies that the current set up of the company is unwittingly optimising on LTV.

The majority of stakeholders fear the concept of change as it would force managers to do more, or less, than they currently do. In the case of this company the major changes came as no surprise, and for some treatments they had been recommended before, but the impact had never previously been quantified. The chart above proves that to optimise spend across many alternative communications requires a delicate balance.

OVERALL CONCLUSIONS

This case study shows that there was an opportunity to make substantial financial gains, or savings, through the use of optimisation technology, in this case using an integer linear programme, to distribute marketing budgets.

In every case the amount of these gains, and the way they are derived through the redistribution of budget across multiple competing destinations, will be different.

We learnt that the scenarios we ran were very dependent on our understanding of both saturation curves and halo effects; if these did not exist then the case for using optimisation technology would be much weaker.

However, understanding saturation curves and halo effects depends either on previous experience or on judgement; in this case we had the advantage of a considerable level of historic performance metrics to draw on.