A STRATEGIC ANALYTICS METHODOLOGY

Marcel van Rooyen[1] and Simeon J. Simoff[2]
[1]*University of Technology, Sydney, Australia*

[2]*School of Computing and Mathematics, University of Western Sydney, Australia*

Keywords: Business intelligence, data-mining project methodology, data-mining, Knowledge Discovery in Databases (KDD), knowledge management, concept drift, Telco retention-management, technology enablement.

Abstract: Businesses are experiencing difficulties with integrating data-mining analytics with decision-making and action. At present, two data-mining methodologies play a central role in enabling data-mining as a process. However, the results of reflecting on the application of these methodologies in real-world business cases against specific criteria indicate that both methodologies provide limited integration with business decision-making and action. In this paper we demonstrate the impact of these limitations on a Telco customer retention management project for a global mobile phone company. We also introduce a data-mining and analytics project methodology with improved business integration – the Strategic Analytics Methodology (SAM). The advantage of the methodology is demonstrated through its application in the same project, and comparison of the results.

1 INTRODUCTION

The use of advanced data-mining to add quantitative rigour and break-through to business decision-making, is becoming the 'industry standard' for companies that aim at strategic advantage (Hirji 2003). Business decision-making is a hypothesis-driven business process consisting of an iterative, intertwined sequence of qualitative and quantitative activities that produce business action (Pyle 2004, pp.35ff, 54ff, 165, 662.) (Hastie, Tibshirani et al. 2001, p.99) (Schön 1995) (Pearce and Robinson 2004) (Liu 2003, p.429). (Van Rooyen 2004, p.85) (Van Rooyen 2005). In order for data-mining to best support business decision-making, its must be methodologically integrated with this business process.

Data-mining project methodologies have been developed to facilitate this integration (Pyle 1999, p.10) (Van Rooyen 2004, p.86) (SAS Institute 2000, p.xi) (Chapman, Clinton et al. 1999-2000, p.3). At present, the two most common data-mining project methodologies are Data-mining Projects Methodology (SDMPM) (SAS Institute 2000), and Cross-Industry Standard Process for Data Mining (Chapman, Clinton et al. 1999-2000). Hirji (Hirji 2003) has also outlined the phases of a methodology, which at the time of the publication was at a conceptual level and had not been tested completely.

Despite the existence and use of these methodologies, organisations are still experiencing difficulty in integrating data-mining and analytics with the business. Van Rooyen showed that one of the main reasons for this is that data-mining project methodologies were data-centric. (Van Rooyen 2005). The reports of several panel discussions at the ACM SIGKDD annual forums on data-mining and knowledge discovery also pointed out the need for better alignment of the output of the data-mining process and business knowledge development activities (Ankerst 2002) (Fayyad, Shapiro et al. 2003). Hirji (Hirji 2003, p.89) describes the problem as relating to strategy, process and technology variables. Kolyshkina and Simoff (Kolyshkina and Simoff 2007) stressed that the potential value of Analytics has not been fully realised or utilised in business settings due to lack of congruency between business issues and analytics targets, and lack of analytics project management. They demonstrated a stage model of the analytics project that provides capabilities for improving business analytics projects (Kolyshkina and Simoff 2007).

The first author considered improving the integration of data-mining with the business through embedding the data-mining process into the business

van Rooyen M. and J. Simoff S. (2008).
A STRATEGIC ANALYTICS METHODOLOGY.
In *Proceedings of the Third International Conference on Software and Data Technologies - ISDM/ABF*, pages 20-28
DOI: 10.5220/0001873300200028

decision-making process. Conceptually the business decision process acts as a "shell" for the data-mining process, feeding input requirements and data, and utilising the output of the data-mining. Hence central to such approach is the structure of the business decision-making process. The Strategic Planning Cycle (SPC) (Pearce and Robinson 2004) is a generic business decision-making process, which is widely used by business, the military, etc. In this paper, we assume the reader has a basic degree of familiarity with the methodology and terminology of SPC in its commercial application.

Van Rooyen (Van Rooyen 2004) presented a comparative study based on the evaluation of the documentation of the CRISP-DM and SDMPM methodologies against SPC. The study indicated that neither of these flagship methodologies offered the required integration with the SPC planning process. These preliminary, theoretical findings justified the need to investigate and reflect on the results of an actual, business application of one of the methodologies. CRISP-DM was chosen because of its availability in the public domain and frequency of use by business. We next present background on the project.

1.1 Background

The test case study aimed at enabling the business decision-making process in a Mobile Telecommunications company which had an unacceptably high customer churn rate of 2.5% per month in their 2nd generation (2G) mobile consumer business. This resulted in lost revenue of tens of millions dollars per annum. Consequently, the Chief Financial Officer (CFO) set the *grand* objective (in SPC terminology) to reduce voluntary customer churn to 1.5%. The 2G management team, including the product managers and the 2G retention manager developed a *grand* strategy of using predictive data-mining in retention marketing, and put a *three-week* delivery time on the project. Tasked with delivering the project, the data analysts decided to use CRISP-DM as project methodology.

As the first author played a key role in the project, in this study we use the *Participatory Action Research* Methodology (Denzin and Lincoln 2003), also known as *action science* or *expert reflection-in-action* (Schön 1995). During the reflection and evaluation phase, the first author (Van Rooyen 2005) evaluated the project's outcomes against the *grand* objective, its pace of progress against business imperatives, and the quality of its marketing content against the marketing principles of Segment, Target and Position (STP) (Kotler 2002). The outcome, progress, and quality gaps were identified and defined in SPC terminology (Van Rooyen 2005), and linked back to the previous evaluation of CRISP-DM (Van Rooyen 2004). During the second stage the first author subsequently experimented with an iterative re-design of the data-mining methodology, guided by SPC principles. Eventually, the initial version of the business-oriented data-mining methodology, evolved into a new methodology, labeled Strategic Analytics Methodology (SAM). This paper presents evaluation outcomes of the pilot CRISP-DM project. The paper then presents the SAM framework, and discusses the utility of SAM in integrating data-mining with the business. The paper also discusses the framework of the SAM, and concludes with the future research direction in the area.

2 OBSERVATIONS AND EVALUATION: PROJECT OUTCOMES USING CRISP-DM

In this section we present our observations and evaluation of the test-case project. The detailed steps of the reflection methodology can be found in (Van Rooyen 2005). The headings in this section are a refinement of Van Rooyen's categories in (Van Rooyen 2004).

2.1 Introducing Expert Business Subject Matter

New retention-management subject matter was introduced to target retention-marketing campaigns at customers whose impending churn has been predicted by data-mining.

However, no innovation was proposed in segmenting the targeted customers on behavioural proxies in the data for value, need, and loyalty (Wedel and Kamakura 2000). The analysts decided to retain the existing retention segmentation based on socio-demographic data cubes despite the fact that there were known problems with this approach, and it was accepted that the segmentation was negatively impacting on campaign response rates. Also, no consideration was given to the use of customer behavioural profiles in tailoring campaign offers. The retention-anagement innovation was therefore limited to the T (targeting) component of the STP marketing principles, restricting the potential improvements in campaign response rates at the outset of the project.

2.2 Formulating Project Scope and Objectives

The defining of the business rules about voluntary churn - and therefore the project scope - was not undertaken by the management team. Instead, the data analysts were left on their own to scope the project. This led the analysts to take an extra step in extracting business rules from the business, which added an extra four weeks to the first phase of the work – an unacceptable delay in the current globally competitive business climate.

The management team had decided at the outset that the new retention campaigns must be delivered within the predicted 3-month time window. This was unrealistic given organisational circumstances we discuss later, and in reality the campaigns took an unacceptable six months to execute.

Also, despite a caution from the analysts about potential lack of evidence, management expressed their determination to continue addressing an anecdotal driver of churn (handset type) in the campaign offer.

No measurement link was established with Return on Marketing Investment (ROMI) as a measure of the project's success.

2.3 Mapping Technique to Data-mining Plan

The data-mining *strategy* selected for prediction (T in STP) was a neural network model. However, the computational models that it offered, and the predicting effects that it generated, were too difficult to interpret and to explain to management. Hence, the strategy therefore failed to produce the needed information and intelligence towards addressing churn drivers in the campaigns.

Further, because no new business subject-matter on segmentation (S) and promotion (P) was introduced earlier, no data-mining objectives could be formulated to re-visit the current segmentation with another statistical technique that offers better explanation, or to profile the segments for campaign-relevant behavioural dimensions.

2.4 Evaluation of Interestingness, Decision-making and Business Action

We observed that the interestingness of knowing each customer's propensity to churn – and model lift - was not sufficiently considered by the retention manager. Instead of, say, halving the number of customers targeted for retention, the manager selected a similar 5% proportion of the 2G customer base for retention management as before. Retention-management costs were therefore not reduced, which diluted the potential lift in ROMI associated with model lift.

Further example of the lack of strategic flavour observed in this analytics project is the mismatch in the 2G customer-analytics retention campaign and the introduction of the recently launched third generation network (3G) acquisition campaign. The campaign manager had incentives to perform in 3G acquisition. The 2G manager did not consider the impact of this on the 2G retention campaigns. Because of this, the 2G retention campaigns were not executed within their 3-month predicted time window, resulting in some at-risk customers receiving their retention offers after they had actually cancelled their service. Also, the campaign manager decided to target at random from the selected 5%. This meant that the 2G retention campaigns were not prioritised towards the most at-risk customers, further diluting the benefits of model lift.

The model produced by the neural network proved difficult to interpret, and no interesting insights were developed relating to how customer value, loyalty, and the degree of product feature and volume use, indicated churn. As a result, the existing promotional content was retained; *all* 2G customers with predictive handsets were offered 2G replacement handsets, irrespective of their degree and volume of use. There were two negative results from this. First, the high cost-base relating to handset promotion was perpetuated, and a potential cost saving from the data-mining approach was forfeited. Second, the existing 2G *retention* offer (with replacement 2G handset) was maintained to high-use 2G customers, despite the fact that they had a disproportional low 2G campaign response rate. There were no insights to stimulate thought about perhaps targeting high-use customers with 3G *migration* offers - since 3G is more suitable to high use and is better featured – and to at least retain revenue from these customers within the organisation, albeit as 3G revenue.

Also, no new insights were developed regarding the role of demographics in churn, and campaign offer components continued to be mapped to demographic descriptors, instead of being updated potential for behavioural drivers as the literature indicated.

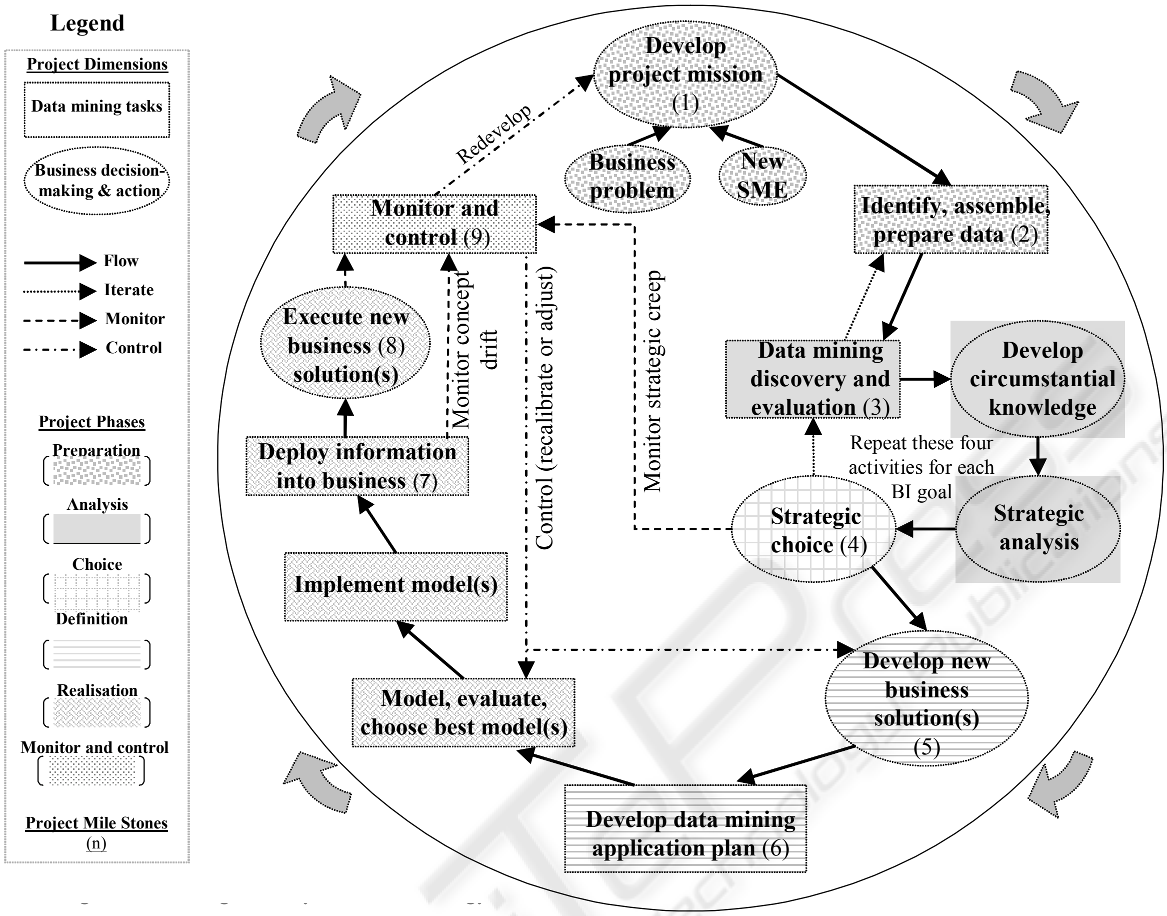

Figure 1: Strategic Analytics Methodology.

2.5 Monitor and Control

The campaign manager continued to monitor campaign response rates, and a 25% improvement was recorded for the CRISP-DM project. The retention manager monitored the percentage of voluntary customer churn, and recorded a similar improvement. However, both results fell short of an expected 50% response rate increase resulting from predictive lift.

There were no plans to monitor the *drift of concepts* that should accompany the proactive, predictive approach to retention management (Van Rooyen 2004, p.94) i.e. monitoring changes in (1) predictive effect scores (2) distance between segments and changes in their behavioural profiles (3) declining model accuracy against actual churners, and (4) the overall predicted risk associated with the churn problem. The business remained locked into reactively responding to falling campaign response rates; that is over time fewer and fewer targeted customers responded positively to save campaigns.

Further, since the calculation and tracking of ROMI was not introduced, there was no return on investment measure for quantifying the benefits that data-mining had brought to retention management in their company.

3 REFRAMING: THE SAM FRAMEWORK

In this section we present the framework of Strategic Analytics Methodology (SAM), and discuss the role of each framework element in the project. The full details are presented in (Van Rooyen 2005, Chapter 5). The framework is shown in Figure 1. The sequence of steps in the framework follows clock-wise, starting at the 12 o'clock position. We follow the six project phases shown in the legend, and highlight the project milestones. The intermediate steps within a project phase are referred in italics.

3.1 Preparation

In the preparation phase, we develop hypotheses about the nature and extent of the business problem or opportunity, its root causes, how data-mining and analytics can enable its solution, and about the actions the business could take. In *Business problem* we define the status quo. The infusion of *New SME* represents domain-specific, new subject matter based on research that challenge the existing boundaries of interestingness, poising the organisation for break-through. In *Develop project mission* the owner of the business problem and the analysts together apply a technique that challenges the current thinking about the business's status quo. The analysts and business agree on what will constitute interestingness, on what the project must deliver for the business action, and on how the analytics will support the project. They then use a guide to convert these hypotheses and agreements into a project mission, the first milestone. The project mission's components include the project scope, its specific business deliverables, interestingness measurement criteria, and the time frames, responsibilities etc. of the project. In *Identify, assemble, prepare data* we identify and source the interesting data, discover information *about* the data quality and interesting signals, and chart how we want to use the data on the project. This constitutes the second milestone.

3.2 Analysis

In the analysis phase, we develop a technical analysis charter and discover interesting information from the data and evaluate it against the interestingness criteria in the project mission. The milestone is interesting information, which supports the project mission. This information is interpreted into circumstantial business insights and knowledge (*Develop circumstantial knowledge*). In *Strategic analysis* we test and redevelop our hypotheses in light of the new insights, and formulate a range of high-level business action options.

3.3 Choice

In this phase, we evaluate the range of action options against organisational and environmental circumstances using a SWOT analysis tool, and *reject* those options that are not realistically actionable. The choices made here are the 4th milestone, since they funnel the project through to the definition of business action.

3.4 Definition

In this phase we develop the functional business actions (*Develop new business solution(s)*) as functional objectives and strategies (e.g. in marketing, operations etc.) for all business enablers (structure, systems, management etc.). We also develop: (i) the key business performance metrics that we will use in monitoring the success of the business action; (ii) the data-mining application plan that will support the business actions (*Develop data application mining plan*). Further, we identify the technical concepts that we model and monitor for drift. Together these constitute the 5th and 6th milestones.

3.5 Realisation

In *Model, evaluate choose best model(s)* we refine or modify the models from *Data mining discovery and evaluation*, and format their informational output for interfacing with other IT applications.

Before the business can *Execute new business solution(s)*, two application layers need enablement. In *Implement model(s)* we build the extract, transform and load (ETL) and analytics layers, and in *Deploy information into business* we enable the business and IT application layers e.g. the customer-relationship management (CRM) application. As the enabling of the business applications, this constitutes the 7th milestone. In *Execute new business solutions*, we execute the business actions through structure, management, IT applications, channels, operating strategies etc. This is the completion of business action and therefore the 8th milestone.

3.6 Monitor and Control

In this last phase of the project we assure the ongoing relevance of the business actions and of their supporting applications (including the operational models) through (i) a process of monitoring the business KPI's and the technical concepts (ii) recalibrating actions and enablers where required; (iii) and monitoring strategic creep.

Monitor and control is an ongoing activity but we have made it the project's 9th milestone. SAM draws on the principles and practice of *concept drift* and statistical process control for the technical components, and on the monitor and control activities of SPC for the business solution. When either the business actions or the supporting enablers become out of control, the project needs to be reformulated, or a new project of higher return is formulated.

4 MOVE TESTING: AN INDUSTRIAL APPLICATION OF SAM

In this section we demonstrate the application of SAM on the same project, and describe the business and technical benefits from this application. The benefits are compared to the CRISP-DM application where relevant.

4.1 Problem Formulation

We used SAM's utilities for defining the business problem in *Business problem* with management participation. Within two one-hour workshops, we refined the business rule for voluntary churn to exclude bad debtors. This later resulted in improving the regression model's predictive accuracy against test cases from 75% to 80%.

4.2 Introducing Expert Business Subject Matter

In *New SME* we introduced the concept of ROMI, and assisted with developing a measure of ROMI suitable for the business. We also linked for the business the concept of predictive lift with the way it could lead to a reduction in the number of customers targeted (T). Further, we agreed with management on: (i) selecting the concept of behavioural segmentation (value, loyalty, needs); (ii) segmenting the targeted customers into 5 statistically relevant segments (to fit the existing segment numbers) by these new measures (S); (iii) and tailoring campaign offers based on behavioural profiles (P).

4.3 Formulating Project Mission

Given the previous re-defining of the churn business rules, the project mission was reformulated to exclude bad debtors. This better focused the analysts on the problem at hand.

4.4 Formulating Technical Charter

Considering the business need for understanding the drivers of churn, we used mapping techniques in *Data mining discovery and evaluation* to lead the project to a decision to build a predictive logistic regression model, instead of using a neural network model (T). This resulted in a model that was interpretable, which in turn facilitated the development of significant insights: (i) the degree and volume of use actually were strong churn drivers; (ii) there were use interactions with certain handset-types; (ii) demographics were insignificant.

These insights supported the case for more statistically valid behavioural segmentation. We therefore decided to use clustering techniques to segment the selected 2.5% customers by behavioural criteria. We found that there were 5 distinct, statistically valid segments (S), with distinct behavioural profiles. One segment (segment 2) consisted of 25% of the targeted customers, and had a distinct handset-type high-use combination, which was well-supported by the regression model.

4.5 Evaluating Interesting Information, Decision-making and Business Action

In the Analysis phase the retention manager now considered a number of things:
(i) a new understanding of the benefits of lift;
(ii) the fact that predictive accuracy had increased after the exclusion of bad debtors;
(iii) as the subject matter about behavior had suggested, we had found strong, statistically valid behavioural churn drivers and segments. The combination of handset type and high-use behaviour had featured, while no demographic factor was driving churn as strong as the behavioural factors drove;
(iv) limited campaign resources were available to 2G retention;
(v) the retention and campaign managers now had ROMI targets to attain. Therefore, during *Strategic choice* they decided collaboratively to target (T) only the most at-risk 2.5% of the 2G customer database. We observed that this halving of targeted customers resulted in the campaigns being executed within their 3-month window to all the targeted customers, eliminating the campaign execution issues observed before.

During the Analysis phase of SAM, the retention manager found segment 2 particularly interesting (S). In *Develop circumstantial knowledge* and

Strategic analysis he considered the impact that the recent launch of 3G had had on the first round of 2G retention campaign execution, and in *Strategic choice* he decided to hand segment 2 over to the 3G acquisition manager for *migration* onto the 3G network (T and P), to a technology that better matched this behavioural profile. This further contributed to attaining the ROMI target, because they did not have to fund any 2G replacement handsets for this segment, and this segment's uptake of 3G was very high.

In light of some evidence that demographics were not driving churn, the retention manager reformulated the retention offer content (P) for segments 1, 3, 4 and 5. Previously, males in the 19 – 25 age group had received a promotional ticket concession to Australian Football League (AFL) matches, based on the assumption that most males in this age group followed AFL. Now, the AFL promotion was offered to all targeted individuals who had a history of requesting sms results of AFL games, irrespective of their demographic. Also, where previously there had been replacement of all 2G handsets irrespective of use level, now 2G handsets were promoted only to individuals with medium-use patterns; individuals with low-use patterns had to buy a replacement handset (recall that high users were *migrated* to 3G).

4.6 Monitor and Control

On the SAM-based project, the response rates increased about 75% over the pre-data-mining basis. This exceeded the business target of 50%, and the 25% improvement on the CRISP-DM-driven project. Notably, there was a 90% uptake of the 3G offer by segment 2 members. The CRISP-DM project brought a 10% ROMI improvement over the base (before CRISP-DM pilot) situation, while the SAM project's ROMI improved 30% over base. The smaller ROMI improvement compared to the campaign response improvements, are attributed to fixed campaign costs remaining from the pre-data-mining mining era; over time ROMI will improve as the organization addresses these costs.

The value of a data-mining approach that supports marketing STP was now well understood by management and proved to the business. The business responded with a request to monitor the extent of the churn risk in the database and assist in formulating responses. We monitored drift in the sum of scored p_values within quintiles with every quarterly database re-scoring, and relating these to change in the campaign response rate. The project iterated through the Analysis, Choice, and Definition phases and we formulated various controls for this phenomenon. For instance, if both sum of p_values and campaign response are in decline (scenario 1), the business solution is to address the drivers of the problem. The first control response is to retrain the model to maintain accuracy. If after a model retrain the sum of p_values and campaign response still are in decline, it indicates that the business solution is also weeding out potential churners, and may be approaching a situation of diminishing returns with retention. The response is to fine-tune the campaign offers (scenario 2). If with model retraining (and in some cases after campaign refining) the sum of p_values keep increasing while the campaign response rates are in decline, then circumstances in the business operating or marketing environments are not represented in the data. The response is to undertake qualitative research to identify these circumstances, and to update the campaigns (scenario 3).

We found that the three scenarios actually manifested themselves in this order over scoring cycles 2-4. In cycle 4, the 2G retention manager was convinced that the point of diminishing returns had been reached with 2G retention, with consensus in the business, that most of the 2G churn risk had been 'weeded'. This was the result of events that were not reflected in the data, namely (1) the business had stopped acquiring 2G customers with the launch of 3G in quarter 1, and (2) by the fourth quarter 3G was cannibalising the mid- and high-user 2G business. Since the revenue and margins in 3G were superior to those in 2G, the business now had sufficient evidence to abandon the 2G retention program.

The concept of interacting behavioural churn drivers was now well understood and proved to the business, and there was a necessity for monitoring any drift in churn drivers, and their interactions, over time. In the second SAM campaign cycle, we started to rebuild the predictive model and the segmentation on a monthly basis, and to monitor drift in effect scores and segment parameters.

One noticeable drift in effects was the diminishing impact of the high-use and handset type combination. We attribute this to the fact that we were gradually moving the high-use customers over to 3G. What did emerge in its place was a combination of over-3-years customer tenure, over-18-month minimum contract plan types, and a 26-34 year-age demographic. It was now apparent, that our own and emerging competitive 3G advertising, was influencing the traditionally more conservative customers in this age group to switch brand or technology. This insight enabled the making of a more proactive 3G *migration* offer to customers who exhibited this behaviour. Their take-up was not 90% as with segment 2 previously, because of the diluting effect of competitive 3G activity.

There also was a noticeable reduction in the importance of segment 2 in the segmentation structure over time, since many customers with that profile were migrating to 3G. Over time the importance of customer tenure and plan duration rose in the segmentation structure. When we experimented with the number of segments, we found that, compared to five segments, four segments actually did not detract much from the statistical significance of the clustering. The business found this attractive, as it meant that they could eliminate one of the five campaigns without a significant impact on response rates. This contributed a further 5 points to ROMI on the SAM project, bringing SAM ROMI to 35% over base.

5 CONCLUSIONS AND FUTURE WORK

We conclude that we have proved the hypothesis that CRISP-DM requires extension in order to integrate deeper analytics and the business decision-making process. The improved results indicate it is possible to improve data-mining methodologies to better integrate with business decision-making and action. We have also proved that the Strategic Analytics Methodology (SAM), which follows a deeper integration of data-mining and business decision-making process, is sufficiently robust to produce beneficial results in a dynamic business environment.

SAM, in its current version met its research objective as industry-driven, academic research by late 2004. At present, a SAM-derived framework is used effectively in process-enabled, industrial data analytics.

Several directions are considered in terms of the future development of SAM, aiming at:

1. further experimentation with better controls, in order to quantify the incremental benefit that SAM contributes to business data-mining compared to existing methodologies;
2. specific development and refinement of SAM to better accommodate the unique integration criteria of not-for-profit and government applications;
3. proving the hypothesis that the existence of a more supportive project methodology like SAM would enhance the uptake and adoption of data-mining and analytics by business and government.

ACKNOWLEDGEMENTS

This research was done under the e-Markets Research Program, University of Technology Sydney. We acknowledge the RTS Scholarship and the support of SAS Institute Australia Pty. Ltd. Thanks to a global mobile-phone company for providing the real-world data and analytics environment. Thanks Dr. Michael J. Yerbury for your grammatical contribution.

REFERENCES

Ankerst, M. (2002). "Report on the SIGKDD-2002 panel - The perfect data mining tool: interactive or automated?" ACM SIGKDD Explorations 4(2): 110-111.

Chapman, P., J. Clinton, et al. (1999-2000). CRISP-DM 1.0: Cross Industry Standard Process for Data Mining. http://www.crisp-dm.org/CRISPWP-0800.pdf, CRISP-DM Consortium. Accessed November 2003.

Denzin, N. K. and Y. S. Lincoln (2003). Strategies of qualitative inquiry. Thousand Oaks, CA, Sage.

Fayyad, U., G. Shapiro, et al. (2003). "Summary from the KDD-03 panel - Data mining: the next 10 years." ACM SIGKDD Explorations 5(2): 191-196.

Hastie, T., R. Tibshirani, et al. (2001). The Elements of Statistical Learning. New York, Heidelberg, Berlin, Springer-Verlag.

Hirji, K. K. (2003). A Proposed Process for Performing Data Mining Projects. Managing Data Mining Technologies in Organizations: Techniques and Applications. P. C. Pendharkar, Idea Group Inc.: 350.

Kolyshkina, I. and S. J. Simoff (2007). Customer analytics projects: addressing existing problems with a process that leads to success. Conferences in Research and Practice in Information Technology, Data Mining and Analytics 2007, Australian Computer Society Inc: pp. 13-20.

Kotler, P. (2002). Marketing Management: Analysis, Planning, Implementation, and Control. International, Prentice Hall.

Liu, X. (2003). Systems and Applications. Intelligent Data Analysis. M. Berthold and D. J. Hand. Heidelberg, Springer-Verlag: 429-442.

Pearce, I. J. A. and J. R. B. Robinson (2004). Strategic Management: Formulation, Implementation, and Control, McGraw-Hill.

Pyle, D. (1999). Data Preparation for Data Mining. San Francisco, Morgan Kaufmann Publishers.
Pyle, D. (2004). Business Modelling and Data Mining. London, Morgan Kaufmann.
SAS Institute (2000). SAS Data Mining Projects Methodology. Cary, NC, SAS Institute Inc.
Schön, D. A. (1995). The Reflective Practitioner: How Professionals Think in Action. London, Ashgate Publishing Limited.
Van Rooyen, M. (2004). An evaluation of the utility of two data mining project methodologies. Proceedings of the 3rd Australasian Data Mining Conference, Cairns, Australia, University of Technology Sydney: 85-97.
Van Rooyen, M. (2005). A Strategic Analytics Methodology. Faculty of Information Technology. Sydney, University of Technology: 350.
Wedel, M. and W. Kamakura (2000). Segmentation: Conceptual and Methodological Foundations. Boston, Kluwer Academic Publishers.