Software review

Business application of forecasting with a campaign management content

Received: 11th August, 2004

Shaun Doyle
is CEO of Cognitive Box, a marketing technology solution provider. He currently spends alternate weeks in the USA and UK helping clients to improve their marketing performance. Prior to setting up Cognitive Box he was VP Intelligent Marketing Solutions at SAS. In this role he worked with various parts of the SAS organisation to develop business-oriented solutions for marketing, in particular the SAS Marketing Automation (MA) Solution and SAS's industry-specific solutions for telco and retail banking. He was founder and chairman of Intrinsic, a campaign management vendor acquired by SAS in March 2001. He is well known in the marketing industry and regularly participates in leading industry conferences. He is also author of a number of papers on the application of database marketing that have been published in the UK and elsewhere.

Abstract This paper explores the use of basic forecasting techniques within a campaign management context. It describes some potential areas where forecasting could be applied and the basic techniques that could be used.

Shaun Doyle
Cognitive Box (Consulting) Ltd, Sams Barn, The Greenway, West Hendred, Wantage, Oxon, OX12 8RD, UK.

Tel: +44 (0)1235 821 353; Fax: +44 (0)1235 861 678; e-mail: shaun.doyle@cognitivebox.com

INTRODUCTION

The author has been involved in the development and deployment of campaign management technologies in a wide range of business sectors. The disappointingly low general quality of standard campaign performance reporting in campaign management applications is surprising, particularly how little advanced analytics have been integrated into this reporting process. If anything, some of the mainstream campaign management software vendors are winding down their development efforts for reporting programs in preference of integrating with Business Objects (www.businessobjects.com) or Cognos (www.cognos.com) reporting suites. This approach is understandable given the desire of most IT organisations to standardise on an enterprise-wide single reporting application, which makes selling individual modules too hard, but these reporting technologies do not have any of the key 'smarts' to deliver the advanced reporting requirements of campaign management. This is particularly true when it comes to forecasting.

This paper explores some of the business applications of forecasting in a campaign management context. It then discusses some of the basic techniques that enable campaign managers at least to start to understand what types of forecasting will be relevant to them.

Definition of forecasting

Forecasting can be defined as the estimation of the value of a variable (or set of variables) at some future point in time.

A forecasting exercise is usually carried out in order to provide an aid to decision making and in planning for the future. Typically, all such exercises work on the premise that if the future can be predicted, behaviour can be modified now in order to gain a better position at a later time.

BUSINESS APPLICATIONS OF FORECASTING

The following section describes a number of business applications of forecasting seen in organisations around the world. The list is by no means comprehensive but it does give a flavour of the type of requirements that exist in the market today.

The following business applications are covered:

— Predicting campaign performance;
— Predicting call centre resource requirements;
— Predicting customer service usage behaviour; and
— Predicting the impact of integrated marketing activities.

Predicting campaign performance

Usually the key requirement here is to predict the volume responses and/or sales of a product or service over time for a particular campaign. Historical response patterns are used to forecast the volume of responses and/or sales for a particular campaign.

In some organisations sales forecasting is done as part of the annual planning process and the resulting individual campaign plans are aggregated to produce an overall annual marketing plan for direct marketing activities. In other cases, the forecasting is done as part of the campaign monitoring process and, based on the initial campaign results, the forecast is used to predict the likely outcome at the end of the campaign period. This approach allows an organisation to take corrective action where the campaign is likely to under- or over-perform.

The key benefits of forecasting in this case are:

— The ability to refine plans to align with corporate goals.
— The ability to understand the dynamic of the business.
— The ability to understand the end performance of a campaign well ahead of completion.
— The ability to take corrective action where campaigns are under- or over-performing.

Predicting call centre resource requirements

In this case, companies normally wish to predict the volume of outbound calls to be made either as part of a campaign or as part of an annual planning process. Information about the availability of resources in the call centre is mapped against the likely call pattern for individual campaigns and the resulting outbound call volumes aggregated to produce a schedule of calls to be made over time. When combined with assumptions about hit rates (percentage of calls that connect with the customer), average call durations, available resource and associated skill sets this information can then be converted into a resource plan by the call centre operations team.

The other requirement in this area is a bit more complex. What the organisation is trying to do is to predict the number of inbound communications that occur as a result of an outbound campaign, then aggregate the results for all of the campaigns in a particular time period to produce an overall resource plan by channel.

Key issues that need to be addressed as part of the forecasting process include:

- The nature and volume of outbound communications.
- Customers' likely response channels.
- 'Lead leakage' in the channel as the enquiry is passed to different parties involved in the sale cycle.

In some cases this information is integrated with channel costs data and forms part of the return on investment (ROI) calculation for the campaign.

The key benefits of using forecasting to predict call centre resource requirements are:

- The ability to refine plans to align call centre resources for outbound activity.
- The ability to plan resources in inbound channels to manage the response from campaigns.
- The ability to understand channel dynamics and, where appropriate, to provide an incentive for customers to use a particular channel.

Predicting customer service usage behaviour

Normally, the key requirement here is to predict the volume or pattern of usage of a service over time at the individual customer level. This information is then used to predict customers' future usage of a service and to develop propositions to either reinforce this behaviour or change the pattern of usage, eg changing a mobile phone customer's price plan.

The key benefits of forecasting in this way are:

- The ability to predict customers' likely behaviour.
- The ability to monitor actual and forecasted behaviour.

In some cases, the forecasting is done at a summary level — eg segment or household; in others the forecasting is done at the customer level and then aggregated to segment or household level.

Predicting the impact of integrated marketing activities

Here, the key requirement is normally to predict the impact of the marketing mix on the performance of a campaign. In most organisations, direct marketing is often integrated with a range of other marketing activities (eg press, advertising, in store promotional activities, all of which have a direct or indirect impact on campaign performance).

One of the most commonly used techniques is market mix modelling (MMM). This technique applies statistical processes to determine:

- The factors that drive sales;
- The relevant importance of each of these factors;
- The ROI for various activities; and
- The optimal mix of spending in each of the activities.

The key steps in MMM are described in Figure 1.

The key benefits of forecasting in this way are:

- The ability to understand the impact of marketing mix on the performance of a campaign.
- The ability to understand how best to allocate marketing resources.
- The ability to simulate changes in the marketing plan.

The above is just a sample of the potential business applications of forecasting in marketing. The following

1. Build a model that analyses and predicts historical sales.
2. Test the predictive ability of the model on a hold-out sample.
3. Refit using all the data and predict the future.
4. Compare actual to forecast sales performance and determine incremental revenue.
5. Apply financial data and determine ROI.
6. Model the influence of individual factors.
7. Simulate the impact of different marketing plans.
8. Develop and deploy the optimal marketing plans.

Figure 1: Key steps in MMM

section explores the basic forecasting techniques available.

TYPES OF FORECASTING PROBLEMS

One way of classifying forecasting problems is to consider the timescale involved in the forecast, ie how far forward into the future we are trying to forecast. The usual categories are:

— Short term
— Medium term
— Long term

The nature of each category will vary according to the situation that is being studied. The categories are relative and will mean different things to different companies. For example, in a mail order organisation, forecasting demand for a product for more than one season may be seen as long term if the usual forecasting timeframe is daily or weekly.

Table 1 shows the timescales associated with business decisions.

Classification is important because different forecasting methods apply in each situation; eg a forecasting method that is appropriate for forecasting sales next month (a short-term forecast) would probably be an inappropriate method for forecasting sales in five years time (a long-term forecast). Note that here the use of numbers (data) to which quantitative techniques are applied typically varies from very high for short-term forecasting to very low for long-term forecasting when dealing with business situations.

BASIC FORECASTING METHODS

Forecasting methods can be classified into several different categories:

— **Qualitative methods:** there is no formal mathematical model, often because the data available is not thought to be representative of the future.
— **Regression methods:** an extension of linear regression where a variable is thought to be in a linear relationship with a number of other independent variables.
— **Multiple equation methods:** there are a number of dependent variables that interact with each other through a series of equations (as in economic models).
— **Time series methods:** a single variable changes with time, its future values are related in some way to its past values.

Qualitative methods

Methods of this type are primarily used in situations where there is judged to be no relevant past data on which a forecast

Table 1: Classification of forecasting decisions

Timescale	Type of decision	Examples
Short term Up to 6 months	Operational	Inventory control Production planning, distribution
Medium term 6 months–2 years	Tactical	Leasing of plant and equipment Employment changes
Long term Above 2 years	Strategic	Research and development Acquisitions and mergers Product changes

can be based. They typically concern long-term forecasting. One approach of this kind is the Delphi technique, which involves asking a body of experts to arrive at a consensus opinion as to what the future holds. Underlying the idea of using experts is the belief that their view of the future will be better than that of non-experts. This technique is often used in marketing, particular for forecasting sales of newly-launched products.

In a Delphi study, the experts are all consulted separately to avoid some of the bias that might result were they all brought together. The answers are assembled in the form of a distribution with comments attached and re-circulated to provide revised estimates. This process is repeated until a consensus view emerges.

Plainly, such a method has many deficiencies, but this may be the best method if the relevant data to support a quantitative approach are lacking.

Regression methods

In linear regression, data are fitted a straight line form to $Y = a + bX$. As a simple example, in real terms, Y is total annual sales, a is total advertising spend; b is the number of marketing communications and X is 9.3.

It is possible to extend the method to deal with more than one independent variable X. Suppose there are k independent variables $X_1, X_2, \ldots, X_k$, then one can fit the regression line $Y = a + b_1X_1 + b_2X_2 + \ldots + b_kX_k$. This extension to basic linear regression is known as multiple regression. Plainly, knowing the regression line enables us to forecast Y, given values for X.

Multiple equation methods

Methods of this type are frequently used in economic modelling (econometrics) where there are many dependent variables that interact with each other via a series of equations, the forms of which are given by economic theory.

Time series methods

Methods of this type are concerned with a variable that changes with time and which can be said to depend only upon the current time and the previous values that it took (ie not dependent on any other variables or external factors). If Yt is the value of the variable at time t, then:

$$Yt = f(Yt{-}1, Yt{-}2, \ldots, Y0, t)$$

ie the value of Y at time t is purely some function, f, of its previous values and time, no other variables/factors are of relevance. The purpose of time series analysis is to discover the nature of f and hence allow us to forecast values for Yt.

Time series methods are especially good for short-term forecasting where, within reason, the past behaviour of a particular variable is a good indicator of its future behaviour, at least in the short term.

A typical example is that of short-term demand forecasting. Note the difference between demand and sales — demand is what customers want, sales is what is actually sold; the two may be different.

The purpose of the analysis is to discern some relationship between the Yt values observed so far in order to forecast future Yt values.

The following are examples of time series methods.

Moving average

One very simple method for time series forecasting is to take a moving average (or a weighted moving average).

The moving average (mt) over the last L periods ending in period t is calculated by taking the average of the Y values for previous periods $t-(L+1)$, $t-(L+2)$, $t-(L+3)$, …, $t-1$, t.

Thus:

$$(mt = [(Yt-(L+1)) + (Yt-(L+2)) + (Yt-(L+3)) + \ldots + (Yt-1) + Yt]/L$$

The forecast for all periods beyond t is just mt (although one usually only forecasts for one period ahead, updating the moving average as the actual observation for that period becomes available).

Single exponential smoothing

One disadvantage of using moving averages for forecasting is that in calculating the average all the observations are given equal weight (namely $1/L$), whereas one would expect the more recent observations to be a better indicator of the future (and accordingly ought to be given greater weight). Also only recent observations are used, whereas perhaps all previous observations should be taken into account.

Exponential smoothing (or, more accurately, single exponential smoothing) gives greater weight to more recent observations and takes into account all previous observations.

If one defines a constant μ, where $0 \leq \mu \leq 1$, then the (single) exponentially smoothed moving average for period t (mt) is given by

$$m_t = \mu Yt + [\mu(1-\mu) \times (Yt-1)] + [\mu(1-\mu)^2 \times (Yt-2)] + [\mu(1-\mu)^3 \times (Yt-3)] + \ldots \quad (1)$$

Here, the exponentially smoothed moving average takes into account all of the previous observations. Compare this with the moving average above where only a few of the previous observations were taken into account. The above equation is difficult to use numerically but note that from (1),

$$mt = \mu Yt + (1-\mu) \times [(\mu Yt - 1) + (\mu(1-\mu) \times (Yt-2)) + (\mu(1-\mu)^2 \times (Yt-3)) + \ldots]$$

ie $mt = \mu Yt + (1-\mu) \times (mt-1)$

Hence the exponentially smoothed moving average for period t is a linear combination of the current value (Yt) and the previous exponentially smoothed moving average ($mt-1$). The constant μ is called the smoothing constant and the value of μ reflects the weight given to the current observation (Yt) in calculating the exponentially smoothed moving average mt for period t (which is the forecast for period $t+1$).

For example, if $\mu = 0.2$ then this indicates that 20 per cent of the weight in generating forecasts is assigned to the most recent observation and the remaining 80 per cent to previous observations.

Note that $mt = \mu Yt + (1 - \mu) \times (mt - 1)$ can also be written $mt = (mt - 1) - \mu \times [(mt - 1) - Yt]$ or current forecast = previous forecast − μ (error in previous forecast), so exponential smoothing can be viewed as a forecast continually updated in the light of the forecast error just made. Methods are available which enable the optimal value of the smoothing constant to be easily determined.

Autoregressive integrated moving average (ARIMA)

More advanced time series forecasting methods exist and are often based on autoregressive integrated moving average (ARIMA) models. Essentially, these assume that the time series has been generated by a probability process with future values related to past values, as well as to past forecast errors. To apply ARIMA models, the time series needs to be stationary. A stationary time series is one where the statistical properties — such as mean, variance and autocorrelation — are constant over time. If the initial time series is not stationary it may be that some function of the time series, eg taking the differences between successive values, is stationary.

In fitting an ARIMA model to time series data, the framework usually used is a Box-Jenkins approach. It does, however, have the disadvantage that whereas a number of time series techniques are fully automatic (in the sense that the forecaster has to exercise no judgment other than in choosing the technique to use) the Box-Jenkins technique requires the forecaster to make judgments. Consequently its use requires experience and 'expert judgment' on the part of the forecaster. Some forecasting packages do exist that make these 'expert choices' for you.

SUMMARY

It is surprising that even basic forecasting has been remiss in most mainstream campaign management applications. A change in focus away from delivering reporting as part of the campaign management solution means that this is unlikely to change in the near future. The use of generic reporting tools, such as Business Object or Cognos, is unlikely to address these issues as these rarely have the required statistical smarts embedded. In the short term it looks like those organisations looking to address these business requirements through forecasting will have to integrate 'best-of-breed' analytical applications — eg SAS (www.sas.com) or SPSS (www.spss.com) into their campaign management or reporting environments. Hopefully, in the future someone will look to add more analytical smarts to campaign management software and create an integrated solution that addresses a real business need.