research.
2007

MRS GOLDEN JUBILEE CONFERENCE
21-23 MARCH 2007
HILTON BRIGHTON METROPOLE

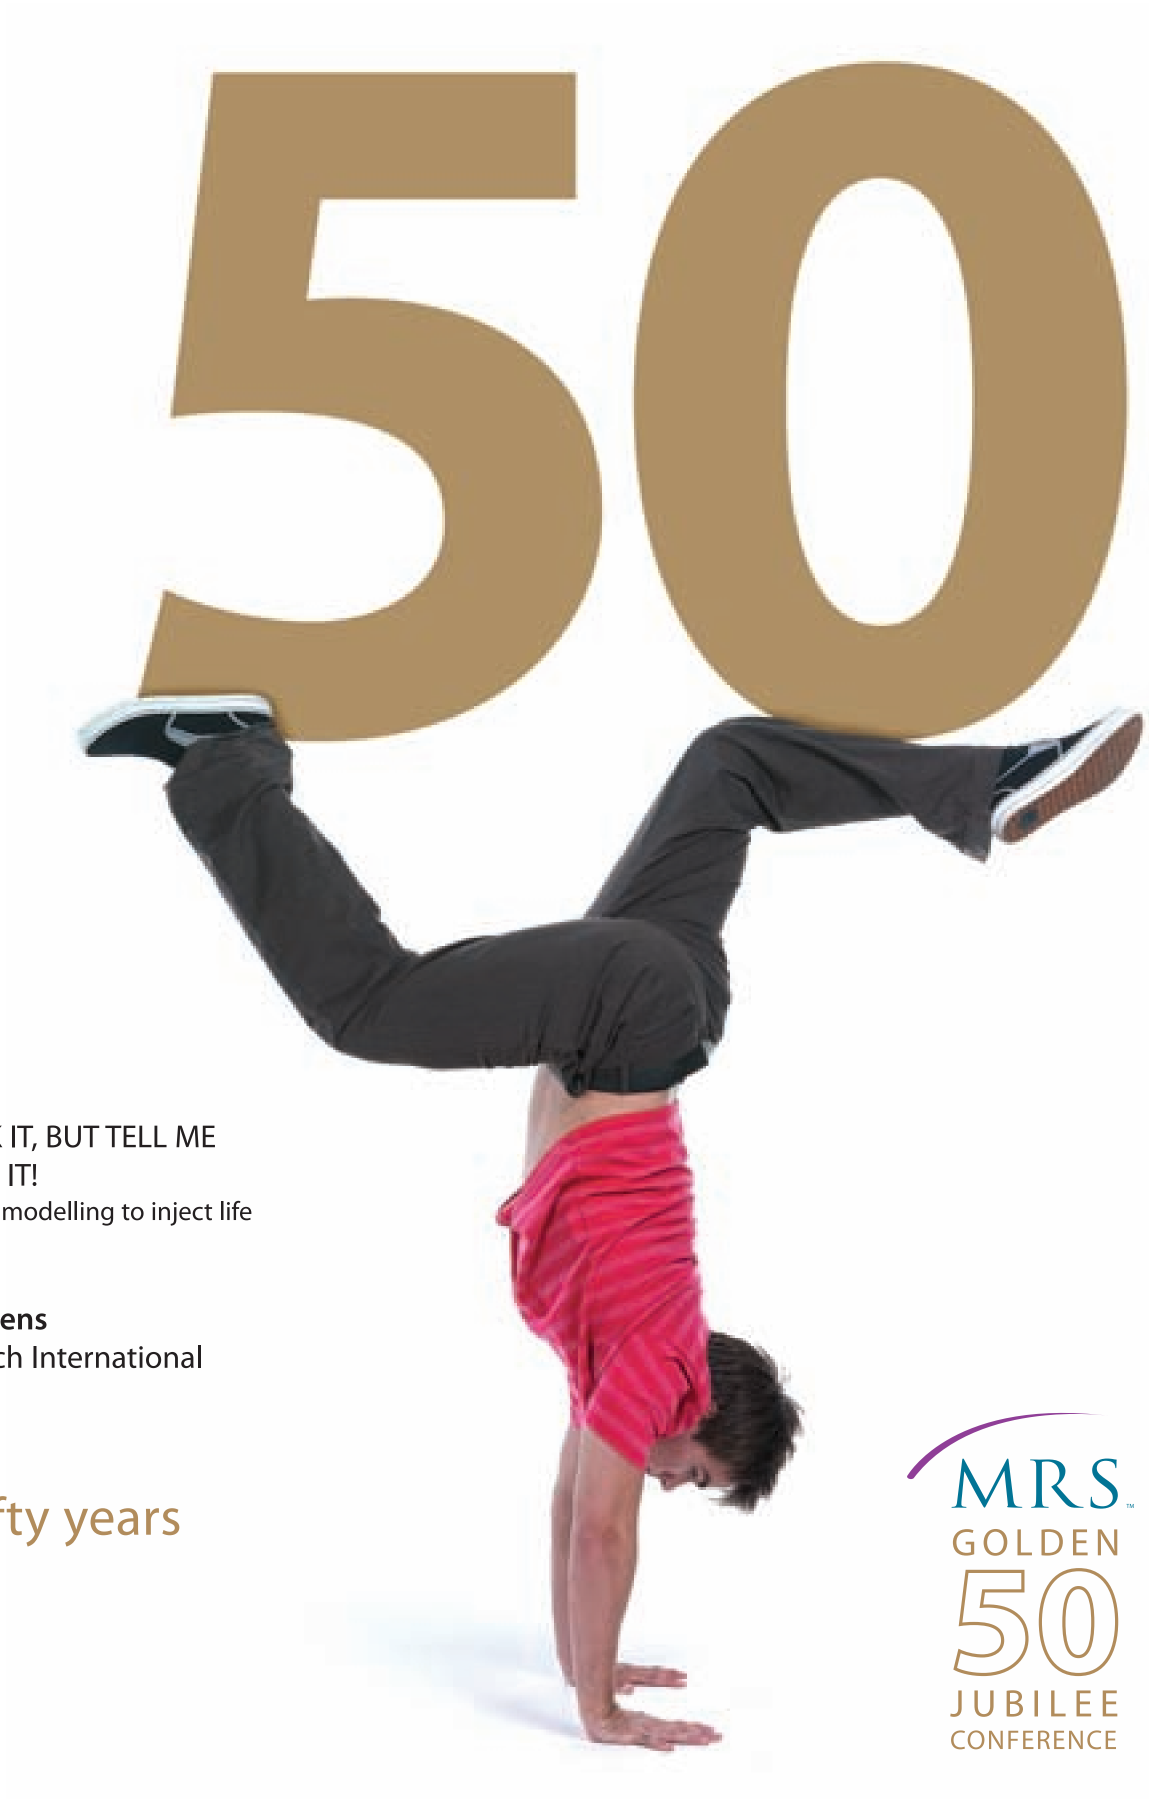

Paper 16

I KNOW WE TRACK IT, BUT TELL ME HOW TO IMPROVE IT!

Using neural network modelling to inject life into research data

Charles Adriaenssens
Consensus Research International

Martin Wootton
RS Consulting

the next fifty years

MRS
GOLDEN
50
JUBILEE
CONFERENCE

research. 2007

I know we track it, but tell me how to improve it!

Using neural network modelling to inject life into research data

Charles Adriaenssens Consensus Research International
Martin Wootton RS Consulting

INTRODUCTION

An ongoing evolution

The regular measurement of key performance indicators, such as customer satisfaction, is an essential part of evaluating how effective a company is performing. There are few companies who do not undertake some form of research to evaluate their performance from this perspective.

In our experience, gained by working with a wide range of companies across a range of industry sectors, we have observed patterns in the way that companies evolve their key tracking programmes. At the outset, a typical company will seek simply to understand the proportion of people who consider themselves satisfied with the company's products or services. Upon receiving this information, a company may then decide to monitor how these measures change over time; equally, they may decide to undertake further research to better understand why scores for certain performance areas are lower than they might have expected.

As these programmes develop and the volume of data collected grows in size, a company might apply more sophistication to their analyses. They might then decide to concoct an 'index' score, represent overall performance across a range of key performance indicators via one single number, or segment customers according to how valuable they are or how they score across the key indicators.

It is at this point that senior managers ask the killer question: ***"I know we track it, but tell me how to improve it!"***

From tracking to action plans

Many companies will typically decide that further progammes of research are needed to understand what improvements should be made, in which areas, and to what degree. We have worked with companies whose logical next step would be to conduct qualitative research among customers who give lower than average scores across one or more key performance indicators. While this approach is valid, it is often both time-consuming and expensive to put into place.

A quicker and more efficient initial approach is to mine the existing data using mathematical modelling techniques, such as regression. Many companies do precisely this in an attempt to distil the relationships between key performance indicators and other performance areas covered by the research programme, in order to understand which areas drive the key performance indicators most strongly. A few companies (albeit, an increasing number) take a further step forward and use a model to create a predictive simulator, which can be used to measure the effect of hypothetical changes to the input variables of a model on key performance indicators.

We have observed that there are flaws with many of the techniques that are commonly used, including those that allow for the construction of a predictive simulator. For instance:

- Many modelling techniques used in this context, including regression, force a linear (straight line) relationship through a data set. Frequently, the 'true' relationship is not linear, but curved.
- Other techniques are not able to take into account the effects of multiple changes across a number of variables – they are limited to looking at one variable's impact in isolation.

Indeed, many of the predictive simulators built around common mathematical modelling techniques have serious practical drawbacks:

- There is often a lack of transparency among some proprietary techniques offered by research companies – companies do not get to see or understand exactly how the model actually works, and how the predictions are derived (the 'black box' syndrome).
- Some require an expert who understands the mathematical details of the model to run the required simulations on behalf of the end-user, which adds further time and cost implications.

A potential solution to the problems of predictive simulators

In late 2005, RS Consulting Group (which includes both RS Consulting and Consensus Research International) partnered with a Danish company, SJP, who work closely with the Niels-Bohr Institute and the University of Copenhagen. They have written and patented software that applies neural network modelling algorithms to large data sets. Working together, we have been able to apply these principles to customer satisfaction and loyalty data from our clients, to client acclaim. We have subsequently become the sole software licensee in the US and Europe.

Via this, we have been able to add a great deal of value to clients' existing tracking data by using the software to model how key performance indicators are impacted by changes in a range of performance attributes. A key advantage of the software is that hypothetical 'what-if…?' scenarios can be run simply, quickly and transparently. We believe this is a significant advantage over other predictive simulator tools that we have seen.

Our paper discusses how we have applied the neural network modelling approach to provide actionable insight for our clients.

APPLYING NEURAL NETWORK MODELLING TO MARKET RESEARCH DATA

Advantages of neural network modelling

Neural network-based models are often likened to the way the human brain works: essentially, they work by creating a network structure connecting all possible points, (typically called 'nodes') together. The initial model is then restructured and simplified over a number of iterations based on the known relationships between node pairs. The model systematically 'learns' from this process and restructures and simplifies further, creating a more mathematically accurate model with each subsequent iteration. The result is a robust model that takes into account the interaction between each and every data point, including outlying data points that may normally be ignored. Thus, it enables non-linear (curved) relationships to be modelled, because all data points are used to calculate the relationship weighting.

As demonstrated by Gronholdt and Martensen (2005) and Lee et al (2005), techniques enabling non-linear relationships to be described - such as neural network modelling - are a truer fit than purely linear techniques (which include regression, structured equation modelling and partial least squares techniques).

While regression, for instance, can only model a relationship across pairs of variables, a neural network model can take into account all input variables, creating the 'network' structure that gives rise to its name. Such multi-dimensionality is the defining characteristic of these models. The benefit of this is that the models generated exhibit an extremely high level of accuracy. This a key benefit to those wishing to use tracking data as the basis of a model as it is this accuracy that enables us to model the impact on one or several key performance indicators of making changes to several variables simultaneously, to a very high degree of confidence.

Building blocks

The software we use to create neural network-based models relies on accurate input in the form of scalar data. We can use both numeric and semantic scales as our input - or a combination of any of these, including scales with any number of points on it.

They key is that the points on the scale are consistently spaced and are sequential (i.e. Likert scales). Numeric scales fit this criteria perfectly, though equally gradated semantic scales (e.g very likely, quite likely, quite unlikely, very unlikely) can be used.

One or more variables in the data set is pre-selected as being the key performance indicator(s) or dependent variables, while the remaining variables incorporated in the modelling process are input factors that might influence the key indicator(s) in some way.

A typical example is the data from a customer satisfaction survey, where a questionnaire may have asked customers to rate their satisfaction across a wide variety of individual service areas, as well as give their overall rating of satisfaction, likelihood of purchasing the product/service again, and likelihood to recommend. In this instance, we might decide to choose overall satisfaction and repurchase likelihood as our key performance indicators. The remaining scale questions, which might have asked about diverse areas as product quality, friendliness of staff, opening times or value for money would then form the variables to be included in the model.

The neural network approach requires that a robust sample size is generated in order to produce output that offers a high degree of reliability. This is because part of the modelling process involves a random tranche of the data (typically 10%) being set aside and used to measure the effectiveness at each iteration of model creation. Therefore, sample sizes lower than 250 sometimes do not yield a model that has a sufficiently high degree of accuracy in order to be confident in the results of scenario modelling.

The number of variables included also impacts the degree of accuracy of the final model, but to a much lesser extent. Assuming a sample size of n=1000, even as many as 50 separate variables used as input would still yield a model that has a high degree of statistical accuracy.

Our approach

In the last twelve months, RS Consulting and Consensus Research International have applied the modelling software to a wide variety of clients' customer data sets. This has enabled us to add real value to data sets that would otherwise have been used only for tracking and basic analysis and make valuable contributions to senior managers' strategic decision-making.

Our approach offers benefits for both the client-side research team and the senior management within client organisations. At the research team level, the key benefit is cost-effectiveness. As most existing scale question sets are suitable as input to the modelling process, no new data collection is necessary. The only costs involved are the processing of the data and the analysis and reporting of the results. Another significant benefit is the ease and speed with which the associated software can be used. This software incorporates a Java-based reader programme, designed to allow end-users to interrogate the model quickly and intuitively.

Quantifying driver strength

Upon construction of the model, the reader can be used to analyse the results of the modelling process. In the first instance, the statistical reliability of the model can be viewed. This expresses the proportion of data correctly predicted when, during the building of the model, the model's output is compared to that of the initial data set aside. Typically models predict more than 90% of data points correctly.

Assuming the user is satisfied with the reliability of the model, they can then look at the extent to which each individual input attribute impacts on each key performance indicator. The rank order and magnitude of driver strength is shown graphically as a stacked bar chart based on the relative impact of a one-point improvement to the score given by each respondent for each input variable on the key performance indicator. An example of this can be seen in figure 1:

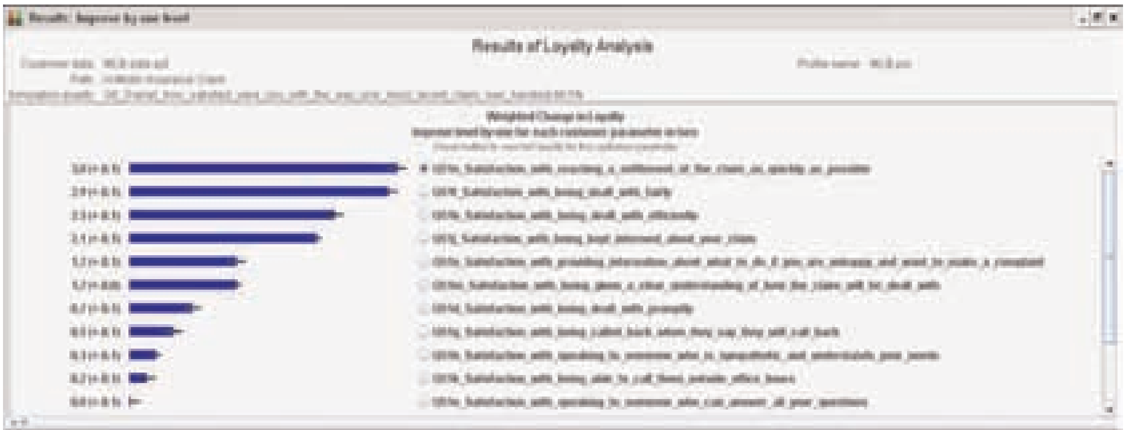

Figure 1: impact on overall satisfaction of a one-step improvement in each input attribute. The rank order and relative driver strength can be seen at a glance.

These data are aggregated findings from Consensus Research International's Motor Claims Benchmarking Survey, which is a syndicated research project run for a number of the main insurers. It shows the relative importance of a range of attributes on motor insurance claimants' rating of their service experiences while making a claim. The driver strength data can be directly compared across attributes to identify those drivers that would have the greatest impact on overall satisfaction - so, for example, reaching the settlement of a claim as quickly as possible would have almost three times the impact than providing information about what to do if claimants want to make a complaint. This output is the starting point for making decisions about which service changes to effect.

Predictive modelling

Aside from quantifying the relative strengths of attributes as drivers of key performance indicators, 'what-if?' scenario modelling can be undertaken with the resultant data set. For most clients, it is this ability that is particularly attractive. The software allows the user to specify hypothetical changes in the level of performance on each input variable, in order to establish how these changes would impact on each key performance indicator. The graphical user interface of the simulator makes this process very straightforward, and produces a visual output that is easy to understand, segment and quantify.

Furthermore, a consistent index calculation enables the effect of a range of different models to be compared directly. This index is simply a representation of the overall distribution of scores for the selected key indicator, and is calculated both before (the current situation) and after improvements are effected. Via this simulator, clients are able to run and interpret simulations themselves, without having to rely on us each time they want to test a hypothesis.

We can illustrate this using a very simple example. One might decide to see what the impact on overall satisfaction would be if you increase customers' scores on one input attribute (for instance, reaching a settlement of a claim as quickly as possible). We may choose to hypothetically increase every customer's rating of this attribute by one step – what would happen if those scoring this attribute 4/10 were to score it 5/10, those scoring this attribute 5/10 were to score it 6/10 and so on. The results of the simulation would describe the frequency distribution of our key indicator, in this case overall satisfaction, both before the hypothetical change and after (using the predictions of the model). By running the simulation, the change in index score can be seen at a glance. This reflects the relative impact of the changes and is illustrated in figure 2:

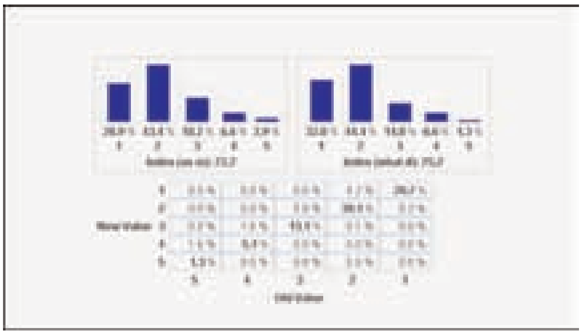

Figure 2: Example of mapping of scores for the key performance indicator before and after a simulated change in an input attribute.

In this example, the key performance indicator (satisfaction) is measured on a 5-point scale, where 1 is the most positive score, and 5 the least positive. We can see that as a result of the modelled change, 4.2% of customers are predicted to rate satisfaction as a 1 rather than 2. Overall, the satisfaction key performance indicator index increases from 72.2 to 75.2.

There is also the ability to look at the compound effect of such hypothetical changes on two key performance indicators simultaneously. Building on the example cited above, one might decide to look at the predicted impact of the same change on overall satisfaction and likelihood to remain a customer simultaneously. By segmenting customers according to their initial scores for these two key performance indicators, this output can be much simplified and made easier to interpret. A before and after visualisation of the satisfaction vs recommendation 'landscape' illustrates how customers' scores change as a result of a hypothetical change. This is summarised in figure 3, where we have classified each customers to one of nine segments according to how they rate satisfaction and loyalty.

This 3x3 matrix is derived by transforming customers' rating of the two key measures (typically given on a 5-, 7- or 10-point scale in the research) into a smaller number of discrete categories. The reason that the data is tranfsormed in this way is merely to provide an easy to understand graphical representation of customers' ratings - the customer 'landscape'.

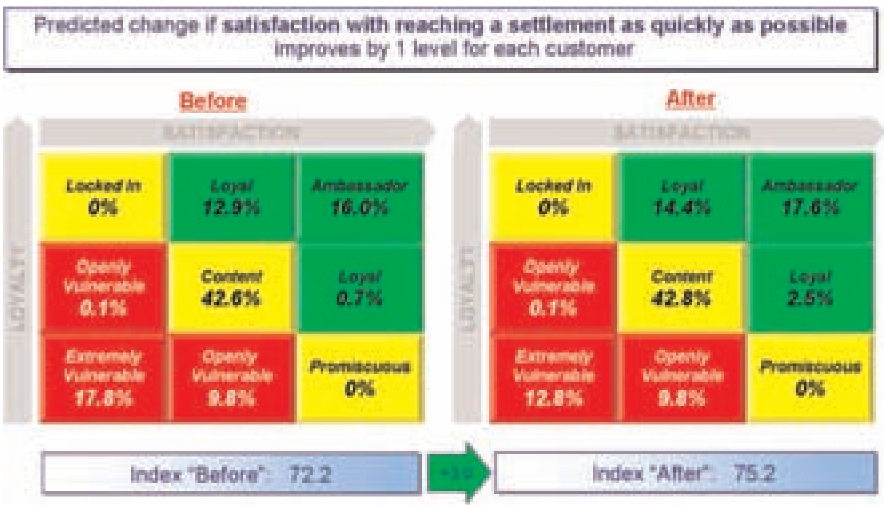

Figure 3: We can model the impact on both satisfaction and loyalty KPIs of improving one or more attributes.

Thus, the changes in the size of each segment as a result of making hypothetical change can be displayed clearly:

In this instance, the proportion of 'Ambassador' customers has increased 1.6% and the proportion of 'Extremely Vulnerable' has decreased by 5%.

The rationale underpinning neural network modelling is based on the need for any business to minimise the number of its customers that sit in the red boxes (representing its vulnerable customers) in the customer landscape, and maximise the number that sit in the green boxes (which represent its loyal/ambassador customers).

For senior management, the ability to model reliably how to stimulate future performance is a vital tool in formulating and validating wider strategies for improving satisfaction, share of wallet, loyalty and many other important company health measures. The predictive ability of the model allows enables us to divine the optimal, most cost-effective strategy for reaching a target performance level for a key indicator.

For example, a company may decide that their target (based on the customer landscape) is to ensure that less than 5% of customers are extremely vulnerable. Our solution would allow them to test any number of changes to one or several input attributes, in order to draw up a list of possible changes that would result in reaching this 5% target. This would enable senior managers to choose the most realistic and cost-effective menu of changes to be taken forward as part of their overall strategy for reaching the target.

The next part of this paper outlines two recent, real-life applications of our solution, showing how specific client organisations have benefited directly from this new approach to analysing their existing customer data.

CASE STUDY 1: IDENTIFYING PRIORITIES FOR SERVICE IMPROVEMENT ACROSS MULTIPLE CHANNELS OF INTERACTION

Channel-specific models

Our first case study comes from the financial services sector, and provides a template for measuring service performance across a range of distribution channels.

Like the majority of major players in this sector, our client interacts with its customers via the full range of available channels: branch, post, telephone and internet. Moreover, its customers use these channels at a number of different stages in the customer 'journey', encompassing: decision making, application, arrangement, post-sales servicing and maturity/closure.

Most customer experience tracking studies have had to deal with this interplay of channel of interaction and event during their evolution and two main types of approach have been adopted:

1. Event-driven sampling – respondents are interviewed soon after they have interacted with the brand on each key event (irrespective of channel of interaction); and

2. Channel-based sampling – respondents are interviewed soon after they have interacted via a particular channel (for one of a broad range of events).

The ideal is to overlay channel with event, although this provides a sampling structure that is typically too complex and costly to consider. Instead, this type of intelligence is commonly generated via straightforward cross-analysis of the resultant data.

This case study follows the channel-based approach, and it involves interviewing customers within a month of them interacting with the brand. Each respondent is asked to rate a broad range of generic attributes relating to service, as well as a number that are specific to the channel via which they recently interacted.

Modelling 'what-if?' scenarios

Running a neural network model based only on those respondents who have interacted via each individual channel enabled our client to identify those attributes for which making improvements would have the greatest effect on their rating of service experience. By establishing the drivers that impacted on each channel individually, we were able to identify the set of improvements – which varied by channel – that would in combination have the greatest overall impact on its service rating.

The key output that enabled them to identify this was a diagram that compared the relative increase in satisfaction that would result by improving each attribute by one level (as shown in fig 1) against customers' current rating of each attribute. Drawing in lines on the axis to highlight those attributes that would bring about relatively large improvements in satisfaction, as well those where the client is currently performing above average, enabled us to categorise each attribute into one of the following groupings: immediate priority, less critical priority, key strength and secondary strength. Via this, and further discussions, the client was able to develop an action plan both for the short- and the longer- term; prioritising initially the 'quick wins' but also identifying longer-term targets.

Figure 4: Template grid for plotting individual satisfaction mean scores against the modelled impact of making service improvements which enables priorities to be identified

As this solution was a relatively sophisticated approach compared to the pre-existing analysis undertaken by our client, one of the barriers we faced was the ability to explain the key outputs to internal stakeholders, as well as how they should interpret these and translate them into action plans. This barrier necessitated a very 'hands on' approach to data dissemination, in which the client was effectively guided through the process of modelling and data interpretation.

In order to facilitate this, we ran a workshop attended by all key staff, at which we explained the implications of the models generated to everyone responsible for the delivery of service performance. This information enabled our client to subsequently identify a set of achievable improvements that would have the greatest impact of service rating for each individual channel, as well as understand the wider picture of similarities and differences between channels i.e gain a holistic view of the overall customer experience.

Benefits for the client organisation

By explaining the priorities for improvement grid, we were able to identify to the client those aspects of service that would increase satisfaction for all customers, and those elements which would only have an impact on individual channels.

By discussions between researchers, the internal research team and the client channel managers, the client was able to develop an action plan which included both immediate and longer-term priorities.

Finally, and perhaps most importantly, the client was able to gauge an understanding of how each customer interaction (event) dove-tailed with wider customer experience.

CASE STUDY 2: HELPING A LEADING PRINTER MANUFACTURER TO FORMULATE A STRATEGY TO MAXIMISE LOYALTY

Setting the scene

Our client, a leading manufacturer of professional printing solutions, had been running a monthly satisfaction and loyalty survey among customers of its high-end solutions. They had amassed a database of around 12,000 surveys over a 2-year period. They were able to report and track satisfaction metrics to a high degree of statistical confidence, and had noted that their satisfaction and loyalty scores were not improving markedly, despite concentrating marketing and account management efforts on the areas of service receiving the lowest mean satisfaction scores.

Having undertaken correlation analysis and regression analysis on the data previously, and having found that the results of this did not differentiate strongly between possible drivers of satisfaction and loyalty, RS Consulting was commissioned to apply our neural network-based analysis to the data.

Using the 6 most recent months' worth of data, we built a neural network model designed to look at how two key performance indicators – overall satisfaction and likelihood of repurchasing – were influenced by 15 other performance attributes:

- Satisfaction with overall support from the sales representative
- Ease of doing business with the company
- Equipment reliability
- Servicing
- Speed of resolving software incidents
- Telephone support for software incidents
- Handling of queries and complaints
- Value for money of equipment
- Ability to fix faults at first attempt
- Offers solutions to fit your needs
- Billing and invoicing
- Ordering process
- Frequency of face to face contact
- Engineer response time
- Keeping you informed
- Customer training
- Quality of software consultants

Initial output

Compared to the results of previous analysis conducted by our client, the output from the neural network model was much clearer to interpret. It immediately showed that the areas of performance that drive satisfaction were distinct from those that drive loyalty. This made logical sense, as the lowest performing area in terms of mean scores – 'billing and invoicing' and 'customer training' – had little impact on satisfaction, which tended to explain why the effort put into rectifying the poor performance in this area was not being translated into increased overall satisfaction or loyalty. Furthermore, we see that 'value for money of the equipment' is a key factor that drives propensity to repurchase, but that it is only a minor driver of overall satisfaction.

The relatively high driver strengths for 'Ease of doing business with the company' was a surprise to the client, however. Subsequent qualitative research focusing on this area demonstrated that the complexity of the client's contracts with customers made it difficult for their customers to deal with them. This affected both satisfaction scores and the propensity of the customers to repurchase, since repurchasing meant having to go through the difficult process of contract writing.

In Figure 5, we have reproduced the resulting driver strengths for both overall satisfaction and repurchase likelihood.

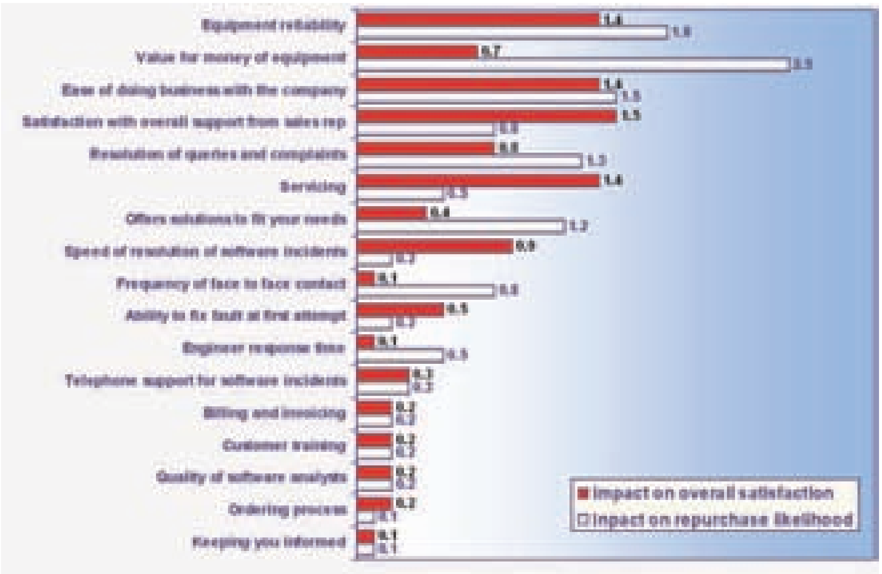

Figure 5: Relative impact of improving each perofmance attribute by one step (e.g. if scores of 6/10 become 7/10 etc) on overall satisfaction and repurchase likelihood

These findings may seem logical in retrospect, but the prior analysis conducted by the client did not give them such a clear message.

Our client then decided that the focus of their remedial action should change to ensure that the most sensitive areas in terms of influencing satisfaction and loyalty are at the centre of any future strategy. With the ability to build on this output and create a predictive model, we were able to simulate the effects of such a change in strategy.

Modelling 'what-if?' scenarios

Knowing that in reality, only a handful of changes could be made, we simulated the effect on increasing satisfaction with only 2 or 3 attributes at a time. Furthermore, making the assumption that action to address performance across an attribute would only stimulate a small rise in attribute rating, we only ran simulations based on a one-point change to the scoring of an attribute (i.e. what if those scoring an attribute 4/10 were to score it one point higher – 5/10).

Simulating permutations of 2 or 3 simultaneous changes of this magnitude enabled us to list a selection of 'packages' of hypothetical changes that the model predicted would have the biggest effect on either, satisfaction, loyalty or both of these. These were then ranked in order of their overall impact. Using the segmentation technique discussed earlier in this document, where we assign customers into one of nine segments according to their overall satisfaction and loyalty scores, we then simulated how the size of each customer segment changes as a result of applying the change 'packages' to the model.

The following scenario was deemed to offer the greatest increase in satisfaction and loyalty:

- '*Servicing*' increased by 1 level for all customers
- '*Ease of doing business with the company*' increased by 1 level for all customers
- '*Handling of queries and complaints*' increased by 1 level for all customers

The predicted result of these changes is a 4.7% decrease in customers categorised as 'potentially lost' (low satisfaction and loyalty), and a 6.7% increase in the number of "ambassador" customers (high satisfaction and loyalty).

This is shown in figure 6:

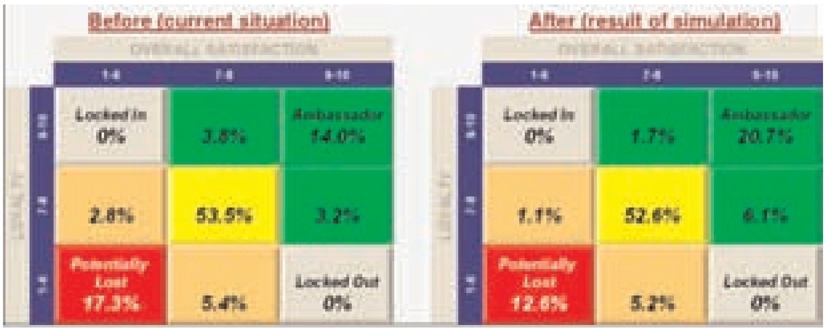

Figure 6: Predicted change in segment size as a result of improvements to 'servicing', 'ease of doing business' and 'handling of queries and complaints'

Benefits for the client organisation

The client was able to use these predictions as the basis of a new strategy for maximising customer satisfaction and loyalty. Targets were set for individual country organisations based on the predictions of the model. Whilte it is too early to tell whether or not the predictions have been borne out by appropriate increases in the latest scores for the key performance indicators, there is a distinct recent upward trend in these scores, steeper than measured prior to the change in strategy.

The final part of this paper will now look at how the modelling techniques described hitherto might be used in a wider context.

WIDER APPLICATIONS

Sub-group modelling

One of the key advantages of the neural network approach is the ability to apply the process consistently to sub-groups of the overall sample quickly and efficiently, thereby enabling companies to embed any key characteristics of their customer base within the resultant output.

Our first case study showed how building different models for each channel of interaction enabled a financial organisation to implement a range of service improvements which differed by channel in order to achieve the maximum possible overall return (measured by increased satisfaction). The same principle could, of course, be applied to any segmentations used by a company to enhance understanding of its customer base: obvious examples include generating unique models for any behavioural, attitudinal or broader lifestyle segment found on its customer database so that differing service levels can be applied to segments to suit their needs, although more tactical characteristics such as customers' current/future value or purchase propensity could be employed to generate a greater impact on profitability.

A more sophisticated approach may well necessitate bringing more data into the modelling process than merely the primary research. This fusion of research, customer data, third party data and sophisticated analytics is very much the growth area of marketing services and we expect much of our future development to be in this direction.

Return on Investment (ROI) models

As discussed earlier, client organisations can use the simulation tool to work out programmes of changes and improvements that would have the most positive impact on key performance indicators. The magnitude of the impact of each hypothetical change can be quantified in terms of its increase on the index score, or in terms of the percentage of customers who would meet a certain threshold of scoring across one or several key performance indicators. The optimal change programme would be one that offers the greatest increase in index score or proportion of satisfied customers, for instance. However, the optimal change may not offer the company the best value – it may prove to be very expensive or time-consuming to implement, or both.

To determine which changes offer the best return on investment, we can apply a relative 'cost' to each performance attribute. This can then be used as an inverse weighting mechanism. Multiplying the magnitude of the impact of each change by the 'cost' weight produces a tempered value, reflecting the true value of a change.

Using the data presented in Case Study 2, we can illustrate how such a calculation works in practice, and how it can be used to provide a simple view of which changes offer the best ROI to the company.

Let us compare three separate performance attributes:

- Equipment reliability
- Ability to fix fault at the first attempt
- Billing and invoicing

Each individual performance attribute is assigned a notional 'cost' which represents the relative monetary and resource effort required to increase performance in that attribute by one step in the scale. Note that these relative costs are determined by the client organisation in an arbitrary manner – it is not necessary to know the exact cost of any one change; rather, it is the relative cost *difference* between performance attributes that is key.

Continuing our example, we might determine that the relative cost differences between the three performance attributes are as follows:

- Relative cost of improving *reliability* = 3.0
- Relative cost of improving *ability to fix fault at first attempt* = 1.0
- Relative cost of improving *billing and invoicing* = 0.5

In other words, it would require three times as much investment to effectuate an improvement in reliability than it would to improve first time fault-fixing to the same degree. Equally, it would require six times more investment to improve reliability than to improve billing and invoicing by the same magnitude.

Next, we invert these cost values, so that the most expensive attributes had the lowest values.

- Relative cost of improving *reliability* = 0.33
- Relative cost of improving *ability to fix fault at first attempt* = 1.0
- Relative cost of improving *billing and invoicing* = 2.0

Let us now remind ourselves of the relative impact scores each performance attribute has on overall satisfaction:

- Index change by improving all respondents' scores for *reliability* by one step = 1.4
- Index change by improving *ability to fix fault at first attempt* by one step = 0.5
- Index change by improving *billing and invoicing* by one step = 0.2

Multiplying the impact scores by the relative cost weighting provides our ROI measure, where the higher the score, the more value offered by the change :

- Value of improving *reliability* = 1.4 * 0.33 **= 0.462**
- Value of improving *ability to fix fault at first attempt* = 0.5 * 1.0 **= 0.50**
- Value of improving *billing and invoicing* = 0.2 * 2.0 **= 0.40**

Assuming the goal is to increase overall satisfaction in the most efficient manner, the best ROI is to be gained from improving the ability to fix faults at the first attempt. Improving reliability may have a greater impact, but the high relative cost of doing so outweighs the benefits. Likewise, improving billing and invoicing would be relatively cheap to do, but doing so would yield little in terms of increased satisfaction scores.

This demonstrates that the results of the simulation tool can be easily translated into truly actionable recommendations, which offer clear guidance on the most efficient and effective way of increasing performance across key performance indicators such as satisfaction, advocacy and loyalty.

Furthermore, we believe that this application represents an extremely attractive and cost-effective ROI model, since it can be implemented using clients' existing survey data and is transparent in its workings.

The bigger picture: creating a consistent reporting mechanism to facilitate the harmonisation of key data across multiple audiences

The original rationale for the development of the neural network algorithm was to provide a means for improving employee satisfaction. Its application to customer satisfaction surveys therefore - particularly in the consumer sphere, where studies typically include a large number of respondents – was a natural one. The origins of the technique, however, mean that it also provides a ready-made 'bridge' that can be used to harmonise the findings of employee and customer satisfaction surveys via a consistent output format (which can be adopted even if data is captured using a different questionnaire and rating scale).

Providing a consistent approach to generating output across complementary research studies is another key trend in research, and neural network modelling is well placed to provide the direct comparisons that clients are seeking. In fact, harmonised employee and customer satisfaction models are only part of the picture, as the need to align senior management decision making with employee satisfaction in order to provide the required service to customers is ultimately driven by the need to deliver the profit demanded by shareholders.

This is the food chain that underpins our entire corporate culture, and having a consistent means of expressing research data across this chain is a key way of enabling companies to ensure that the objectives and motivations of all four audiences are aligned. In the wider context, we feel that employing consistent modelling techniques is one of the tools that can be used to facilitate this alignment.

Furthermore, in a corporate environment when economies of scale as well as developments in marketing and related activities require the management of multiple brands across multiple channels, this consistency is imperative as it enables direct comparisons to be made between product areas, between a range of (sub- and parent-) brands, and across multiple distribution channels.

CONCLUSION

We firmly believe that neural network modelling represents a forward progression in the evolution of how companies use their satisfaction data. Since this type of approach has been proven to be more accurate than linear modelling, we forecast that awareness and confidence in these techniques will grow, and that more and more organisations will strive to pro-actively model optimal strategies for improving key performance indicators - and from there develop an approach to maximising profitability.

Meanwhile, the need to provide consistent performance metrics across multiple channels of interaction while managing a portfolio of brands will demand an approach that is flexible enough to be applied to a range of data structures, while at the same time generating feedback for each individually. By using a neural network approach in a logical and consistent manner to generate actionable output, research organisations will be able to ensure that the ever more demanding insight needs of their clients can be satisfied.

References

Gronholdt, Lars and Anne Martensen (2005), Analysing Customer Satisfaction Data: A comparison of regression and artificial neural networks, International Journal of Market Research, 45, pp.121-130.

Lee, Carl, Tim Rey, James Mentele and Michael Garver (2005), Structured Neural Network Techniques for Modelling Loyalty and Profitability, Proceedings of SUGI 30 Conference, Pennsylvania, PA. USA. Available at: www2.sas.com/proceedings/sugi30/toc.html.