



WHITE PAPER

Essential Guide for Predicting Customer Churn

Exacaster:::

www.exacaster.com

Contents

INTRODUCTION	3
WHAT IS CUSTOMER CHURN?	4
IMPORTANCE OF CUSTOMER SEGMENTS	6
WHAT TO PREDICT?	7
PROACTIVE VS REACTIVE CHURN PREDICTION STRATEGIES	9
SELECTING ALGORITHMS FOR CHURN PREDICTION	10
PREPARING CUSTOMER DATA FOR CHURN PREDICTION	11
TRAINING ALGORITHMS AND MEASURING PERFORMANCE	12
MEASURING RETENTION PROCESS IMPACT	13
THE BUSINESS CASE FOR CHURN PREDICTION	14
ABOUT EXACASTER	15



INTRODUCTION

Each week, Exacaster runs hundreds of churn prediction and propensity models scoring more than 20 million consumers in various countries globally. After five years of extensive R&D, we believe we have distilled the science and art of churn prediction and identified the most effective techniques for carrying out this crucial service to our clients.

Our goal is to help our clients improve their customer retention by interpreting big data into a language they can understand and deploy into tactical strategies.

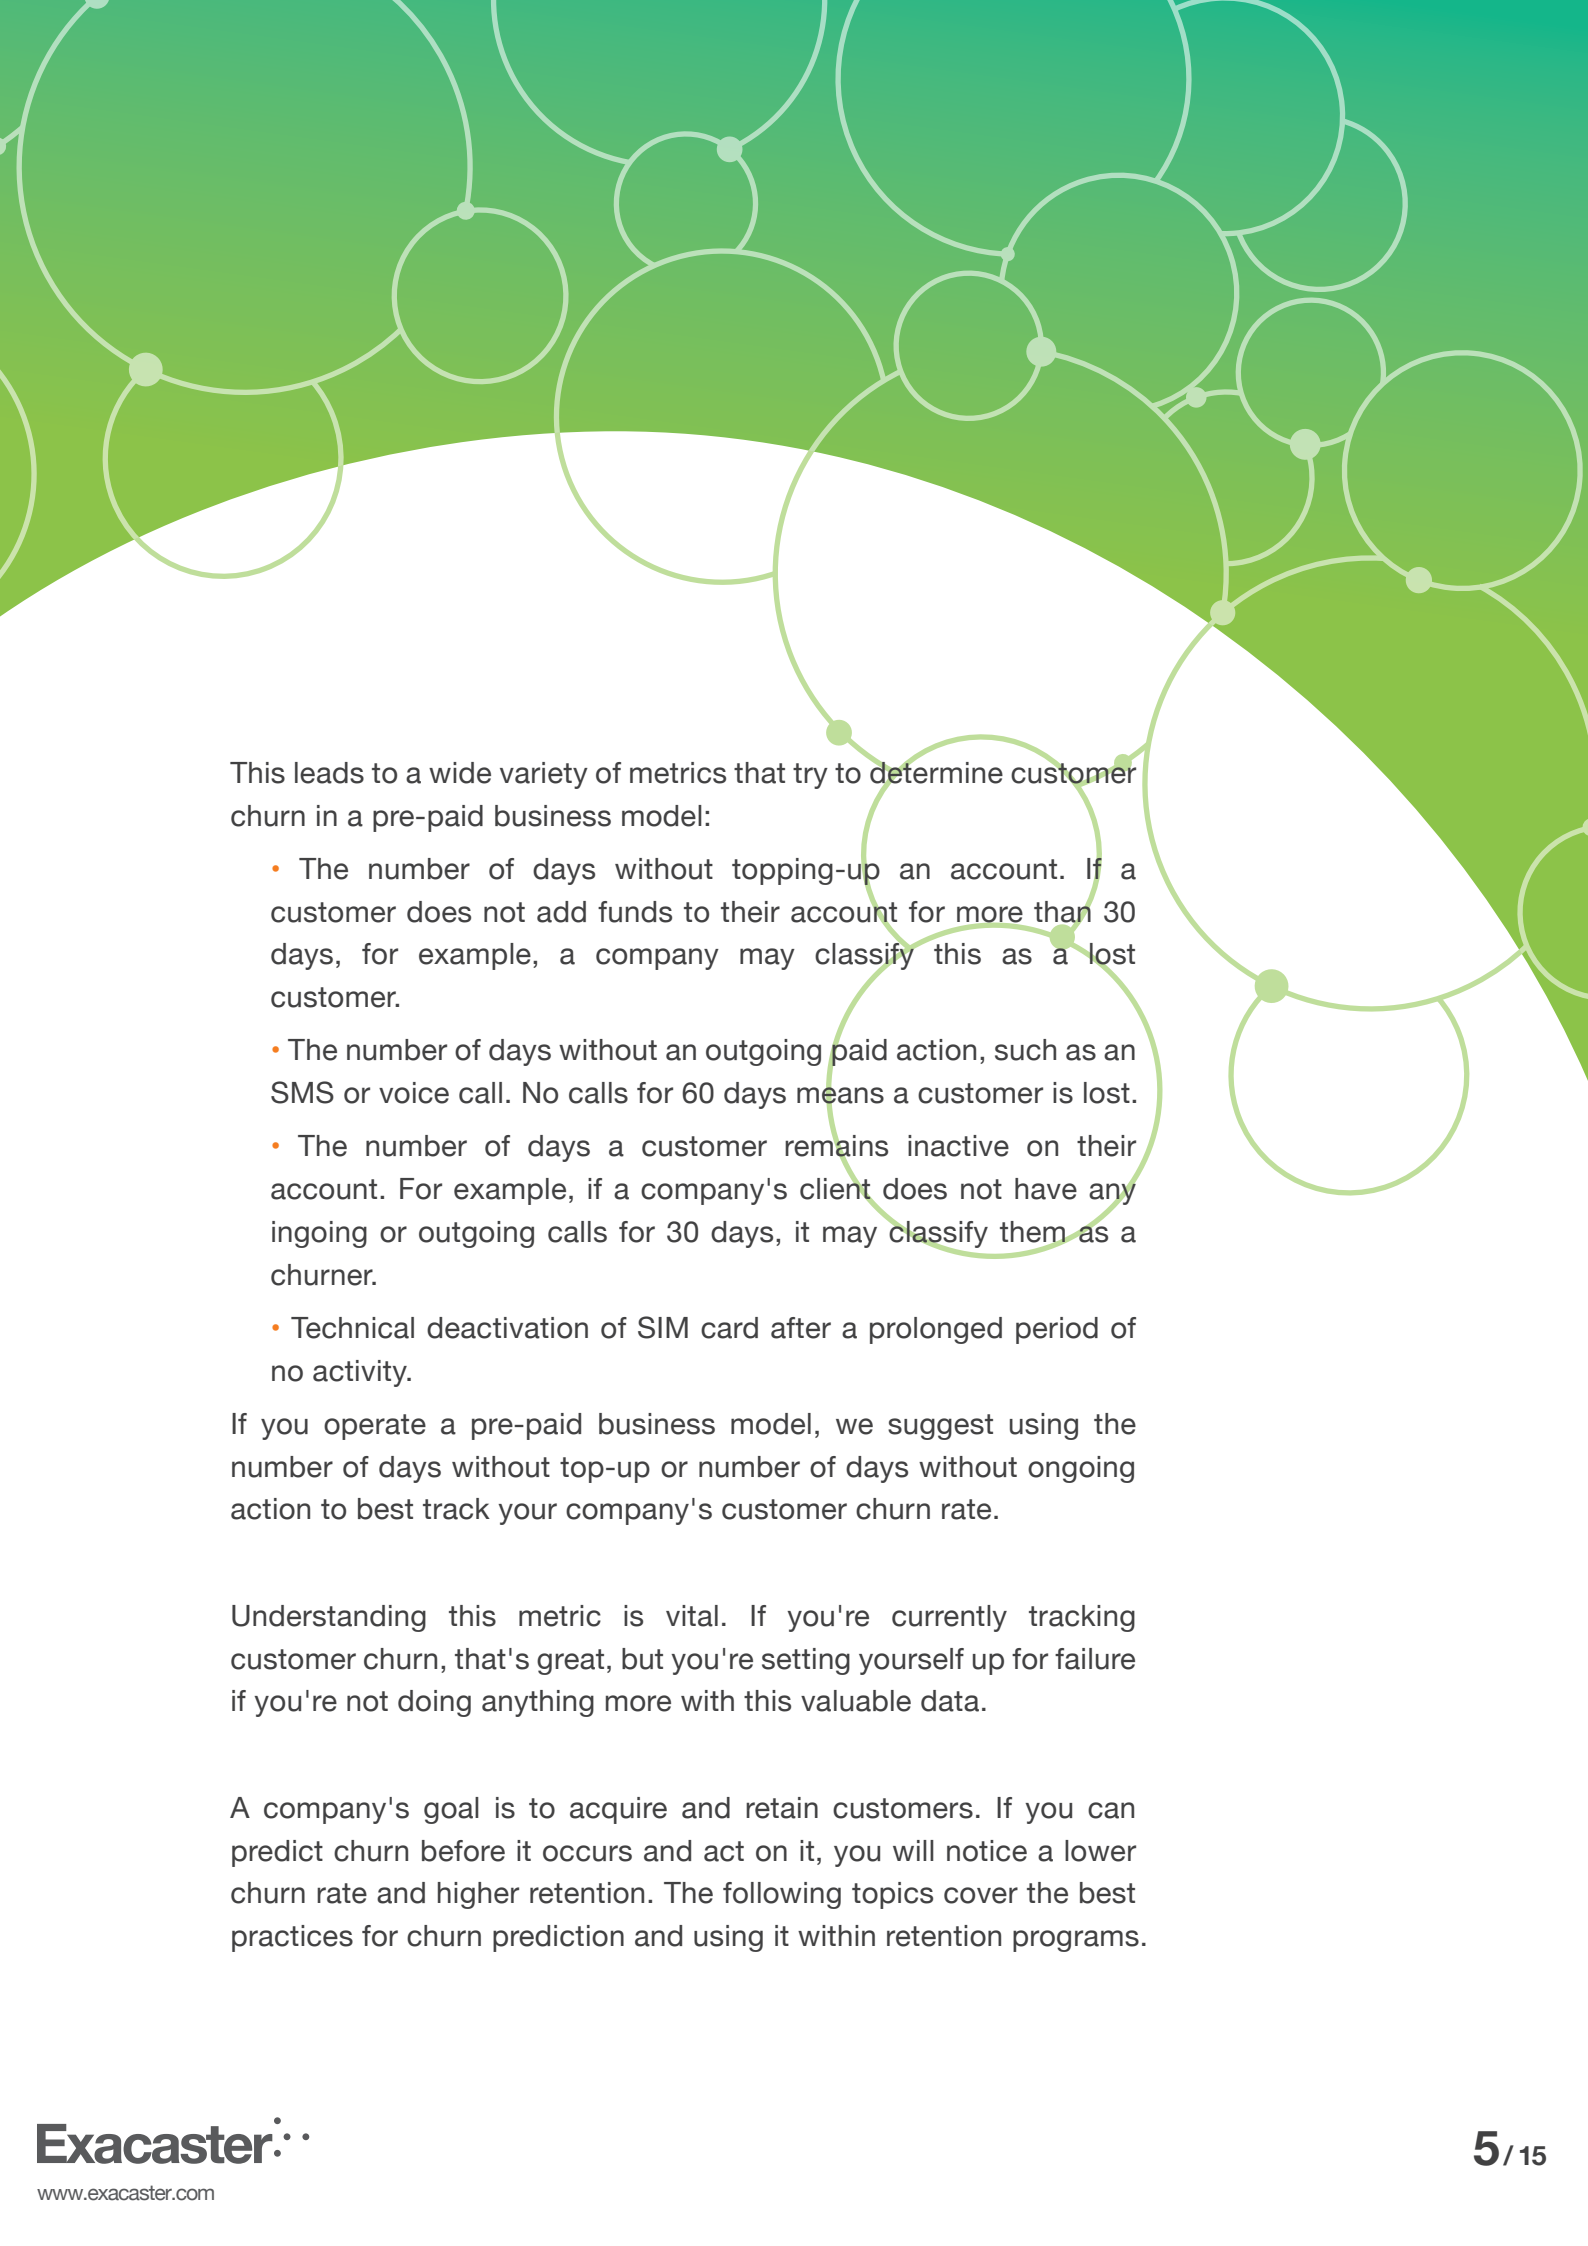
In this report, we share our best practices for using churn prediction and propensity models so that you and your company can improve your customer retention, too.

WHAT IS CUSTOMER CHURN?

Customer churn is a fancy word for losing a consumer. Customer churn affects all businesses. We happen to specialize in telecommunications, hence our material focuses on pre-paid and post-paid billing telecom companies.

Post-paid telecom service providers can easily understand churn rate because they work with contracts. When a customer terminates their contract, the company knows its lost a customer and can use this information to calculate its churn rate, as well as try to prevent this contract cancellation from happening.

With a pre-paid telecom service, churn rate is harder to measure because the customer does not sign a formal service agreement. This model uses a pay-as-you-go method and customer retention must be monitored closely by the company using many possible definitions of when a customer should be considered lost.



This leads to a wide variety of metrics that try to determine customer churn in a pre-paid business model:

- The number of days without topping-up an account. If a customer does not add funds to their account for more than 30 days, for example, a company may classify this as a lost customer.
- The number of days without an outgoing paid action, such as an SMS or voice call. No calls for 60 days means a customer is lost.
- The number of days a customer remains inactive on their account. For example, if a company's client does not have any ingoing or outgoing calls for 30 days, it may classify them as a churner.
- Technical deactivation of SIM card after a prolonged period of no activity.

If you operate a pre-paid business model, we suggest using the number of days without top-up or number of days without ongoing action to best track your company's customer churn rate.

Understanding this metric is vital. If you're currently tracking customer churn, that's great, but you're setting yourself up for failure if you're not doing anything more with this valuable data.

A company's goal is to acquire and retain customers. If you can predict churn before it occurs and act on it, you will notice a lower churn rate and higher retention. The following topics cover the best practices for churn prediction and using it within retention programs.

IMPORTANCE OF CUSTOMER SEGMENTS

Observing and tracking your customers' behaviors is the first step towards developing an efficient churn prediction model. Roughly 20% of your customers create the majority of your company's value according to the Pareto Principle. As a consequence, you should at least split your customer base into three value segments:

High-Value
Customers

Medium-Value
Customers

Low-Value
Customers

These simple segments are the foundation of your consumer retention strategy. According to our data, High-Value Customers have churn rates that are five to ten times lower than the Low-Value Customers. You cannot combine the data from these two segments to build a single retention campaign or predictive model because you will develop skewed results.

Businesses with the best results create separate churn prediction models for each customer segment. This ensures accuracy, and provides additional benefits for your firm, including:

- **Better understanding of your client base.** Knowing how your customers are grouped improves other business strategies outside of retention.
- **Pin-pointed data.** If you work with a data scientist or an analyst, they can allocate their time more efficiently if you focus on just one customer segment.
- **Clear focus for activities.** Thinking about one customer segment at a time simplifies implementation of CRM activities.

We take segmentation to the next step at Exacaster by using behavioral customer base segmentation (please see our white paper on this topic) and apply churn prediction algorithms to every segment separately. With our additional process, we can assist your team's Customer Base Manager (CBM) with segment-specific retention models and techniques leading to more focused offers for your client base.



WHAT TO PREDICT?

A practical use of churn prediction looks like this:

- Know your Current Churn Rate
- Develop a Desired Goal for Reducing that Number
- Have your CBM focus on your Current Customer Base that Show Similar Patterns to Those that Disengaged with Activities Designed to Re-Engage.

Sounds simple, right? Not really - the devil is in the details.

As you recall, the difference between post-paid businesses and pre-paid businesses is a contractual agreement between the consumer and the company. Predicting churn in the post-paid scenario is more straightforward than the pre-paid scenario because you're forecasting that your customer will terminate their contract (with or without a set expiration term).

In the pre-paid business model, companies need to monitor many more customer actions and come up with a better churn definition to predict successfully, because the right churn definition is not obvious at all.

Let's see what happens if we adopt a typical reporting definition of 60-day inactivity indicator as our churn measure. Using this definition, data scientists will produce very accurate churn predictions. However, once CBM tries to contact these customers with retention campaigns, he or she will immediately realize something is wrong. The customers are already inactive. Our models are predicting what is obvious: it finds customers who are no longer active for 30 days and simply predicts they will be no longer active for 60 days. So, the churn definition is wrong for our practical purpose.

Thus, we realize that churn prediction here needs a more specific definition.

We call this our *prediction target*. This should be an early indicator of specific consumer behaviors that lead to churn. For example, if a customer did not have any outgoing activity for the past five days, she is more likely to be disengaged moving forward. We will classify her as a potential churner. Therefore, customers with inactivity for five days is one of potential prediction targets. For new customers such an indicator may be the fact that this customer did not top up the account for a second time within first two weeks.

Thus, a sufficiently skilled analyst can find a good definition that starts to work. However, many businesses fail to realize that such definitions vary within customer groups and all the way down to individual customer's usage pattern. High-Value Customers share the behavioral aspect of high volume consumption, which leads to increased spend on your services. When we dig deeper into this usage, we notice that usage patterns within the same segment of customers can differ significantly.

- One group of customers may use the service heavily every day. Therefore, a drop off of seven days without any outgoing activity would be a red flag for data scientists and possibly indicate churn.
- In contrast, the second group of customers may recharge their account each Sunday and spend all of their cash on international calls that day. For the rest of the week, they remain inactive. This inactivity is normal and will not lead to churn.

Thus, an important second conclusion must be made: If you try to use a single prediction target definition for every usage pattern, you will fail to predict churn accurately.

Thus, when your company defines your desired prediction target, you should spend the biggest effort on this. It's the single most important thing to get right.

Exacaster always develops personalized churn target definitions. We analyze each customer and define *typical* behaviors and *non-typical* behaviors that are an early indicator of churn on an individual customer level. We use these indicators as churn prediction targets.



PROACTIVE VS REACTIVE CHURN PREDICTION STRATEGIES

Once we help you develop an effective churn prediction target, the next step is a strategy buildout. We focus on proactive and reactive churn prediction strategies to combine efforts and prepare for every situation.

It's typical for us to see between 10-30% of a company's customer base past the churn prediction target threshold at any given moment in time. Consumers that are normally active every day, for example, now have seven-day gaps without any outgoing activities. These customers have already changed their behavior, and we do not even need to predict their churn – it's obvious they are on the way out.

The inactive clients, who you will potentially lose, are grouped into the *reactive churn* prediction strategy. This segment of customers will churn in the future, and you should try to retain them, but your marketing efforts will likely not see a high rate of success. As a result, an aggressive retention strategy is needed to quickly gather whatever can be regained from this group of customers.

For the majority of your customer base, typically 70-90%, who still show regular activity on their accounts, Exacaster helps you design sophisticated models to detect a change in behavioral patterns as early as possible. We pinpoint customers that exhibit risky behaviors that may lead to churn, which is defined as our *proactive churn* prediction process.

Here you must develop a less aggressive retention strategy for these consumers as they are active and your aim is to help these customers stay engaged. We recommend building your strategy for proactive retention *after* you have a solid strategy in place for reactive retention. The reason is that a proactive retention strategy requires more creativity in offers and flexibility in execution. For example, you may decide to recommend a rate plan change that aligns their usage better with your product, help set up an automated top-up process and try out a new service free of charge.

SELECTING ALGORITHMS FOR CHURN PREDICTION

There are two standard preferences for crunching the big data collected during the churn prediction process. These are:

- Bias towards predictive models clarity
- Bias towards predictive models accuracy

At Exacaster, we leverage both. Our R&D department performed extensive research on more than 100 algorithms, discovering different configurations to find the more efficient and effective model for churn prediction.

Our research determined that the family of Random Forests algorithms consistently delivers better, stable and accurate results in comparison to widely used algorithms such as Neural Networks (not including Deep Learning Neural Networks), Decision Trees, Linear Regression, Logistic Regression, and Support Vector Machines.

Our algorithms have additional

benefits including the ability to approximate variable importance; strength against multiple statistical issues like outliers in data, multicollinearity, and unbalanced data sets; and ability to maneuver vast data sets with easy calculation scaling across multiple servers.

Other data scientists may argue that any algorithm could produce very accurate predictions given adequate time to find a suitable parameterization. However, our team's experience shows that the effort required to clarify superb churn prediction accuracy using Random Forest is, by far, smaller compared to the other algorithms, resulting in faster implementation time.

PREPARING CUSTOMER DATA FOR CHURN PREDICTION

Delivering a solid churn prediction strategy centers around well-designed metrics that accurately track customer behavior.

A typical mistake in propensity modeling is when a scientist does not focus full effort in the design of customer metrics because they are too concerned about selecting the best algorithm. Commonly used metrics that produce low-value churn predictions include:

- Customer Age
- Gender
- Number of Calls
- SMS Messages
- Data Usage

These metrics represent a static picture of your customer, and you cannot produce an accurate prediction model with this data.

Our team uses hundreds, sometimes thousands, of specially crafted metrics that identify customer context and behavioral changes. For example, we measure the percentage of primary contacts that recently churned and what their voice, SMS, and data usage trends were in the final weeks before you lost them as a client.

This contextual information dramatically increases churn prediction accuracy. Since we have a preferred algorithm, Random Forests, our goal is to gather a plethora of contextual metrics from your data about customers. Having a vast amount of data allows Exacaster's team to build robust, accurate churn predictions.

As a result, we're able to automate the churn prediction process and demonstrate best-in-class accuracy.

TRAINING ALGORITHMS AND MEASURING PERFORMANCE

Algorithm training is an incredibly technical process. Looking past the specifics and understanding this method on a conceptual level is an essential step for you to validate the prediction quality.

In theory, when a data scientist builds a churn prediction algorithm, they must split all the data into three blocks:

- Training Set (50% of the data, used for algorithm training)
- Testing Set (25% of the data, used for results testing)
- Validation Set (25% of the data, used for results validation)

Exacaster uses a slightly different process for algorithm training. Our ultimate goal is creating robust, production oriented, high-performing churn prediction algorithms that run automatically without supervision. As a result, we take the algorithm validation process to the next level.

Instead of pushing a sample data set from a training data set and preparing it for validation, Exacaster's team validates algorithm performance with hidden data. For example, if we train churn prediction models with January 2016 data, we verify the algorithm performance with February 2016 data. This additional step gives us a clear understanding if our models are stable enough in the real world.

For churn scoring, we retrain algorithms every day, week or month, depending on prediction cycles. This additional process ensures that our algorithms are always up to date, by accounting for the latest market changes.

MEASURING RETENTION PROCESS IMPACT

Once we help you prepare a list of customers with the biggest churn risk score, retention efforts should start.

It's extremely important but counterintuitive to set aside a small control group for each campaign (this group does not receive a retention offer and you will do nothing to keep them). However, a control group provides you with a valuable baseline that helps measure two key factors:

- How effective are our churn prediction algorithms?
- How effective are our retention strategies?

When a company begins to track and understand these two metrics, management of your proactive retention strategy becomes simplified and based on fact, not speculation. Without this measurement, you have no real data to understand if your retention process produces any value. In fact, it is not rare to find retention processes that are too aggressive and, as a result, destroy value. Nobody measured it, so nobody knows.

THE BUSINESS CASE FOR CHURN PREDICTION

By now, the main outlines of the business case to operate churn prediction should be clear. It can be distilled into three major value drivers:

- **Understand the real health of your customer base at any moment:** A very sharp engaged / at risk / disengaged customer definition that considers individual customer differences arms CBMs with key decision information.
- **Stop wasting money on misdirected retention:** Stopping reactive retention programs that mistakenly target customers who are, in fact, fully engaged is a huge money saver.
- **Start working on targeted retention well in advance:** Churn prediction shows the CBM the at-risk group, which can be addressed with specific pro-active retention programs. Talking to the customer before your SIM card is in the trash bin and gives much better response rates to your retention activities.

ABOUT EXACASTER

Each day Exacaster monitors more than 20 million consumers of telecommunications services and retail shoppers searching for churn patterns and value increase campaign scenarios. We use every available data point you can muster and predict customer behavior automatically using world-class predictive modeling techniques.

Companies of almost any size can enjoy benefits of our advanced predictive analytics due to economies of scale achievable in the cloud.

Contact us at:

hello@exacaster.com

www.exacaster.com