Institute of Actuaries of Australia

An Actuarial Approach to Optimising the Trade-Off Between Media and Price Promotions

Prepared by Adam Driussi, Caroline Stevenson and Tony Davis

Presented to the Institute of Actuaries of Australia
Biennial Convention 23-26 September 2007
Christchurch, New Zealand

This paper has been prepared for the Institute of Actuaries of Australia's (Institute) Biennial Convention 2007. The Institute Council wishes it to be understood that opinions put forward herein are not necessarily those of the Institute and the Council is not responsible for those opinions.

© Driussi, Stevenson & Davis

The Institute will ensure that all reproductions of the paper acknowledge the Author/s as the author/s, and include the above copyright statement:

The Institute of Actuaries of Australia
Level 7 Challis House 4 Martin Place
Sydney NSW Australia 2000
Telephone: +61 2 9233 3466 Facsimile: +61 2 9233 3446
Email: actuaries@actuaries.asn.au Website: www.actuaries.asn.au

© Driussi, Stevenson and Davis

An Actuarial Approach to Optimising the Trade-Off Between Media and Price Promotions

© Driussi, Stevenson and Davis

1. Introduction

Marketing decisions have historically been categorised as falling into one of four controllable categories, otherwise known as the 4 Ps of marketing – product, price, place and promotions. Achieving a successful business outcome requires an appropriate balance of time and money invested in each of these four areas.

As businesses become increasingly customer centric, many organisations have felt that their overall marketing strategy (encompassing all of the 4 Ps) is too important to be only considered by the marketing function. In many cases, the marketing function focuses primarily on marketing communications, with marketing strategy in aggregate being considered by a wider group. In this paper, we are referring to marketing strategy in the wider sense when we discuss marketing or marketers.

In a world where there is increasing pressure on marketers to become more accountable for their expenditure, the scrutiny on the effectiveness of spend in each separate area, and in total, will only continue to increase. This highlights the need to become efficient in considering the marketing budget in total, i.e. for carefully balancing the investment across different marketing activities. Marketers are increasingly turning to actuaries and financial experts to help respond to these increasing demands.

Despite the necessity for balancing the investment in each marketing dimension, it is not uncommon to find different people responsible for each of the individual dimensions, whether it is for structural, cultural or technical reasons. This often leads to a disconnect between the planning done for each of these dimensions, with each being “optimised” relatively independently.

The biggest issue in this regard occurs where the investments made in promotion and the development of pricing strategies are considered independently. This paper considers integration of these two key pillars of marketing and discusses approaches to analysing the effectiveness of the two, both separately and in combination.

1.1. Traditional analysis of promotional expenditure

Promotional expenditure alone is often a very large component of company expenditure. In fact, total media expenditure in Australia over the last year was in excess of $10 billion, with more than 50 companies spending in excess of $25 million per annum.

© Driussi, Stevenson and Davis

Despite its importance in budgetary terms, the effectiveness of advertising expenditure is rarely measured using rigorous, evidence based approaches. Traditionally, measurement of the effectiveness of advertising expenditure has focused on qualitative measures and brand tracking, rather than measuring the actual impact on sales. While we recognise that longer term brand building, to maintain a certain level of baseline sales, is an important function of some promotions, this paper examines the effectiveness of media in generating more immediate sales (or enquiries) above this baseline level.

Simple 'source code' analysis represents an attempt to improve on qualitative measurement, by effectively attributing sales outcomes to the effect of a single channel. Usually, the channel credited with the response is identified either through the telephone number called, or channel used to respond, or by asking the respondent what prompted them to contact the company.

While this methodology can be very useful in understanding the effectiveness of advertising within a channel, source code methods fail to take account of multi-media effects. A response is attributed to one channel only, without taking into account the influence of other media the respondent may have been exposed to. In a mature market characterised by increasing proliferation of media channels, as shown in Table 1.1 below, such an approach will almost always give a biased view of the effectiveness of each media channel in the mix.

Table 1.1 – Media Landscape 2000 to 2008

Medium	2000	2008
TV Channels	50	100+
Pay TV Penetration	23%	35%
Radio Channels	220	250+
Cinema Screens	1554	1800
Newspapers	630	600
Consumer Magazines	725	800
Internet Penetration	48%	90%

Whilst source code analysis can be useful in understanding what is occurring within a channel, its limitations make it difficult to identify appropriate strategies across combinations of channels. Other methodologies are more suited to identifying these across channel effects.

When considering the traditional measurement of multi-media effects, it is still true that many marketers are still struggling with the conundrum raised by Lord Leverhume quite a number of years ago - "I know half my advertising expenditure is wasted, I just don't know which half".

© Driussi, Stevenson and Davis

One approach to help answer this question is the use of econometric techniques to better understand the effectiveness of different media channels. The spectrum of econometric modelling is broad, ranging from relatively simple modelling techniques, to approaches that analyse past experience with a high degree of detail, obtaining full value from the granularity of information available.

Most companies start with relatively simple econometric approaches to analysing media effectiveness. While these techniques represent a step forward, they are still subject to a number of shortcomings, as discussed in more detail in section 1.3 below.

Simplistic econometric models can be helpful in obtaining an initial understanding of the major influences on advertising effectiveness, but often suffer from a lack of predictiveness. They are usually unable to fully describe the impact of different media channels and rarely incorporate the full range of price factors and other non-media factors in as much detail as required. This typically makes them unsuitable for use in developing detailed media strategies or for forecasting the effect of the strategies considered.

It is worth recognising that many companies are subject to a variety of impediments to introducing new approaches to measurement. These could be cultural resistance to new measurement techniques and a comfort with entrenched management information, legacy management information systems that are difficult to change, or a reliance on planning and analysis provided by their agencies.

1.2. Interaction between price and promotion

Companies with separate departments managing advertising spend and price (or revenue/margin) per sale can have substantial inefficiencies in their marketing budgets. In this situation it is not uncommon for the company to be spending large amounts of money on promotional advertising and separately making a large investment in setting appropriate pricing rules while still not achieving a good sales volume outcome.

A good example of where this can go wrong can be found in retail banking, when one considers products such as credit cards or personal loans. Pure promotional optimisation would be focused only on increasing the number of credit card/personal loan applications, regardless of the quality of the applicant. The credit risk and pricing areas, on the other hand, are concerned with optimising the profitability from this book, and will therefore look to set interest

© Driussi, Stevenson and Davis

rates and decline conditions in such as way as to ensure that only applicants of a particular quality will get through the application process.

If increased promotional expenditure only manages to generate applications from poorer quality leads, then most of the additional leads are likely to be either declined or priced out of the market – a situation where large marketing investments are being made, for very limited return.

It is not uncommon to find that the individual KPIs of the 'promotional' team and the 'pricing' team have been set so as to encourage this sort of outcome. As a result, a business may achieve a 'successful' marketing outcome and a 'successful' pricing outcome without achieving a 'successful' business outcome overall. In fact, in some cases, the structure of the business may not be suitably designed to enable the realisation that there is a problem – perpetuating inefficient investment in price and promotions.

Where price promises are used in advertising, pricing strategies set without taking this into account can also give unexpected and negative outcomes. The actual price advertised can have a significant impact on the profile of leads attracted to the company, which in turn can make a substantial difference to the bottom line.

Again, a good example of this effect can be found in personal loans, in this case using a UK example. Legislation in the UK requires any provider who advertises a set APR (annual percentage rate) for a personal loan to offer an interest rate at least as good as this APR to two thirds of its customer base. This typically leads to a profile of offered interest rates similar to those shown in Figure 1.1 below, where interest rates increase gradually for the first two thirds of the customer base, and then jump sharply for the last third of customers (when ranked by risk or some other measure of desirability to the bank).

© Driussi, Stevenson and Davis

Figure 1.1 – UK Personal Loans: Distribution of Interest Rates Offered

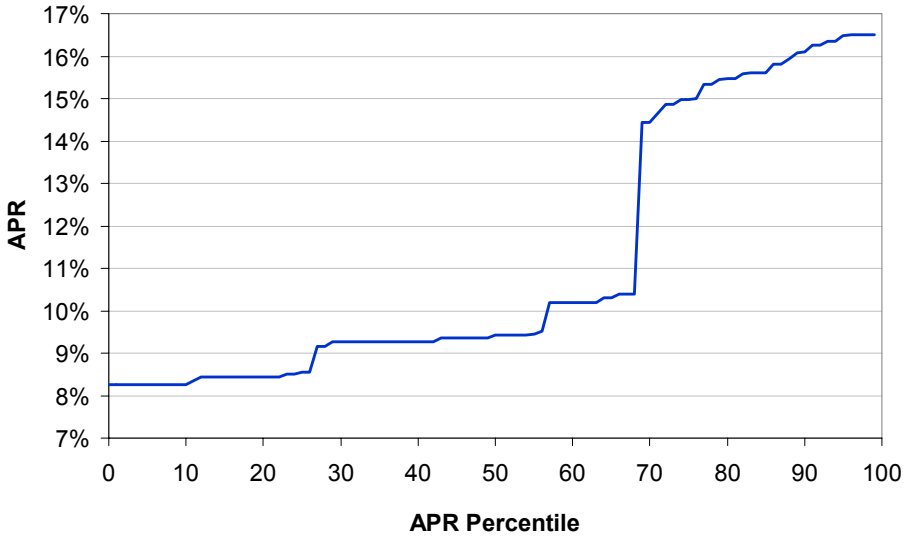

Usually, the customers who are slightly below the two thirds cut-off point are subsidised by more desirable customers, in an attempt to keep the advertised APR low. If a pure pricing exercise were conducted, the bank might decide to increase rates for these customers, to improve profitability in this area. This should allow it to offer lower interest rates to the most desirable customers, hopefully increasing conversion in this segment and theoretically improving the outcome for the portfolio overall.

However, this relationship can break down in a situation where the APR is used as part of an advertising message. By increasing rates for, say, customers in the 50th to 67th percentile, the bank will need to advertise a higher APR to comply with the legislation. This in turn influences the profile of potential customers who apply for a loan, often biasing it towards higher risk and less desirable customers. This can serve to counteract the gains that should be made due to the changes in pricing, potentially leading to a worse outcome for the portfolio overall.

These simple examples highlight the need to understand promotional expenditure not just in isolation, but also in the context of the pricing strategy of a given company.

© Driussi, Stevenson and Davis

1.3. Simple econometric approaches to analysing promotional expenditure

Before we consider the impact of price and promotion in combination in more detail, it is worth considering 'traditional' simple econometric approaches that are often used by marketing agencies to understand the effectiveness of media expenditure.

Simple econometric approaches typically utilise a multiple linear regression approach to understand the effectiveness of different media channels. Analysis is typically done using weekly or fortnightly observations at a national or state level, over a time period of anywhere between one and three years. This gives anywhere between about 30 and 150 individual data points for use in the regression analysis.

Although econometric models will seek to find the relationship between the response variable (for example, applications or sales) and a range of explanatory factors including channel, price and external effects, the low volume of data can make it very difficult to identify all of the drivers of the experience. Often, media campaigns are run concurrently across multiple channels, further decreasing the modeller's ability to isolate the impact of each channel individually.

The difficulty of identifying the contribution of each explanatory factor is exacerbated by the fact that there are so many different influences on the response variable. Often, many of the non-media effects are ignored, or can only be incorporated in a very simplistic way. Consider the example shown in Figure 1.2 below.

Figure 1.2 – Simple Econometric Modelling Example

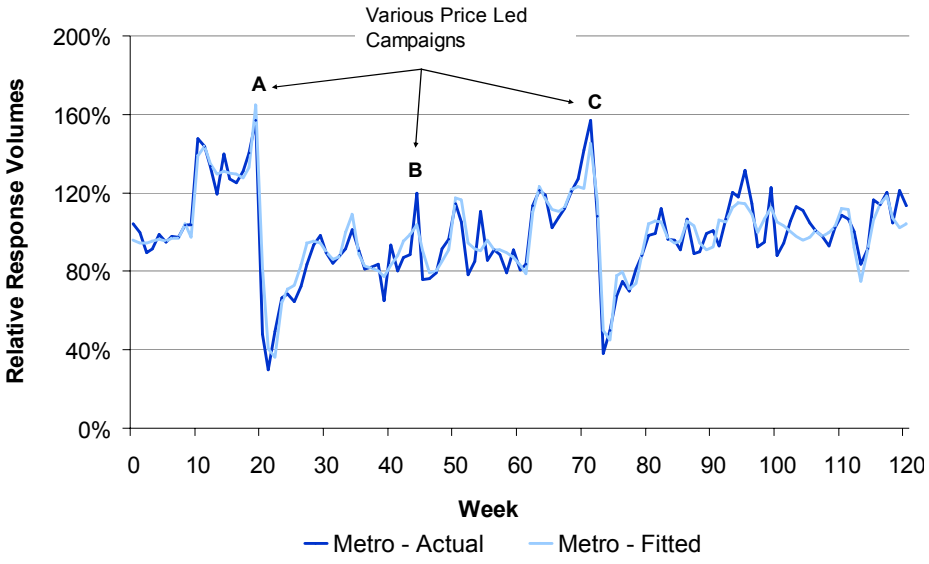

© Driussi, Stevenson and Davis

In this instance, a simple time series based approach might attribute the peaks seen at points A, B and C to a series of price led campaigns. Such simple approaches cannot effectively separate the impact of customer characteristics, other media or external factors from the price impact, and hence can lead the user to draw potentially incorrect conclusions about the impact of price (or indeed media or other factors). If this is used to inform pricing or media strategy, the results could be highly inaccurate and potentially value destroying for the business.

Even where it is possible to discern that a relationship between the response variable and a given factor exists, the lack of modelling data can make it difficult to determine the exact shape of the relationship. Consider the example given in Figure 1.3 below.

Figure 1.3 – Relationship Between Television Spend and Leads

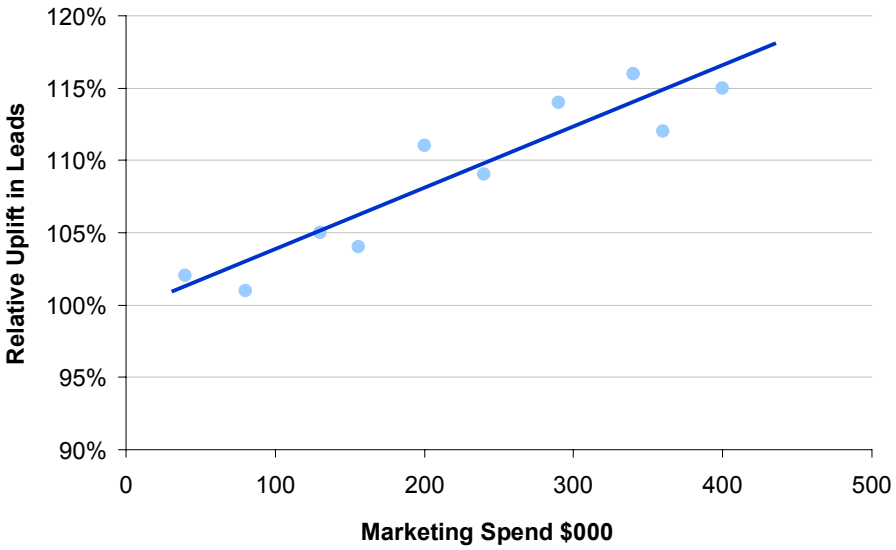

While a strong positive relationship between television advertising and leads generated appears to exist, it is not possible to definitively identify the shape of this relationship. In particular, we cannot identify whether a straight line or a curved relationship is more appropriate or at what point the marginal return from television spend diminishes. This can have a significant impact on the media strategy selected. Consider Figure 1.4 below, which shows the same data points as seen in Figure 1.3 above, but presents an alternative relationship between television advertising and leads, that could also be supported by the data given.

© Driussi, Stevenson and Davis

Figure 1.4 – Alternative Relationship Between Television Spend and Leads

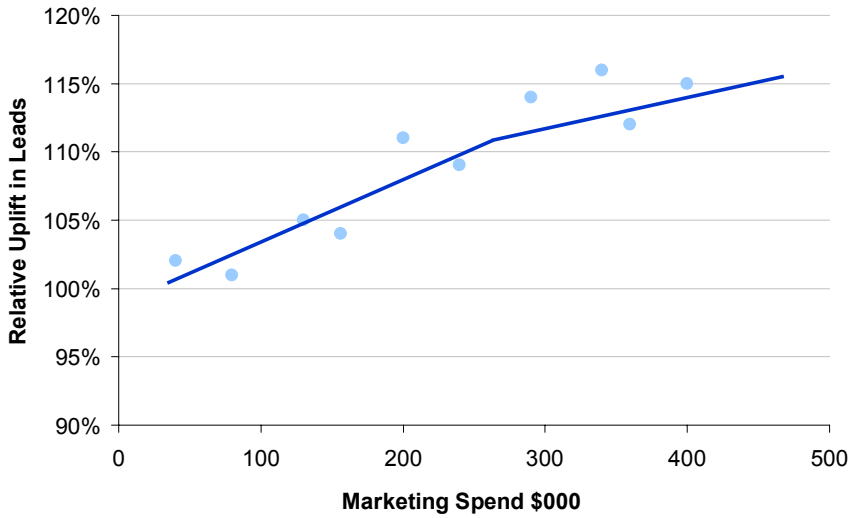

If the straight line relationship shown in Figure 1.3 is correct, then increasing advertising expenditure beyond $250,000 will continue to generate increased returns at the same level per dollar as expenditure below this point. However, if the relationship between television advertising and leads is more like the line shown in Figure 1.4, then the benefit gained by increasing advertising beyond this level will be smaller. Simple econometric approaches based on limited data cannot accurately identify the exact shape of these relationships, and hence many of the media strategies developed utilising these models do not perform according to expectations.

Another major weakness that exists in many simple econometric approaches is the fact that the regressions performed attempt to identify the relationship between leads and media spend, rather than media exposure. Media exposure can be measured in a variety of ways, but gives a truer picture of the audience for a given advertisement than spend does. Often, the cost of an individual unit of exposure is not constant over time, which implies that the audience one reaches for a given dollar spend in advertising (and hence the potential effectiveness of the advert) is not constant over time. An example of this can be seen in Figure 1.5 below, which shows the relative cost per TARP[1] for television advertising by month of the year.

[1] TARPs stands for Target Audience Rating Points and is a way of measuring population weighted impacts (exposure) on television

© Driussi, Stevenson and Davis

Figure 1.5 – Relative Cost per TARP by Month

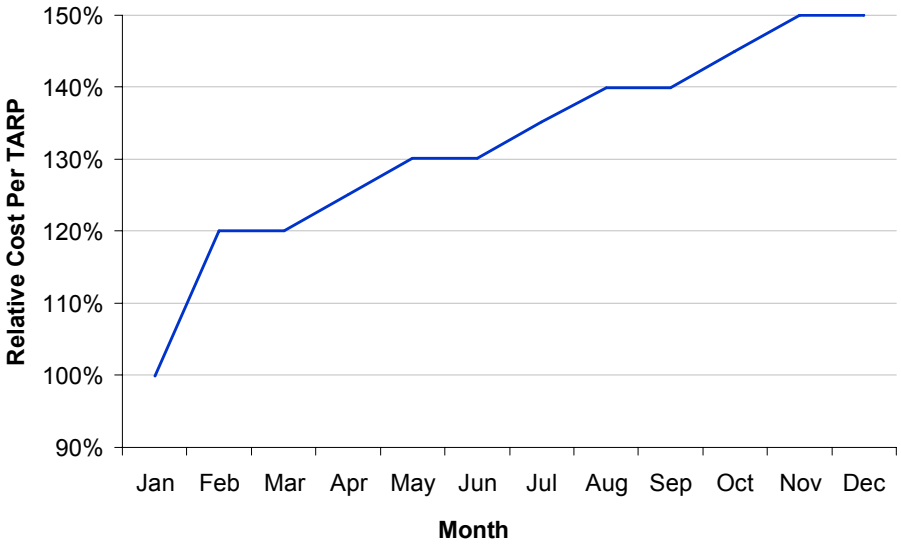

We clearly see an upward trend in the cost of television advertising as the calendar year progresses. If we model the effectiveness of television using spend, then the conclusions drawn may differ depending on whether we advertised early or late in the year. For the same dollar expenditure, an advert run earlier in the year will be deemed more effective than one run later in the year, simply because we reach a smaller audience but have not taken this into account in the modelling approach.

Even if the modelling is standardised by the average costs per month, there can still be inaccuracies in spend based modelling due to:

- Short term fluctuations in the cost of media where the agency takes advantage of tactical buying opportunities;
- 'Events' that drive unusual media costs in particular markets (e.g. state and federal election campaigns, Olympic Games, etc);
- The issue of 'bonus' exposure issued by the media owner (where surplus inventory is available and provided on a free of charge basis).

An exposure based approach allows for the differences in reach over the time period, and will draw more correct conclusions about the effectiveness of advertising for each channel.

© Driussi, Stevenson and Davis

The weaknesses in over simplistic econometric approaches arising out of the lack of data and the prevalence of spend based modelling are substantial. As discussed above, these techniques are typically unable to properly isolate the impact of different media channels and other factors and are hence generally quite poor at predicting future outcomes. Therefore, development of a detailed marketing strategy needs to be based on more information than is available from these type of approaches.

1.4. Overcoming the weaknesses of simple econometric approaches

In an environment where marketers are under increasing pressure to provide justification for marketing budgets and expenditure, simplistic econometric approaches to analysing and determining appropriate advertising expenditure or pricing strategies are not sufficient. Media and pricing strategy needs to be developed using robust, evidence-based approaches that are able to explain the impact of different channels and factors, in order to provide good predictiveness into the future.

The remainder of this paper presents a more rigorous approach that can be used to not only understand the impact of different channels and factors, but also to optimise expenditure, obtaining the best sales/profit outcome for different levels of media expenditure and price.

Chapter 2 discusses an improved, more granular approach to media modelling. We then discuss how this approach can be used to simulate media strategies and determine an optimal allocation of media budget in chapter 3.

We further discuss the influence of price and customer value, and how this can be incorporated into the media strategy in chapter 4. We then consider some initial thoughts on extending the media optimisation problem to concurrent optimisation of media and price in chapter 5.

It is important to recognise that all of the techniques discussed in this paper utilise past data to predict future experience. Thus, although they form a helpful guide for use in media planning, and can be extremely valuable in ensuring that the bulk of the media budget is utilised efficiently, they should not be used without appropriate practitioner judgement and experience. For example, future markets and opportunities may be different to those available in the past, necessitating experienced based judgements to determine suitable allowance for these variations.

© Driussi, Stevenson and Davis

The approach presented in this paper is not trying to automate the production of media schedules. Rather, they are designed to give advertisers an additional tool in planning, to obtain greater certainty about the majority of the media spend. By having greater confidence in the results expected under, say, 80% of the existing budget, advertisers have the flexibility to trial new approaches that might otherwise never have been found, thereby allowing greater creativity than might be possible without the use of these tools.

1.5. Case study

Throughout the paper we have included a case study at the end of each section which gives an example of how the approach works in practice, using a general insurance company to demonstrate the concepts.

Whilst we have amended actual results and data to protect client confidentiality, broad details of the insurer are as follows:

- An established consumer brand with substantial annual marketing spend;
- Predominantly writes two classes of business – motor insurance and householders insurance – in all states of Australia;
- Advertising has historically consisted of a combination of both product specific executions (both motor insurance and home insurance) and brand executions;
- Have used a wide range of media channels historically, with a particularly heavy emphasis on television and press;
- Historically there was a disconnect between the individuals responsible for pricing (the actuaries) and the individuals responsible for promotional activities (the marketing team).

The insurer was keen to answer questions such as:

- What is the effect of each media channel both independently and in the mix?
- How are channels inter-dependent and what is the value of the 'halo effect' (i.e. the impact of advertising one product on the sales for another product)?
- What impact do brand promotional activities have on immediate sales volumes?

© Driussi, Stevenson and Davis

- How does competitive activity influence sales volumes and what should we do when certain competitors advertise?
- What is the minimum spend required to deliver a certain sales target (or alternatively the maximum number of sales achievable with a set budget)?
- What is the response tail for each media – and how does this affect optimal flighting patterns for a campaign?[2]
- What is the optimal media mix by channel, region and week/month?
- Which regions have higher conversion rates and profitability and what impact does this have on the optimal media strategy?

[2] The response tail of a media channel allows the practitioner to understand the pattern of effectiveness of a media channel, both during exposure on the channel and after it ceases. Flighting refers to the phasing of media deployment, e.g. constant levels of media exposure over a sustained period or bursts of media exposure followed by absence from market. Flighting also considers variations in the combination of channels used.

© Driussi, Stevenson and Davis

2. Media Modelling

2.1. Modelling Approach

Many of the problems associated with simple econometric modelling can be overcome by applying a more rigorous, actuarially based approach to media analysis. We have found an approach based on using generalised linear modelling (GLM) techniques to model applications or sales using accumulated exposure to various types of media much more effective in predicting responses.

Benefits of greater granularity of exposure

Central to the approach is expressing both the exposure data and the response data in a much more granular fashion: modelling responses at postcode, CCD or even household level – as opposed to a national or state level. This has the effect of giving a much larger volume of points for use in modelling the effectiveness of each channel, thereby giving far greater certainty about the shape of the media effect, as is shown in Figure 2.1 below.

Figure 2.1 – The Benefits of Additional Data Points

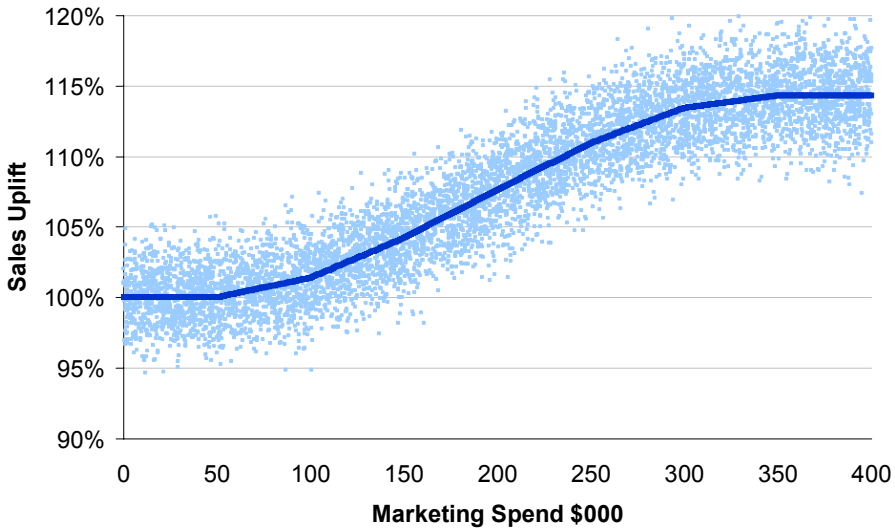

Whilst Figure 2.1 is only an illustration and in practice there is considerably more variance in the underlying data points, the additional data from modelling at a more granular level does give the modeller significantly greater confidence in the underlying shape of each factor.

© Driussi, Stevenson and Davis

By applying this more granular approach, we are able to identify the 'pure effect' of both broadcast and direct response media channels (as well as non media factors). This represents the impact of increasing advertising in that particular channel, while holding everything else constant. An example pure effect chart for television is shown in Figure 2.2 below.

Figure 2.2 – Pure Effect of Television TARPs over the Last 7 Days

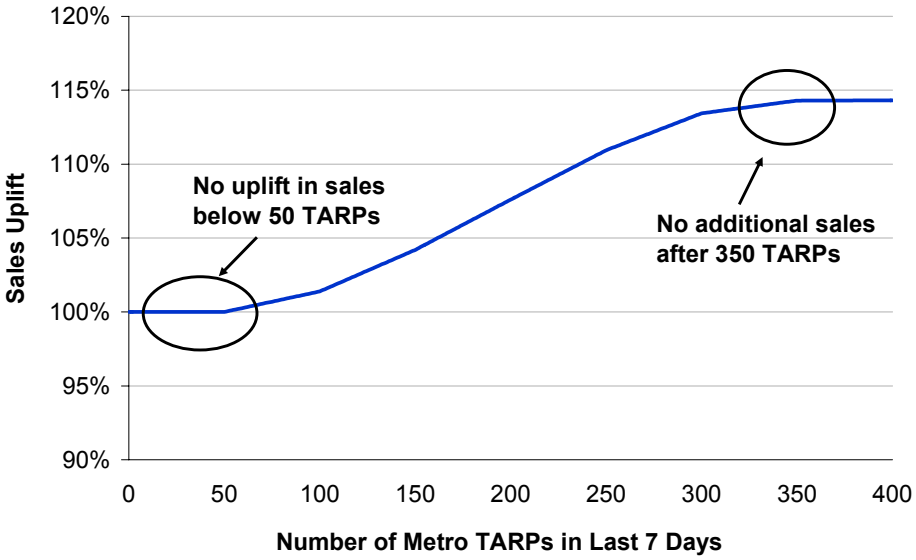

This chart shows how the relative response (shown on the vertical axis) is expected to increase when we increase television advertising. For example, when we advertise at a level of 350 TARPs on metropolitan television, this gives us almost fifteen percent higher expected sales than if we advertise at a level of 50 TARPs.

What is particularly interesting in looking at this chart is the fact that there is no discernable uplift in response for advertising at level of 50 TARPs or below in a seven day period. We also reach a point of diminishing returns once advertising reaches a level of approximately 350 TARPs in a seven day period, with limited further gains being possible for advertising spend beyond this level. From a planning perspective, we want television advertising to be at a level somewhere between these two points, so the ability to identify each pure effect accurately is an extremely valuable input into strategy determination.

The key point here is that the exact shape of each media factor is bespoke to each advertiser and each channel. What works in one industry or for one brand does not necessarily apply to a competitor or another industry. As an example,

© Driussi, Stevenson and Davis

mortgage providers advertising on a real estate website could reasonably expect a stronger response from each advertisement than an insurance company advertising on the same website.

What is required to model at this more granular level?

In order to obtain these additional data points, we need a method for accurately determining media exposure and response at a granular level.

The process of estimating media exposures at a postcode, CCD or even household level can be time consuming and complex. Publicly available exposure measures are typically published at a higher level than required – for example, television ratings are typically published by state for each of metropolitan and regional areas, usually split by demographic factors.

For other 'above the line' channels such as press or magazines the challenge is even greater. With literally hundreds of publications, we need to allocate readership statistics for each individual publication by postcode, CCD or household. As an example, the media exposures generated by postcode will obviously be very different for the Sydney Morning Herald in New South Wales compared with exposures generated by the Herald Sun in Victoria.

For 'below the line' channels such as direct mail or outbound telemarketing, it tends to be easier to determine the footprint of historical media activities. Typically companies or their agencies will retain a file of households mailed or called which can either be used at an individual household level or summarised to a CCD or postcode level as appropriate.

It is particularly difficult to estimate exposures for some promotional activities at a granular level – for example, sponsorship, PR and outdoor advertising.

It is also possible to identify and consider differently media which causes the response, which we might call 'catalyst' media, compared to media through which response is channelled, which we might call 'harvest' media (for example the yellow pages or the brand's own website).

Once the 'footprint' of each individual media channel/publication has been estimated it is possible to construct a dataset which estimates the historical exposures generated by each individual advertisement over time. This, for example, could then be summarised to estimate the total exposures in each postcode or CCD at a weekly level.

© Driussi, Stevenson and Davis

The next step is to determine the level of response at a granular level. Depending on the industry, this process can range from being straight forward to relatively complex. For example, for insurers, prospective policyholders would almost always supply a postcode at the time of quotation – enabling responses (in this case quote volumes) to be easily summarised by day and postcode. For FMCG brands, however, it is often more difficult to estimate sales volumes by postcode historically.

Specific modelling challenges

There are a number of challenges associated with applying GLMs in this context that do not arise, at least to the same extent, in many other applications relying on these modelling techniques.

The biggest difficulty arises due to high degrees of correlation between different media channels. These correlations arise due to the fact that most media campaigns utilise multiple channels at the same time.

Differences in media footprints and media costs over time will mean that, at a postcode level, there will be distinct combinations of media exposure in different areas. While this variation is sufficient to enable the identification of separate effects, high levels of correlation between different effects remain. Due to these high levels of correlation, media modelling needs to rely on a combination of judgement and consistency tests (over time or region) to a much greater extent than many other GLM models.

Because of the way the modelling datasets are constructed and the high levels of correlation in the data, any hold out samples need to be created with care. The hold out sample cannot be just as a random selection of observations as in other situations, because it is possible to have a model that fits such a holdout sample well, but fits any other dataset poorly.

Out of time testing is also important, but can present particular challenges. For example, it is much more common to find exposure levels outside the range of the levels tested historically, especially for competitor factors and channels that are new, or have had only limited budget allocated to them in the past. The greater the difference between the new experience and the levels seen historically, the greater the potential degradation in predictiveness of the results. If the changes are too substantial, it will be necessary to recalibrate the underlying media model.

It is also import to understand substantial differences in creative and advertising formats, particularly for media channels that have only been used to a limited

© Driussi, Stevenson and Davis

extent historically. For channels that are not heavily used, it is not uncommon to see particular levels of media exposure occurring only in one or two states, where particular test campaigns may have been run. In this case, it can be difficult to separate state effects, or even the test campaign effects if a substantially different offer or creative was used, from the effects of the media channel in question. Understanding these types of potential issues is an important component of building successful media models.

Optimal definition of the modelling problem

It is generally easier to build successful media models if the response modelled is as immediate and as close to the start of any sales pipeline process as possible. It is preferable to model initial enquiries (such as an application or quote) rather than eventual sales because the initial enquiry is the pure outcome of the media activity. Whether or not the enquiry eventually converts to a sale will depend on a range of additional factors, including the price, competitive position, product quality and the customer's overall experience of the sales process. Whilst important, all of these factors will dilute the ability of the modeller to identify the ability of media to drive initial enquiries, and should typically be considered separately.

We therefore recommend modelling initial enquiries wherever possible. The likelihood of converting to a sale can be overlaid on enquiries, to still enable targeting of outcomes that will prove most beneficial to a company's bottom line. This is discussed in more detail in section 4 of this paper.

Where possible, we recommend modelling at the most granular level possible. If it is possible to model at household level, then the ability of this style of modelling to drill into the efficacy of below the line activities in particular is increased. In order to do this, it is necessary to have a modelling set of all potential households that might be interested in purchasing the product marketed – not just a data set of those households who already have the product.

Large banks or telecommunications providers often have access to this type of information due to their extensive existing customer databases (covering large sections of the population, from which a very high proportion of product purchases originate) or lists of households where coverage would be available. Similar benefits can be obtained if we are interested in understanding the efficacy of media activity in helping to move existing prospects through a sales pipeline process.

© Driussi, Stevenson and Davis

In this situation, we are able to match direct mail activity directly to the household that received it, rather than just looking at total mailing activity per head of population in a postcode. This allows the incorporation of a whole range of additional factors about direct mail activity, such as the frequency of contact and time since the last contact, all of which will allow us to better understand the impact of direct mail activity on sales enquiries.

The importance of building models at a granular level

In addition to allowing efficient use of additional data, there are a number of other key reasons why we believe that it is critical to construct media effectiveness models at a granular level.

Consider the example of a major bank advertising a product such as a credit card or mortgage. There are a range of factors which vary at a postcode or CCD level which we would expect to impact on the level of responses in any individual postcode or CCD. These might include:

- The existing penetration of a core product such as a transaction or savings account – individuals with a core product such as a transaction or savings account will be much more likely to apply for a credit card or a mortgage irrespective of any media activity – and hence we would expect postcodes or CCDs with a high penetration rate of such products to have naturally higher levels of credit card or mortgage applications. Existing customers are also likely to be more predisposed for a given brand, which can enhance the effectiveness of any media activity deployed;

- Branch footprint of the bank (and their competitors) – we would typically expect the presence of a branch in a postcode or an adjacent postcode to increase the likelihood of response, all other things being equal;

- The socio-demographic characteristics of the individuals living in each postcode – certain socio-demographic groups are more or less naturally inclined towards particular brands and this has a significant influence on the level of response in each postcode.

For some industries (e.g. telecommunications) there is an additional element of serviceability – whereby the product may not be available for certain postcodes or CCDs or there could be limited availability.

Failing to take these factors into account, which all typically vary by postcode or CCD, can mean that our media models would incorrectly ascribe the impact of these variables to media factors – whether it is over or under stating the impact of each type of media. Indeed, one of the key insights of interest in the process

© Driussi, Stevenson and Davis

is typically to identify regions in which different types of media might be expected to have an increased or reduced effect.

To the extent that data is available to construct models at an individual household level it becomes possible to incorporate actual product holdings and demographics as opposed to product penetration or average demographics in the postcode or CCD. This enables the models developed to be even more accurate.

2.2. Case Study

We now revisit the case study introduced in section 1.5 to illustrate the types of insights achievable when utilising the techniques described above. By building a suitable model of the effectiveness of different levels of media exposure, we were able to start answering the questions posed by our client about the effectiveness of individual channels and the impact of competitive activity on their expected quote volumes.

What is the effect of each channel both independently and in the mix?

The modelling outcomes identified some clear opportunities for improving the effectiveness of the media exposure. One example of this can be seen by considering the pure effect of television advertising that is shown in Figure 2.3 below.

Figure 2.3 – Pure Effect of Television Advertising over the Last 7 Days

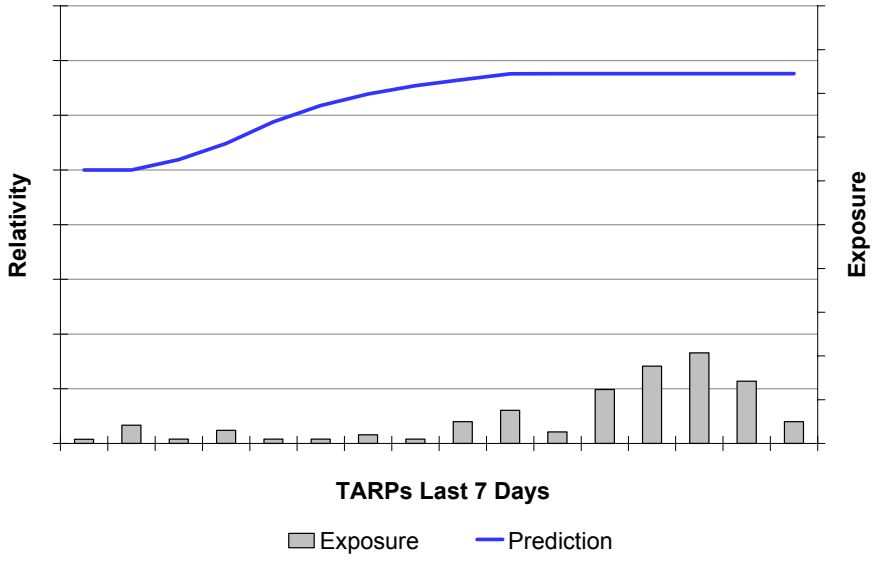

© Driussi, Stevenson and Davis

In this case, we have also overlaid the historical utilisation of television, which is represented by the grey bars on the chart (read relative to the right hand axis). Historically, a large proportion of the advertising had been at levels in excess of the point of diminishing returns. This suggested some immediate improvements in media deployment – decreasing television expenditure below historical levels.

How are channels inter-dependent and what is the value of the 'halo effect' (i.e. the impact of advertising one product on the sales for another product)?

What impact do brand promotional activities have on immediate sales volumes?

We were also able to identify that the short term uplift in quotes achievable by conducting high levels of home product specific advertising was significantly more than that achievable using similar levels of brand advertising, as is shown in Figure 2.4 below.

Figure 2.4 – Pure Effect of Product and Brand Television TARPs

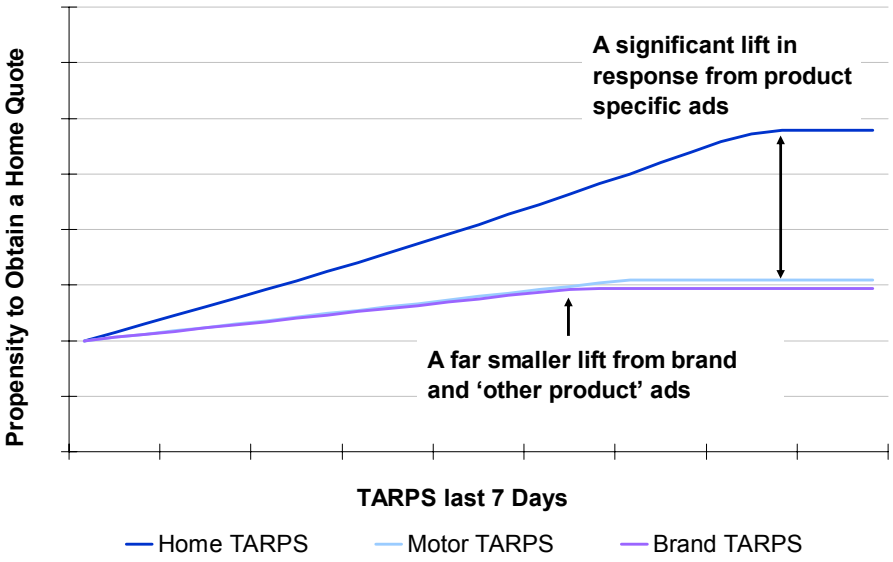

Interestingly enough, the impact of motor specific advertising on home quotes was similar to the effect of brand advertising. From a short-term sales perspective, there was therefore no real advantage to be gained by conducting brand specific advertising as opposed to product specific advertising, in this case.

© Driussi, Stevenson and Davis

Note that this ignores the longer term impact of building brand strength in helping to maintain baseline sales. This effect was analysed separately and was also considered in strategic decisions made about media deployment going forward.

How does competitive activity influence sales volumes and what should we do when certain competitors advertise?

The insurance company found that advertising by its competitors had variable effects on its own quote volumes. As shown in Figure 2.5 below, advertising by competitor B, a market leading company, decreased quote volumes for the insurance company. However, advertising by competitor A, a newer entrant, actually increased quote volumes for the insurer in question. Advertising by the newer entrant actually prompted a more general shopping response, as potential customers in the market benchmarked competitor A's premiums against more established players in the market.

Figure 2.5 – Pure Effect of Competitor Television Advertising

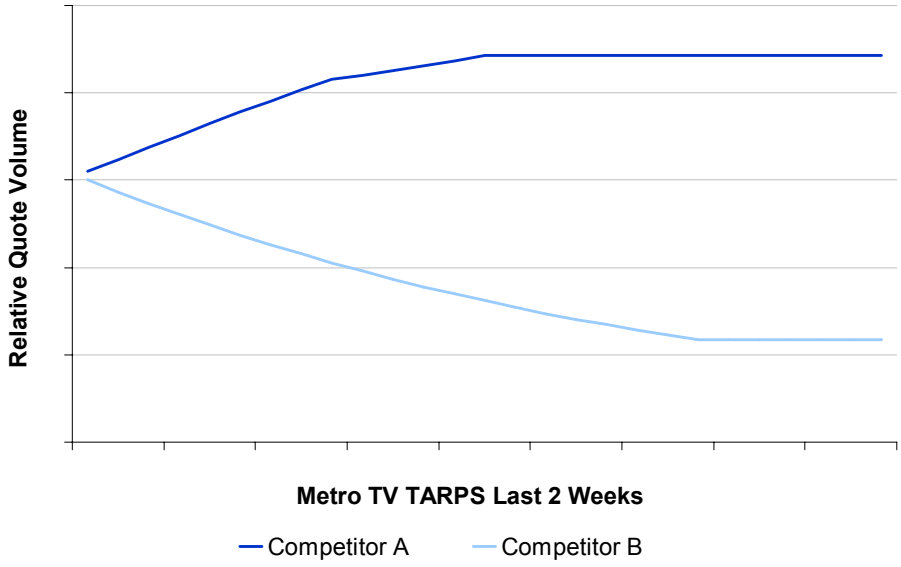

In the case of this insurer therefore, the strategic response to competitor advertising differed depending on whether it was competitor A or B advertising.

© Driussi, Stevenson and Davis

3. Simulating the Outcome of Media Scenarios

3.1. Scenario Simulation

The discussion to date has focused around determining the relationship between media exposure and enquiries or sales. The next step is to incorporate these findings into the planning process and the development of media strategies going forward.

While it is important to model the effectiveness of different media channels using exposure data, any media strategy that is developed will ultimately need to be expressed in terms of spend levels for different channels. To do this, we need to combine the pure effects with exposure footprints and cost estimates for each channel for the planning time horizon – i.e. what does it cost to buy each alternative level of media exposure for the proposed period?

These inputs can then be used to simulate the expected outcome of any given media strategy, thus providing a valuable tool for media planners to use in assessing alternative strategies prior to implementation.

The first step is to specify a range of alternative media strategies that are to be considered within the simulation algorithm. This algorithm is the result of the modelling phase discussed in section 2. The media strategy would typically be specified at the same level at which the advertiser or their agency currently specifies a proposed media schedule (potentially down to a spot level).

Figures 3.1 below sets out an example of a simple high level proposed media schedule, in this case investing approximately $370,000 over an 8 week period in three media channels – Metropolitan Television, Metropolitan Press and Online.

© Driussi, Stevenson and Davis

Figure 3.1 – High Level Proposed Media Schedule

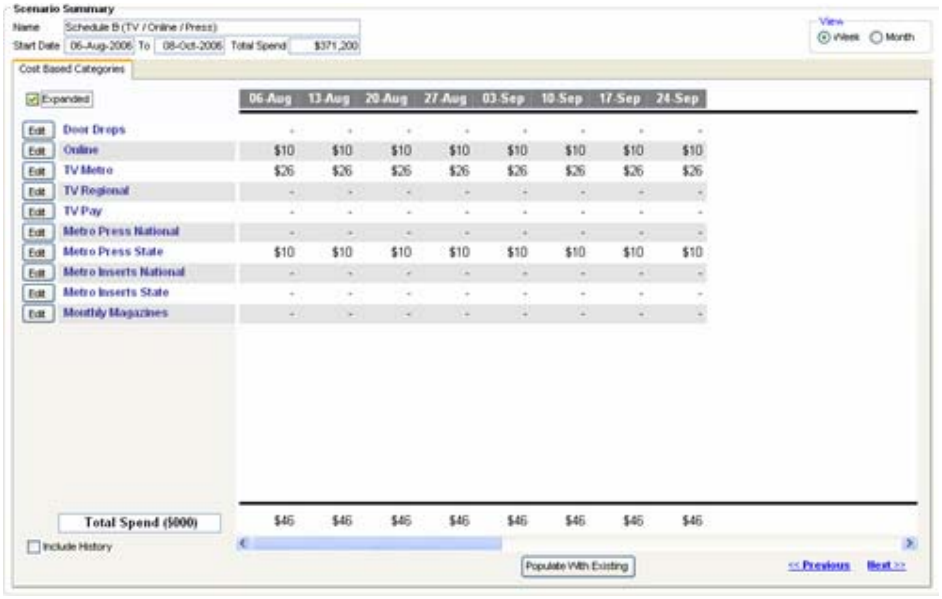

As we drill down further into the Metropolitan Press spend in Figure 3.2, we can see that it consists of a series of advertisements in each of the Daily Telegraph (Monday to Friday), Daily Telegraph (Saturday) and the Sydney Morning Herald (Monday to Friday), each with a slightly different cost.

Figure 3.2 – Metropolitan Press Schedule

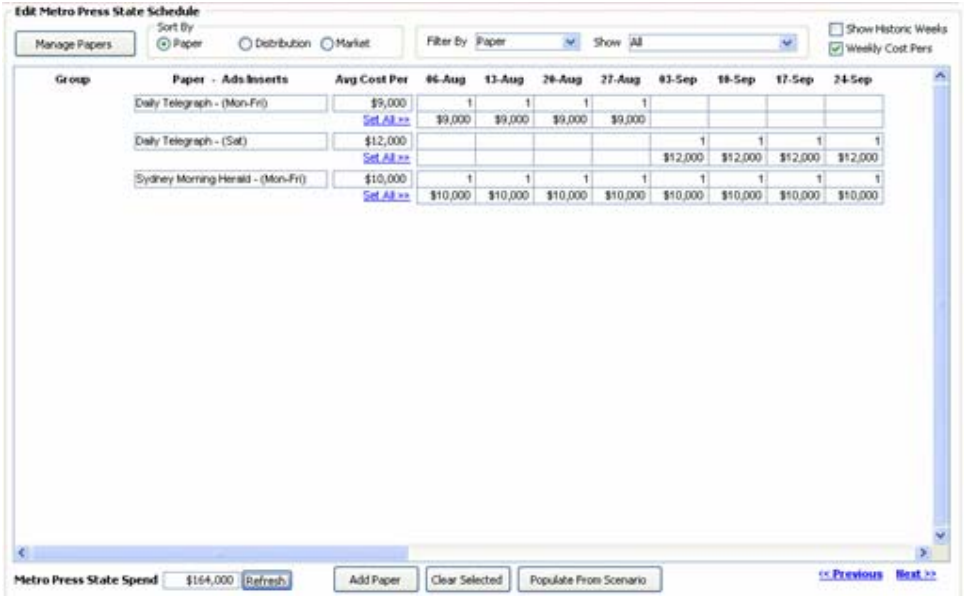

© Driussi, Stevenson and Davis

In practice the cost of each type of media can vary depending on the time of year, the volume of media being purchased, how far in advance the media is purchased and any particular deals that the advertiser or agency can negotiate.

Each of these individual advertisements will generate different levels of exposure depending on the footprint of the publication (see example below in Figure 3.3), the total readership of the publication and the size of the advertisement.

Figure 3.3 – Footprint of the Sydney Morning Herald

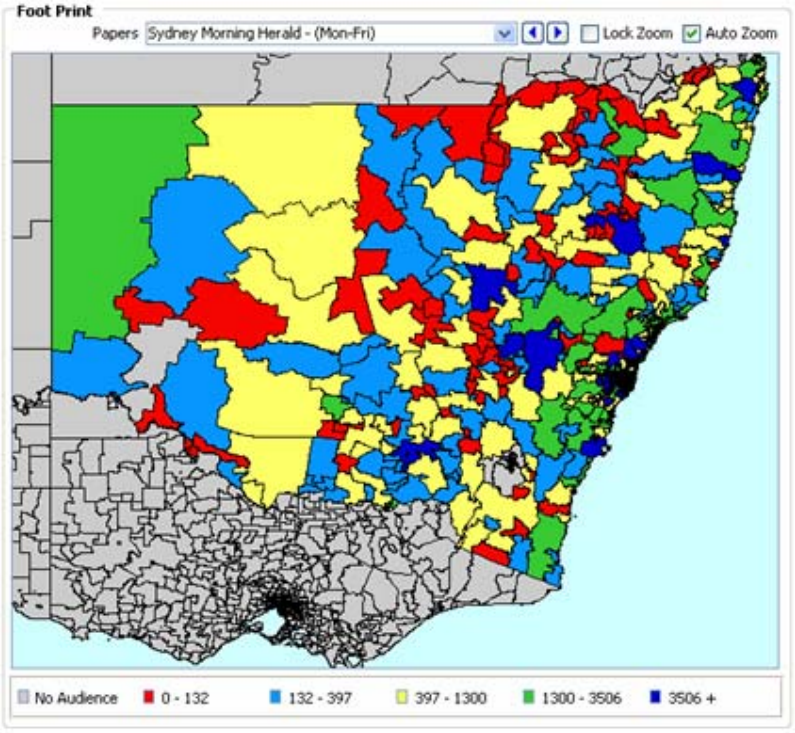

Ideally proposed media strategies should draw on the findings discovered during the modelling process, while still being valid strategies that can actually be purchased in the marketplace. As such, initial scenarios should ideally be jointly developed by the advertiser, the agency and the team conducting the modelling.

Providing scenarios are constructed sensibly, it then becomes simple to compare the effectiveness of alternative media strategies – for example the impact of increasing or decreasing investment in a particular media channel. A simplified example output is shown in Figure 3.4 below, comparing the spend versus expected sales trade off of four alternative media scenarios.

© Driussi, Stevenson and Davis

Figure 3.4 – Assessing the Results of Alternative Media Scenarios

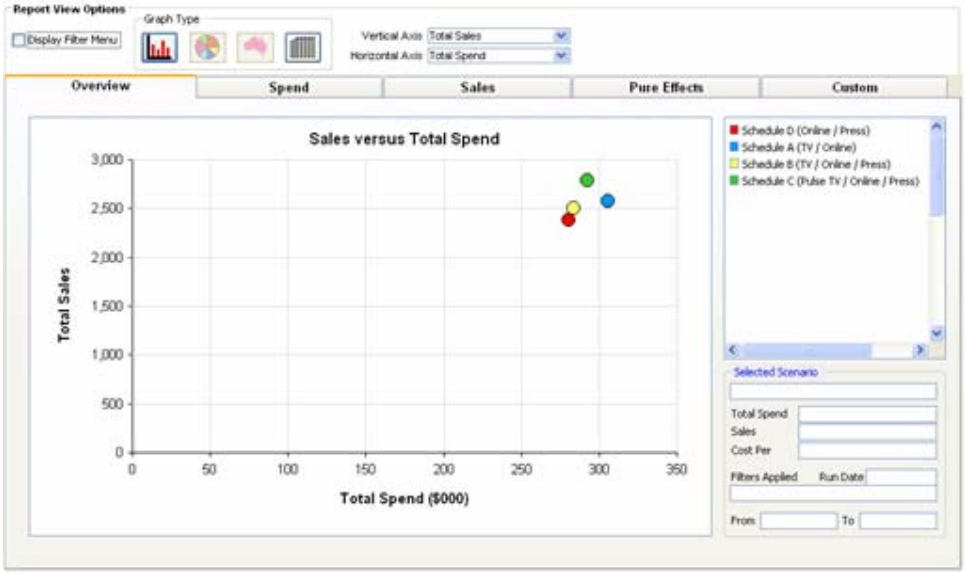

3.2. Importance of ongoing active media planning

The simulation process described above should not be seen as a replacement for the detailed knowledge developed through working in media planning. Rather, it should be seen as a useful tool that forms part of the overall planning process, but still needs to be interpreted carefully, with an understanding of both its role and its limitations.

The simulation process is particularly useful for assessing the efficacy of different 'across channel' allocations. Analysis of performance within a channel can be conducted to the extent that enough past information is available. As an example, this may involve analysing the effectiveness of radio advertising by day part or comparing the efficacy of television peak and off-peak exposures. However, the level of analysis possible within each channel in this style of approach is likely to be relatively high level, so it is advisable to supplement the macro-channel planning done here with more granular planning tools that assess the suitability of different strategies within each channel.

There will always be some factors that will disrupt the effectiveness of the media prediction. Major changes in any of the following factors have the potential to reduce the accuracy of the predicted outcome:

- Overall budgets (particularly periods of inactivity resulting from budget cuts);

© Driussi, Stevenson and Davis

- Channels utilised;
- Advertising formats and creative themes;
- Pricing and short-term promotional offers;
- Target profiles; and
- The overall market place including legislative changes or changes to the competitor set.

Small variations in these types of inputs can usually be allowed for by judgemental overlays from experienced channel planners. If the types of changes listed above have occurred previously, then the modelling, which relies on past data, should also be able to capture a large component of these effects. The extent to which this is not the case can be determined by continual monitoring of the ongoing predictiveness of the simulation process (comparing actual outcomes to those predicted using the simulation algorithm and the actual strategy implemented). If this monitoring indicates significant degradation in the accuracy of the predictions, then a recalibration of the underlying models may be required.

3.3. Media Optimisation

Having developed a process for accurately simulating the outcomes of alternative media scenarios, it becomes feasible to answer a number of key strategic questions in relation to media planning. These questions might include:

- What is the optimal deployment of a given media budget to generate the maximum number of enquiries possible with this level of spend?
- If I want to generate a certain number of enquiries, what is the lowest level of media expenditure that will allow me to achieve this?
- How does this vary for different geographic areas or times of the year?
- Given my chosen level of media spend, what flighting patterns should I employ for each media channel?

Rather than going through the labour intensive process of specifying thousands of alternative media scenarios to try and answer these questions, it is possible to use mathematical optimisation techniques to automate this process. The result is an efficient frontier, which sets out the best possible trade-off between media spend and expected enquiry levels (in this case) as shown in Figure 3.5 below.

© Driussi, Stevenson and Davis

Figure 3.5 – Optimal Trade Off Between Media Spend and Expected Enquiries

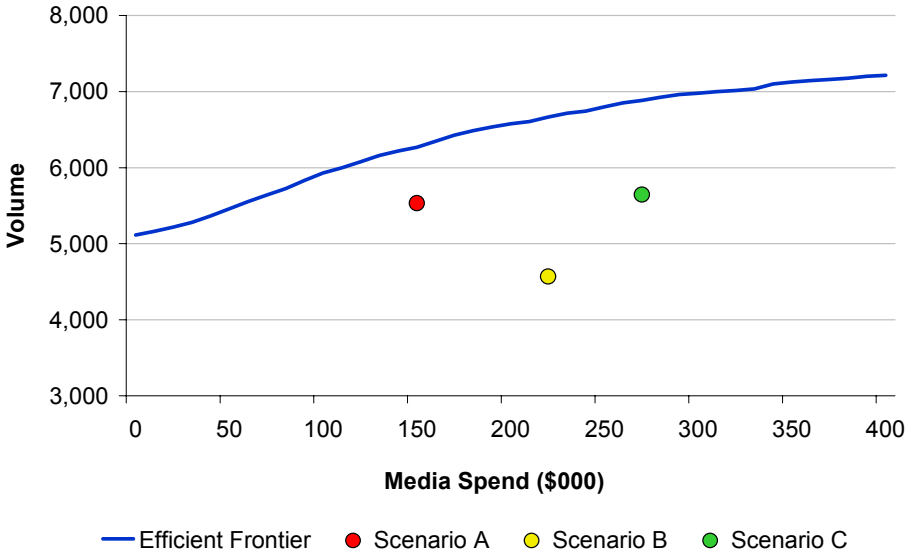

Each point on the frontier represents a different media strategy (including media deployment across channel, geography and time). In each case, the media strategy in question gives the optimal spend/volume outcome i.e. it is not possible to get the same number of enquiries for a lower level of spend, or to get more enquiries for the same level of spend with the current assumptions made around factors such as strength of offer, competitor activity and mix of media within each channel.

It is important to understand that the media optimisation is not just a theoretical exercise but needs to operate within the real world, incorporating real world business constraints and assumptions. In practice, it may not be possible to purchase the optimal schedule due to the availability (or lack thereof) of media or because of pre-existing commitments. Therefore, any media optimisation needs to operate within pre-defined parameters that constrain channel expenditure to within pre-determined bounds.

Changes in any of the constraints or assumptions, will typically serve to shift the efficient frontier either up or down as shown in Figure 3.6 below.

© Driussi, Stevenson and Davis

Figure 3.6 – Impact of Constraints and Assumptions on the Efficient Frontier

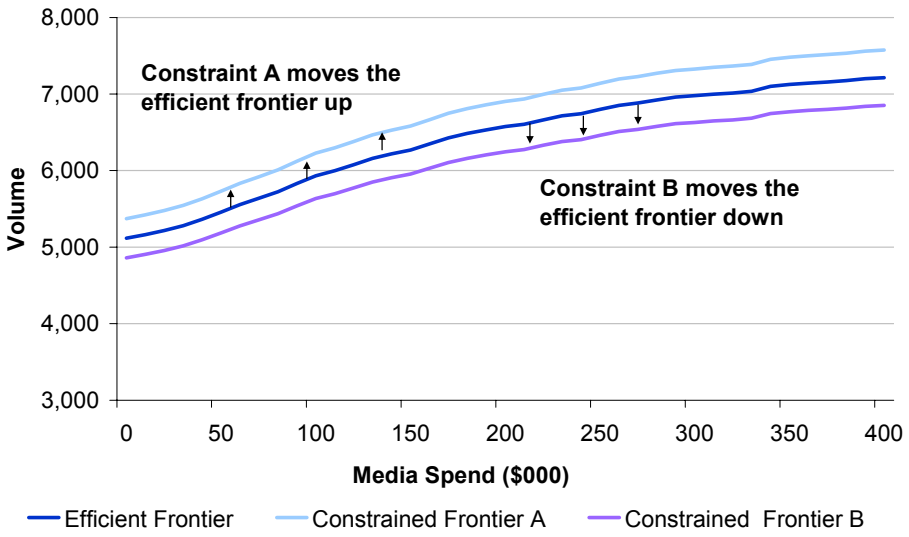

The most common constraints which are likely to shift the efficient frontier downwards might include imposing minimum or maximum investments in any particular channel. There are a range of reasons why this may occur in practice including:

- Having already committed to buying certain levels of exposure in certain channels – either because the media was pre-purchased earlier in the year or because the business is conducting a trial;
- Not having sufficient time to invest in a particular channel – for example, investing in monthly magazines may not be possible with one week's notice;
- Having physical limits or rules on how much certain channels can be utilised – for example, there may be business rules around the level of cross-sell activity that can be undertaken in a particular period or a physical limit of how much acquisition direct marketing can be carried out without mailing poor prospects.

Business 'constraints' or assumptions which might serve to shift the efficient frontier upwards include:

- Increasing the strength of offer associated with the campaign;
- Assuming different levels of market size – for example, the total demand for credit if modelling credit card or mortgage applications;

© Driussi, Stevenson and Davis

- Assuming different levels of competitor media activity;
- Assuming different levels of competitor pricing;
- Assuming cheaper media costs per unit of exposure.

By understanding the sensitivity of the results to each different assumption or constraint we can obtain an understanding of the importance of the constraint or assumption in question, which provides another valuable input into the planning process.

3.4. Case Study

In section 2.2 we discussed some of the findings derived using media modelling techniques on the past advertising experience of the insurance company considered. Although there were some quick wins available from the initial consideration of the pure effects (such as maximum levels of exposure in particular channels), development of a comprehensive media strategy required careful consideration of the effectiveness and cost of all channels in the potential mix.

What is the minimum spend required to deliver a certain sales target (or alternatively the maximum number of sales achievable with a set budget)?

What is the optimal media mix by channel, region and week/month?

These questions were answered for an upcoming campaign by utilising the simulation and optimisation techniques described earlier in this chapter. The insurer initially determined the best possible expected quote outcome for a range of different media spends to produce an efficient frontier similar to that shown in Figure 3.5 above.

Each of the points on this frontier represented a different media strategy that had been optimised to give the best results for a given level of spend by varying media deployment across channel, geography and week.

In this case, the insurer had a fixed budget for a 12 week campaign, which they wished to optimise. The initial optimisation process allowed the insurer to determine the highest number of expected quotes achievable for that budget.

In addition to determining the media mix which generated the highest expected number of quotes, the simulation process also predicted quotes at an individual

© Driussi, Stevenson and Davis

postcode level as shown in Figure 3.7 below for Sydney (relative quote volumes in one month, standardised for the population in each postcode).

Figure 3.7 – Expected Relative Quote Proportions for the Sydney Region

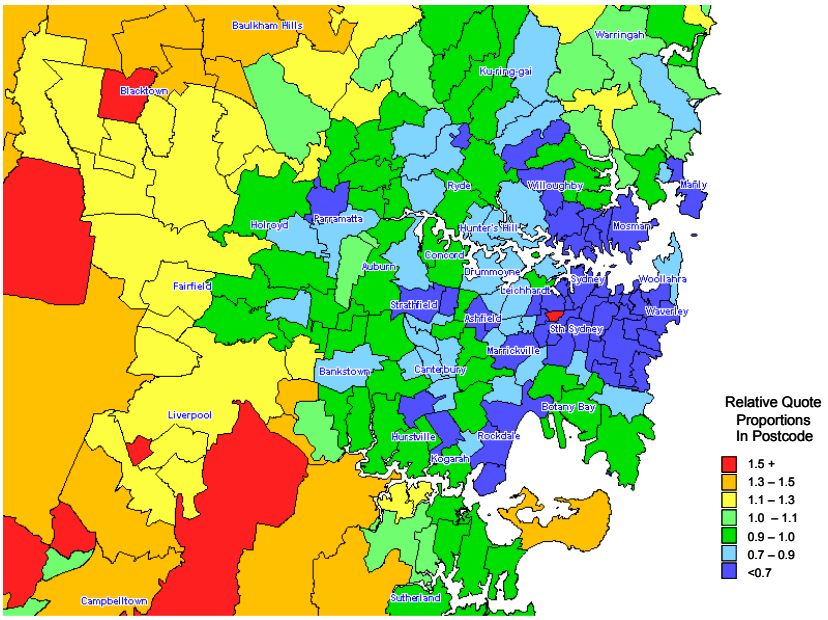

What is the response tail for each media – and how does this affect optimal flighting patterns for a campaign?

Simulation techniques were also used to help the insurer understand the response tail for each media channel.

In the modelling phase described in chapter 2, exposure in some media channels impacted quote volumes for as short a period as one week (e.g. online), whilst exposure in some others impacted quote volumes for up to 8 weeks (e.g. magazines).

Figure 3.8 below sets out the simulated response tails of each media. In this case we assumed constant levels of exposure in each channel over a 4 week period. The chart shows what proportion of lift from that activity we would expect in each week.

© Driussi, Stevenson and Davis

Figure 3.8 – Response Tails of Different Media Channels

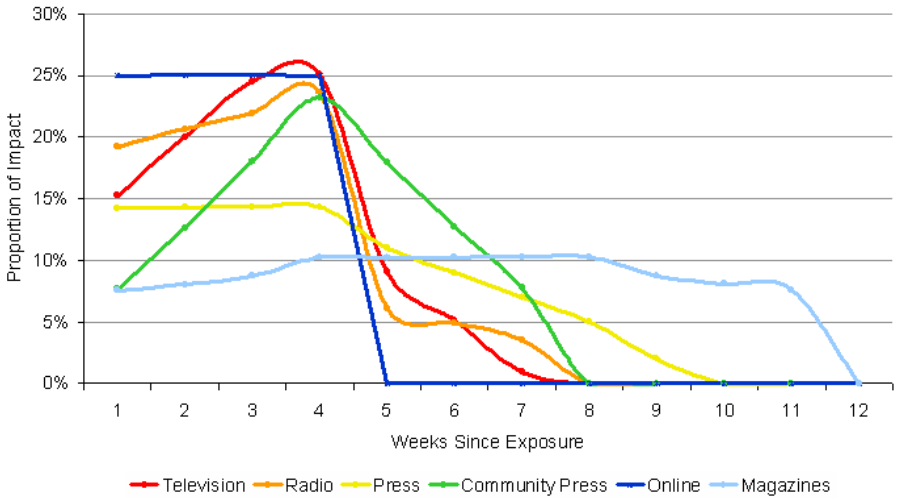

As one would expect, channels such as online had the most immediate effect, while channels like monthly magazines have a much longer residual effect. These types of results form a very useful input into determining appropriate campaign flighting.

For example, if the insurer wants to very quickly generate a short term boost in sales, then online advertising is the most effective way of meeting this objective. However, if the insurer wishes to maintain a prolonged response, online is less effective, as it requires a constant level of advertising. Press and magazines are much more suitable for delivering a sustained response, although the level of lift in short term enquiries is lower.

The availability of these simulation techniques has enabled our client to incorporate the modelling results into its day to day planning process, allowing it to make both macro level decisions about spend levels and test and compare a smaller number of specific media strategies to make more short term tactical decisions.

© Driussi, Stevenson and Davis

4. Incorporating Customer Value and Price

As discussed in section 1.2, it is important to not just consider media expenditure in isolation, but to also consider the impact of price and the quality of the enquiries generated. This is explored further in this chapter.

Influence of price factors on enquiries

At the simplest level, price factors can be incorporated into the media modelling process in the same way that any other non-media factors are included. This works best in industries or for products where pricing is largely advertised at a portfolio level, and the price charged does not depend heavily on the individual customer characteristics.

The types of outputs that might be developed using this approach can be illustrated with an example from the pay television industry. In this case, the product pricing is quite simple relative to most industries with pricing levers consisting of:

- The monthly charge for the basic channel package, with no price variation depending on individual customer circumstances;
- An installation charge – which can vary in individual campaign periods;
- Campaign offers related to a period of free service – for example, two or three months free.

Figure 4.1 below shows the uplift in new sales that was achievable under a variety of different price led campaigns. Whilst the details of the actual campaigns have been de-identified in this example, we can see how understanding the impact of each individual offer on new sales volumes enabled the company to develop a deeper understanding of how price and media could each be used to influence sales volumes going forward.

© Driussi, Stevenson and Davis

Figure 4.1 – Impact of Price Led Campaigns on Relative Sales Volumes

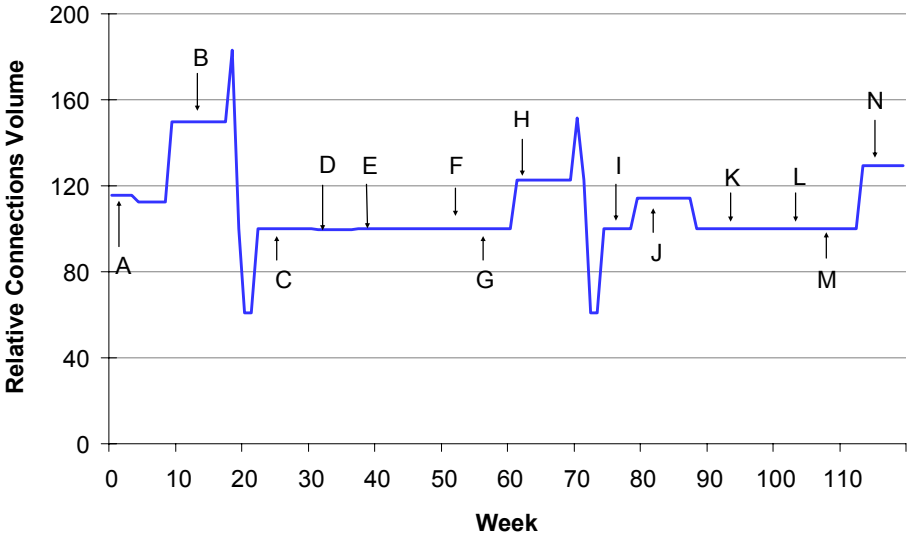

In industries where price setting depends on individual customer characteristics to a greater degree, price factors can be incorporated by looking at aggregate price statistics for a particular postcode or CCD. These could represent average price levels, or some measure of overall competitiveness, based on the demographics of the postcode in question.

Where modelling is done at an individual household level, it may be possible to improve on this methodology, due to the greater amount of detail available in this instance. In this case we can use the customer characteristics to 'guess' at a possible range of prices for use in the modelling process.

Assessing the quality of the response

Understanding how offered prices influence enquiries is only half the story. It is equally important to understand how the selected combination of media and offered price influences the quality of the response.

The transition to considering value, rather than just volume, can be achieved by giving each response a ranking, or value score. For example, this could be based on:

- Expected conversion rates of potential customers, taking into account competitiveness and/or acceptance criteria at a postcode/CCD/household level; or

© Driussi, Stevenson and Davis

- The expected profitability (or lifetime value) of a potential customer, calculated at a postcode/CCD/household level, incorporating all known demographics of the customer in question;

- The type of response elicited, where potential customers have the ability to respond in more than one way. For example, a provider of retirement solutions was looking to generate appointments with its financial planners, to help generate future sales. However, not all individuals who respond to the advertising are necessarily willing to go straight to an appointment, so the company developed a series of mail outs and seminars to provide additional information to its potential customers. Thus, new enquiries could result in an appointment, a seminar attendance or the sending of an information pack. Each of these initial contacts has a different probability of eventually converting to a sale, which can be used to apply a different value weighting in each case.

Rather than modelling pure enquiry volumes, we can now overlay the value measure for each response and assess the various media strategies considered relative to the expected total value generated.

This allows us to re-express the efficient frontier to show the trade-off between spend and expected value, as shown in Figure 4.2 below.

Figure 4.2 – Optimal Trade Off Between Media Spend and Expected Value

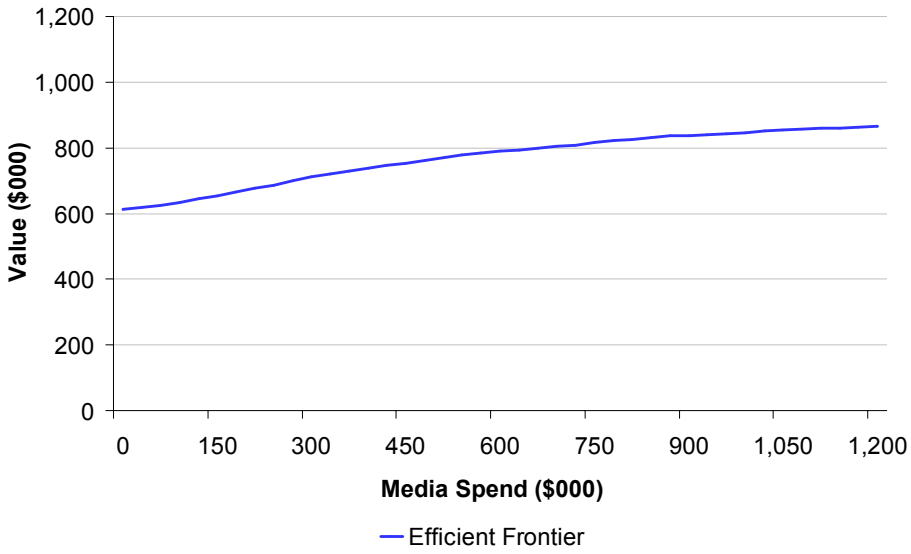

There are a number of benefits of overlaying value in this manner. By being able to make informed decisions about the spend/volume/value trade-off, a company can set marketing strategies that match current business objectives.

© Driussi, Stevenson and Davis

This may mean optimising either value or total enquiry volumes at different times, depending on whether the current strategic objective is focused more around growing market share or value harvesting.

We can also now directly compare the cost of each media strategy with the expected value generated from that strategy – enabling us to start considering metrics such as return on media investment. Consider Figure 4.3 below where we re-express the efficient frontier shown in Figure 4.2 to show the return on media investment, or the expected additional value generated as a proportion of the additional media spend needed to produce it.

Figure 4.3 – Return on Investment: Expected Value as a Proportion of Spend

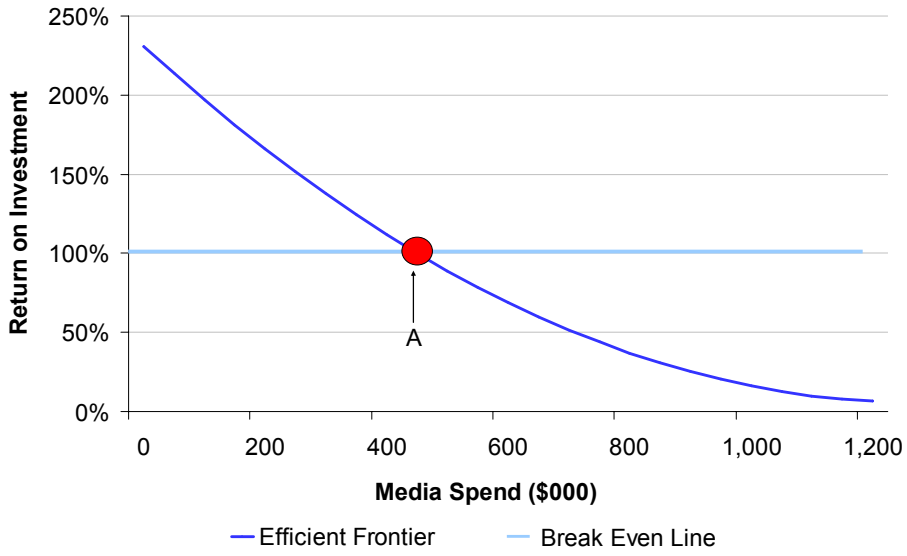

Any strategies to the right of the point where the break even line crosses the frontier are not generating sufficient additional value to justify expenditure in excess of that associated with strategy A. The set of media strategies that are considered should be limited to all practically achievable strategies that lie along the efficient frontier to the left of, and including, strategy A.

4.1. Case Study

We now return to the case study discussed at the end of each of the preceding chapters. So far, we have just discussed the generation of an optimal number of quotes for a given level of spend. However, the insurance company in question was interested not only in quote volumes, but also in actual sales results and the associated profitability of those sales.

© Driussi, Stevenson and Davis

Which regions have higher conversion rates and profitability and what impact does this have on the optimal media strategy?

As described in section 3.4, the outcome of the initial simulation process was an expected number of quotes in each postcode for each alternative media strategy.

In this case, it was possible to overlay a number of measures of value on each postcode to determine the expected value of each alternative strategy:

- Expected conversion rates by postcode – i.e. the likelihood that the quote converts to a sale by postcode. Conversion rates varied from anywhere between 10% for a postcode to 60% depending on the number of competitors in the state, the relative strength of the brand in that region and the competitiveness of the premium offered in that postcode;

- The expected profit associated with each quote, calculated at a postcode level. With a reasonable degree of cross subsidisation existing in the pricing by postcode across Australia, there were some postcodes (and broad regions) with very different levels of expected profitability.

This translation from pure quote volumes to expected profit allowed the insurance company to refine its evaluation of different media strategies, selecting the strategy that gave it the best result for its bottom line. Indeed the optimal media mix by geography changed quite significantly as a direct result of overlaying expected conversion rates and profitability – with certain regions generating high quote volumes but low conversion and expected profit.

© Driussi, Stevenson and Davis

5. Simultaneous Optimisation of Media and Price

The discussion in chapter 4 focused on incorporating a fixed price level for each customer. The next step is to consider varying the price level for each customer at the same time, effectively conducting simultaneous optimisation of media and price. There are a number of challenges associated with conducting a combined optimisation of this kind, not all of which have been solved at this stage. Therefore, this chapter discusses some initial ideas on how this problem might be addressed, but does not attempt to present a full solution to the issue.

In order to be able to conduct a simultaneous price and media optimisation, we need the ability to estimate the price that would be charged for each potential applicant. We also need to have sufficient information about all past enquiries to allow us to estimate the price elasticity, or price sensitivity of each enquiry.

These two pieces of information are then combined in the optimisation algorithm, enabling us to predict both the number of enquiries that we expect to generate with a given pricing and media strategy, and how likely these enquiries are to convert to a sale at the price specified. The optimisation algorithm then simultaneously varies the media spend and pricing (with price reductions being offset against lower media expenditure) to optimise the expected sales/profitability outcome overall.

The challenge in most cases arises because we do not have sufficient information to allow us to estimate the price that would be charged (and the associated price elasticity) for every initial enquiry. Even if prices are set at a relatively aggregated segment level, accurate estimation of price elasticity usually requires detailed information about the potential customer's characteristics.

Therefore, a combined optimisation is most easily conducted for a company that is able to do household level analysis. The greater richness of information that is available when modelling is conducted at this level allows the inclusion not only of demographic information, but also of information about the household's past interaction and existing relationship with the company (if any).

Extension of the approach to CCD level modelling becomes more feasible in industries where price setting has to remain at a broader segment level. However, the nature of CCD level modelling implies that the elasticity models can only be applied at quite an aggregated level, which will decrease the value allocated to price discounting, and hence overweight the allocation to media spend (rather than price discounting) in a standard optimisation algorithm.

© Driussi, Stevenson and Davis

While the approach described above gives a potential broad outline that could be applied to the combined optimisation problem, numerous technical issues still remain to be resolved. The potential benefits of determining an appropriate solution are enormous, so further research and investment in this area is warranted.

Further research in this area is currently underway, through continuous testing of new methods in collaboration with experts in other fields of strategy and marketing. The end result is very significant, building both efficiency and performance across the marketing value chain.

© Driussi, Stevenson and Davis