Institution: The University of Edinburgh/Heriot-Watt University (Maxwell Institute)

Unit of Assessment: B10, Mathematical Sciences

Title of case study: Uplift modelling for improved customer targeting

1. Summary of the impact

Research at the Maxwell Institute led by Radcliffe from 1996 onwards has developed new statistical models of the response of customers to targeted marketing. Traditional customer targeting misallocates resources by failing to estimate the change in the probability of customer behaviour that results from a given marketing action. This results in three kinds of waste: treating customers for whom intervention is ineffective, failing to treat customers for whom it would be effective, and treating customers for whom the intervention is counterproductive. The new models, known as uplift models, predict the change in behaviour, allowing lower target volumes, larger changes in customer behaviour, and suppressing counterproductive interventions. Uplift modelling has been commercialised in the form of software and consulting services from 2000: it is the core of the software Portrait Uplift sold by Pitney Bowes since 2010. The research has therefore had a major economic impact on Pitney Bowes and earlier companies selling uplift software and services, and on their customers which include US Bank and phone operators T-Mobile Austria and Telenor.

2. Underpinning research

Background. Traditional approaches to targeted marketing have applied statistical and machine-learning methods in a rather simplistic manner, resulting in suboptimal (and sometimes counter-productive) performance. The basis for most targeted marketing is a predictive model which is fitted to data on a population that has been subject to a marketing intervention (a 'treatment'). The model attempts to classify population members as either responders or non-responders, or in some cases to predict the size of response. Typical responses are purchases, renewals or (in the negative) account closures. Typical models of this form consider the output O (e.g. purchase/non-purchase) as a binary event and given treatment T and customer covariates x, typically combining geodemographic information with behavioural characteristics, to model the probability $\Pr(O \mid T; x)$. Similarly, for continuous-valued response S, the conditional expectation $E(S \mid T; x)$ is modelled. Traditional targeting methods usually select customers with high values of $\Pr(O \mid T; x)$ or expected response size $E(S \mid T; x)$, where the threshold may be simply determined by volume (e.g. target the top 30%) or may be selected to maximise the expected return on investment.

Uplift modelling. In contrast, uplift modelling focuses on the incremental impact of marketing by modelling quantities that reflect *change* in behaviour, such as $P(O \mid T; x) - P(O \mid T'; x)$ or $E(S \mid T; x) - E(S \mid T'; x)$ where T' denotes non treatment. Although it had been recognised best practice to maintain a control group to allow *assessment* of the incremental impact of a marketing initiative, prior to uplift modelling, we are aware of no attempts to target on the basis of modelled incremental impact.

The research programme led by Radcliffe from 1996 onwards applied a range of statistical approaches including generalised linear models and generalised additive models along with approaches from machine learning such as decision trees (e.g. CART, CHAID or C5). Key developments included identifying and correcting the traditional mis-formulation of the targeted marketing problem and developing a suite of increasingly sophisticated methods for building uplift models, which directly model the *change* in behaviour exhibited by individuals in the treatment and control groups. The core of the current method, the *significance-based uplift tree* [1-3], for the case of a binary outcome, models the probability of conversion as $p_{ij} = \mu + \alpha_i + \beta_j + \gamma_{ij}$ where α quantifies the effect of the treatment, β quantifies the effect of the split and γ is the interaction term. This is solved with a regression using as a split criterion (for a greedy binary decision tree) the square of a *t*-statistic for the significance of the γ_{TR} parameter, corresponding (without loss of generality) to the strength of the interaction between the treatment (T[reated]) and the split side (R[ight]).

As well as constructing models, the team has developed measures of performance of uplift models; in particular, the qini measure [4] is a rank-based statistic formed by generalizing of the widely used gini coefficient to the case of uplift. More recently a family of *moment of uplift* measures have been developed by Radcliffe and Mesalles Narajo [5]. Models have been built for particular customer data sets provided by companies and used to guide the targeting of future campaigns, in which effective performance has been verified.

Attribution. N. J. Radcliffe was with the Maxwell Institute (MI) from 1995 to 1998 and has remained a MI visiting professor since; he was also a director of Quadstone Limited – a spin-out company from the University of Edinburgh (UoE) – from 1995-2008; since 2008, he has run Stochastic Solutions Limited. Several members of Quadstone staff also made significant contributions to the work, including D. Signorini, T. Harding and P. Surry. UoE postgraduates who worked on Uplift Modelling under Radcliffe's supervision include P. Surry (Ph.D. student, graduated 1998), D. Hofmeyr (M.Sc. student, 2010-11), O. Mesalles Narajo (M.Sc. student, 2011-12).

3. References to the research

Although information on and results from uplift modelling have been published in industry-relevant publications and conferences, until recently the details of the algorithms were considered commercially confidential by Quadstone Limited. Full details of the core algorithm have only recently been published in [1]. Further details on implementation methods and performance measures are reported in MSc theses supervised by Radcliffe [3,4].

References marked (*) best indicate the quality of the research.

[1]* Radcliffe, N. J. and Surry, P. D., Real-World Uplift Modelling with Significance-Based Uplift Trees, submitted to *Data Mining and Knowledge Discovery* (2012). Available at http://stochasticsolutions.com/pdf/sig-based-up-trees.pdf.

[2] Radcliffe, N. J. and Surry, P. D., Differential Response Analysis: Modeling True Responses by Isolating the Effect of a Single Action, *Credit Scoring and Credit Control IV* (1999). http://www.maths.ed.ac.uk/~mthdat25/uplift/cscc99-1

[3] Hofmeyr, D., An Application of Genetic Algorithms to Uplift Modelling. *M.Sc. Thesis, Department of Mathematics and Statistics, University of Edinburgh* (2011). http://www.maths.ed.ac.uk/~mthdat25/uplift/HofmeyrDavid-1

[4] Radcliffe, N. J., Using control groups to target on predicted lift: Building and assessing uplift model, *Direct Marketing Analytics Journal*, Direct Marketing Association Analytics Council, 14–21 (2007). http://www.maths.ed.ac.uk/~mthdat25/uplift/dma2006-3-1

[5] Mesalles Narajo, O., Testing a New Metric for Uplift Models, *M.Sc. Thesis, Department of Mathematics and Statistics, University of Edinburgh* (2012). http://www.maths.ed.ac.uk/~mthdat25/uplift/MesallesNaranjoOscar-1

4. Details of the impact

The research has had an economic impact for the companies commercialising uplift modelling through software sales and consultancy, and on their customers who have improved the cost-effectiveness of their marketing investments. Uplift modelling has been a core part of the consulting and software solutions marketed by Quadstone Limited from c. 2000 onwards. Quadstone was acquired by Portrait Software, in December 2005 for £3.5M. Pitney Bowes Software then acquired Portrait Software (including Quadstone) in 2010 for £44M. The Uplift Software continues to be a key part of the analytical software and services delivered by Pitney Bowes today, now marketed as Portrait Uplift [6-7]. While the precise impact and results of uplift modelling are in many cases not shared publicly by Quadstone's customers, in some cases they are. For example, US Bank (the fifth largest commercial bank in USA as of 2010) and Telenor (the world's 7th largest mobile phone operator) have both published case studies discussing the results in some detail. These are:

US Bank. The bank traditionally used 'straight response modelling' to target sales of various products including HELOCS (Home Equity Line of Credits, i.e. mortgage-backed loans). Traditional response models performed so poorly, in some cases, that with the typical 30% cutoff, they achieved no incremental sales (compared with the control group) at all or a small negative uplift. [Text removed for publication]. Work presented at Predictive Analytics World showed large improvement when uplift modelling was used. [Text removed for publication]

Telenor. The published case study [10] shows how, by using uplift modelling for its customer retention programme, Telenor reduced the rate of customer defection. [Text removed for publication].

Uplift modelling software from (now) **Pitney Bowes** is used by dozens of financial services, telecommunications and other major companies in the US, UK and mainland Europe. See [11] for further details. An example is **T-Mobile** Austria who have been using the software since 2009. A senior expert in their Consumer-Customer Insights division made the statement: 'with the use of the uplift modelling approach T-Mobile Austria has successfully optimized big retention campaigns; this has not only reduced communication costs in direct marketing activities but also had a significant uplift in contribution margin as an effect of targeting only segments which should be "moved" by simultaneously avoiding common side effects in pro-active targeting customers.' [12]

Other applications. In addition to these direct impacts of the research through Quadstone Limited and Pitney Bowes, Uplift Modelling has, after a slow gestation, started to be recognised more widely as a powerful method for increasing marketing efficiency in areas such as demand generation (cross-selling, up-selling, deep-selling) and customer retention (where campaigns with significant negative effects are not uncommon). The following examples illustrate this. **SAS** (the world's largest private software company, and the leading provider of statistical software) now includes an Incremental Response Node in its Enterprise Miner 7.1 product, which implements some form of uplift modelling. Similarly **KXEN**, another analytics company, lists uplift modelling as a capability of its InfiniteInsight Explorer [13]. The ideas underlying uplift modelling have now been diffused broadly and adopted in models that, although not directly traceable to the original research, have most likely been influenced by it. An example is the 2012 Obama campaign which used 'persuasion' models for each state (equivalent to uplift models) to decide who to target [14].

5. Sources to corroborate the impact (indicative maximum of 10 references)

[6] See http://www.portraitsoftware.com/products/portrait-uplift-optimizer for a description of the Portrait Uplift Software.

[7] The crucial importance of the research to this product can be confirmed by a former Vice President at Pitney Bowes.

[8] See http://www.portraitsoftware.com/newsandevents/press-releases/portrait-software-and-us-bank-present-predictive-analytics-world-2009 for a report on the presentation.

[9] The impact of Uplift Modelling at US Bank can be confirmed by Vice President of Marketing Analytics at US Bank.

[10] See http://www.pbinsight.com/assets_microsite/resources/files/telenor-cs.pdf for a report on the benefits of Uplift Modelling for Telenor.

[11] The page http://www.portraitsoftware.com/uplift-modeling/who-uses-uplift-modeling describes some users of Portrait Uplift.

[12] The use of Uplift Modelling by T-Mobile Austria and the statement can be confirmed by a Senior Expert, Consumer-Customer Insight, T-Mobile Austria.

[13] A description of KXEN's use of Uplift Modelling is given at http://www.kxen.com/blog/2012/01/uplift-modeling-with-kxens-infiniteinsight/

[14] The page http://www.thefiscaltimes.com/Articles/2013/01/21/The-Real-Story-Behind-Obamas-Election-Victory.aspx describes the 'persuasion' model used in the 2012 Obama campaign.

Note: should links to web pages be broken, please use the website http://www.maths.ed.ac.uk/~mthdat25/ to access pdf versions of the pages