Practical Methods for Predicting Customer Retention

Cherkashin, Alexander and Sakhadzhi, Vladislav and Guliev, Ruslan and Bolshunova, Elena

T2 Mobile LLC

15 October 2024

Practical Methods for Predicting Customer Retention

Aleksandr Cherkashin[1], Vladislav Sakhadzhi[2], Ruslan Guliev[3], Elena Bolshunova[4]

[1]*T2 Mobile LLC, Moscow, Russia*, a.cherkashin@t2.ru
[2]*T2 Mobile LLC, Moscow, Russia*, vladislav.sakhadzhi@t2.ru
[3]*T2 Mobile LLC, Moscow, Russia*, ruslan.guliev@t2.ru
[4]*T2 Mobile LLC, Moscow, Russia*, elena.bolshunova@t2.ru

October 15, 2024

Abstract. This study examines methods for analyzing and forecasting the retention of active subscribers in the telecommunications industry using various criteria for subscriber activity. The results demonstrate that the retention dynamics of an active subscriber base can be effectively modeled using a decreasing power function. This allows for medium-term forecasting based on initial subscriber activity data. However, it is important to note the potential limitations in the effectiveness of the proposed approach for long-term forecasting, associated with changes in subscriber churn dynamics over time.

Keywords: subscriber base, customer retention, customer churn, power law, power function, telecommunications, LTV.

JEL Classification: C53

Introduction

Known research in the field of subscriber churn prediction primarily focuses on the use of machine learning methods and big data analysis (Verhelst et al., 2021; Verbeke et al., 2012). These approaches often require significant computational resources and large amounts of data. At the same time, they are typically aimed at identifying factors influencing subscribers' propensity to churn (Ribeiro et al., 2023; Jain et al., 2021) or predicting the behavior of specific subscribers in order to influence them. This study focuses on analyzing the behavior of subscriber groups, as well as on the inverse of churn: the retention of an active subscriber base. The study demonstrates that the retention of a subscriber groups follows a power law. This allows predicting the retention of a group of subscribers, which is necessary, among other things, for economic modeling tasks and estimating the total revenue from a subscriber – LTV (Lifetime value) (Gupta & Zeithaml, 2006). Before moving on to modeling the subscriber retention, it is necessary to discuss practical issues of choosing criteria for subscriber activity and constructing a retention curve.

1. Signs of subscriber activity

The analysis and forecasting of active subscriber retention largely depend on the choice of criteria used to determine subscriber activity. This chapter will examine various approaches to defining subscriber activity and their impact on assessing base retention. The following classification of activity criteria is proposed:

1. legal criteria;
2. financial criteria;
3. criteria for actual service usage.

Each of these categories has its own advantages and limitations, which will be discussed in detail in the following sections.

1.1. Legal criteria for activity

Legal criteria for activity are based on the formal status of contractual relationships between the company and the subscriber. The main indicator in this category is the presence of an active service agreement. In the field of mass service provision, customer activity and the existence of a valid contract with them are generally loosely related. There is a significant delay between the actual cessation of service use by a subscriber and the legal termination of the contract. Many mobile operators, including T2, provide for automatic termination of the contract after a long period of inactivity. For example, according to T2 terms, the contract is terminated if the subscriber does not perform any paid actions for 180 days with a zero or negative balance. In some areas, such as retail or one-time services, legal relationships may be formalized for each individual transaction, making this criterion inapplicable for long-term retention analysis. Accordingly, the use of legal criteria for analyzing subscriber retention has significant limitations.

Despite limited applicability in the telecommunications industry, legal criteria may be relevant in other economic sectors or legal systems where the formal status of a contract is more closely linked to actual service usage.

1.2. Financial criteria for activity

Financial criteria for activity are based on various aspects of financial interaction between the subscriber and the telecommunications operator, including revenue generation, subscriber payments, balance status, and service-related costs. Using financial indicators as activity criteria has a number of obvious advantages. First and foremost, they directly reflect the economic value of the subscriber to the company, which is particularly important from an economic modeling perspective.

However, despite this advantage, using financial criteria to assess subscriber retention is associated with several significant limitations. One of the key factors affecting the reliability of financial criteria is the dependence on the company's accounting policies. The moments of revenue recognition can vary significantly depending on the rules adopted in the organization. For example, with an advance payment for annual service, revenue can be recognized at once or distributed over the entire paid period, which will show completely different subscriber activity. Such payment for services a year in advance can create either false constant activity, even if the subscriber actually stopped using the services after a few months of payment, or, conversely, show a complete lack of activity beyond the month of advance. Neither will provide reliable information about the subscriber's real behavior, which will not allow using this data for forecasting. This is also relevant in the context of various marketing promotions and special offers that may encourage subscribers to make long-term prepayments.

The problem of residual funds in inactive subscribers' accounts deserves special attention. If a subscriber who has stopped using the services has funds left in their account, their gradual write-off to revenue may show false activity. This phenomenon can lead to systematic distortion of data on subscriber retention.

Balance replenishment, as a potential criterion for activity, on the one hand, does not depend on the company's accounting policy, which is its undoubted advantage. However, it is also not necessarily synchronized with the real use of services: a subscriber may deposit a large amount and not use the services for a long time, which will lead to a distortion in the assessment of their activity.

Using costs as an activity criterion faces the same problems as using revenue, but is exacerbated by a more complex system of cost accounting and allocation. Moreover, costs are usually associated with service usage, which is more efficient to measure directly, as will be discussed in more detail in the next chapter on criteria for actual service usage.

Considering all of the above, it can be concluded that despite the importance of financial indicators for the economic evaluation of the subscriber base, their use as criteria for subscriber activity requires an extremely cautious approach. These limitations can lead to significant distortions in the analysis and forecasting of the dynamics of active subscriber base retention. To overcome the described limitations, it is recommended to carefully analyze the company's accounting policy before using financial criteria. When interpreting financial indicators of subscriber activity, it is necessary to take into account the specifics of the business model and the peculiarities of service provision. Financial criteria, although important for the overall assessment of the economic efficiency of the subscriber base, cannot be considered a universal and sufficient tool for analyzing subscriber retention. Nevertheless, financial criteria can be effective for solving narrowly focused tasks, for example, analyzing the payment of subscription fees in tariffs with fixed monthly payments.

1.3. Criteria for actual service usage

Criteria for actual service usage represent the most objective and direct method of assessing subscriber activity in the telecommunications industry. Unlike legal and financial criteria discussed earlier, these indicators are based on real consumption of communication services, which ensures high accuracy and relevance of data on subscriber behavior. The main advantage of these criteria is their independence from the company's accounting policies and other managerial factors. The use of communication services creates a direct load on the network infrastructure, reflecting real physical processes that are recorded at the moment of their occurrence. This allows obtaining up-to-date data on subscriber behavior without the time delays characteristic of financial indicators.

In the context of a mobile operator, criteria for actual usage cover a wide range of services, including voice calls, data transmission, and text messaging. However, not all types of network activity are equally relevant for assessing subscriber retention. For example, technical traffic generated by a device to maintain connection with the network cannot be considered as an indicator of subscriber activity (although it does indicate device activity). When developing a system of criteria for actual service usage, it is necessary to consider the specifics of the provided services and the peculiarities of their consumption.

Despite the obvious advantages, criteria for actual service usage also have certain limitations. The main disadvantage is the potential gap between actual service usage and their economic value for the company. A subscriber may actively use services included in a package without generating additional revenue. On the other hand, a subscriber inactive according to these criteria may regularly pay for services, which is important from a financial perspective. For example, a subscriber may have a positive balance and regularly make payments, but not use the services for various reasons. Such a situation can arise, for instance, when a subscriber has multiple SIM cards from different operators or when keeping a number, for example, for registration on online services. To overcome this limitation, one can, for example, evaluate the actual difference between "paying" and "service-using" subscribers, and take it into account in modeling. Another way is to use a mixed approach that combines criteria for actual service usage with selected financial indicators.

The choice of a specific criterion or combination of criteria should be made taking into account the specifics of the business, analysis goals, and characteristics of the subscriber base under study. There is no universal solution suitable for all situations, and each company needs to develop its own approach to defining and measuring subscriber retention. Within the framework of this article, we will use two specially developed criteria for subscriber activity:

1. Criterion A – considers only some service usage by the subscriber.
2. Criterion B – combined, includes both service usage and financial signs of subscriber activity.

2. Methods for calculating and constructing the retention curve

The subscriber retention curve allows visualizing and quantitatively assessing the dynamics of active subscriber retention over time. It is a graphical representation of the proportion of subscribers remaining active over a certain period after initial activation (usually moment of sale). In the context of the telecommunications business, "active" means the continued use of the company's services by a subscriber. Construction and analysis of retention curves serve two main purposes: diagnosing the current state and creating a basis for building predictive models.

A retention curve can be constructed using two different methods. These methods will be discussed in more detail below, their advantages and disadvantages will be defined, as well as their relationship to each other.

2.1. Retention curve of a fixed group of subscribers

The first method calculates and constructs the retention curve of a fixed group of subscribers. The method is based on observing a fixed group of subscribers who connected at a certain point in time and allows tracking the dynamics of activity of the same group of users throughout the entire observation period, representing a cohort analysis (Fader & Hardie, 2009; Zhang & Chang, 2021).

The process of constructing a subscriber retention curve includes the following stages:

1. defining the initial group of subscribers who connected during the selected time period (usually a calendar month);
2. tracking the activity of this group according to the chosen activity criterion throughout the study period;
3. calculating the proportion of active subscribers from their initial number for each time interval.

Mathematically, retention for each period can be expressed by the following formula:

$$R_i = \frac{A_i(T, c)}{G_T} \tag{1}$$

where:

- R_i – is the retention rate for the i-th period;
- $A_i(T, c)$ – is the number of active subscribers (according to activity criterion c) in period i from the number of subscribers activated in period T;
- G_T – is the number of initial activations in period T;
- $i \in [1; +\infty]$ – is the period for calculating retention;
- T – is the fixed initial period (usually the initial activation or sales period);
- c – is the subscriber activity criterion.

To illustrate this method, let's consider an example in Figure 1 of calculating subscriber retention according to criterion A for a group of subscribers who connected to the T2 network in April 2023:

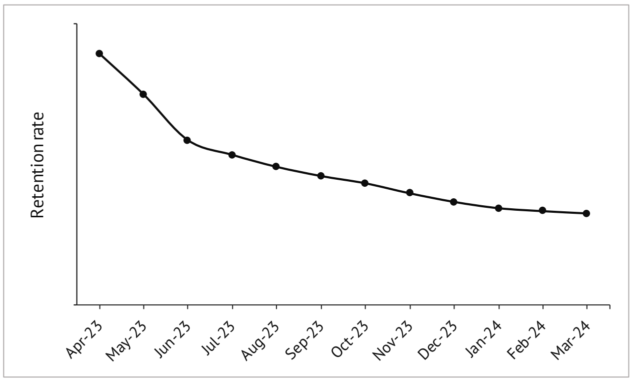

Figure 1. A-retention curve of subscribers activated in April 2023.

The vertical axis should have a range from 0 to 1 (from 0% to 100%) and is usually presented on a linear scale. Within the article, to prevent disclosure of the Company's actual data, the range has been deliberately changed and value labels have been removed in this and subsequent figures, while the linear scale of the axis is maintained. Given the above, Figure 1 can be read as follows: for example, during the first month (April 2023), the proportion of connected subscribers remaining active according to criterion A was close to 100%, during the second month – to 90%, in the third month – let's say about 75%, and so on.

It's important to note that to construct a retention curve for a certain period, for example, 12 months, it's necessary to have data on a group of subscribers who connected at least 12 months before the analysis moment. At the same time, the retention rate for the first month for this curve will be calculated based on subscriber activity data from 12 months ago.

2.2. Retention curve by subscriber cohorts

An alternative method is to construct a retention curve by subscriber cohorts. A cohort is a group of subscribers united by initial SIM activation month. The age of a cohort is measured in full calendar months from the month of activation. Accordingly, this approach is based on the analysis of the current retention of subscriber groups of different ages.

The process of constructing a retention curve by cohorts includes the following stages:

1. dividing current active subscribers into groups (cohorts) by "subscriber age", for example, expressed in months;
2. calculating retention for each cohort, starting with the youngest;
3. constructing a curve where each point represents the retention of the corresponding cohort at the current moment.

To illustrate, let's consider an example:

- Cohort 1 (age 1 month): subscribers whose initial activation was a month ago. Retention is calculated as the proportion of active subscribers in this cohort at this month to the total number of SIM activations a month ago.
- Cohort 2 (age 2 months): subscribers who activated two months ago. Retention is calculated similarly, but relative to activations from two months ago.

Mathematically, retention for each cohort can be expressed by a similar formula:

$$R_j = \frac{A(j,c)}{G(P(j))} \tag{2}$$

where:

- R_j – is the retention rate of the j-th subscriber cohort;
- $A(j,c)$ – is the number of active subscribers in the j-th cohort, according to activity criterion c;
- $G(P(j))$ – is the number of SIM activations in period $P(j)$;
- $P(j)$ – is the initial activation period corresponding to cohort j;
- $j \in [1; +\infty]$ – is the cohort number;
- c – is the subscriber activity criterion.

Note that unlike the retention curve of a fixed group of subscribers, the value of the denominator changes for each cohort j. Here we see the currently observed retention of different groups of subscribers, which were first activated at different time.

2.3. Relationship between both curves

The retention curve by cohorts can be represented as a combination of values from several fixed group retention curves. Schematically, this relationship is shown in Figure 2:

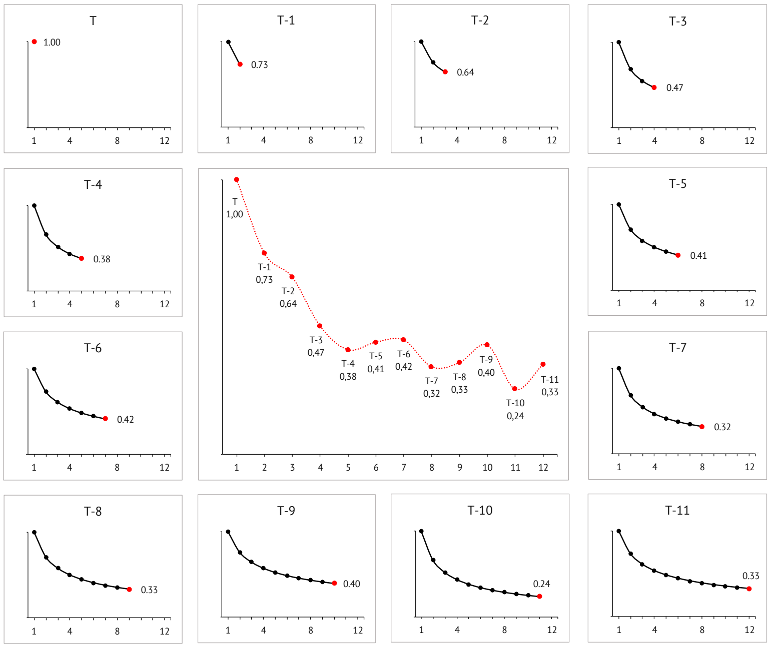

Figure 2. Relationship between the retention curve by cohorts and fixed group retention curves. Data randomly modeled for illustration.

In this diagram, T represents the current (or last available for analysis) month. The small graphs show retention curves for fixed groups of subscribers activated in different periods (from T to T-11). The large graph in the middle represents the retention curve by cohorts, where each point corresponds to the last point of the respective retention curve: the first point (T – retention for the first month) corresponds to the last point of the small graph T, the second point (T-1 – retention for the second month) corresponds to the last point of the small graph T-1, and so on. Thus, the retention curve by cohorts consists of a sequence of retention values for the first, second, and subsequent months of life of different subscriber groups, observed at the current point in time.

Note that the retention curve by cohorts generally repeats the retention dynamics of the curves from which its points are composed, while the dynamics are "noisy" due to various factors associated with the fact that it represents the retention of different subscriber groups.

2.4. Comparative analysis and choice of retention curve construction method

Economics is a field of knowledge forced to operate in constantly changing external conditions. Empirical laws and dependencies from the field of economics are not constant. Data and models tend to become outdated quickly, losing relevance. This forces a constant search for balance between the amount of data and its relevance (in the opinion of the authors of this study, this is analogous to the well-known bias-variance tradeoff in statistics and machine learning). On the one hand, using a larger volume of historical data can increase the statistical stability of the model. On the other hand, older data may not reflect current market realities, leading to increased bias in estimates. In our view, when choosing a method for constructing a retention curve, one should be guided by the purpose of the analysis (operational control, or long-term forecasting, or research tasks) and the specifics of the business and market in which it operates.

The retention curve of the fixed group of subscribers reflects the retention of the same group of activations throughout the entire observation period. The curve is always strictly non-increasing, which corresponds to the logic of the subscriber churn process. However, it has low operability: constructing a long-term curve requires a long observation period. As a result, by the time a complete curve is obtained, the retention data for the initial periods may lose relevance. For operational assessment tasks, it is necessary to see, on the one hand, the most up-to-date data, and on the other hand, the overall picture as a whole. This is made possible by the retention curve by cohorts. This curve always presents the latest available data for each period of subscriber life as it reflects the current state of the subscribers at different stages of their life cycle. However, because the retention curve by cohorts is built on different groups of activations, it is subject to high fluctuations associated with seasonality and all other factors affecting the quality of activations in different periods. This creates, among other things, the possibility of local increase in the curve, which is counter-intuitive as it contradicts the logic of the churn process.

In more theoretical tasks, in studies of the influence of various factors on subscriber retention or in studies comparing different activation channels with each other, the retention curve of the fixed group of subscribers turns out to be preferable. It provides a cleaner picture of the retention dynamics of a specific group of subscribers, which can be critically important for identifying fundamental patterns and trends. Since this study is devoted to developing methods for predicting subscriber base retention, the retention curve of the fixed group of subscribers will be used further (referred to simply as the "retention curve"). In practical work, when making forecasts themselves or monitoring the quality of sales, a choice can be made in favor of the cohort survival curve.

It is necessary to briefly mention the construction of survival curves based on averaged data, which is often used in practical work. In the presence of seasonal fluctuations and market anomalies, averaging over several periods allows smoothing the dynamics and excluding the influence of one-time factors. At the same time, the application of averaging inevitably leads to the use of older data, which can negatively affect the relevance of the analysis.

3. Correspondence of subscriber retention to the power law

To identify stable patterns in subscriber behavior, it is desirable to study retention curves over significant time intervals. In this study, we analyzed the retention curve of subscribers who activated T2 network in November 2015, with subsequent observation of their activity according to criterion B over 105 months (8.5 years) until July 2024 (Figure 3).

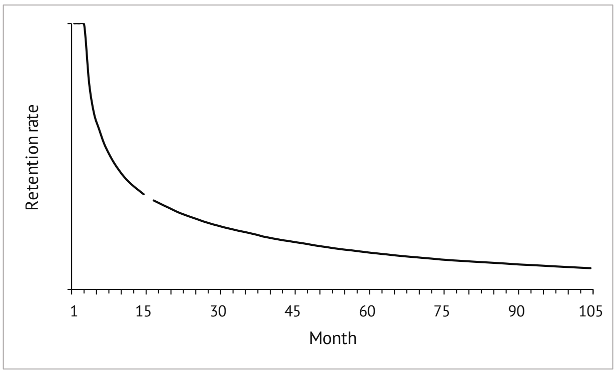

Figure 3. B-retention of subscribers activated in T2 in November 2015.

Note: Data on active subscribers for February 2017 is missing, which appears as a gap on the graph.

Visual analysis of the graph demonstrates a monotonic and smooth nature of the retention curve, which allows us to hypothesize about the possibility of its approximation by a relatively simple mathematical function.

To test the proposed hypothesis, four types of functions were considered:

- power: ax^b;
- exponential: ae^{bx};
- logarithmic: $aln(x) + b$;
- linear: $ax + b$.

Nonlinear functions were linearized, the parameters of the functions were determined by the least squares method (LSM). The approximation results are presented in Figure 4 and Table 1:

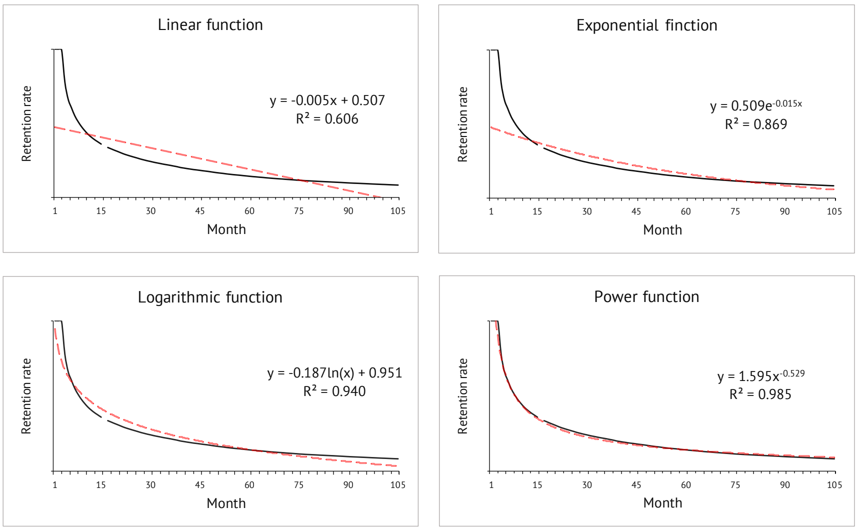

Figure 4. Approximation of the B-retention curve by various functions.

Analysis of the results shows that the power function demonstrates a qualitative and best approximation of the retention curve with a coefficient of determination (R^2) = 0.985.

Table 1

Evaluation of retention approximation accuracy across various functions

Function type	Function	R^2
Power	$y = 1.595x^{-0.529}$	0.985
Logarithmic	$y = -0.187\ln(x) + 0.951$	0.940
Exponential	$y = 0.509e^{-0.015x}$	0.869
Linear	$y = -0.005x + 0.507$	0.606

To verify the obtained results, we performed an approximation of retention curves according to criterion A for various groups of subscribers activated in the company's network at different periods and through different sales channels. The results are presented in Figure 5 and Table 2:

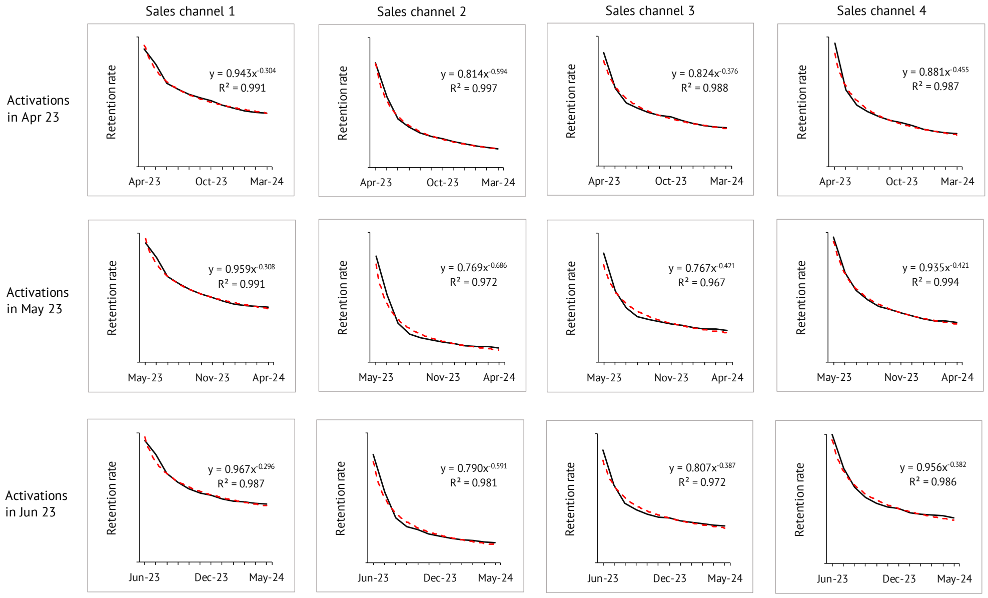

Figure 5. Approximation of A-retention curves by a power function for various sales channels and activation periods.

Analysis of the data in Table 2 confirms the high stability of the results of approximation by a power function. In all considered cases, R^2 exceeds 0.96, which indicates a good correspondence of the model to empirical data regardless of the sales channel and subscriber activation period.

Table 2

Evaluation of power function approximation for various sales channels and periods

Sales channel	Activation month	Function	R^2
	Apr 23	$y = 0.943\,x^{-0.304}$	0.991
1	May 23	$y = 0.959\,x^{-0.308}$	0.991
	Jun 23	$y = 0.967\,x^{-0.296}$	0.987
	Apr 23	$y = 0.814\,x^{-0.594}$	0.997
2	May 23	$y = 0.769\,x^{-0.686}$	0.972
	Jun 23	$y = 0.790\,x^{-0.591}$	0.981
	Apr 23	$y = 0.824\,x^{-0.376}$	0.988
3	May 23	$y = 0.767\,x^{-0.421}$	0.967
	Jun 23	$y = 0.807\,x^{-0.387}$	0.972
	Apr 23	$y = 0.881\,x^{-0.455}$	0.987
4	May 23	$y = 0.935\,x^{-0.421}$	0.994
	Jun 23	$y = 0.956\,x^{-0.382}$	0.986

The obtained results allow us to formulate the following observations and hypotheses:

1. Universality of the power function: high accuracy of approximation of retention curves by power functions for various activity criteria, subscriber base segments,

and time periods may indicate the presence of a fundamental pattern in the process of subscriber retention and churn.
2. Dependence on initial conditions: variations in the parameters of the approximating power function for different subscriber groups may reflect exclusively the influence of initial conditions (sales channel, activation period) on long-term retention dynamics.
3. Resistance to external factors: preservation of the shape of retention curves, despite the influence of a large number of various changing factors (service quality, subscriber retention activities, competitors' activities to attract the base). Perhaps this indicates some dynamic equilibrium of the competitive market between the company's efforts to retain subscribers and factors contributing to their churn.
4. Power functions are widely encountered in describing patterns in various areas of the surrounding world: physics, astronomy, biology, geology, sociology, linguistics, psychology, economics, etc. (Andriani & McKelvey, 2007). Newman's study (Newman, 2005) examines power laws in such areas as physics, biology, and economics, as well as their manifestations in the form of Pareto distribution and Zipf's law. Power functions describe the following patterns in economics and finance: distribution of income and wealth, company sizes, stock market returns, trading volumes, indicators of international trade, etc. (Gabaix, 2009). The prevalence of power dependence in describing subscriber retention raises a number of fundamental questions about the nature of the processes that determine subscriber behavior. For example, in a study by J. Staddon (Staddon, 1978), an attempt was made to justify the emergence of power dependencies in the behavior of living organisms in response to various stimuli.

The high efficiency of approximating subscriber retention curves by power functions opens up opportunities for predicting subscriber base dynamics. The next chapter will consider the possibilities of applying the obtained results for building predictive models and discuss the limitations of the proposed approach.

4. Extrapolation of retention and its use as a forecast

The high accuracy of approximating retention curves by power functions, identified in the previous chapter, opens up prospects for their use in forecasting. Extrapolation of these functions may allow estimating subscriber base retention over long time intervals without the need to wait for actual data. This is essential for economic modeling and evaluating the effectiveness of subscriber acquisition activities. Forecasting subscriber retention plays a key role in calculating such important economic indicators as expected subscriber lifetime (Customer Lifetime) and total revenue from a subscriber (LTV – Customer Lifetime Value) (Kumar, 2014). These metrics are critically important for comparison with subscriber acquisition costs (SAC – Sales/Subscriber Acquisition Cost) and, consequently, for making informed decisions about launching, continuing, or terminating marketing campaigns to grow the subscriber base (Krstevski & Mancheski, 2016).

To assess the effectiveness of extrapolating power functions in predicting retention, the following analysis was conducted:

1. data on subscriber retention according to criterion A were used, for subscribers activated in April 2023, over a period of 12 months;

2. forecast models were built based on the extrapolation of power functions using different numbers of initial data points (from 2 to 11 months);
3. the average monthly absolute forecast error for each model was calculated as the average monthly absolute difference between the predicted and actual values (only forecast periods were used for calculation);
4. a graph was plotted showing the dependence of the average monthly absolute error on the number of actual data points used in the extrapolation.

The results of the analysis are presented in Figures 6, 7 and in Table 3.

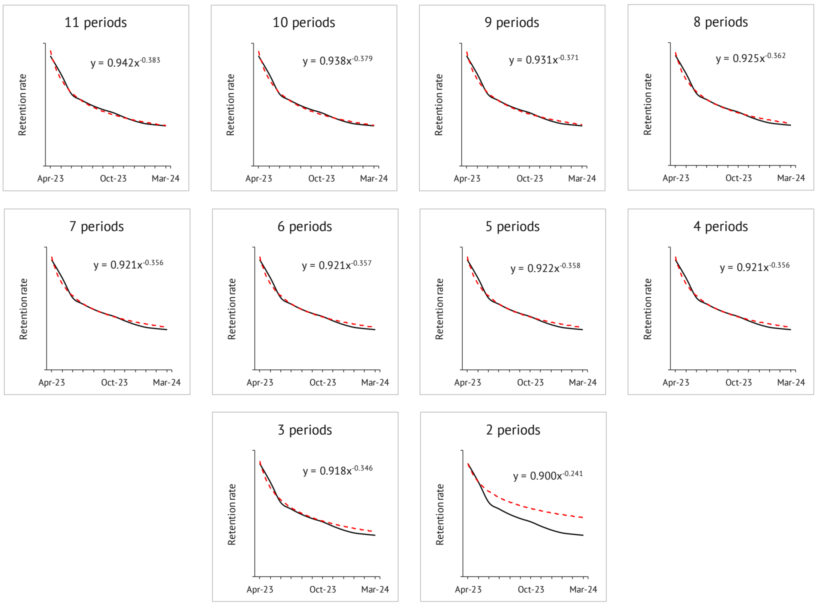

Figure 6. Comparison of actual A-retention with forecasts based on extrapolation of power functions.

A sharp improvement in forecast quality is observed when using three actual data points, after which the forecast quality shows only a slight improvement. A small anomalous increase in forecast error observed in the middle of the studied range (5-7 months) is related to the chosen metric for assessing forecast quality and the specifics of the particular activation month. It can be assumed that, on average, the error will monotonically decrease with an increase in the number of actual data points used.

Table 3

Parameters of approximating functions and forecast errors

Actual periods	Forecast periods	Function	Avg. abs. residuals
11	1	$y = 0.940x^{-0.383}$	0.004
10	2	$y = 0.938x^{-0.379}$	0.009
9	3	$y = 0.931x^{-0.371}$	0.015
8	4	$y = 0.925x^{-0.362}$	0.020
7	5	$y = 0.921x^{-0.356}$	0.021
6	6	$y = 0.921x^{-0.357}$	0.017
5	7	$y = 0.922x^{-0.358}$	0.015
4	8	$y = 0.921x^{-0.356}$	0.014
3	9	$y = 0.918x^{-0.346}$	0.020
2	10	$y = 0.900x^{-0.241}$	0.113

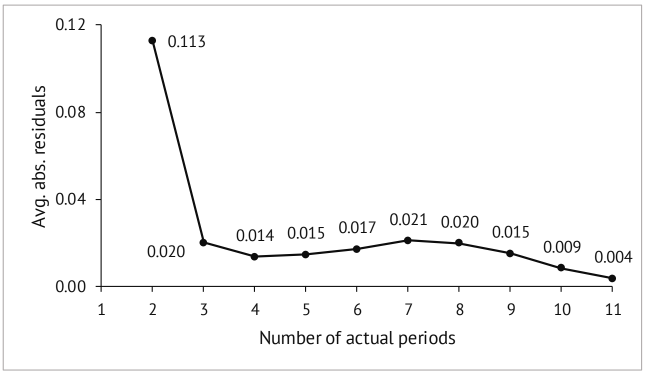

Figure 7. Dependence of the average monthly forecast error on the number of actual data points used.

The aim of this analysis was to propose a general approach to assessing forecast quality, to demonstrate the applicability of forecasting based on a limited set of actual data, and to propose a method for determining the necessary amount of actual data for its construction. The fact that the quality becomes sufficiently high when using just a few actual values, after which it only slightly improves, has important practical significance. For example, only the first three retention values can be used for a fairly accurate estimate of retention in the future. Of course, for practical application, this conclusion needs to be further validated on a wider set of data.

It should be noted that it is not necessary to take exactly the first few values. This method provides an opportunity to assess the "hypothetical past" – that is, to conduct a retrospective analysis and evaluation of alternative scenarios. This can be useful in cases where it is required to analyze the effectiveness of already conducted actions aimed at subscriber retention.

In practical tasks, it may be advisable to build such a forecast based on the retention curve by cohorts, which better captures the current dynamics of subscriber base retention. However, given its heterogeneity, a larger amount of actual data may be required for forecasts of the required accuracy.

The conducted research demonstrates the high potential of using power function extrapolation for predicting subscriber base retention. However, despite the revealed effectiveness, this approach has certain limitations that require careful consideration. The next chapter will analyze these limitations, in particular, the question of long-term applicability of power functions for modeling retention will be discussed.

5. Limitations of applicability

As demonstrated in Chapter 3, power functions of the form ax^b show high accuracy in approximating the actual retention of the subscriber base. Chapter 4 considered the possibility of using these functions for building forecasts. However, despite their effectiveness, it is evident that the high quality of forecasts based on the extrapolation of power functions has certain limits of applicability. From a mathematical point of view, as $x \to +\infty$ a function of the form ax^b, with $a > 0$ and $b < 0$, asymptotically approaches zero, never crossing the X-axis:

$$\lim_{x \to +\infty} a\,x^b = 0, \quad \text{при } a > 0, b < 0 \tag{3}$$

This means that in forecasts based on such functions, retention will never reach zero value. However, based on general considerations, it is obvious that such an assumption cannot correspond to reality. The retention of a group of subscribers is an aggregated indicator of individual states, each of which at any given time can only take two values: 1 (subscriber is active) or 0 (subscriber is inactive). It is logical to assume that over time, the number of active subscribers will steadily decrease until the last active subscriber remains. After their departure, the retention of the entire group will inevitably become zero – an event that will never be reflected by a function of the form ax^b.

Moreover, subscriber base retention is affected not only by factors related to customers' desire to use the company's services but also by objective limitations such as the physical ability to use communication services. Subscribers are living people (with the exception of the segment of devices with SIM cards, which also have a limited service life). Inevitably, there will come a time when each subscriber will physically be unable to continue using the services.

These considerations point to a potential limitation in the applicability of power function extrapolation for long-term forecasting purposes. For a more detailed analysis of this limitation, let's refer to the retention graph according to criterion B and its approximation by a power function, presented in Figures 3 and 4. Upon careful examination of the graph area corresponding to later periods, it can be noticed that, despite high indicators of overall approximation quality, with increasing time, there is a tendency for systematic overestimation of actual retention by the approximating curve. Interestingly, in the first half of the graph, the opposite situation is observed – the function slightly underestimates the actual retention. This can be explained by the fact that the coefficients of the approximating function are selected in such a way as to minimize the overall error along the entire length of the curve.

For a more vivid demonstration of the retention overestimation effect, let's construct an approximation using only 3 initial actual values and extrapolate retention for the entire period.

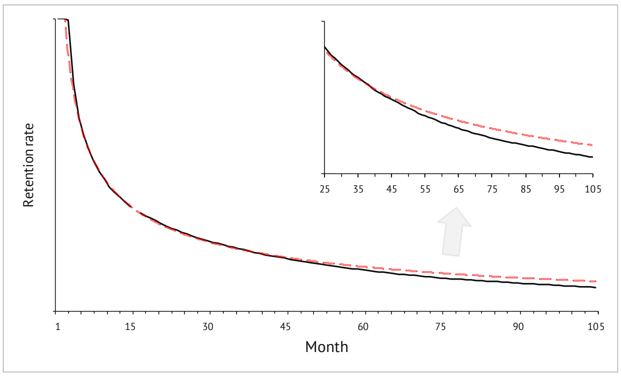

Figure 8. Extrapolation of B-retention based on initial actual values.

Note: Due to the peculiarities of criterion B, the very first values give a poor representation of the subscriber's real activity, so actual values for the 5th, 6th, and 7th months were chosen for extrapolation. The deviation in the first few months is also explained by the specifics of the chosen activity criterion.

The results (Figure 8) demonstrate that the curve constructed based on a limited set of initial data more accurately describes the actual retention in the first half of the observation period. However, now the tendency for divergence between predicted and actual retention as the age of the subscriber group increases is even more pronounced. Such an observation allows us to hypothesize about the emergence of additional factors affecting subscriber base retention over time. As mentioned above, one such factor could be the physical age of subscribers, whose influence on the use of communication services presumably increases with time.

Based on the survival curve of large mammals [Deevey, 1947; Pearl & Miner, 1935], including humans (type I curve in Figure 9), we can make an assumption about the general form of the subscriber retention curve, taking into account the survival factor.

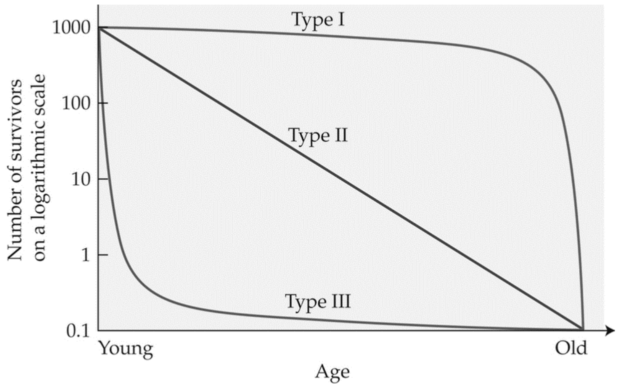

Figure 9. Three main types of survival curves for animal populations in the wild according to E. Deevey Jr. [Hill et al., 2021]

If we consider the subscriber retention curve as the probability of maintaining their activity as a subscriber, then the survival factor should gradually modify the curve in accordance with the type I curve. Then, the general form of the subscriber retention curve should have an S-shape, consisting of two sections separated by an interval of indefinite duration (Figure 10). The first section reflects the dynamics of retention due to factors related to service quality and loyalty. The second section demonstrates the influence of the survival factor:

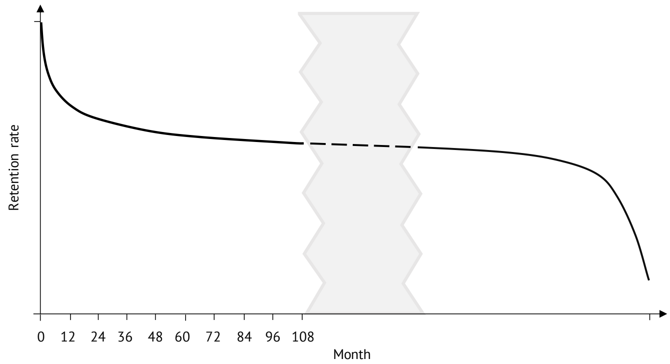

Figure 10. General view of the S-shaped subscriber retention curve.

It is possible that this variant of the curve with areas of deceleration and acceleration could be a general form of the physical survival curve of organisms. Possibly, if comparable observation start points are selected, all three types of curves are variants of such a general curve, and the differences arise in the location of the deceleration and acceleration areas, as well as in the life period that usually receives attention.

The study of the long-term influence of the age factor and other demographic parameters on subscriber base retention represents a direction for further research. Such analysis could contribute to the development of more accurate long-term forecasting models and help determine the general form of the subscriber retention curve throughout their entire life cycle. Nevertheless, as shown earlier, for short-term and medium-term forecasting tasks, this factor can be neglected.

Conclusion

This study proposes a classification of subscriber activity criteria and discusses issues of their practical application. Combined indicators of actual service usage, sometimes in combination with financial indicators, appear to be the most universal for modeling subscriber base retention tasks.

Two methods of compiling a retention curve are considered, and their preferred areas of application are determined. For analyzing the actual results, it is often preferable to use the cohort retention curve, as it contains more up-to-date data. For tasks of researching the shape of the retention curve and factors influencing it, for comparing different curves with each other, it is preferable to use the retention curve of the fixed group of subscribers, as it reflects the pure actual dynamics of activity.

It is shown that the dynamics of subscriber base retention corresponds to a power law, namely curves of the form ax^b ($a > 0$ and $b < 0$). This fact can be used to build forecasts of subscriber base retention based on simple extrapolation, using, in our case, just three to four actual values. The ability to build simple and accurate forecasts knowing only the initial retention values finds great applicability in economic modeling tasks and in assessing the full aggregate expected effects of subscriber activation, such as average subscriber lifetime and LTV.

At the same time, long-term forecasts based on the extrapolation of power functions have limitations, which are presumably associated with the increasing influence of other factors over time, such as the physical age of the subscriber. It can be assumed that the general form of the retention function should have some combined form, in which the initial part of the curve has a more power-like appearance, and the final part differs from it.

This study leaves the following questions open, each of which requires separate study:

1. Why is subscriber retention so well described by power functions?
2. Why does base retention (and, accordingly, churn) look so uniform and, as if, dependent only on initial conditions? Is this a manifestation of the dynamic equilibrium of efforts of all participants in a competitive market?
3. What could be the general form of the long-term subscriber retention curve?

References

Andriani, P. & McKelvey, B. (2007). Beyond Gaussian averages: Redirecting international business and management research toward extreme events and power laws. Journal of International Business Studies, 38(7), 1212-1230

Deevey, E.S., Jr (1947). Life tables for natural populations of animals. The Quarterly Review of Biology, 22, 283-314.

Fader, P. S. & Hardie, B. G. (2009). Probability models for customer-base analysis. Journal of interactive Marketing, 23(1), 61-69.

Gabaix, X. (2009). Power Laws in Economics and Finance. Annual Review of Economics, 1, 255-293.

Gupta, S. & Zeithaml, V. (2006). Customer metrics and their impact on financial performance. Marketing Science, 25(6), 718-739.

Hill, R. W., Sleboda, D. & Millar, J.J. (2021). Youth in the study of comparative physiology: insights from demography in the wild. Journal of Comparative Physiology B, 191, 1-16.

Jain, H., Khunteta, A. & Srivastava, S. (2021). Telecom churn prediction and used techniques, datasets and performance measures: a review. Telecommunication Systems, 76(4), 613-630.

Krstevski, D. & Manchenski, G. (2016). Managerial Accounting: Modeling Customer Lifetime Value – An Application in the Telecommunication Industry. European Journal of Business and Social Sciences, 5(1), 64-77.

Kumar, V. (2008). Costumer lifetime value: The path to profitability. Now Publishers.

Newman, M. E. J. (2005). Power laws, Pareto distributions and Zipf’s law. Contemporary Physics, 46(5), 323-351.

Pearl, R. & Miner, J. R. (1935). Experimental studies on the duration of life. XIV. The comparative mortality of certain lower organisms. The Quarterly Review of Biology, 10, 60-79.

Ribeiro, H., Barbosa, B., Moreira, A. C. & Rodrigues, R. G. (2023). Determinants of churn in telecommunication services: a systematic literature review. Management Review Quarterly, 74(3), 1327-1364.

Staddon, J. (1978). Theory of behavioral power functions. Psychological Review, 85(4), 305-320.

Verhelst, T., Caelen, O., Dewitte, J. C., Lebichot, B. & Bontempi, G. (2021). Understanding Telecom Customer Churn with Machine Learning: From Prediction to Casual Inference. Artifival Intelligence and Machine Learning, 1996, 182-200.

Verbeke, W., Martens, D., Mues, C. & Baesens, B. (2011). Building comprehensible customer churn prediction models with advanced rule induction techniques. Expert Systems with Applications, 38(3), 2354-2364.

Zhang J. Z. & Chang, C. W. (2021). Consumer dynamics: Theories, methods, and emerging directions. Journal of the Academy of Marketing Science, 49(1), 166-196.