REDUCE CUSTOMER CHURN USING CHURN MODELLING

With Cliently

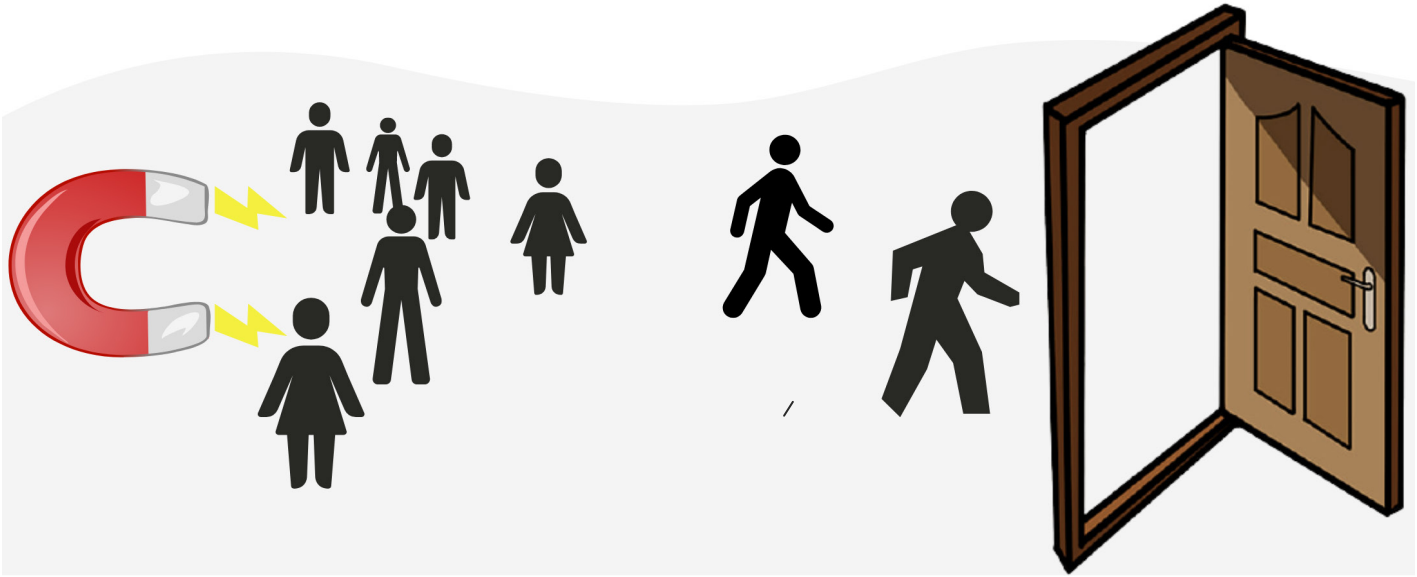

TABLE OF CONTENTS

01

INTRODUCTION

INTRODUCTION

As subscription services have risen to dominance in recent years, many companies have begun to realize that they need a way in which they can predict customer churn in advance of it happening in order to be able to do something to reduce it. Today, this problem is much easier to solve using Machine Learning with its ability to predict what will happen next all without the bias of human judgement or faulty intuition. We also have the ability to predict when customers will churn (using survival analysis) and what actions we can take to reduce the chances of churn.

02

WHAT IS CUSTOMER CHURN AND REVENUE CHURN?

WHAT IS CUSTOMER CHURN AND REVENUE CHURN?

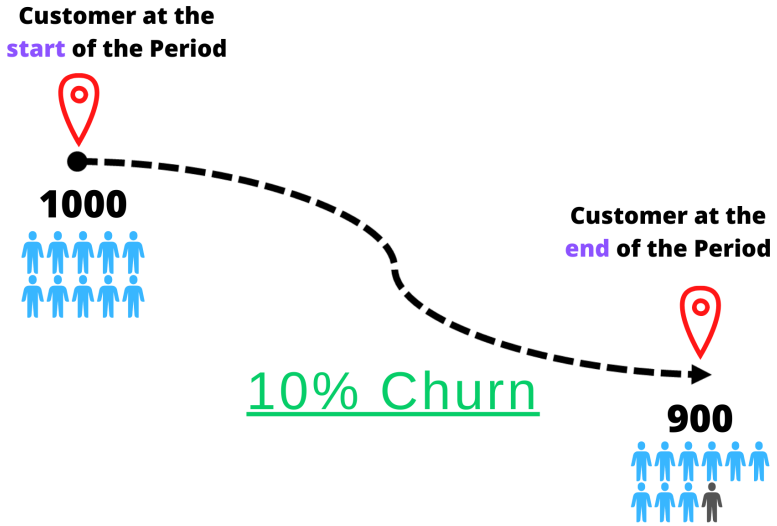

Customer churn can be defined as the people that unsubscribe to your product or services in a certain time frame.

REVENUE CHURN

Revenue churn is the direct measure of lost revenue as a result of customer churn and downgraded subscriptions . It can be measured in terms of lost monthly recurring revenue (MRR)or annual recurring revenue (ARR)

The causes of revenue churn can be one of the following:

- Lost contracts or cancellations
- Downgrades to a cheaper plan
- Customer acquisition by a competitor
- Customer going bankrupt

Except for the last point, all the above causes are in your control and largely depend on the quality of service you provide to your customers.

CUSTOMER CHURN VS REVENUE CHURN

In order to evaluate the strength of your business, you have to take into account both customer churn and revenue churn. Customer churn would tell you how good you are at retaining customers whereas Revenue churn would show how good you are at retaining your revenues.

Customer churn is more relevant in the beginning when every customer starts with their initial purchase of your product. With time, each customer starts generating different revenues. This is where the Revenue churn calculation becomes more relevant for each case.

03

WHY CHURN IS A BIG PROBLEM

WHY CHURN IS A BIG PROBLEM

Significant dollar amounts are spent by companies to acquire customers. When a customer churns, or leaves, you lose business thus impacting your recurring revenue. There are many reasons why keeping existing customers are more important than new client acquisitions. For example, the cost of acquiring a new customer is far more than retaining an existing one. In fact, it's between 5 and 25 times more!

Customer relationships also become more valuable over time, as operating costs to serve them decline and their inclination to purchase increases. Loyal clients may also refer new customers through their social networking interactions. In concrete terms, research by Bain & Company found that just a 5% increase in customer retention improved profits by over 25% in some sectors. Another McKinsey report estimated that reducing customer churn could increase earnings for a company by as much as 9.9 percent.

LOST REVENUE

The problem confronting management is that it is very difficult to determine which subscribers left the company and for what reasons. 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.

“Companies lose $1.6 trillion per year due to customer churn!”

According to the Forrester, it costs “5 TIMES MORE to acquire new customers than it does to keep an existing one.”

FUTURE REVENUE

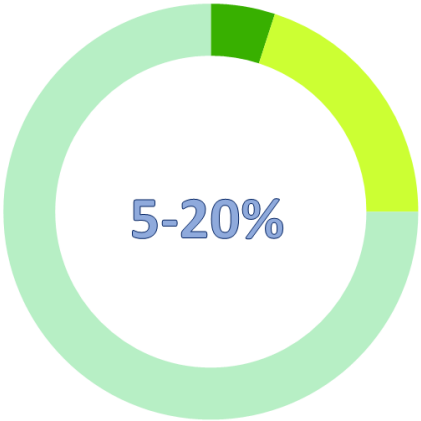

Probability of selling to a
New Prospect

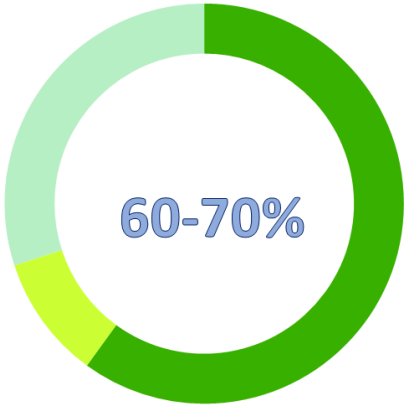

Probability of selling to a
Existing Customer

According to Gartner, a staggering 80% of a company's future revenue will come from just 20% of its existing customers.

Meanwhile, Marketing Metrics claims that the probability of selling to an existing customer is 60-70%, and only 5-20% to sell to a new prospect.

04

HOW DO COMPANIES KNOW WHICH CUSTOMERS ARE at risk of churn?

HOW DO COMPANIES KNOW WHICH CUSTOMERS ARE AT RISK OF CHURN TODAY?

Identifying at-risk customers is one of the most popular churn tactics for B2B companies. In fact, 35% of B2B of organizations have used this tactic to successfully reduce customer churn.

Most companies are identifying at risk customers by looking for red flags; monitoring community forums and third-party review sites; and through continuous communication with customer success teams.

MOST COMPANIES USE ONE OF THESE WAYS

Looking for Red Flags -

- Customer Feedback forms like NPS are collected and they can give a sense if a customer is upset or not. Unhappy customers are likely to churn. The problem with this is that many people choose not to fill any information or fill wrong information.

- Tickets, issues raised by customers - If they are too many and indicate problems, it means customers won't likely stay.

- Negative Sentiment expressed in email, phone calls, chat etc.

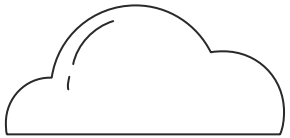

Monitoring Community Forums and Third-Party Review Sites -
Community forums and third-party review sites are terrific resources for companies. If you notice the same users bashing your products again and again, then you know these people are at a higher risk of churn than those who are posting positive comments.

The same goes for third-party reviews. The people who give your company the lowest ratings are the ones who are most likely to leave your business.

Continuous communication with Customer Success Team-
This is especially true for big accounts, where individual representatives are assigned a few big-ticket customers. They monitor these accounts and reach out to users who are unhappy. That way, your team can immediately clear any roadblocks for customers instead of having them reach out to your support team.

WHAT COMPANIES CAN DO AFTER IDENTIFYING CUSTOMER CHURN

Once the companies have identified at-risk people, then they may offer incentives or give those customers more attention. A recent survey found these that companies frequently use these feature importance characteristics -

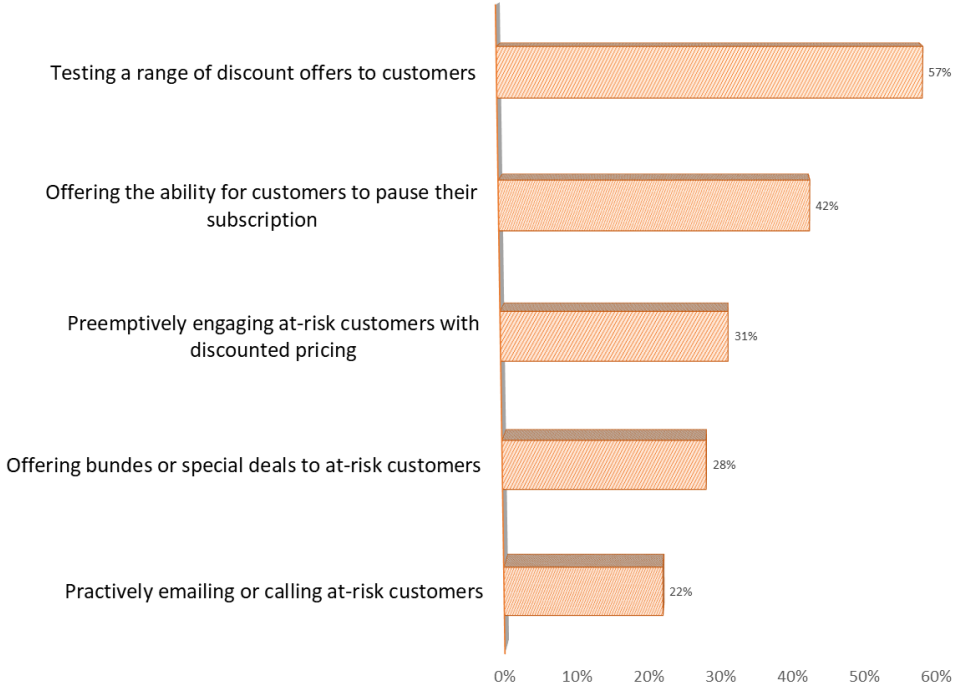

05

WHAT COMPANIES ARE MISSING TODAY

Customer Churn Modelling

WHAT COMPANIES ARE MISSING TODAY

Customers vary in their behaviours and preferences; this is where churn modelling is most useful. Customer Churn Modelling helps you identify which of your customers are most likely to stop engaging with your business. The models take historical data of customers as their source of input and the result is an output probability of churn for everyone.

The customers likely to cancel contracts or subscriptions can then be offered customized incentives designed to deter cancellation. Analysis can also provide insight into why clients are leaving, and what they need from you in order to continue engaging with your brand over time.

Detecting causes of churn that lie within the sphere of influence of the company also enables it to eliminate them in the future. For instance, the analysis may reveal that factors as different as long wait for customer support and bad billing procedures are the root causes of churn. Another example can be, especially in product-based companies, there could be an issue with the product that is blocking customers' ability to come back easily. This could result in product improvements (or potential new features) to act on to prevent attrition.

06

HOW TO GET STARTED

HOW TO GET STARTED

We should do a basic analysis upfront to decide which users should even be considered in the churn analysis.

Some questions to answer internally are -

1) If someone used a product on a free trial or a very limited time, should they be taken into churn analysis?

2) What is the threshold or a period, after which a customer should be taken into attrition study?

3) What is organization's definition of churn?

4) Often times you will have monthly, quarterly, and annual subscriptions. So, are all these cohorts different? Do you want to have four separate churn models? Or one model for all? Or a model for the subscription with most customers?

5) What about involuntary churn or factors that are not in your control? For example, somebody moving out of the country, and your services are not available there. Are these people to be included in the study?

6) Do you want to take low or negative profitability people into your study?

Whatever assumptions, you decide to come up with, it will play a role in creating your data for Machine Learning and hence the results would be different too.

07

DATA CREATION

for Churn Modelling

DATA CREATION FOR CHURN MODELLING

It is a good idea to conduct analysis and build models for market segments that exhibit some broad commonality. The minimum data required to predict churn is simply customer transactional details and the date or time of their last and first interaction with your company. While this may not be incredibly detailed, it would allow you build basic level predictions.

The reality is that adding additional data on top of this minimal data set is recommended and highly encouraged. The more the sources and data included, the better the churn predictions will be. Since ML model performance and the quality of insights generated depend on the quality of data, you will also want to make sure that all data points are presented in a consistent form suitable for building the models.

INFO YOU SHOULD BE COLLECTING

There are 5 types of information you should be collecting and using:

- Customer demographics such as their basic information like job, industry, age, gender, location, education, income, etc.
- Contractual data like the terms of the contract, duration of contract, etc.
- Billing and usage data such as usage patterns of your product by the customer
- Events-type data such as the dates when there is significant event in the customer's life
- Customer interactions such as phone and email transcripts

08

OUR APPROACH

OUR APPROACH

Our team of data scientists have multiple meetings with the marketing and sales teams of companies we work with. During the back-and-forth discussions, we develop a full picture of our clients' needs.

We then build unique churn prediction and propensity models for our clients. A customized and unique predictive AI model is developed analysing historical data over large periods of time. This is tailored to each clients' needs.

Further, once the infrastructure is set up, you receive real-time predictions. Our models are learning real-time from data and the moment they flag somebody as a risk for churn, notifications are sent.

OUR PREDICTIVE CHURN METHODOLOGY PART I

Understand

- Build and understanding of business problem, create objectives and determine how the churn predictions will be used.

Define

- Develop a definition for churn including input from stakeholders. For instance, who should be included in churn modelling and at what point they considered to be churners?

Identify Data

- Which data will you want to use? At a minimum, we require Customer ID and their last transaction date. Your data will then be synced in one place within our app.

Conversion to one-row-per-customer data

- Often data is transactional, and we have to form it into a one-row-per-customer format in order to proceed with Machine Learning.

Data Cleaning

- Our algorithms clean the data, mostly making it homogenous data. This process includes deleting duplicate or irrelevant observations, correcting misspellings, filtering unwanted outliers, handling missing data, and standardizing/normalizing data.

Feature Engineering

- Much of the current information cannot be used as it is in AI algorithms. Algorithms understand numeric information. For example, the date column has to be broken down into day of the month, the value of the the month, year, number of days from the nearest holiday etc. for Machine Learning to understand. Another example is of text column, which has to be converted into sentiment score, subjectivity score, text topics. Cliently's Automated Feature Engineering automatically do it for you, bringing down project duration from weeks to a mere few days.

OUR PREDICTIVE CHURN METHODOLOGY PART 2

Building Machine Learning models	• Our autoML code builds dozens of models and finds the one with the highest accuracies on the test dataset. We use many new algorithms such as Catboost, XGBoost, LightGBM, as well as traditional models like Random Forest, Neural Networks, and Ensembles.
Hyperparameter Tuning	• The same kind of machine learning model can require different constraints, weights or learning rates to generalize different data patterns. These measures are called hyperparameters and must be tuned so that the model can optimally solve the machine learning problem.
Predictive Probability of Churn	• We have a probability of churn for every customer now. Anybody with a probability of more than 0.5 is considered as churned or in jeopardy of churning.
Insights	• We tell you which variables are most important in determining churn and how you can reduce churn in the long term overall.
AI triggers	• We tell you the best possible action to take for every customer that will enhance chances of retention.
Predicting the time of churn	• Knowing probability isn't enough, we also need to know when exactly customers will churn. This is done using survival analysis.
Iterate	• We iterate the effectiveness of models. Over time, we see if the predictions are right and if they are missing and by how much.
Deploy	• Ensuring a continuous churn prevention strategy - whenever any new or old customer is in risk of churn, our dashboards notify you and suggest action plans.

IN THE END, THIS IS WHAT YOU WILL HAVE FOR EACH OF YOUR CUSTOMERS

Churn Scoring GPS

Dontae Little
Vice President of Marketing
XYZ Inc.

Company Revenue >$6M

Probability of churn > 0.8

Dontae has still $40,000 future worth left

Dontae will churn in next 6 months

Why is this customer churning?

Dontae dislikes long waiting for support & is unhappy with pricing plans.

Action plan

Since she has a high return in future, prioritize her queries and make a customzied pricing plan according to her needs. Needs urgent escalation.

09

COMPARING OUR ACCURACIES TO OTHER

autoML solutions

OUR ACCURACIES ARE HIGHER

Our accuracies are very high because of our algorithms.

From day one, our focus has been on data cleaning and feature engineering. The kind of information we create from current information is unparalleled. We have 50,000 lines of code to clean multiple types of data, take various permutations, consider a number of combinations into account, and create more variables from the current ones.

We have learned best practices from Data Science competitors in Kaggle and have implemented them into ours. Not to mention, our autoML product performs in the top 1% of rankings in every competition including Kaggle

HOW WE COMPARE

A few automl solutions like google automl, Microsoft azure, aws sagemaker and datarobot could be good solutions if you have data scientists and engineers to set up infrastructure. They may be more cost-effective as well.

We have compared 100+ datasets and 85% of the time, we performed better than these four autoML platforms in terms of accuracies (weighted f1, mcc, rmse)!

How do we achieve higher accuracies? It stems from the intensive data cleaning and feature engineering we discussed in the above sections. If we were just to build models on raw data, our accuracies would be similar.

10

THE NEXT STEPS

after churn modelling

THE NEXT STEPS

Merely uncovering the probabilities of finding churn is just the starting point. As an executive, your next questions should be:

- Okay, now that we know the probabilities of churn and when they will be churning, what do we do next? How do we reduce that churn? Tell us the action plan to reduce it immediately and then the long-term strategy too.
- When exactly are the customers going to churn? Two people may have 0.8 probability of churning. But one may be churning within next year and other could be churning within the next five years. So, the strategies for both of them should be completely different.
- There are thousands of customers predicted to churn. Where do we apply our focus? Is there a way to prioritize? If we have to offer any incentives, to whom should we make incentives?

BEST POSSIBLE ACTIONS TO TAKE TO REDUCE CHANCES OF CHURN

Once you have that churn model in place, the job is only half complete. The next big thing is to take action to reduce churn. But where to start? What's the best way to tackle a potentially large pool of potential churners?

The next possible action can easily be determined. There are 2 ways to go about it – learning from the past churn management campaigns and the other is causal modelling.

Using AI, campaigns will be suggested that have the highest conversion rate for retention. It's using a combination of factors to segment each client based on demographics, engagement, previously sent messaging, and additional info within your data.

Let's take a closer look with some examples…

LEARNING FROM PAST CHURN MANAGEMENT CAMPAIGNS

Let's say you had these four campaigns in the past which you initialiated just before cancellation of contracts:

Action Plans (Campaign to stop Churn)
Phone call
20% Discount
2 months of free usage
Offer Extra Features

So now we will build machine learning models to understand what type of action plan will work on what type of customers the best. The simplest model can look like this.

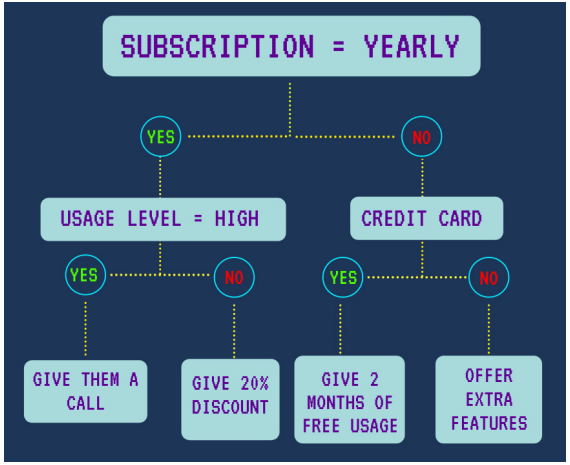

BEST POSSIBLE ACTIONS TO TAKE TO REDUCE CHANCES OF CHURN

CAUSAL MODELLING

The other way to determine best action plan (to reduce churn) is by discovering the most important attributes differentiating between churn and non-churners. Then we look at every variable one by one to produce a plan. This requires detailed study and investigation.

At Cliently, we give you action plans and recommendations to reduce immediate and long-term churn. We do this using Explainable AI or Causal modelling. It focuses on determining which factors effectively influence outcomes. This is significant in that it helps determine which kinds of churn reduction strategies are most effective.

SOME OF THE SNIPPETS FROM OUR ACTION PLANS - FEATURE IMPORTANCE

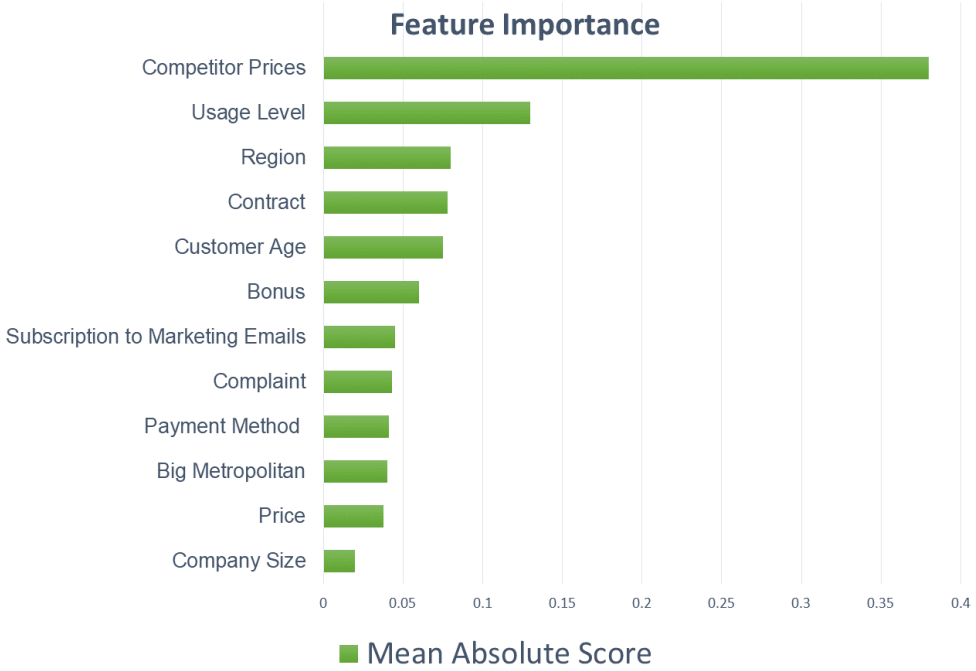

Feature importance charts can tell us which variables are creating churn.

BREAKDOWN OF INDIVIDUAL VARIABLES

A breakdown of important individual variables can help us understand how change values affect conversion.

For instance, looking at graph below, we can understand the type of contract and its effect on churn. The month-to-month contract sees the highest churn, so we would want to move customers to two-year contracts as much as possible.

SHAP PLOTS

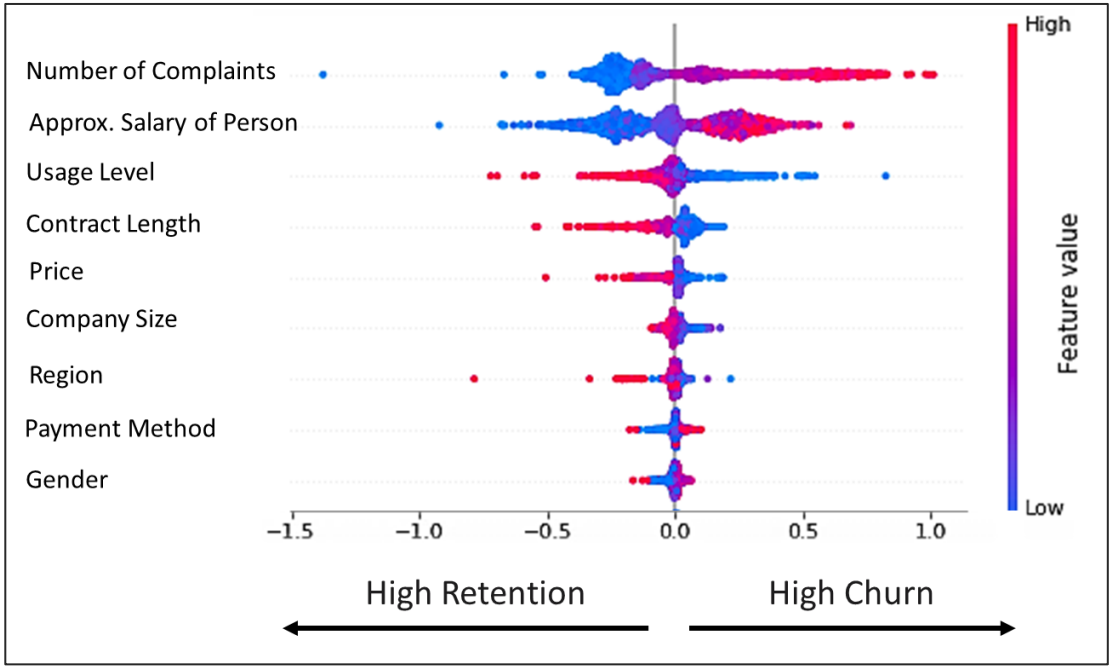

Shap plots can help us understand how increasing or decreasing variables affect conversion.

This graph indicates that as the number of complaints increase, the churn also increases. A counterintuitive insight is, the higher the salary of the customer, the more the chances of attrition.

WHAT-IF SCENARIOS

We give you **what-if** dashboards to play with to see if a customer behaved differently or if you intervened at certain step, what would've been their new probability of churn.

Why is this important?

This helps marketing teams understand:

- What they could have done differently to reduce or avoid churn for a specific customer.
- What can be done differently to increase the chances of the retention of current specific customers.

Having these answers will be especially more valuable for high LTV customers.

A SNIPPET OF WHAT-IF SCENARIOS

Here, the value of one means churn. and on the right-hand side, as you change the values of variables, you will see new probabilities of churn.

WHEN ARE MY CUSTOMERS GOING TO CHURN?

Knowing the probability of churn and the action plan to reduce churn, puts you well in the advantage, but in today's world, it isn't enough. That is when focus must be placed on an additional type of statistical modelling called Survival Analysis to predict the time at which a customer will churn.

Survival analysis is a series of statistical methods that deals with an outcome variable of interest such as time to an event. This is unlike a typical regression problem where we might be working with a continuous outcome variable (e.g., housing prices) or a classification problem where we simply have a discrete variable (e.g., Class I or Class II). In survival analysis, the outcome variable has both an event and a time value associated with it.

The analysis addresses questions such as "How long would it be before a particular event occurs?". Traditionally, survival analysis is used in medicine to study things like a drug having the ability to prevent a disease from occurring (event) and how long it can prevent a disease for (time).

SURVIVAL ANALYSIS

Survival analysis is applicable in situations where we can define:

- A ‘Birth’ event: in this case, this is the date a customer enters a contract with a company.
- A ‘Death’ event: in this case, ‘death’ occurs when a customer does not renew a subscription.

While the usual models are appropriate for binary classification problems (Will the customer churn or not?), survival analysis goes beyond that to answer when the customer will churn. If we do not expect a customer to churn right now, it does not automatically imply that they will never churn. If not now, there is a high probability that they might churn later. In such cases, survival analysis can throw light on such ‘censored’ customers by predicting their survival (or churn) at various points of time in the future. The ability to deal with ‘censorship’ in data makes survival analysis a superior technique to other regression and tree-based models.

SURVIVAL ANALYSIS - CENSORING

An important concept in survival analysis is “censoring”. Censoring occurs when the start and/or end of a measured value is unknown. For example, suppose our historical data includes records for the two customers below. In the case of customer X, we know the precise duration of the subscription because the customer churned in July 2020. For customer Y, we know that the contract started 24 months ago and is still active, but we do not know how much longer they will be a customer. This is an example of right censoring because we do not yet know the end date for the measured value. Right censoring is what we most commonly see with this type of survival analysis

Customer	Subscription Date	Subscription End Date	Subscription Duration	Is Subscription Active?
X	April 2017	July 2020	39 months	No
Y	October 2019		24 months	Yes

As illustrated above, we could move forward with a duration of 24 months for customer Y, but this would lead to underestimating survival time. This problem is alleviated when using survival analysis since censoring is considered.

OTHER TYPES OF CENSORSHIP

Other types of censorship include Left-Censorship and Interval-Censorship. The former refers to the scenario where we do not know when an individual became our customer because his ‘birth’ event happened before our observation time period.

In contrast, Interval-Censorship occurs when data is collected at specific time intervals that do not constitute a continuous observation period. Thereby, some life/death events could have missed out from being represented in the data.

SURVIVAL ANALYSIS TECHNIQUES

Survival Analysis can be done through various techniques such as:

- Life tables
- Kaplan-Meier analysis
- Survivor and hazard function rates
- Cox proportional hazards regression analysis
- Parametric survival analytic models
- Survival trees
- Survival random forest

Once you have figured out the best survival analysis techniques, you can easily answer questions like -

- When will 25% of my customers churn?
- In the next 3 years, which product or contract type will see the most churn?
- Which customers are expected to churn immediately in the next 1 year, 2 years, 3 years, etc.?

Finally, the most important metric from survival analysis is mean residual lifetime. The time remaining until a customer churn is called the mean residual lifetime. This is derived from the models. We won't be going into the mathematics of it as it can be a white paper in itself 😅. But here are some of the papers for reference - University of Maine paper and NCSU paper.

SURVIVAL CURVES

The Kaplan-Meier estimate is the simplest way of computing survival over time. The survival curve can be created assuming various situations. It involves the computing of probabilities of an event at a certain point in time and multiplying these successive probabilities by any earlier computed probabilities to get the final estimate.

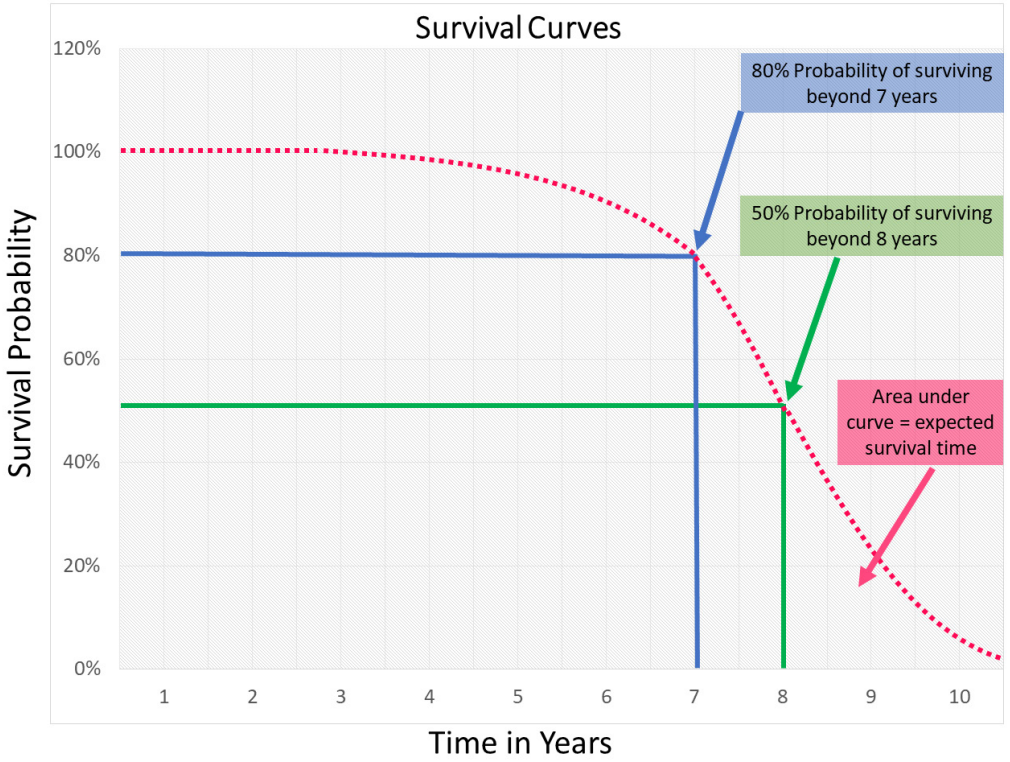

WHEN TO PRIORITIZE AMONG CHURNERS?

Should we contact everyone and offer incentives? Many businesses make the mistake of taking those who scored the highest (i.e., are most likely to churn) and targeting them.

Decide on whom you want to target. Not every customer will come back, nor do you have the time and resources to reach out to everyone. Advanced teams go one step further and use customer future value metrics in churn marketing campaigns. In other words, you should only spend time and resources targeting those churners or potential churners who will not only respond positively but also offer good returns in future.

> "If I offer an incentive to customers most likely to churn, they may not leave the company, but will it be profitable for me? The traditional method is focused on reducing churn, but we contend the goal should be maximizing profits, rather than only reducing churn.“
>
> - Sunil Gupta

CUSTOMER FUTURE VALUE

Customer Future Value is simply the future profits expected from customers until they churn. It is determined by 3 key factors:
(1) The length of time the customer will maintain a relationship with the firm from today onwards.
(2) Expected profits from this relationship in future.
(3) The discount rate used to determine the present value of the customer revenue stream.

Another big thing we like to do with our clients is to give them a list of churners (churned in the past) who have a high future value remaining. So, the marketing teams can now try to reach out to them when they are beyond a threshold future value (dollars).
We call these lost opportunities. Our experience has seen 20-30% of these high worth churners coming back when contacted personally.

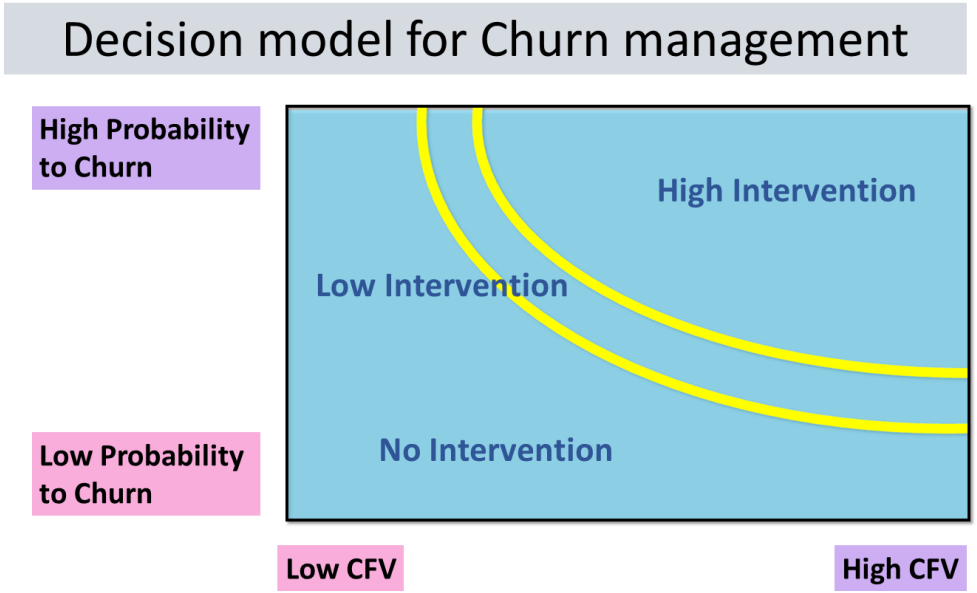

II

CONCLUSION

CONCLUSION

Predicting and preventing customer churn will not only save your company a lot of money on acquiring new customers, but also represent a huge additional potential revenue stream for your business.

Each company who globally develops customer service programs needs to have a line of sight to predict customer churn to reduce costs of non-conformance and ensure customer loyalty while remaining competitive. Early signals from AI and Machine Learning are very powerful and uncannily accurate. Our data science team at Cliently continues to explore the intricacies of churn modelling to accurately measure when and why your customers churn.

Attrition analyses offer some of the greatest return on investment for companies looking to expand their data-driven marketing strategies. The data needs are modest and approachable for nearly all firms, and the effect of saving marketing budget by targeting only those customers most likely to end their relationship with the firm and tailoring marketing decisions to proactively respond to customer concerns can be significant. The potential gains are augmented even further with the choice of the appropriate model and explainable strategies.

To stay motivated, just keep in mind that 82% of companies claim that retention is cheaper than acquisition and even a small 2% increase in retention can lower costs by as much as 10%.

THANKS!

Do you have any questions?

info@cliently.com

www.cliently.com