Paper 084-30

A Proportional Hazards Approach to Campaign List Selection

By Charles Manahan, Ph. D.
Cingular Wireless, LLC.
Atlanta, GA

ABSTRACT:
Churn, or customer loss, is a serious problem in the wireless telecommunications industry. Hence, churn reduction is a priority goal in the wireless industry, and marketing campaigns are frequently done to reduce this problem. But, how well do they succeed, and what can be done to make them more successful? There are three traditional parts to any marketing campaign, the offer, the creative and the list. This paper addresses developing predictive models using SAS/Enterprise Miner® and SAS/STAT® to give lists with a higher response rate.

SAS® software is sufficiently powerful to do both analysis of past campaigns and produce predictive models. The results of a campaign on both churn and contract renewal were analyzed using a life table approach. This gave the response shape that a predictive model should emulate. The proportional hazards approach was used to model the predicted response to future campaigns. Models treating churn and contract renewal as competing events using both proc logistic and proc neural were developed and evaluated.

This is an intermediate level presentation and audience should have knowledge of base SAS® and conceptual familiarity with regression.

INTRODUCTION:
One of the churn reduction initiatives at Cingular is a contract renewal campaign directed to customers who will be coming off contract in the next 90 days. Initially, offers were directed to all customers who met minimum financial criteria. However the results from this approach were mixed. There is a fraction of customers who are sufficiently unhappy with their service that they will churn within the first two months after coming out of contract. The results showed that the fraction of these “immediately churning” customers who go elsewhere upon contract expiration was only minimally impacted by these campaigns. In fact the treated customers churn at a slightly higher rate than do the control customers. (See Figure1 – Comparison of rate for churn event in treated vs. control customers). Figure 1 is for the cohort whose contracts expired in April of 2004. Several inferences about contacting customers who are about to go off contract can be made from the graph.

1. Contacting these customers may actually increase churn slightly.
2. Contacting these customers will increase contract renewal, but this increase is partially offset by a slight subsequent drop in contract renewal below the control rate after the offer expires
3. Looking at the last few months of contract renewal one could conclude that you get more renewals be not contacting them at all; however, this is not the case. There was an additional campaign that drove the contract renewal (and apparently a slight increase in churn) in the last few months. The later campaign did not keep the original control group.

The ideal would be to contact only those customers who renewed who would not have otherwise have done so and ignore the potential churners. This study was done to identify the most profitable customers to contact.

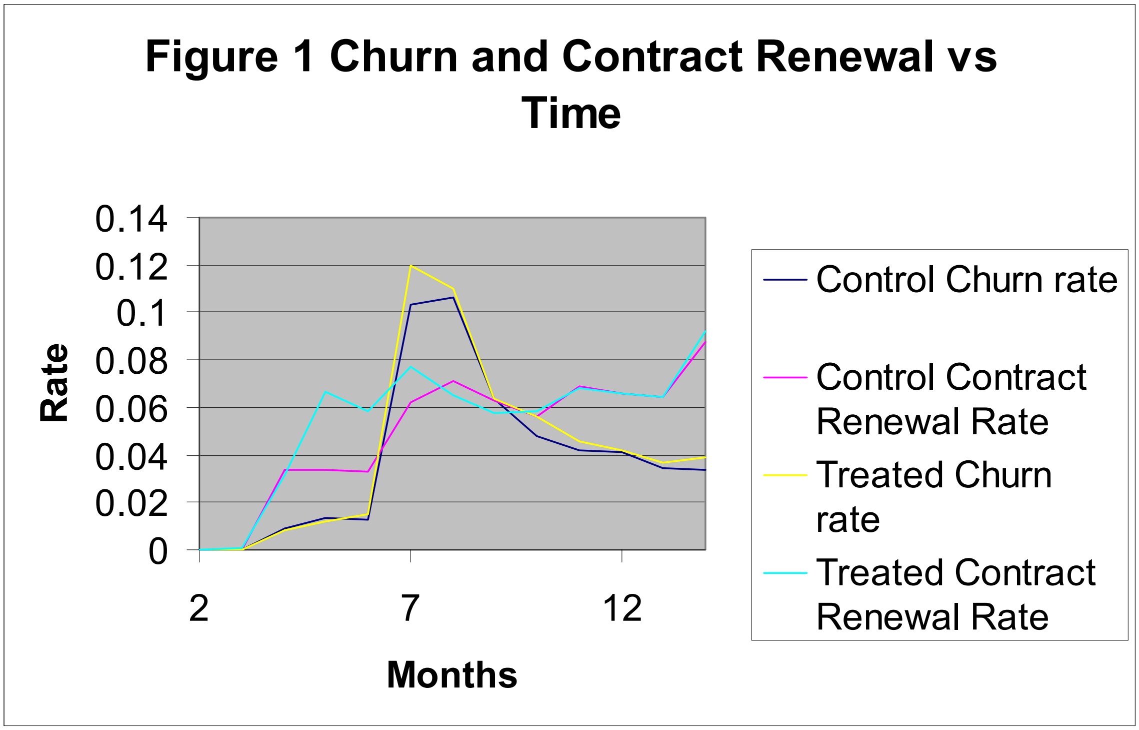

However the campaigns did have a long term effect in increasing the number of customers under contract and hence reducing the fraction of long tern churn in the treated group. The overall positive campaign result was a significant increase in the number of customers who chose to extend their contracts in the treated group over the fraction who extended their contracts in the control group. (See Figure1 – Comparison of rate for contract renewal event in treated vs. control customers)

It was decided to use a proportional hazards model to give an insight into the churn and contract renewal that occur as customers come off existing contracts. Specifically, it was decided to use a proportional hazards approach to model the effect of Cingular's monthly retention campaigns for the second quarter of 2004. These were chosen because the selection parameters were similar, and the customer response was typical.

METHODS:

Data sets were defined using Nov 30, 2003 as an arbitrarily defined left truncation date and Dec 31, 2004 as the censoring date. The data structure included one observation per month per subscriber with a month value of "t" going 0 to "n" with 0 being November 2003 and incrementing "n" by one each month. Cingular Wireless data for customers with contract expirations in the 2nd quarter of 2004 were chosen as the source for data used to train and validate the models.

Two competing hazards were defined: Churn and Contract Renewal. These were assumed to be mutually exclusive categories so as to simplify the model, although there is some churn that occurs after contract renewal. After this first approach a second pass was made with churn as a single hazard with contract renewal as a covariate and with contract renewal as a single hazard. This did not give any particular improvement in the model's ability to predict churn. Data sets in the following format were set up in SAS. (See Fig.2)

FIGURE 2 DATA STRUCTRUE FOR ANALYSIS

Obs	srv_accs_id	g	t	tot_accs_chrg_amt	tot_air_chrg_amt	call_tot_qty	treatment	tot_roam_amt	min_tot_qty	rplan_change
1	328	0	0	49.99	50.40	330.00	test	0.00	1114.00	0
2	328	0	1	49.99	0.35	249.00	test	0.00	903.00	0
3	328	0	2	49.99	0.00	135.00	test	0.00	354.00	0
4	328	0	3	49.99	6.65	236.00	test	0.00	726.00	0
5	328	0	4	49.99	0.00	200.00	test	1.29	554.00	0
6	328	0	5	49.99	0.00	203.00	test	0.00	693.00	0
7	328	0	6	49.99	0.00	215.00	test	0.00	526.00	0
8	328	0	7	49.99	0.00	177.00	test	0.00	472.00	0
9	328	0	8	49.99	0.00	160.00	test	0.00	510.00	0
10	328	0	9	49.99	0.00	220.00	test	0.00	691.00	0
11	328	0	10	49.99	0.00	231.00	test	0.00	563.00	0
12	328	1	11	0.00	0.00	246.00	test	0.00	860.00	0
13	362	0	0	76.99	0.00	590.00	control	0.00	1992.00	0
14	362	0	1	76.99	0.00	646.00	control	1.29	2746.00	0
15	362	0	2	76.99	0.00	696.00	control	1.29	3358.00	0
16	362	0	3	76.99	0.00	838.00	control	0.00	4387.00	0
17	362	0	4	76.99	0.00	1015.00	control	0.00	3296.00	0
18	362	0	5	76.99	0.00	853.00	control	1.29	3351.00	0
19	362	0	6	76.99	0.00	938.00	control	0.00	3496.00	0
20	362	0	7	76.99	0.00	791.00	control	0.00	3414.00	0
21	362	2	8	76.99	14.82	907.00	control	2.58	2755.00	1
22	363	0	0	29.99	0.00	36.00	test	0.00	75.00	0
23	363	0	1	29.99	0.00	35.00	test	0.00	88.00	0

SAS code to produce this structure is not included here since this code will depend on the data source

The variable “g” is the target variable. The value of “g” is the end event with churn represented by 1, contract renewal represented by 2 and censored represented by 0. “t” is the month relative to the left truncation date. In the single hazard models the data structure was essentially the same, but the target variable was only allowed to assume two values. This data structure can be analyzed in either SAS STAT or Enterprise Miner. Both were used.

Cubic splines were defined in the logistic model to help with the curve fitting. Knots were defined at quartile points, and knots were defined for every month of the study using both stepwise selection and a simple quartile assumption to determine the most relevant knots.

The SAS code to do this is relatively simple with the key being the selection of the LINK = GLOGIT option on PROC LOGISTIC See Figure3. The SAS/STAT manual gives all the details on how to code this. This is a typical run. A number of relevant variables were tested. The csb: variables are splines.

FIGURE 3 PROC LOGISTIC CODE

```
ods output parameterestimates = pe2;
dm 'clear lst';
proc logistic data = churn5;
class cat4 cat1 cat2 treatment;
model g(ref = '0') =  tot_accs_chrg_amt t csb1 csb2 csb3 csb4 csb5 cat2
cat4 cat1 tot_air_chrg_amt  call_tot_qty  treatment
tot_roam_amt  min_tot_qty   t*tot_accs_chrg_amt
t*min_tot_qty  rplan_change            /link = glogit;
run;
```

RESULTS:

The results from these approaches approximated the overall shape for churn immediately after contract expiration reasonably well, although the model showed churn decay to be slower than was the actual. Figure4 is the best fit to the time response curve of a number of models using PROC LOGISTIC and different variables to predict customer behavior.

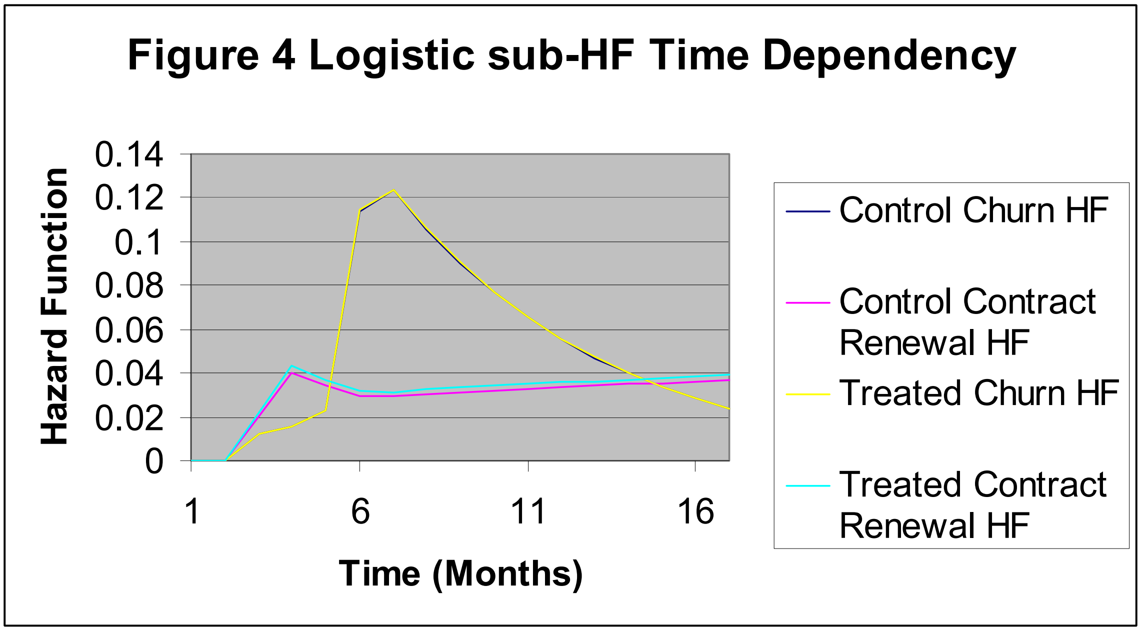

However, the contract renewal was not as good a fit to the measured response (Figure 1) since the large difference between treated and control seen in the actual was not modeled.

Neural nets are considered more flexible curve fitting models than regression, so neural net was tried to attempt a better fit to the observed churn and contract renewal rates. PROC NEURAL from SAS/Enterprise Miner was used to generate a neural net model using the code in Figure 5.

FIGURE 5 PROC NEURAL CODE

```
proc dmdb data = train  out = dtrain dmdbcat = dctrain;
var t tot_air_chrg_amt call_tot_qty rplan_change tot_roam_amt  tot_chrg_amt
 missing_bd_flag tt_change;
class treatment cat1 cat2 cat4 tech_type  age_cat g (desc);
run;

proc neural data = train dmdbcat = dctrain;
input  treatment cat1 cat2 cat4 age_cat tech_type /level = nom;
input   t tot_air_chrg_amt call_tot_qty rplan_change tot_roam_amt
tot_chrg_amt
        missing_bd_flag tt_change /level = int;
target g /level = nom;
archi mlp hidden = 6;
nloptions absconv = 0 absgconv = 1e-16;
train tech=quanew;
code file = '~/april_ooc_neural_score_10_21.txt';
run;
```

Neural net gave the following approximation (Figure6) of the time dependency of the hazard function for churn and contract renewal.

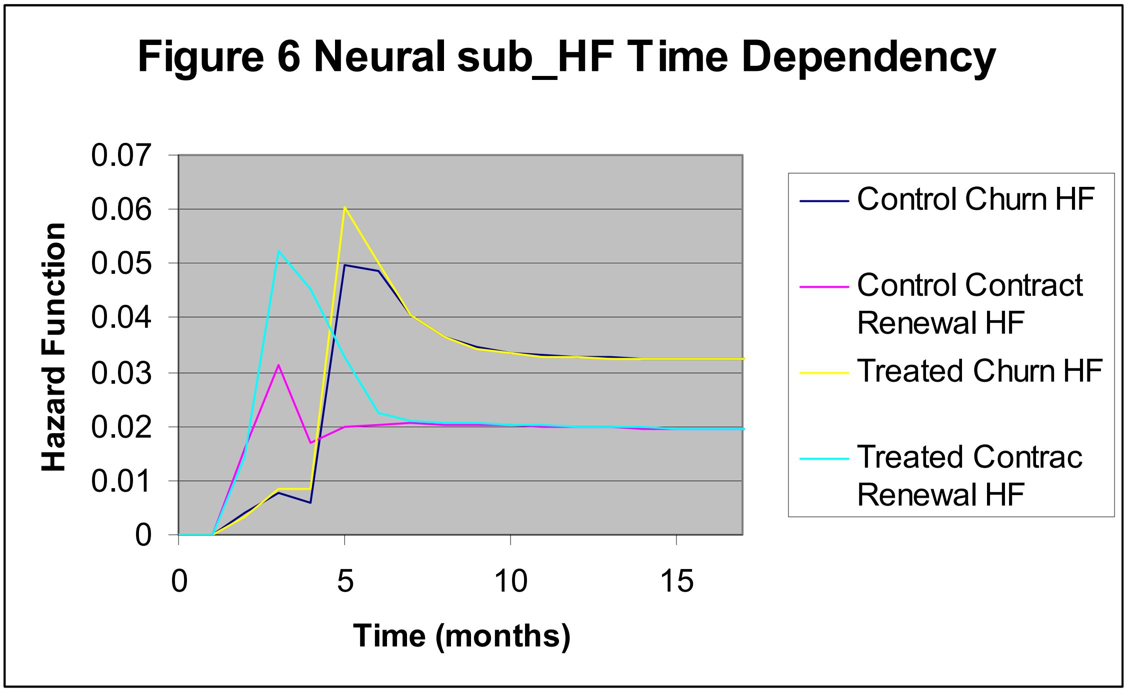

Neural net gave the best modeled approximation of the actual response shape for both churn. See Figure7.

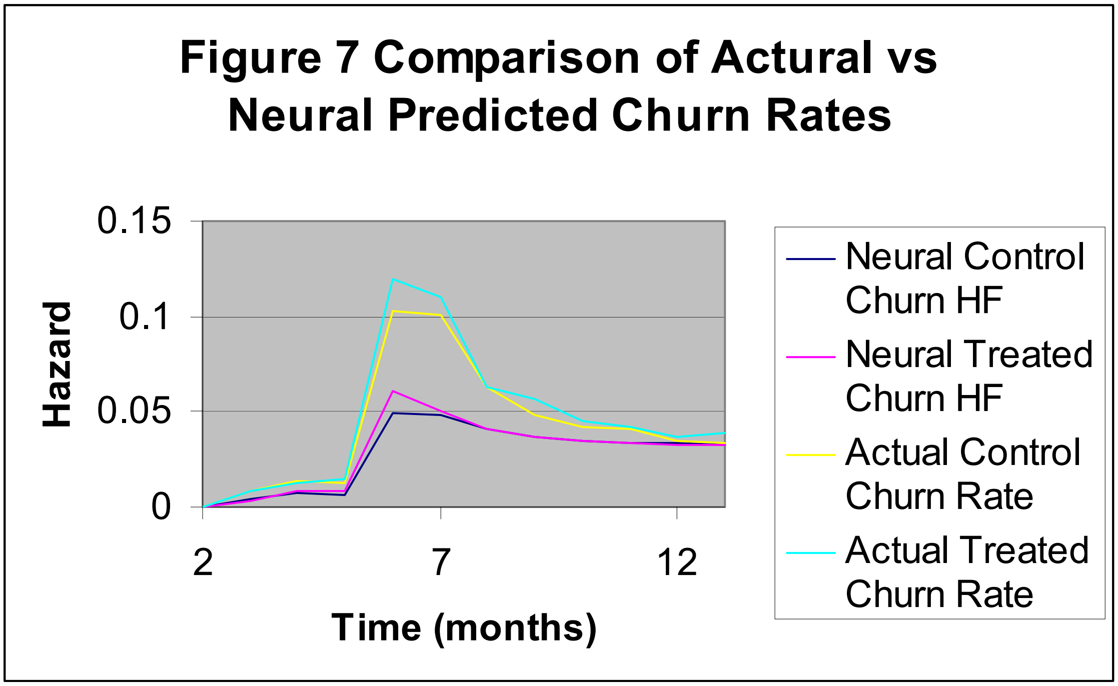

The shapes are quite similar, although the neural model lacks the magnitude of the actual data.

Next the models were evaluated using a validation data set. June expiration cohort as well as the April expiration cohort only were used as validation data sets ability somewhat from non-existent to low.

The posterior probabilities were calculated for the validation data using the model developed on the training data averaged for the months of interest – those months just prior to and immediately after contract expiration - to give a mean posterior probability. Code to do this can be found in Potts' course book for the SAS Survival Data Mining Course. The predictive abilities of the models were evaluated visually by the standard marketing percent captured response chart and by the numeric KS (Kolmogorov-Smirnov) score.

The specific model from Figure7 while giving a promising hazard vs. time plot had virtually no predictive ability for churn as measured with the standard percent captured response chart, showing a line that varied slightly above and below the random response diagonal and gave a KS score of 0.02. The predictive ability for contract renewal was slightly better with a KS score of 0.18, but still in the "poor predictor" range.

The model was rerun using additional predictor variables such as feature codes (e.g. unlimited nights and weekends, unlimited mobile to mobile, phone technology type and how recently the customer had purchased a new phone.) This improved the predictive ability of the model giving a modest lift as can be seen in Figure8.

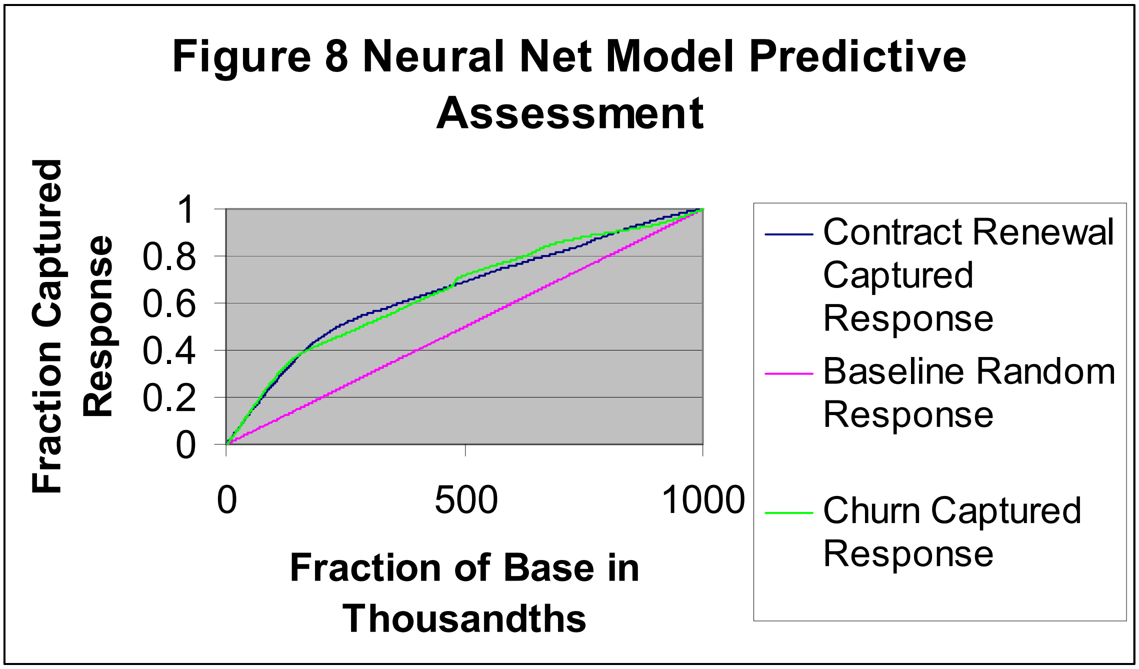

The additional predictors gave a KS = 0.29 for contract renewal and KS = 0.26 for churn.

Various regression models were also evaluated. These gave KS scores somewhat lower than the scores for the neural net models. With KS scores ranging from 0.16 to 0.22. Generally the contract renewal prediction was somewhat better than the ability of the model to predict churn.

CONCLUSIONS

The PROC NEURAL models gave the best predictive scores. At that, the best of the models were not highly predictive with KS scores generally in the 0.25 to 0.31 range. However, there are several significant churn predictors that we currently do not have readily available that we would like to add. Specifically network experience at the mobile level is something that we would like to add to our models, when this becomes available

Further there is no variable on the "richness" of the offer. It has been observed that more financially attractive offers gather more takers, so adding some sort of financial incentive predictor would most likely improve the predictive value of the models. Proportional hazards using either proc logistic or proc neural is another approach in addition to the traditional binary logistic regression to modeling campaign results and generating predictive scores to weed out unlikely responders and to include probable acceptors. Cingular™ Wireless LLC has traditional logistic regression churn models which do a significantly better job of predicting churn. With tuning the proportional hazard models should be able to approximate this level of predictive ability.

TRADEMARK CITATION:

REFERENCES

Potts, W., 2003, Survival Data Mining Predictive Hazard Modeling for Customer History Data, SAS Institute, Cary, NC

SAS Institute,1999, SAS/STAT User's Guide, Version 8, SAS Institute, Cary, NC

ACKNOWLEDGEMENTS

The author would like to express his appreciation to the following:

Stephen Butler, MKIS Director, Cingular™ Wireless, LLC for making the data available and allowing these results to be presented.

Will Potts, Data Miners, Inc., for an excellent introduction to this approach in the SAS® Survival Data Mining Course

Jerry Musial, Cingular Wireless, LLC. for proofreading the draft.

CONTACT INFORMATION:

Charles Manahan
Cingular Wireless, LLC.
Glenridge Highlands Two
5565 Glenridge Connector
Atlanta, GA 30342