



Predicting Churn

*Analytical Strategies for Retaining Profitable
Customers in the Telecommunications Industry*

Table of Contents

Introduction	1
The Price of Churn	1
The Raw Costs	1
The Challenge for Management	2
Information Delivery in the Telecommunications Industry.....	2
The Challenge for Information Technology Management	2
“Think strategic, start focused”	2
Churn Prediction	4
Business Goals	4
Data Definition	5
The Analysis Process	5
Deployment and Review.....	6
Warehousing.....	6
Outlook	6
Case Study	6
Process Flow	6
Input Data	8
Modeling Techniques.....	8
Assessment	9
Summary	9

Introduction

Customer acquisition and retention is a concern for all industries, but it is particularly acute in the strongly competitive and now broadly liberalized telecommunications industry. For the marketing departments of newcomer companies, the major short- to medium-term issue is likely to be attracting new customers. However, for the incumbent operators and the more mature market entrants, retaining profitable customers is the number one business pain.

The information delivery systems of many telecommunication companies are reaching a maturity level that allows them to make the step from simple query and reporting of past cancellations towards the semi-automatic creation of predictive models. Companies can assure a more constant flow of revenue and higher profit margins through targeted activity such as the fine-tuning of services and promotional messages.

This white paper introduces some of the terms and concepts involved in the process of accurately predicting which customers are likely to deliver high value and at the same time are likely to exhibit a high propensity to change suppliers.

The Price of Churn

The Raw Costs

Churn is the common denominator in the world's liberalized telecommunications industry. It now costs European and US telcos close to US\$4 billion each year, and the global cost of customer defection may well approach a staggering \$10billion.¹ Annual churn rates of 25 to 30 percent are the norm, and carriers at the upper end of this spectrum will get no return on investment on new subscribers.

Why? Because it typically takes three years to pay back the cost (approximately US\$400 in the United States and US\$700 in Europe) of replacing each lost customer with a new one (customer acquisition).

In the European and Asian markets in particular, the number of new market entrants is adding to the churn phenomenon. In Europe, 30 new telcos entered the market in 1998, seeking the 15 percent market share that analysts say they will need to survive. The growth in the number of subscribers has eased this situation in the past, but as market growth slows and average revenue per user declines, we are likely to see an increase in predatory activity.

¹ *Telecommunications Online*, February 1999.

The Challenge for Management

The problem confronting telcos' management is that it is very difficult to determine which subscribers left the company and why. It is therefore even more difficult to predict which customers are likely to leave the company, and more difficult still to devise cost-effective incentives that will persuade likely "churners" to stay.

Churn is such a massive problem that it affects other aspects of customer relationship management, such as customer acquisition. A manager must ask himself, "Am I recruiting the right people or are they likely to churn before I have made a return on my investment?" "How is churn affecting the lifetime value of my customer base?" and "Can we get a complete view of our customer information, so that we can profile likely churners?" The answer to these questions depends largely on having the right information delivery solutions in place.

Information Delivery in the Telecommunications Industry

The Challenge for Information Technology Management

Telcos are among the biggest users of information technology (IT) systems, yet their IT departments tend to be focused on meeting day-to-day operational goals. In many cases the technology is not yet in place to support the complex requests for information from the sales and marketing departments that must address the issue of churn. Also they may lack the expertise to support complex data mining and analytical and predictive tasks. The volumes of data that are needed to undertake such tasks are huge and sometimes difficult or impossible to access and consolidate using conventional operational system tools.

"Think strategic, start focused"

"Think strategic, start focused" is a good motto for anybody who wants to build an information delivery solution.² You must always keep the big picture in mind, because information delivery projects are invariably designed to solve business problems relating to the organization's overall strategy. An information delivery project therefore requires the support of senior management. (Clearly, this is not the case with systems that merely process data and migrate it from one machine to the next.)

On the other hand, it is foolish to try to solve everything in one go. Such an approach will suck in untold resources, delay return on investment and put the project managers under increasing strain. Business executives who must answer to shareholders while fending off the competition understand this message. Every investment in data warehousing and data mining needs to show a prompt return on investment.

² SAS *Rapid Warehousing Methodology*, SAS White Paper (43093US_0501), SAS, Cary, NC.

Now for the good news: SAS has done it all before, many times over. Experience tells us that the only feasible way to build an enterprise information delivery solution is with support from the top down and action from the bottom up, adding one information delivery application after the other.

Everybody involved in an information delivery project must have in mind the clear vision of an integrated information architecture all along the way. As Rob Mattison points out, the main goal to envision is the alignment of the telecommunications value chain, the organizational structure and the architecture of the IT systems, with the value chain taking the leading role (see Figure 1).³ At the start is the creation of a product and the acquisition of the right to distribute it. A network then has to be built and maintained and the customers must be properly billed and managed.

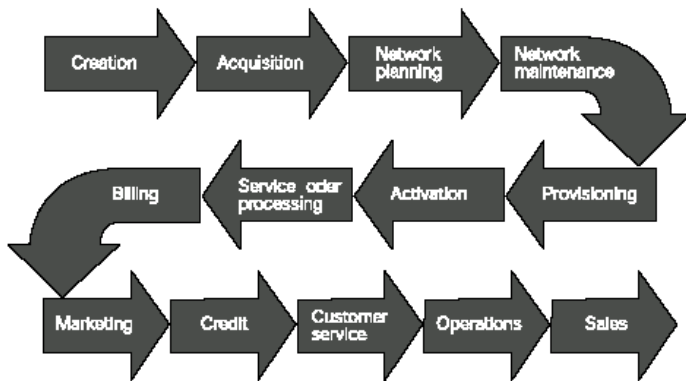


Figure 1: The Value Chain in the Telecommunications Industry

The IT systems in most companies already align with the value chain to some degree. What can usually be found is that two main systems make up the largest part of the IT infrastructure: the switching system and the billing system, corresponding to the network side and the business side of the value chain. The records that the switching system keeps on the call details are passed to the billing system for summarization and the issuing of bills.

Both these systems are essentially operational systems. However, they often have been stretched beyond their original purpose to serve as platforms for information delivery. So, on the network side, traffic monitoring or even capacity planning functionality may have been added, while the billing system may have turned into something like a customer management system. If that is so, then on both sides of the house possibly hundreds of users will run reports directly against what used to be and still is supposed to function as an operational system. This won't work in the long run.

What is really needed is an information delivery layer, a *data warehouse*. The warehouse allows the switching and the billing system to work more efficiently, because they don't have to cater for so many user requests. Maybe even more importantly, such a buffer layer insulates the business users from any changes in the organizational or IT structure of the switching and billing operations. All the information needed for their purpose will always be found at the same place, processed and enhanced in a way appropriate for their specific purposes.

³ *Data Warehousing and Data Mining for Telecommunications*, by Rob Mattison, Artech Computer Science Library, Norwood, MA.

Adding this extra degree of freedom is what data warehousing is all about. It is this freedom that makes it possible and advisable to build the data warehouse step-by-step from the bottom up. It would be wasteful to re-enter unnecessary constraints up-front to limit the business user's freedom to cater for his needs. It is rarely possible to anticipate information needs very long in advance and this is especially true when data mining comes into play. *Data mining* is a process that is meant to generate new knowledge, which in turn will lead to new processes and data requirements. To a certain degree this circle is open-ended. As a consequence, the nature of the relationship between IT and business in the area of information delivery systems must change to be one of constant co-operation and adaptation, and must not be inhibited by a spirit of once-and-for all system implementation.

Data warehousing and mining can add value at a number of different places along the telecommunications value chain. Consolidating and analyzing customer information in the marketing department for retention or acquisition purposes, for example profiling profitable customers, and designing promotional campaigns has surely been the most prominent application area so far. Another important value proposition is located in the credit area, where it is necessary to evaluate a customer's tendency not to pay the bill. The results of such an analysis can then be used to limit the customer's access to certain services or to reject applications of risky customers in the first place. Closely related to this are fraud detection applications. Other applications concern the optimization of customer service as well as network traffic monitoring and capacity planning.

Customer retention strategy is seen by many telcos and other companies as today's most pressing business pain and is therefore often chosen as the first application area of data mining. Not all companies, however, are yet at the point where they want to start with predicting future cancellations. Some are still fully focusing on acquisition. Others are so far only able to generate lists of churners and to get a feeling of the size of the problem.

Yet others are already able to carry out simple analyses and queries that describe the cancellations that have occurred in the past more in detail. Over the past year many companies have embarked on data mining projects that let them learn from the past and predict each customer's likelihood to cancel in the future.

Churn Prediction

Business Goals

The primary goal of churn analysis is usually to create a list of contracts that are likely to be cancelled in the near future. The customers holding these contracts are then targeted with special incentives designed to deter cancellation. At a more sophisticated level, the telecommunications company attempts to detect the reasons for an expected cancellation, because this information may help customize the offer. Detecting causes of churn that lie within the sphere of influence of the company also enable it to eliminate them in the future. For example, analysis may reveal that factors as different as inadequate billing procedures and connection quality are the root causes of churn.

Data Definition

The statistical unit for churn analysis is most often not a customer but a contract. In other words, propensities of cancellation are calculated on a per contract rather than a per customer basis. The main reason for this is that many important predictor variables, for example the length of time since a contract has been signed or the time left until the end of the obligation period, are associated with contracts rather than customers. Also, even though a customer may hold several contracts, usually each of these contracts contributes to revenue. However, it should be remembered that mailings and other follow-up actions target customers, not contracts, so there must be some post-analysis processing to summarize the predictions for customers as individuals.

It is a good idea to conduct analysis and build models for market segments that exhibit some broad commonality. Building a model means finding rules that relate customer attributes (input variables) to the likelihood of the churn event coded in the target variable. The customer attributes typically considered in a churn analysis can be broadly categorized into four types:

- Customer demographics.
- Contractual data.
- Technical quality data.
- Billing and usage data.
- Events-type data.

The most commonly used historic variables are:

- The time a customer spends on air.
- The number of calls.
- The revenue.

The Analysis Process

The process of analyzing the data follows the SAS Sample, Explore, Modify, Model, Assess (SEMMA) methodology.⁴

⁴ From *Data to Business Advantage: Data Mining, SEMMA Methodology and the SAS System*, SAS White Paper, SAS, Cary, NC.

Deployment and Review

Deployment means taking the churn management process out of the limited realm of the analyst. As a first step it is often necessary to make the information that defines and results from the analysis flow available to a broader group of business users. These reports can include a depiction of the process flow itself, the different assessment charts, the decision tree structure, the coefficients of the logistic regression or the details of the pre-processing.

Before launching the prevention campaign, it is generally considered necessary to verify the accuracy of the model at least once by comparing the predictions with the actual cancellations of the most recent month.

Warehousing

There are a number of data management tasks that need to be executed at regular intervals when a churn management solution has been deployed. Most importantly, the usage and revenue data from the operational billing and switching systems needs to be passed to the warehouse, where it is aggregated and appended to the churn analysis data set.

Outlook

It usually is not just a customer's likelihood to cancel that makes her or him a good target for a prevention campaign. The value that a customer will bring to the company in the future — usually a projection of past revenue — plays a major role in designing a campaign in such a way that it brings maximum return on investment. Analysis should also take into account the likelihood that a targeted customer can be deterred from canceling. All three qualities — churn likelihood, prevention likelihood and value — need to be combined to decide on campaign targets.

Case Study

Process Flow

A simple process flow will illustrate how churn analysis can be implemented with the help of SAS Enterprise Miner software. The flow concentrates on the main aspects of the SEMMA methodology.

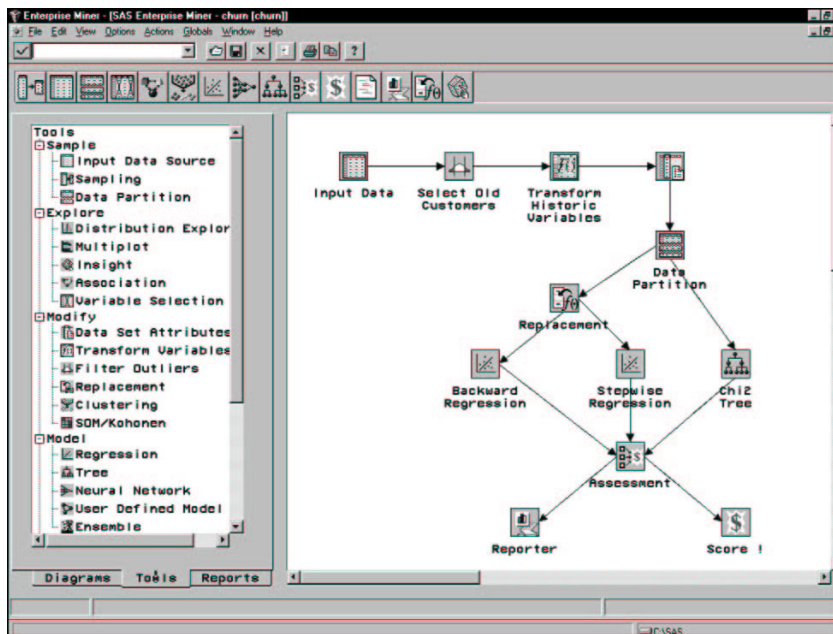


Figure 2: SAS Enterprise Miner

In figure 2, the workspace containing the process flow diagram is in the right-hand window, and the tools palette is in the left-hand window.

The flow starts in the upper left corner with the Input Data Source node, where the data is read in, model roles and measurement levels are assigned, univariate summary statistics are calculated and univariate distribution charts can be viewed. The data set discussed will be described more closely in the following section.

A subset of customers is selected, namely all customers whose contract started at least six months ago. The Filter node is used for this purpose.

Some simple transformations of historic variables are then done in the Transform Variables node. Instead of using all the detail available for the modeling, it was decided to retain only the values for the last month and summarize the values of the other months by building the average of the monthly totals. Additionally, flags are created that indicate if the usage was zero in the last month. A Data Set Attributes node is used afterwards to set the model role and the measurement levels of the newly created variables.

The data is then split in half with the Data Partition node to obtain a training data set and a validation data set that is later used to choose the best model.

The Replacement node then substitutes missing values with sensible estimates. Three models are built, two of which are regression models that only use complete observations as input. The two Regression nodes are therefore connected to the Replacement node.

Two alternative variable selection methods are experimented with, namely backward and stepwise selection. The third model is a decision tree. The Tree node is not connected to the Replacement node, since the tree algorithm can handle missing values.

The quality of the three models is then assessed and compared using the Assessment node.

The reporter node creates a hypertext (HTML) documentation of the analysis and the Score node creates SAS code that can be used for assigning churn propensities on any system independently of SAS Enterprise Miner.

Input Data

The input data set should typically be organized to include names, model roles, measurement levels and labels of the variables. Variables would typically describe the contract and the respective customer, together with variables containing the usage and revenue in the relevant time period (totals and more detailed information).

Modeling Techniques

A decision tree finds optimal if-then rules that split the customers, variable by variable, into ever finer segments in such a way that the terminal segments (“leaves”) contain either a very high or a very low proportion of churners. Regression models try to find optimal coefficients for a linear equation that relates the input variables to the likelihood of the churn event.

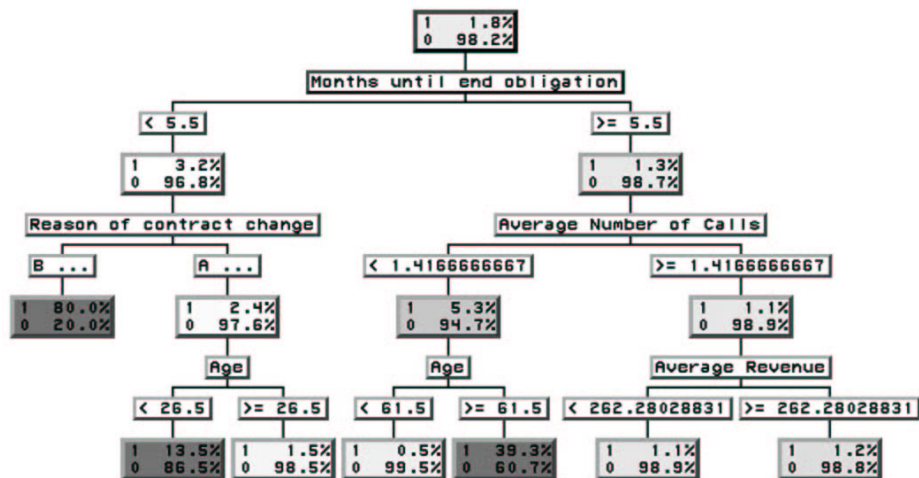


Figure 3. Decision Tree Results

Only the first three of six levels are shown in Figure 3. Dark shading represents a high proportion of churners in the corresponding segment.

Assessment

The quality of the models needs to be assessed and compared in order to pick the model that gives the highest business value.

After modeling, the customers in the validation data set are sorted by the churn propensities assigned by each model. Equally sized groups (in this case containing two percent of the customers each) are then built and the proportions of actual churners in each group determined.

Summary

Churn is widely recognized as a major threat to telecommunications service providers. If they can retain the best customers, telcos can increase overall profitability. To do this, they must first of all identify which customers are likely to churn. What are their characteristics? What can be done to incentive them to stay loyal?

This white paper discussed some of the requirements of a methodology and software solution that will help telcos to answer these questions by exploring and mining customer data.

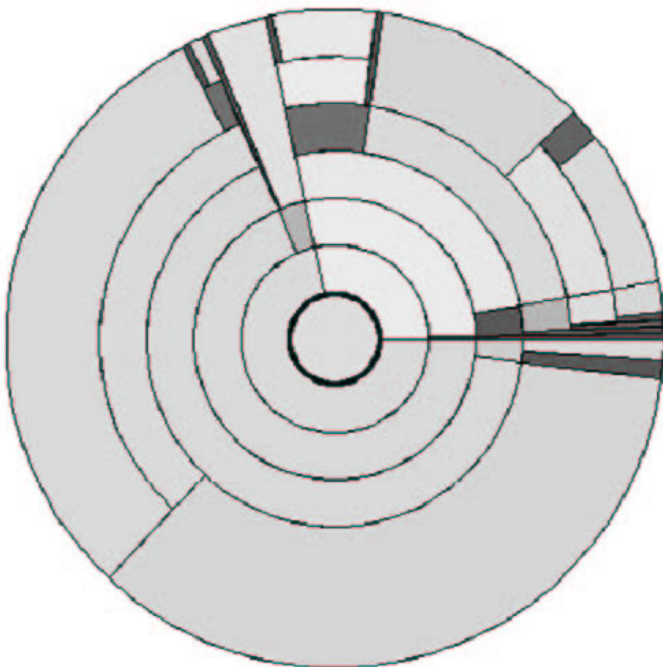


Figure 4: Decision Tree Results Shown in Ring View

In figure 4, dark shading represents a high proportion of churners in the segment.

Predicting Churn

A clear methodological approach and powerful enabling software technology are, however, only two of the essentials for a successful churn modeling exercise. The traditional virtues of science, patience and curiosity will eventually also have their share, since data mining is an iterative process — answers to one set of questions lead to more interesting and specific questions.

The many possible routes that a data mining analysis can take require a solution that can harness this degree of complexity without limiting the freedom and creativity of the analyst. SAS Enterprise Miner software solves this dilemma by synthesizing the world-renowned statistical analysis and reporting system of Institute with an easy-to-use graphical user interface that can be understood and used by business analysts as well as quantitative experts.



World Headquarters
and SAS Americas
SAS Campus Drive
Cary, NC 27513 USA
Tel: (919) 677 8000
Fax: (919) 677 4444
U.S. & Canada sales:
(800) 727 0025

SAS International
PO Box 10 53 40
Neuenheimer Landstr. 28-30
D-69043 Heidelberg, Germany
Tel: (49) 6221 4160
Fax: (49) 6221 474850
www.sas.com