Omega xxx (xxxx) xxx

Contents lists available at ScienceDirect

Omega

journal homepage: www.elsevier.com/locate/omega

Profit-driven churn prediction for the mutual fund industry: A multisegment approach

Sebastián Maldonado[a,d], Gonzalo Domínguez[b], Diego Olaya[c,*], Wouter Verbeke[e]

[a] Department of Management Control and Information Systems, School of Economics and Business, Universidad de Chile, Diagonal Paraguay 257, Santiago, Chile
[b] Facultad de Ingeniería y Ciencias Aplicadas, Universidad de los Andes Monseñor Álvaro del Portillo 12455, Las Condes, Santiago, Chile
[c] Data Analytics Laboratory, Faculty of Social Sciences & Solvay Business School, Vrije Universiteit Brussel Pleinlaan 2, Brussels 1050, Belgium
[d] Instituto Sistemas Complejos de Ingeniería (ISCI), Chile
[e] Faculty of Economics and Business, KU Leuven, Naamsestraat 69, B-3000 Leuven, Belgium

ARTICLE INFO

Article history:
Received 23 December 2019
Accepted 25 November 2020
Available online xxx

Keywords:
Profit metrics
Churn prediction
Mutual funds
Analytics
Finance

ABSTRACT

This paper proposes a novel approach for profit-based classification for churn prediction in the mutual fund industry. The maximum profit measure is redefined to address multiple segments that differ strongly in the average customer lifetime values (CLVs). The proposed multithreshold framework for churn prediction aims to maximize the profit of retention campaigns in binary classification settings. The multithreshold framework is empirically tested on data from a Chilean mutual fund company with varying and heterogeneous individual CLVs. Our results demonstrate the virtues of the proposed approach in achieving the best profit when compared to other metrics. Although presented in the context of investment companies, our framework can be implemented in any churn prediction task, representing an important contribution for decision-making in business analytics.

© 2020 Elsevier Ltd. All rights reserved.

1. Introduction

The current scenario for the service industry, which is characterized by strong competition and mature markets, is forcing companies to build strong relationships with their customers [37]. Enhancing loyalty is a profitable decision since it has been estimated that attracting new customers is five to six times more expensive than retaining an existing customer [17,18,37].

The churn prediction problem can be addressed via supervised learning, casting the task into a binary classification problem. Machine learning can be used to predict which customers are more likely to attrite [43]. Given the wide variety of predictive techniques, the usual strategy for choosing the correct one is constructing several models and assessing their predictive performance using statistical measures. However, this approach has recently been criticized since it distances itself from the main goal of the learning process, which is taking profitable action for customers via retention campaigns [43,44]. In this sense, goal-oriented metrics have been developed in recent years for addressing the churn prediction problem, computing the expected profit of a retention campaign [20,44].

Profit metrics for churn prediction are designed to identify the right customers to target in order to lead to profitable retention campaigns. Under this premise, customers with large CLVs are more profitable when retained in contrast to those with smaller CLVs. The latter group is not too attractive for customer retention even when their churn probabilities is high. Although the importance of considering individual CLVs has been acknowledged in most of the studies in the area (see e.g. [36,44,47]), the first profit metrics developed did not account for heterogeneous CLVs, they averaged them. Recently, Óskarsdóttir et al. [36] computed individual CLVs but considered a single decision threshold for a given classifier. This approach does not allow to distinguish between the most valuable customers and the rest, ending up aggregating customers in such a way that the use of a single average CLV is almost equivalent to the use of individual CLVs.

Although churn prediction has been studied in a wide variety of domains, such as telecommunications [8,45], banks and other financial institutions [25,26,37,41], and higher education (as student dropout) [1,12], this topic has not been addressed for predicting customer attrition in the mutual fund industry.

Churn prediction for mutual funds is different from traditional churn prediction, e.g., in the telco industry. First, attrition strongly

* Corresponding author.
E-mail addresses: sebastianm@fen.uchile.cl (S. Maldonado), gdominguez@miuandes.cl (G. Domínguez), diego.olaya@vub.be (D. Olaya), wouter.verbeke@kuleuven.be (W. Verbeke).

https://doi.org/10.1016/j.omega.2020.102380
0305-0483/© 2020 Elsevier Ltd. All rights reserved.

Please cite this article as: S. Maldonado, G. Domínguez, D. Olaya et al., Profit-driven churn prediction for the mutual fund industry: A multisegment approach, Omega, https://doi.org/10.1016/j.omega.2020.102380

depends on the state of the portfolio: customers are more likely to withdraw their funds when they are down. Second, customer lifetime values (CLVs) are very heterogeneous since they depend on the mutual fund performance indicators. Finally, there is a high risk of waking up "sleeping dogs" or "do-not-disturbs": contacting customers that churn because of the retention campaign, who would not have churned otherwise. These customers may act in the exact opposite way that was intended if they realize that their funds are down because of the campaign.

In this study, we adapt profit metrics for predicting churn in the mutual fund industry. The main contributions of this paper are as follows:

- A case study based on data from a Chilean mutual fund company is presented, in which the cost and benefits of applying retention programs are discussed.
- We present a novel multithreshold, multisegment framework for addressing heterogeneous CLVs in the design of profit metrics. This scheme can be used in any churn prediction domain in which the benefits and costs of a retention campaign can be estimated.
- The effectiveness of the proposal is assessed, demonstrating that it leads to the largest average profit when compared with well-known profit metrics for churn prediction.
- The key input for the profit metric is discussed, that is, the amount of the incentive related to the retention campaign. A sensitivity analysis is performed to provide a tool for an adequate design of the campaign and addressing the possible risks.

The remainder of this paper is organized as follows: Section 2 discusses prior studies that are relevant for our proposal. The proposed framework for profit-driven churn prediction in the mutual fund industry is introduced in Section 3, which discusses relevant aspects of this application. Section 4 reports our results on a real-world dataset from a Chilean mutual fund company. Finally, Section 5 summarizes the main conclusions and potential future developments.

2. Prior work

This section is structured as follows: the mutual fund industry is described in Section 2.1, while recent developments in profit-based classification for predicting customer attrition are discussed in Section 2.2.

2.1. The mutual fund industry

A mutual fund is a financial product composed of a pool of money collected from various investors, consisting of a portfolio of bonds, stocks, and other securities. Mutual funds were created in the Netherlands due in response to a crisis that occurred in the years 1772-73, and the concept was designed as a means for small investors to have an option to diversify their portfolio [39].

When comparing mutual funds with a direct investment option, some advantages and disadvantages can be identified [38]. One of the main advantages is that economies of scale can be achieved [14]: a participant with a low initial capital investment may widen the range of their portfolio, thereby lowering their financial risk. Another advantage is that these investments are managed by professional investors called fund administrators [11]. In contrast, those who invest in these funds must pay several fees and charges, which can affect the overall returns[32].

There are many types of mutual funds: open-end funds, closed-end funds, and unit investment trusts. Furthermore, exchange-traded funds (ETFs) are popular investment vehicles that are related to open-end funds and to unit investment trusts. Open-end funds can be further classified according to the main investment of the fund, as follows:

- Money Market: These funds invest in money market instruments. These are regulated funds that are generally known for keeping instruments of high quality in a short term.
- Bond Funds: They consist of debt instruments in public offer values and fixed incomes.
- Equity Funds: They consist of those who invest a great percentage of their portfolio in variable income instruments.
- Mixed Funds: They invest in both types of funds, bond and equity funds.

This research focuses on open-end funds. Most mutual funds belong to this type of fund. Towards the end of the year 2018, there were 118,978 open-end funds, which accounted for $46.7 trillion dollars, an amount that is a bit lower than the previous year, when it reached $49.3 trillion dollars [23].

Analytics and operational research (OR) have been shown to be useful in improving decision-making in the mutual fund industry. For example, portfolio selection decisions can be enriched with investor preference data elicited from questionnaires [4]. Mathematical modeling can be used to approximate utility optima via compromise programming (CP). Additionally, machine learning methods have been used for predicting the performance of equity mutual funds. For example, Indro et al. [24] utilizes artificial neural networks (ANNs) for this task, outperforming traditional linear models. In the same line, forecasting approaches have been used for predicting stock market returns [33] and the price of commodities [35], which are important inputs for the mutual fund industry. Another OR approach that has been used in the mutual fund industry is data envelopment analysis (DEA), which has been successful as a fund diversification approach to find an adequate compromise between risk and returns [16,29].

2.2. Profit-driven churn prediction

Several churn prediction approaches have been developed and reported in the literature. Most strategies, including the one presented in this study, perform single-period predictions in order to anticipate whether a customer will churn in the following period [6]. This problem can be framed as a binary classification task, which has been solved by employing conventional machine learning algorithms such as decision trees [50], logit classifiers [7,34], k-nearest neighbors [13], neural networks [10], or support vector machines (SVMs) [45]. Alternatively, recent studies have tackled the churn prediction problem with the current trends in artificial intelligence: deep learning [49], big data/Hadoop ecosystem [41], and social network analysis via graph theory [8].

Formally, the estimation of customer attrition consists of predicting a binary label, $Y \in \{0, 1\}$. The task is to accurately identify customers with a high propensity to churn. Hence, a classification model, $\mathcal{C}$, is trained to return a probabilistic outcome $s \in [0, 1]$. Following the notation proposed in [47], at a threshold t, a new customer described by a vector of characteristics $\mathbf{x}$, is predicted to be a churner ($y = 0$), if $s \leq t$, or classified as non-churner ($y = 1$) otherwise. It is worth noting that, by convention, $y = 1$ is used for churners and $y = 0$ for non-churners in many machine learning studies.

The results of the classification problem can be summarized in a confusion matrix that displays the proportions of correct (diagonal) and incorrect (off-diagonal) predicted test cases (see Table 1). These proportions vary according to the chosen threshold. The prior probabilities of class class 0 (churner) and 1 (non-churner) are represented as π_0 and π_1, respectively. Moreover, the cumulative density functions for churners and non-churners at threshold t are indicated as $F_0(t)$ and $F_1(t)$, respectively.

Table 1
Confusion matrix.

		Predicted	
		Churner	Non-churner
Actual	Churner	$\pi_0 F_0(t)$	$\pi_0\big(1-F_0(t)\big)$
	Non-churner	$\pi_1 F_1(t)$	$\pi_1\big(1-F_1(t)\big)$

The main goal of predicting customer churn is to perform a targeted retention campaign on the fraction of customers η with the highest risk of attrition. Contacting a customer leads to a cost of f, while the incentive offered has a monetary value of d which is only incurred if the customer accepts the offer. It is expected that all false would-be churners (customers incorrectly labelled as churners) accept the incentive, since their intention has never been to churn. However, only a fraction γ of the contacted would-be churners, i.e., correctly classified as churners, would accept the incentive and choose to remain in the company [47]. The latter group of customers have the largest impact on the profit since retaining them imply gaining their customer lifetime value (CLV). The CLV tends to be significantly larger than the campaign costs f and d. This framework considers that no benefits are gained from: the fraction of customers that is not contacted, the retained false would-be churners, and the fraction $(1-\gamma)$ that leaves despite the retention campaign [47]. This leads to the following formula for *the average profit of a retention campaign* given a threshold t:

$$P_{\mathcal{C}}(t;\gamma,CLV,\delta,\phi) = CLV(\gamma(1-\delta)-\phi)\pi_0 F_0(t) - CLV(\delta+\phi)\pi_1 F_1(t), \quad (1)$$

where $\delta = d/CLV$ and $\phi = f/CLV$ are the two dimensionless parameters introduced by Verbraken et al. [47].

Verbeke et al. [44] proposed the maximum profit criterion (MPC) that maximizes the profit $P_{\mathcal{C}}(t;\gamma,CLV,\delta,\phi)$ by optimizing t. The MPC is defined as follows:

$$\text{MPC} = \max_{\forall t} P_{\mathcal{C}}(t;\gamma,CLV,\delta,\phi), \quad (2)$$

and considers the parameters in Eq. (1) deterministic. Moreover, the fraction of customers that is targeted given model predictions, η_{MPC}, and the optimal threshold t^* is calculated as follows:

$$\eta_{MPC} = \pi_0 F_0(t^*) + \pi_1 F_1(t^*). \quad (3)$$

The expected maximum profit criterion (EMPC) [47] is an extension of the MPC metric. It allows the introduction of uncertainty in the estimation of profit by assuming that γ is a beta distributed random variable. Formally, the EMPC is the maximization of the expected value of the profit, and is defined as follows:

$$\text{EMPC} = \int_{\gamma} P_{\mathcal{C}}(t^*(\gamma);\gamma,CLV,\delta,\phi)\cdot h(\gamma)d\gamma, \quad (4)$$

where $t^*(\gamma)$ and $h(\gamma)$ are the optimal threshold that leads to the maximum value for $P_{\mathcal{C}}(t;\gamma,CLV,\delta,\phi)$, and the probability density function associated to the random variable γ, respectively.

Subsequently, from the EMPC, we can compute the fraction of customers to be contacted η_{EMPC} as follows:

$$\eta_{EMPC} = \int_{\gamma} [\pi_0 F_0(T(\gamma)) + \pi_1 F_1(T(\gamma))]\cdot h(\gamma)d\gamma. \quad (5)$$

The objective of the MPC and the EMPC is to determine the best model for churn prediction by assessing the performance of several classifiers. The intuition behind the profit evaluation framework is that conventional evaluation techniques such as accuracy, sensitivity and AUC, are not goal-oriented metrics, and hence suboptimal to evalute classification approaches from a business-driven perspective [44,47]. Therefore, this study extends the profit metric to the prediction of churn in the mutual fund industry.

Note that the profit-driven framework has been extended to other classification domains in business analytics such as credit scoring [46]. Moreover, the intuition behind profit metrics has been applied to other predictive analytics tasks. Examples of these are the model and feature selection for SVM classification [30], and the design of novel classification approaches that maximize profit during model training [21,42]. For the latter extension, profit-driven classifiers has been proposed by adapting existing models, such as decision trees [21], logistic regression [21], and robust optimization models for maximum-margin classification [31].

Several studies have emphasized the importance of using individual CLVs rather than a single value (generally the average CLV); see [2,3,36]. Although the framework by Verbeke et al. [44] allows taking on a distribution for the CLV, this has not been done empirically in most EMPC studies. The main drawback is the lack of datasets that include the CLV of each individual. Therefore, our proposal represents an important contribution since experiments with individual CLVs are performed, demonstrating the importance of a multisegment, multithreshold approach. The idea that firms want to optimally allocate limited resources to heterogeneous segments to maximize profit has been discussed on simulated data for general-purpose migration models using, e.g., Markovian processes [27]. We extend this idea to real-world setting for churn prediction in the mutual fund industry incorporating profit metrics.

3. Proposed framework for customer attrition from investment firms

The proposed profit-based framework for churn prediction in the mutual fund industry is introduced in this section. The main goal is to seize the–sometimes extreme–differences among CLVs that are observed in this industry by developing a multisegment, multithreshold approach. In particular, we extend the MPC and EMPC measures for effectively addressing individual CLVs. The proposal is presented in two parts. First, the profit metric extensions are depicted in Section 3.1 as generic approaches for churn prediction. Subsequently, relevant aspects of the mutual fund industry and its relation with the profit measures are discussed in Section 3.2.

3.1. A multithreshold approach for profit-based classification

Our proposal considers a trained classifier $\mathcal{C}$ that produces a probabilistic output for each customer. We define a set of segments $\mathcal{I}$ in which its cardinality $|\mathcal{I}| = q$ is set a priori. Each customer is assigned to a segment $i \in \mathcal{I}$ by obtaining q-quantiles of the CLV distribution, i.e., the fraction of $1/q$ customers with the highest CLVs are assigned to the first segment ($i = 1$), and the ones between the cut points $1/q$ and $2/q$ are assigned to the second segment. Each segment also has a corresponding threshold $\tau_i \in [0,1]$: given a new sample $\mathbf{x}$ that belongs to the i-th segment, if $s \leq \tau_i$, then this customer is classified as a churner; otherwise, $\mathbf{x}$ is classified as a non-churner. The multisegment, multithreshold framework is illustrated in Fig. 1.

Following [47], the parameters that are required for implementing the profit of a retention campaign are the cost of contacting a customer f, the monetary cost of the incentive d, the CLV, and the fraction γ of the would-be churners that accept this incentive and remain with the company. This latter parameter can either be deterministic (as in the MPC measure) or stochastic (a beta-

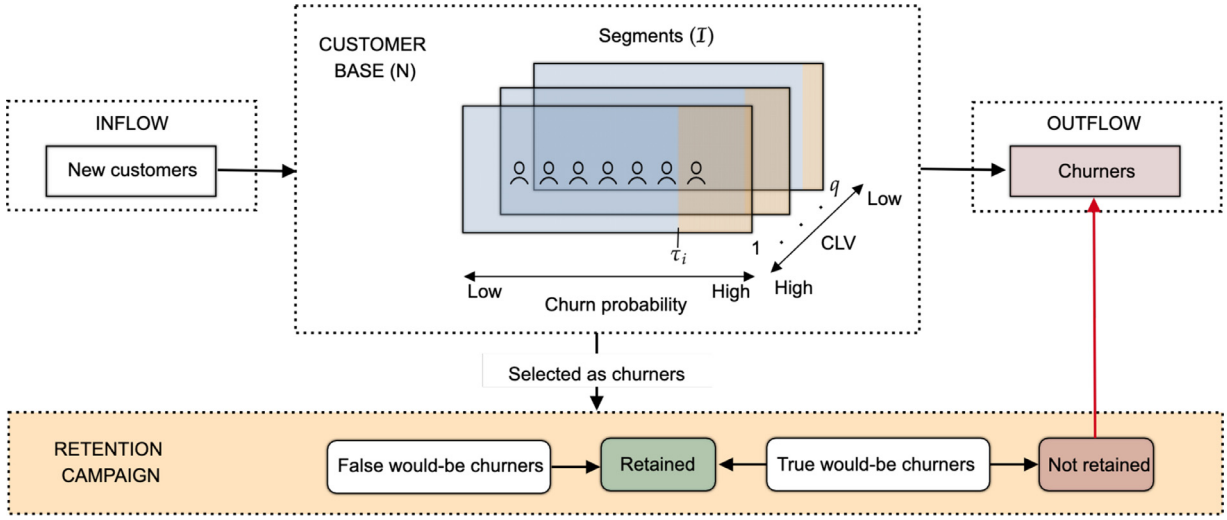


Fig. 1. Multisegment, multithreshold framework: 1) Customers are assigned to a segment $i \in \mathcal{I}$ based on their CLV, 2) the retention campaign is applied to customers whose $s \leq \tau_i$, 3) among the targeted customers (highlighted in light orange), all false would-be churners are retained, whereas only a fraction of the true would-be churners are retained.

distributed random variable, as in the EMPC measure). The average profit of a classifier is as follows:

$$P_{\mathcal{C}}(\boldsymbol{\tau}; \gamma, \mathbf{CLV}, \delta, \phi) = \frac{\sum_{i=1}^{q} [CLV_i(\gamma(1-\delta_i)-\phi_i)\pi_{0,i}F_{0,i}(\tau_i) - CLV_i(\delta_i+\phi_i)\pi_{1,i}F_{1,i}(\tau_i)]}{q} \quad (6)$$

where the parameters $\pi_{0,i}$ and $\pi_{1,i}$ are the prior probabilities of class 0 (churner) and 1 (non-churner), respectively, and $F_{0,i}(\tau_i)$ and $F_{1,i}(\tau_i)$ are the cumulative density functions for churners and non-churners for a given threshold τ_i of a given customer from segment i, respectively. Additionally, CLV_i represents the average CLV for segment i, and $\delta_i = d/CLV_i$ and $\phi_i = f/CLV_i$. Finally, $\boldsymbol{\tau}$ and $\mathbf{CLV}$ are the vectors of thresholds and average per-segment CLVs, respectively.

The number of segments is, in principle, a decision variable in our approach. Note that companies usually have predefined customer segments in terms of CLVs or other variables. In such cases, our model remains valid. However, q is not longer a decision variable, as the segment-wise CLVs are computed based on arbitrary segments rather than on CLV quantiles.

The proposed *maximum segment profit criterion* (MSPC) extension corresponds to maximizing Eq. (6) with respect to τ_i, $i = 1, \ldots, q$:

$$\text{MSPC} = \max_{\forall \boldsymbol{\tau}} P_{\mathcal{C}}(\boldsymbol{\tau}; \gamma, \mathbf{CLV}, \delta, \phi). \quad (7)$$

The fraction of customers to target η_{MSPC} given the optimal set of thresholds $\boldsymbol{\tau}^*$ is computed as follows:

$$\eta_{MSPC} = \frac{\sum_{i=1}^{q} \left(\pi_{0,i}F_{0,i}(\tau_i^*) + \pi_{1,i}F_{1,i}(\tau_i^*)\right)}{q}. \quad (8)$$

Subsequently, the proposed *expected maximum segment profit criterion* (EMSPC) is obtained as follows:

$$\text{EMSPC} = \int_{\gamma} P_{\mathcal{C}}(\boldsymbol{\tau}^*(\gamma); \gamma, \mathbf{CLV}, \delta, \phi) \cdot h(\gamma) d\gamma, \quad (9)$$

where $\boldsymbol{\tau}^*(\gamma)$ and $h(\gamma)$ are the optimal thresholds and the probability density function for γ, respectively. In this study, we take the same values that Verbraken et al. [47] considered for the parameters α and β. These are the inputs required for the beta distribution.

Finally, the fraction of the contacted customers η_{EMSP} that results from using the optimal set of thresholds $\boldsymbol{\tau}^*$ is given by:

$$\eta_{EMSP} = \frac{\sum_{i=1}^{q} \int_{\gamma} \left[\pi_{0,i}F_{0,i}(\tau_i^*(\gamma)) + \pi_{1,i}F_{1,i}(\tau_i^*(\gamma))\right] \cdot h(\gamma) d\gamma.}{q} \quad (10)$$

Note that our proposal involves training a single model for a given classifier. When the model is applied in an independent dataset for prediction purposes, the multisegment profit metric is computed as an aggregation of the individual CLVs. Hence, our approach has similar running times, as other evaluation methods. However, the additional operations consists of simple arithmetical computations. Alternatively, multiple classification models could be considered; one for each segment. Nevertheless, using a single classifier allows to leverage bigger datasets which typically yields a better classification model and reduce the risk of overfitting.

3.2. Important considerations for the mutual fund industry

In this section, we discuss the parameters of the proposed metrics and their relation with the mutual fund industry. The individual CLV is computed as the net present value of the future revenues generated by a customer j and the costs of serving him/her [5,43]:

$$CLV_j = \sum_{t=1}^{T} \frac{(Revenue_{j,t} - Cost_{j,t})r_t}{(1+r)^t} \quad (11)$$

where r is the discount rate and T is the time horizon. Following the framework designed by Safari et al. [40] for the banking industry, we consider $T = 3$ years, r equals the inflation rate, and the direct costs about the accounts $Cost_{j,t}$ to be negligible. Parameter r_t is the retention rate for a given period, i.e., the probability of a customer staying in the company at time t [36].

According to the company that provided the data, the commission that is charged can be viewed as a percentage of the total amount of money invested by a customer on a given fund. This percentage cannot be disclosed due to confidentiality reasons. The $Revenue_{j,t}$ parameter is computed based on the commissions paid by customers.

The variability in terms of money invested by customers leads to differences across customer segments in terms of CLVs. Hence, the importance of using a multisegment approach in the mutual

Table 2
Dataset variables.

Source	Category	Variables
Demographics		Age, gender, marital status, tertiary studies, and city of residence.
Behavioral		Recency-redemption, frequency-redemption, monetary-redemption, recency-repurchase, frequency-repurchase, monetary-repurchase, and customer seniority.
	Investor profile	Type of investment and company survey.
Financial	Investor portfolio	Percentage of short-term funds, cumulative return of the funds in the week before churn is predicted, cumulative return of the funds in the month before churn is predicted, percentage of funds invested in United States of America, and Percentage of funds invested in Chile.

fund industry. The average CLVs for all segments obtained with our data will be illustrated in the next section. Once the individual CLVs are computed, the q-quantile strategy described in the previous section is applied for creating the segments.

Unfortunately, the literature on churn prediction for the banking industry does not provide guidelines for tailoring the remaining parameters: the cost of the incentive d associated with the marketing campaign, the cost of contacting a customer f, and the rate γ of the would-be churners that accept the incentive and remain with the company. Therefore, we align with the existing literature on profit metrics for telecommunications to set these parameters. Additionally, the company does not have data from previous retention campaigns that would allow us to estimate these parameters empirically. The design of customized retention campaigns is a topic that falls outside the scope of our study. Particularly, value-adding services and products such as personalized counselling and incentives are offered by financial entities to improve customer experience and satisfaction. Subsequently, these interventions can be evaluated and further optimized by setting a randomized control trial or using suitable methods to estimate individual treatment effects from observational data [19,28,48].

We fixed f to a value very close to 1€, which was used in the MPC and EMPC measures [44,47], and $d = \overline{CLV}/20$, where $\overline{CLV}$ is the average customer lifetime value. This relation for d was also proposed for the MPC and EMPC measures [44,47]. Finally, we made the same assumptions for parameter γ indicated in Verbeke et al. [44] and Verbraken et al. [47] for the MPC and EMPC measures, respectively.

Note that profit metrics aim at identifying the right customers to target in order to lead to profitable retention campaigns. In this sense, the only aspect of the mutual fund company's profit that is affected by this decision is the income due to successfully retained customers (the CLVs) and the campaign costs. There are many other relevant aspects that affect a mutual fund company's profitability, such as the returns of the investments, but these elements are not part of the decision scheme associated with the predictive model for customer attrition.

4. Experimental results

The proposed framework for assessing churn prediction models in the mutual fund industry is applied to data from a Chilean company. This section is organized as follows: Section 4.1 provides a description of the dataset. The experimental settings are presented in Section 4.2. Subsequently, the main results are summarized in Section 4.3. Finally, we perform a sensitivity analysis for the most relevant parameters considered in the proposed framework. The main insights obtained from these experiments are reported in Section 4.4.

4.1. Dataset description

The dataset contains information regarding (1) customer demographics, (2) customer behavior with the company in terms of redemption and repurchase activities, and (3) financial indexes related to the funds themselves (see Table 2). For customer demographics, we consider the following variables: age, gender, marital status, tertiary studies, and city of residence as a dichotomous variable.

Table 3
Numbers of customers, variables, and churn rate for all datasets.

	one-month churn	three-month churn
Nº customers (train)	2263	2305
Nº customers (test)	2205	2263
Nº attributes	24	24
Churn rate	6.33%	17.41%

The customer behavior data summarize information regarding the customers' seniority and redemption and repurchase activities based on the following variables: recency-redemption, frequency-redemption, monetary-redemption, recency-repurchase, frequency-repurchase, monetary-repurchase, and customer seniority. The variables were calculated by following the RFM (recency, frequency and monetary) framework [43]. *Recency* is computed as the number of months that have passed since the customer's last activity. *Frequency* is obtained as the number of activities performed by the customer over the past 12 months. Finally, *monetary* is computed for each customer as the average value of all redemption and repurchase activities.

The third source of information corresponds to indexes and other financial variables related to the performance of the funds and the profile of investors. The daily returns are monitored and can be used for predicting whether a customer will churn in the following months based on the average value of its share in a given month. Additionally, three important financial indexes were used as predictors: the Standard & Poor's 500 (S & P 500), the US Dollar to Chilean Peso (USD/CLP) exchange rate, and the IPSA (Selective Price of Shares Index). The IPSA is the main stock market index in Chile and is formed by 40 stocks with the highest average annual trading volume in the Santiago Stock Exchange.

Two datasets were obtained for the experimental evaluation, resulting from considering two different target variables. For the first dataset, churn is predicted for a one-month period (March 2018) using the customer base from February 2018. The second dataset consists of the customer base from January 2018, in which churn is predicted on a three-month period (February 2018 to April 2018). For both datasets, a whole year was considered as customer history for creating interaction variables, such as the RFM framework.

The relevant information for each dataset is summarized in Table 3.

4.2. Experimental setup

Following the seminal studies on profit metrics [44,47], we performed a comprehensible empirical analysis, evaluating and comparing the performance of various well-known classifiers using both profit-based and statistical measures in order to illustrate the virtues of the proposed metrics.

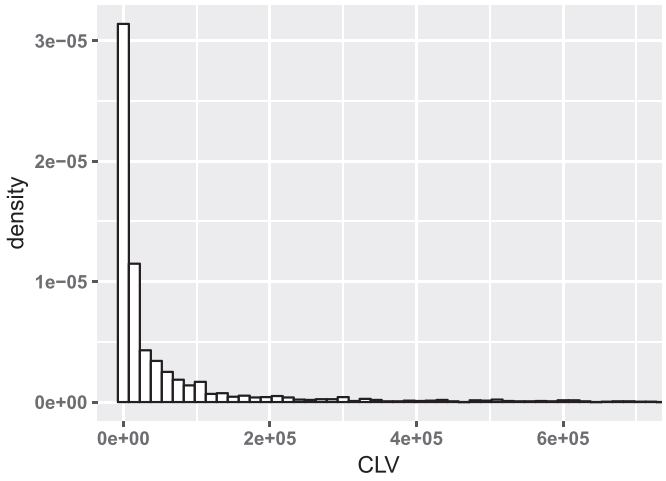
Ten-fold cross-validation is used in the training set for parameter tuning. Several metrics are computed for each fold (area under the curve (AUC), accuracy, sensitivity, specificity, MPC, EMPC, and the proposed MSPC and EMSP), and their average values are reported. AUC was used for model validation to avoid favoring any of the profit metrics.

The presence of class imbalance in both datasets may cause predictive models to favor the majority class. Therefore, we apply SMOTE oversampling [9]. This strategy creates synthetic examples from the minority class (the churners) by calculating the k nearest neighbors of each instance and interpolating between its neighbors. We use the traditional setting that creates $2n_1$ new examples, where n_1 is the cardinality of the minority class in the training set, and $k = 5$ [9]. It is important to note that profit metrics alleviate this issue by providing an adequate framework to balance the benefits and costs of classification. In churn prediction, the benefit of correctly predicted churners (minority class) outweighs the misclassification costs since it implies a gain, the CLVs, which are estimated to be about 20 times more valuable than the campaign cost.

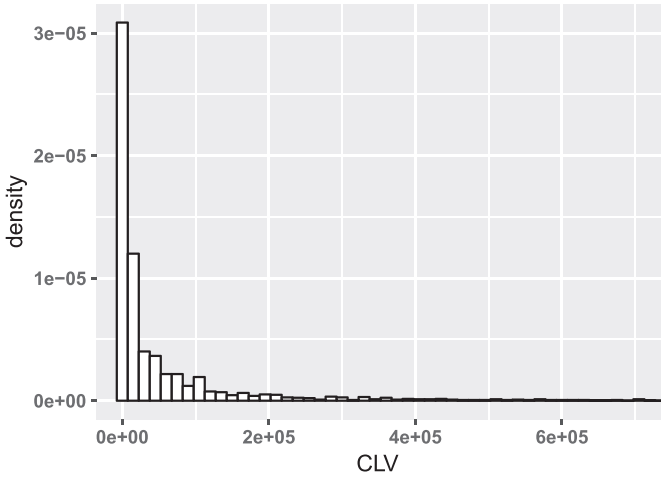
Based on the default configuration of the CARET toolbox for the R software, the following classification methods were studied:

- Standard logistic regression (Logit). This technique does not require parameter tuning. A backward variable elimination process was performed using the Akaike information criterion (AIC). The AIC criterion is a commonly used approach for variable selection in generalized linear models, as it is directly linked to the maximum likelihood function. This criterion also finds a good compromise between goodness of fit and model simplicity, discouraging overfitting and improving out-of-sample prediction [51].
- k-nearest neighbors (k-NN). Odd numbers were studied between 5 and 41 using cross-validation. The final values for this parameter on both datasets are $k = 35$ (one-month churn) and $k = 39$ (three-month churn).
- CART decision tree (CART). Classification and regression trees (CARTs) are a well-known strategy for constructing decision trees. The complexity parameter (CP) is set via cross-validation. The following are the optimal values: $CP = .0092$ (one-month churn) and $CP = 0.0045$ (three-month churn).
- Random forest (RF). For this method, 500 trees were constructed for each classifier. The number of attributes considered for the construction of each tree was set via cross-validation, exploring in the values $\{2, 4, 6\}$. The optimal number of attributes was set to 2 in both datasets.
- Artificial neural network (ANN). Single-layer neural networks were constructed using 500 epochs. We use the default of one hidden layer as it is sufficient for the large majority of problems according to the universal approximation theorem [22]. For both datasets, the number of hidden units and the weight decay parameter were set to 23 and 0.1, respectively, using cross-validation.
- Support vector machine (SVM). Soft-margin SVM classifiers were implemented using a Gaussian kernel. The radial basis function (RBF) or Gaussian kernel is a well known default kernel, that is flexible for modeling nonlinear problems. This kernel function has shown the best predictive results in several empirical studies for classification [15]. The following values for C and γ were used: $C \in \{0.001, 0.01, 0.1, 0.25, 0.5, 1, 2, 4, 8, 16, 32, 64, 128\}$, $\gamma \in \{0.001, 0.01, 0.1, 0.25, 0.5\}$. The best configurations were $C = 2$ and $\gamma = 0.001$ and $C = 5$ and $\gamma = 0.001$ for the one-month and three-month churn datasets, respectively.

For the profit measures, we consider the parameters proposed by Verbeke et al. [44] and Verbraken et al. [47]. The probability γ is set to 0.3 for the MPC and MSPC, while for the EMPC and EMSP, it is assumed that it follows a beta distribution with $\alpha = 6$ and $\beta = 14$ as shape parameters. We compute the average and segment-based CLVs in Chilean Pesos (CLP, CLP/€ exchange rate equals 806.75 on October 24, 2019). The parameter f is set to 1000 CLP (1.24 €), which is very close to the value in Verbeke et al. [44] and Verbraken et al. [47] ($f = 1$€). The retention parameter r_t is set to 1 for the entire time horizon based on the assumption that a customer that accepts the incentive will remain loyal to the company.



(a) 0ne-month dataset



(b) Three-months dataset

Fig. 2. Distribution of the CLV for the One-month and Three-months churn datasets.

Based on our data, the average customer lifetime values $\overline{CLVs}$ in the test set are 65,582.98 CLP and 65,414.55 CLP for the one-month and three-month churn datasets, respectively (approximately 81.5€). Panel A and B of Fig. 2 display the CLV distribution for the one-month and three-months churn datasets, respectively. The red lines represent the exponential distribution fitted for these variables. As it can be observed, the CLVs follow an exponential curve. This is somewhat expected given our estimation for the CLVs as a percentage of the amount invested. The investment amounts follow an exponential curve similar to other monetary variables used in risk assessment, such as the income of the investor.

Table 4
Descriptive statistics of the CLVs.

Dataset	Min.	Q1	Med.	Mean	Q3	Max.
One-month churn	387	1626	9756	65,583	50,300	5,489,297
Three-months churn	387	1625	9553	65,415	48,967	5,489,297

Table 5
Test performance and fraction of contacted customers for all classification methods. One-month churn dataset.

Metric	Logit	k-NN	CART	RF	ANN	SVM	$\overline{X}$
AUC	70.64	66.21	63.90	71.77	70.77	69.26	**68.76**
Accu.	82.99	63.51	75.66	87.46	79.96	63.08	**75.44**
Sens.	34.25	64.38	34.93	25.34	40.41	65.75	**44.18**
Spec.	86.29	63.32	78.42	91.76	82.63	62.90	**77.55**
MPC	9.78	0.00	7.92	24.19	3.96	5.82	**8.61**
η_{MPC}	0.22	0.00	0.26	6.12	0.13	0.09	**1.14**
MSPC	405.3	299.4	229.7	353.2	412.3	388.5	**348.08**
η_{MSPC}	28.03	15.15	50.01	37.29	32.73	27.70	**31.82**
EMP	21.81	1.67	8.39	51.15	13.70	9.07	**17.63**
η_{EMP}	2.63	0.49	0.38	3.21	2.22	0.97	**1.65**
EMSP	408.8	323.2	248.4	366.6	416.7	394.1	**359.64**
η_{EMSP}	3.62	5.39	7.96	5.61	3.78	5.27	**5.27**

In addition, Table 4 summarizes some descriptive statistics of the CLVs for the one-month and three-months churn datasets. The differences between the median and the mean and the distribution of the quartiles clearly shows that the data is skewed to the right or positively skewed. It is important to note that the distribution of the CLVs is likely to vary from application to application. In telecommunications, for example, such a high variance among the customers is not expected, and the CLVs may follow another data distribution. Further studies on the data distribution of the CLVs are encouraged as future work in the final section of this study.

The parameter d is rather complex to set without having information on prior retention campaigns in the mutual fund industry. Therefore, we explore a wide range of values: $d \in \{0, \overline{CLV}/20, \overline{CLV}/15, \overline{CLV}/10, \overline{CLV}/5, \overline{CLV}/3\}$. Finally, the set of segments q is $q \in \{2, 3, 4, 5, 8, 10, 11, 15\}$.

4.3. Results summary

In order to validate the assumption that the net value of the mutual fund correlates with the churn behavior, we computed the average churn rate and the average return of the funds one week before the month of the churn assessment for the 24 months available for this study (January 2016 to December 2017). We observed a positive non-significant correlation between the churn rate and the return (Pearson's ρ=0.103, p-value=0.63, N=24), indicating that investors tend to wait for above-average market conditions before withdrawing their investments.

We observed, however, an interesting pattern: the churn rate increased in periods with abnormally low or high returns. This indicates that some investors assume their losses and wait for positive market conditions, while others prefer to withdraw their funds in periods of risk and uncertainty. We computed the absolute value of the returns minus their average. A large value for this variable implies a large deviation from the average returns in the 24-month period. The Pearson's correlation between the churn rate and this new variable was 0.319 (p-value=0.06, N=24), which is significant for $\alpha = 0.1$. This pattern can be confirmed at individual level since the variable used for this analysis (return of the funds one week before the month of the churn assessment) was one of the most relevant ones in the various classification models considered in this study.

Tables 5 and 6 report the performance on the test set for each classification technique for the one-month and three-month churn datasets, respectively. In addition to the profit metrics discussed in this study (MPC, MPSC, EMP, and EMSP), traditional performance metrics such as the area under the curve (AUC), accuracy (Accu.), sensitivity (Sens.), and specificity (Spec.) were included for completeness.

We present the results for $d = \overline{CLV}/20$, which is the same proportion used in the studies by Verbeke et al. [44] and Verbraken et al. [47]. The MSPC and EMSP measures presented in these tables are obtained with $q = 10$ and $q = 3$ for the one-month and three-month churn datasets, respectively. Additionally, the proportion of contacted customers η is reported in Table 7. For each table, the average of each metric is presented in the last column and emphasized in bold.

Tables 5 and 6 show that the proposed multithreshold approach is more beneficial than the single-segment framework in terms of the generated profit. In the one-month churn dataset, the average MSPC and EMSP are 40 and 20 times larger than the MPC and EMP, respectively. Comparably, in the three-month churn dataset, the MSPC and EMSP are 2.4 and 2.1 times larger than the standard profit metrics. Both results demonstrate that targeting customers according to the multisegment approach is optimal because this framework takes the differences in customer value into account. The dramatic improvement in terms of profitability resulting from the multisegment approach is because more customers with a high

Table 6
Test performance and fraction of contacted customers for all classification methods. Three-month churn dataset.

Metric	Logit	k-NN	CART	RF	ANN	SVM	$\overline{X}$
AUC	68.76	67.68	64.49	72.33	69.25	70.73	**68.87**
Accu.	56.43	64.07	57.05	66.24	58.77	64.12	**61.11**
Sens.	74.62	58.63	76.65	64.72	72.08	65.23	**68.65**
Spec.	52.59	65.22	52.92	66.51	55.97	63.88	**59.51**
MPC	698.80	639.47	697.20	836.11	753.78	773.83	**733.20**
η_{MPC}	30.80	43.53	50.11	46.18	29.74	47.64	**41.33**
MSPC	1,755.1	1,717.7	1,697.0	1,811.5	1,745.2	1,804.8	**1,755.22**
η_{MSPC}	50.32	53.14	56.98	31.71	57.65	45.59	**49.23**
EMP	830.2	769.2	803.2	970.4	862.3	913.1	**858.06**
η_{EMP}	37.09	40.92	44.12	40.14	35.33	37.45	**39.18**
EMSP	1,775.1	1,737.6	1,724.9	1,817.4	1,769.0	1,818.5	**1,773.75**
η_{EMSP}	28.49	29.99	30.12	27.03	28.80	27.69	**28.69**

Table 7
Fraction of contacted customers for all classification methods. One- and three-month churn datasets.

Dataset	Metric	Logit	*k*-NN	CART	RF	ANN	SVM	$\overline{X}$
One-month	η_{MPC}	0.22	0.00	0.26	6.12	0.13	0.09	**1.14**
	η_{MSPC}	28.03	15.15	50.01	37.29	32.73	27.70	**31.82**
	η_{EMP}	2.63	0.49	0.38	3.21	2.22	0.97	**1.65**
	η_{EMSP}	3.62	5.39	7.96	5.61	3.78	5.27	**5.27**
Three-month	η_{MPC}	30.80	43.53	50.11	46.18	29.74	47.64	**41.33**
	η_{MSPC}	50.32	53.14	56.98	31.71	57.65	45.59	**49.23**
	η_{EMP}	37.09	40.92	44.12	40.14	35.33	37.45	**39.18**
	η_{EMSP}	28.49	29.99	30.12	27.03	28.80	27.69	**28.69**

Table 8
Fraction of contacted customers per segment for various *d* values. Random forest as classification method. One-month churn dataset.

Segment	CLV_i	d					
		0	$\overline{CLV}/20$	$\overline{CLV}/15$	$\overline{CLV}/10$	$\overline{CLV}/5$	$\overline{CLV}/3$
1	448,002	84.7	49.2	38.6	30.7	10.2	1.7
2	102,952	54.9	4.1	3.1	2.1	1.2	0.9
3	51,714	22.1	2.4	1.4	0.8	0.2	0.0
4	26,158	20.5	0.0	0.0	0.0	0.0	0.0
5	13,245	0.5	0.4	0.4	0.4	0.0	0.0
6	6625	0.1	0.0	0.0	0.0	0.0	0.0
7	3504	0.0	0.0	0.0	0.0	0.0	0.0
8	1617	0.0	0.0	0.0	0.0	0.0	0.0
9	776	0.0	0.0	0.0	0.0	0.0	0.0
10	424	0.0	0.0	0.0	0.0	0.0	0.0
	η_{EMSP}	18.3	5.6	4.3	3.4	1.2	0.3

Table 9
Fraction of contacted customers per segment for various *d* values. Random forest as classification method. Three-month churn dataset.

Segment	CLV_i	d					
		0	$\overline{CLV}/20$	$\overline{CLV}/15$	$\overline{CLV}/10$	$\overline{CLV}/5$	$\overline{CLV}/3$
1	183,917	98.3	81.1	77.4	65.3	22.4	3.2
2	11,059	26.9	0.0	0.0	0.0	0.0	0.0
3	1111	0.0	0.0	0.0	0.0	0.0	0.0
	η_{EMSP}	41.7	27.0	25.8	21.8	7.5	1.1

CLV are targeted, even when having a low probability to churn in relation to the other segments. This fact will be discussed further in the following section.

In addition, Tables 5 and 6 also indicate that choosing the classifier with the best performance depends on the evaluation metric. In the one-month churn dataset, for example, random forest is the best performing algorithm with respect to MPC and EMP. However, in terms of the multithreshold approach, the ANN algorithm is more suitable for targeting. For the three-month dataset, the RF classifier outperforms other techniques in terms of the MPC, EMP, and MSPC but slightly not for the EMSP.

The average profit measures are larger in the three-month churn dataset (Table 6) than in the one-month churn dataset (Table 5). The traditional measures AUC and accuracy suggest that classification performance is better in the latter. This can be explained by the fact that the three-month churn dataset has a larger churn ratio (17.41% versus 6.33%), leading to campaigns that consider a large fraction of customers.

Last, Table 7 suggests that the large profit achieved with the MSPC and EMSP measures does not necessarily imply that more customers are contacted. In the three-month churn dataset, for example, the average expected profit of the campaign increases from 858.06 (single segment) to 1,773.75 (multiple segments), but the average proportion of contacted customers decreases from 39.18% to 28.69%.

4.4. Sensitivity analysis and discussion

In this section, we perform sensitivity analyses on the parameters that are relevant to estimating the EMSP. The random forest classifier is chosen for performing the experiments since it generally achieves the best performance among all other classification algorithms. Tables 8 and 9 show, from left to right, the chosen segments q, the per-segment customer lifetime value CLV_i, and the optimal proportion of targeted customers within each segment for increasing values of d (cost of the retention campaign). In addition, the total fraction of contacted customers η_{EMSP} for a given d value is reported at the bottom of each table. The number of segments is ten ($q = 10$) and three ($q = 3$) for the one-month and three-month datasets, respectively.

Tables 8 and 9 show that, as expected, our approach prioritizes targeting valuable customers, which are customers within segments with a large CLV_i (top segments). For example, when the cost of the retention campaign is set to zero in the three-month churn dataset, optimally 98.3% (33.3% of the customer base), 26.9% and 0% of the customers in the first, second and third segments are contacted, respectively; see Table 9. Consequently, as the retention campaign becomes more expensive, only customers in segments with large CLVs are targeted, and the fraction of contacted customers decreases.

Note that our approach is not equivalent to segmenting first and then computing the MPC and EMP measures for each segment. Such an strategy fails to prioritize customers with large CLVs, since it treats each segment independently. For $d = \overline{CLV}/20$ in Table 9, for example, our approach targets only customers with large CLVs and ignores those from the remaining segments. A cluster-first-then-classify approach would not be able to identify that the premium segment should be prioritized over the remaining two.

Next, Figs. 3 and 5 illustrate the classifiers' EMSP variation when increasing the cost of the retention campaign for the one-month and three-month churn datasets, respectively. In addition, Figs. 4 and 6 show for each classifier the optimal fraction of customers to target, η_{EMSP}, at different values of d, for the one-month and three-month churn datasets, respectively.

Figs. 3 and 5 show that all classifiers perform similarly for increasing values of d. When the cost of the retention campaign is low, the EMSP and η_{EMSP} are nearly equal for all algorithms. This leads to a very important conclusion: the major gains in terms of profit achieved using the ESMP measure are mostly due to its ability to direct the retention campaign to customers who are more valuable to the company, i.e., multithreshold and multisegment framework, rather than because of the variability in the performance of the classifiers. Nonetheless, some differences are observed for the results of the one-month dataset, as well as when d becomes larger. As expected, both the profit and the size η_{EMSP} decrease as the retention campaign becomes more expensive. We do

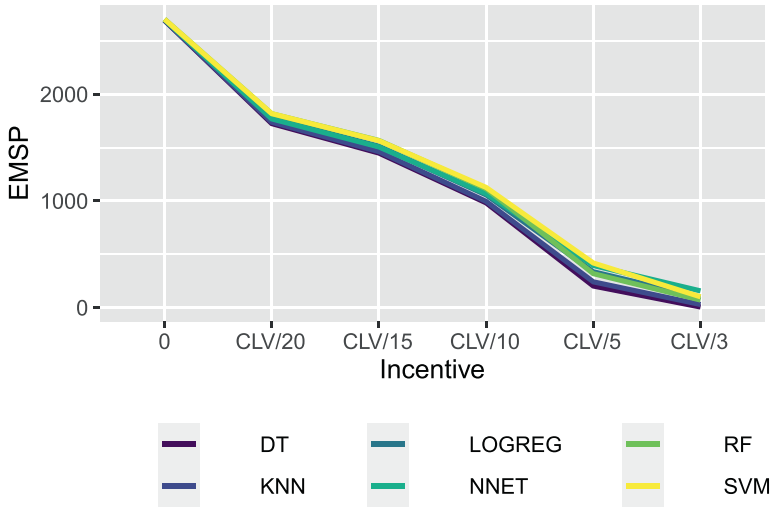


Fig. 3. Average ESMP for an increasing value for the incentive d. All classification approaches. One-month churn dataset.

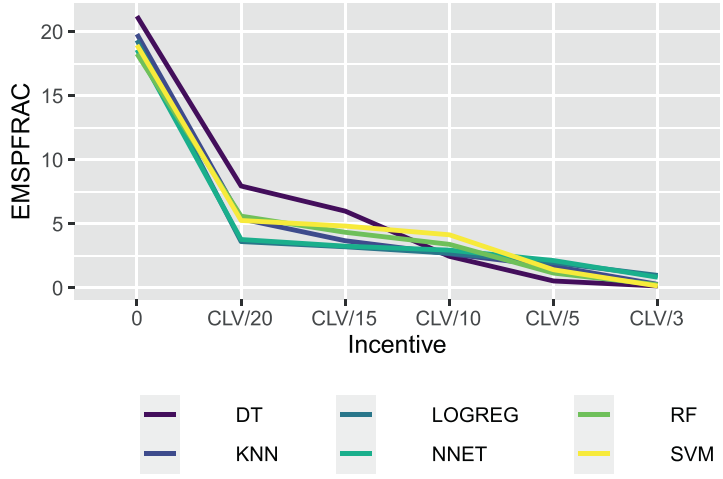


Fig. 4. Average η_{EMSP} for an increasing value for the incentive d. All classification approaches. One-month churn dataset.

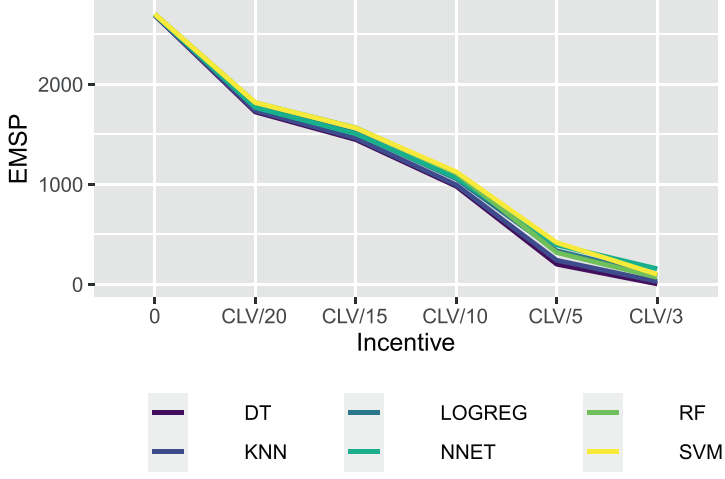


Fig. 5. Average ESMP for an increasing value for the incentive d. All classification approaches. Three-month churn dataset.

not report the results of the MSPC as they are analogous to those of the EMSP.

Tables 10 and 11 report the EMSP and η_{EMSP} measures for the six classifiers in increasing number of segments for the one-month and three-month churn datasets. The cost of the retention campaign, d, is set to $\overline{CLV}/20$, as suggested by the literature on profit measures (see, e.g., [44,47]). The largest EMSPs and their corresponding η_{EMSP} are highlighted in bold. It is observed that there is not a clear relationship between the EMSP and the number of segments. For example, the largest EMSPs are achieved for most of the

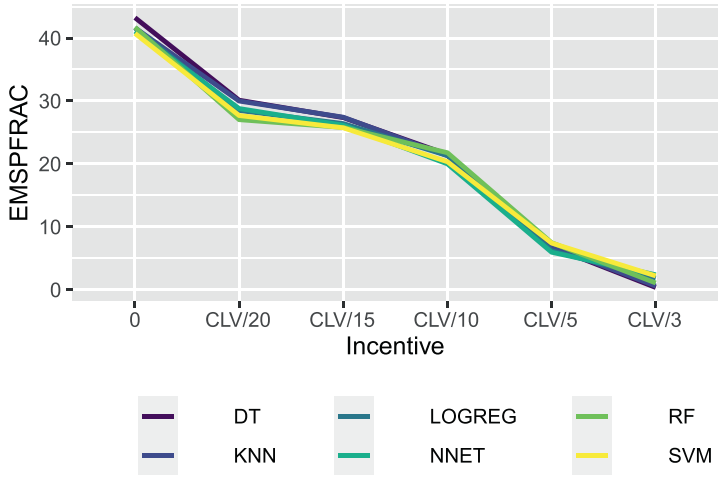


Fig. 6. Average η_{EMSP} for an increasing value for the incentive d. All classification approaches. Three-month churn dataset.

Table 10
Expected profit (EMSP) and fraction of contacted customers (η_{EMSP} in brackets) for various q values (number of segments). All classification methods. One-month churn dataset.

Number of segments (q)							
Method							
	3	4	5	8	10	11	15
Logit	253.4 (8.1)	192.2 (5.7)	274.5 (6.5)	301.7 (4.2)	408.8 (3.6)	385.6 (3.5)	346.1 (3.0)
k-NN	220.6 (9.4)	131.0 (6.4)	220.1 (6.2)	263.0 (5.1)	323.2 (5.4)	306.0 (5.6)	194.7 (4.8)
CART	98.0 (5.7)	60.1 (4.5)	101.0 (7.9)	178.2 (8.4)	248.4 (8.0)	261.4 (7.4)	165.4 (4.0)
RF	271.7 (5.5)	207.9 (5.2)	270.0 (6.9)	310.9 (6.2)	366.6 (5.6)	354.5 (5.5)	279.6 (4.5)
ANN	319.0 (9.5)	224.8 (7.0)	326.3 (6.2)	357.3 (4.1)	**416.7** (**3.8**)	395.0 (3.6)	350.1 (4.3)
SVM	221.9 (9.0)	148.2 (7.0)	235.4 (8.6)	289.4 (6.6)	394.1 (5.3)	396.1 (4.8)	281.0 (5.3)

Table 11
Expected profit (EMSP) and fraction of contacted customers (η_{EMSP} in brackets) for various q values (number of segments). All classification methods. Three-month churn dataset.

Number of segments (q)							
Method							
	2	3	4	5	8	10	11
Logit	1,561.4 (36.9)	1,775.1 (28.5)	1,565.9 (21.8)	1,624.8 (19.6)	1,585.3 (16.9)	1,649.5 (16.8)	1,572.1 (17.9)
k-NN	1,556.2 (40.2)	1,737.6 (30.0)	1,541.6 (21.8)	1,625.3 (18.2)	1,604.9 (15.9)	1,658.1 (17.0)	1,590.5 (16.5)
CART	1,518.1 (38.6)	1,724.9 (30.1)	1,492.6 (23.6)	1,532.5 (21.7)	1,514.8 (18.3)	1,580.5 (18.1)	1,526.5 (17.2)
RF	1,650.4 (37.3)	**1,817.4** (**27.0**)	1,563.7 (21.6)	1,618.6 (18.5)	1,567.1 (17.4)	1,646.9 (17.6)	1,576.8 (18.3)
ANN	1,570.6 (35.6)	1769 (28.8)	1,544.2 (22.6)	1,602.4 (19.6)	1,557.9 (17.9)	1,632.9 (18.3)	1,564.6 (18.6)
SVM	1,603.1 (35.7)	1,818.5 (27.7)	1,575.8 (21.8)	1,644.7 (19.6)	1,600.6 (17.1)	1,654.9 (17.0)	1,580.2 (18.3)

classifiers when the number of segments is set to 8 and 3 in the one-month churn dataset and three-month churn dataset, respectively. Moreover, the proportions of customers to target (in brackets) generally tend to decrease for a larger number of segments, but this does not necessarily lead to a reduction in the EMSP.

Based on our sensitivity analysis, we suggest using a validation procedure for setting the number of segments as another hyperparameter to be tuned, while the individual CLVs should be obtained according to the company data. It is also advisable to consider different scenarios with varying campaign costs. This facilitates the evaluation of different campaign scenarios in terms of costs and fraction of contacted customers such that the optimal proportions of customers are targeted at a reasonable cost. Furthermore, our experiments demonstrate that using traditional evaluation techniques for choosing a classifier does not necessarily maximize the profit of a retention campaign, emphasizing the importance of a segment-based profit metric for performance evaluation (see Tables 5 and 6).

5. Conclusions

This paper presents a novel profit-driven framework for churn prediction. The main methodological contribution is the extension of the MPC [44] and EMP [47] measures for assessing the profit of a retention campaign to be able to handle multiple segments of customers. A multithreshold approach is proposed for this task, which allows contacting a different proportion of customers in the various segments. In this way, we are able to target profitable customers that are likely to churn.

The main applied contribution is the design of a framework for assessing attrition in the mutual fund industry. Unlike traditional churn prediction tasks, such as in telecommunications and banking, the investment sector is characterized by heterogeneous CLVs that depend on the amount invested. Furthermore, attrition depends not only on the sociodemographic and behavioral variables associated with the customers but also on the mutual fund performance indicators. The proposed framework is able to adequately capture the heterogeneity of the customers in terms of CLVs while incorporating the various sources of information as covariates for the classifier.

Furthermore, the particularities of the mutual funds are mainly considered in the variables included in the classifiers (see Section 4.1) and the method to estimate the CLVs (see Section 3.2). Our key contribution in comparison with the existing literature on profit metrics, is the segment-wise decision scheme that distinguishes and prioritizes customers based on the predictions of the classifier and customers heterogeneity, i.e., their value to the company. In addition, we demonstrate the potential advantage of our approach on data from the mutual fund industry, a field relatively unexplored by the majority of scholars on profit metrics.

Experiments are performed on churn prediction datasets obtained from a mutual fund company. When using the same cost/benefit parameter configuration, our proposed MSPC and EMSP approaches achieve an average profit that is up to 40 times larger than their single-threshold counterparts (MPC and EMP). This result demonstrates the virtues of a multisegment, multi-threshold approach when the CLVs are heterogeneous. Our framework not only provides novel profit metrics for assessing the average performance of a retention campaign but also provides a tool for deciding the suitable incentive based on the budget of the campaign, the segments in which the customer belongs, and the fraction of contacted customers.

This study opens several opportunities for future developments. For example, a multicampaign strategy can be designed by considering different segment-based incentives d_i rather than a unique value. It is expected that customers with larger CLVs would require a larger incentive. Another future research direction is uplift modeling, which is designed for targeting those customers that are more likely to stay when contacted by a marketing action. This approach not only avoids targeting "lost causes", i.e., customers that will churn irrespective of whether they are targeted, but also prevents "waking up do-not-disturbs", i.e., customers that are unlikely to churn unless targeted by a retention campaign. This is of the utmost importance for the mutual fund industry because retention campaigns in tough economic times can actually increase the churn rates. The data of this study do not include information on retention campaigns. Therefore, it is not possible to evaluate potential efforts in these two directions. However, the company that provided the data may be able to collect data from retention campaigns that would allow us to validate proposals related to this topic.

Declaration of Competing Interest

The authors declare that they have no conflict of interest.

CRediT authorship contribution statement

Sebastián Maldonado: Conceptualization, Methodology, Validation, Formal analysis, Writing - original draft, Writing - review & editing, Supervision, Project administration, Funding acquisition. **Gonzalo Domínguez:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing - original draft, Writing - review & editing, Visualization. **Diego Olaya:** Conceptualization, Software, Validation, Resources, Data curation, Writing - review & editing, Visualization. **Wouter Verbeke:** Conceptualization, Validation, Writing - review & editing, Funding acquisition.

Acknowledgements

The authors would like to thank José Santomingo from FOL Agencia de Valores for providing the necessary information for this research. The first author gratefully acknowledges financial support from CONICYT PIA/BASAL AFB180003 and FONDECYT-Chile, grants 1160738 and 1200221.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.omega.2020.102380.

References

[1] Alkhasawneh R, Hargraves RH. Developing a hybrid model to predict student first year retention in stem disciplines using machine learning techniques. Journal of STEM Education: Innovations and Research 2014;15(3):35.
[2] Bahnsen AC, Aouada D, Ottersten B. A novel cost-sensitive framework for customer churn predictive modeling. Decision Analytics 2015;2(5):1–15.
[3] Bahnsen A.C., Aouada D., Ottersten B.E.. Ensemble of example-dependent cost-sensitive decision trees. CoRR; 2015b.
[4] Ballestero E, Pla-Santamaria D. Selecting portfolios for mutual funds. Omega (Westport) 2004;32(5):385–94.
[5] Berger PD, Nasr NI. Customer lifetime value: marketing models and applications. Journal of interactive marketing 1998;12(1):17–30.
[6] Blattberg R, Kim B, Neslin S. Database marketing: analyzing and managing customers. New York: Springer Science+Business Meida, LLC; 2008.
[7] Burez J, Van den Poel D. Handling class imbalance in customer churn prediction. Expert Syst Appl 2009;36(3):4626–36.
[8] Calzada-Infante L, Óskarsdóttir M, Baesens B. Evaluation of customer behavior with temporal centrality metrics for churn prediction of prepaid contracts. Expert Syst Appl 2020:113553.
[9] Chawla NV, Hall L, Bowyer K, Kegelmeyer W. Smote: synthetic minority oversampling technique. Journal of Artificial Intelligence Research 2002;16:321–57.
[10] Chen Z-Y, Fan Z-P, Sun M. A hierarchical multiple kernel support vector machine for customer churn prediction using longitudinal behavioral data. Eur J Oper Res 2012;223(2):461–72.
[11] Chevalier J, Ellison G. Risk taking by mutual funds as a response to incentives. Journal of Political Economy 1997;105(6):1167–200.
[12] da Costa FJ, de Souza Bispo M, de Faria Pereira RC. Dropout and retention of undergraduate students in management: a study at a brazilian federal university. RAUSP Management Journal 2018;53(1):74–85.
[13] Datta P, Masand B, Mani D, Li B. Automated cellular modeling and prediction on a large scale. Artif Intell Rev 2000;14:485–502.
[14] Dermine J, Röller L-H. Economies of scale and scope in french mutual funds. Journal of Financial Intermediation 1992;2(1):83–93.
[15] Fernández-Delgado M, Cernadas E, Barro S, Amorim D. Do we need hundreds of classifiers to solve real world classification problems? The journal of machine learning research 2014;15(1):3133–81.
[16] Galagedera D, Roshdi I, Fukuyama H, Zhu J. A new network dea model for mutual fund performance appraisal: an application to u.s. equity mutual funds. Omega (Westport) 2018;77:168–79.
[17] Gandomi A, Zolfaghari S. Profitability of loyalty reward programs: an analytical investigation. Omega (Westport) 2013;41(4):797–807.
[18] Ganesh J, Arnold M, Reynolds K. Understanding the customer base of service providers: an examination of the differences between switchers and stayers. J Mark 2000;64(3):65–87.
[19] Gutierrez P, Gérardy J-Y. Causal inference and uplift modelling: A review of the literature. In: International conference on predictive applications and APIs; 2017. p. 1–13.
[20] Hand D. Measuring classifier performance: a coherent alternative to the area under the ROC curve. Mach Learn 2009;77(1):103–23.
[21] Höppner S, Stripling E, Baesens B, vanden Broucke S, Verdonck T. Profit driven decision trees for churn prediction. European Journal of Operational Research 2019. In press
[22] Hornik K. Approximation capabilities of multilayer feedforward networks. Neural networks 1991;4(2):251–7.
[23] ICI. 2019 Investment company; 2019.
[24] Indro D, Jiang C, Patuwo B, Zhang GP. Predicting mutual fund performance using artificial neural networks. Omega (Westport) 1999;27(3):373–80.
[25] Kaya E, Dong X, Suhara Y, Balcisoy S, Bozkaya B, et al. Behavioral attributes and financial churn prediction. EPJ Data Sci 2018;7(1):41.
[26] Keramati A, Ghaneei H, Mirmohammadi SM. Developing a prediction model for customer churn from electronic banking services using data mining. Financial Innovation 2016;2(1):10.
[27] Klein R, Kolb J. Maximizing customer equity subject to capacity constraints. Omega (Westport) 2015;55:111–25.
[28] Künzel SR, Sekhon JS, Bickel PJ, Yu B. Metalearners for estimating heterogeneous treatment effects using machine learning. Proceedings of the national academy of sciences 2019;116(10):4156–65.
[29] Lin R, Li Z. Directional distance based diversification super-efficiency DEA models for mutual funds. Omega 2019:102096. In Press
[30] Maldonado S, Flores A, Verbraken T, Baesens B, Weber R. Profit-based feature selection using support vector machines - general framework and an application for customer churn prediction. Appl Soft Comput 2015;35:740–8.
[31] Maldonado S, López J, Vairetti C. Profit-based churn prediction based on minimax probability machines. Eur J Oper Res 2019.
[32] Malhotra DK, McLeod RW. An empirical analysis of mutual fund expenses. Journal of financial research 1997;20(2):175–90.
[33] Mallikarjuna M, Rao RP. Evaluation of forecasting methods from selected stock market returns. Financial Innovation 2019;5(1):1–16.
[34] Neslin S, Gupta S, Kamakura W, Lu J, Mason C. Defection detection: measuring and understanding the predictive accuracy of customer churn models. Journal of Marketing Research 2006;43(2):204–11.
[35] Oral E, Unal G. Modeling and forecasting time series of precious metals: a new approach to multifractal data. Financial Innovation 2019;5(1):3.
[36] Óskarsdóttir M, Baesens B, Vanthienen J. Profit-based model selection for customer retention using individual customer lifetime values. Big Data 2018;6(1):53–65.
[37] Van Den Poel D, Lariviere B. Customer attrition analysis for financial services using proportional hazard models. Eur J Oper Res 2003;157(1):196–217.
[38] Pozen R, Hamacher T. The fund industry: how your money is managed. 2nd edition. Wiley Finance; 2015.

[39] Rouwenhorst KG. The origins of mutual funds. Tech. Rep.. Yale ICF Working Paper No 04–48; 2004.
[40] Safari M, Tive M, Babania A, Hesan M. Analyzing the applications of customer lifetime value (clv) based on benefit segmentation for the banking sector. Procedia - Social and Behavioral Sciences 2014;109:590–4.
[41] Shirazi F, Mohammadi M. A big data analytics model for customer churn prediction in the retiree segment. Int J Inf Manage 2019;48:238–53.
[42] Stripling E, vanden Broucke S, Antonio K, Baesens B, Snoeck M. Profit maximizing logistic model for customer churn prediction using genetic algorithms. Swarm Evol Comput 2018;40:116–30.
[43] Verbeke W, Baesens B, Bravo C. Profit driven business analytics. Wiley; 2017.
[44] Verbeke W, Dejaeger K, Martens D, Hur J, Baesens B. New insights into churn prediction in the telecommunication sector: a profit driven data mining approach. Eur J Oper Res 2012;218(1):211–29.
[45] Verbeke W, Martens D, Baesens B. Social network analysis for customer churn prediction. Appl Soft Comput 2014;14:431–46.
[46] Verbraken T, Bravo C, Weber R, Baesens B. Development and application of consumer credit scoring models using profit-based classification measures. Eur J Oper Res 2014;238(2):505–13.
[47] Verbraken T, Verbeke W, Baesens B. A novel profit maximizing metric for measuring classification performance of customer churn prediction models. IEEE Trans Knowl Data Eng 2012;25(5):961–73.
[48] Wager S, Athey S. Estimation and inference of heterogeneous treatment effects using random forests. J Am Stat Assoc 2018;113(523):1228–42.
[49] Wang C, Han D, Fan W, Liu Q. Customer churn prediction with feature embedded convolutional neural network: an empirical study in the internet funds industry. Int J Comput Intell Appl 2019;18(01):1950003.
[50] Wei C, Chiu I. Turning telecommunications call details to churn prediction: a data mining approach. Expert Syst Appl 2002;23:103–12.
[51] Yamashita T, Yamashita K, Kamimura R. A stepwise aic method for variable selection in linear regression. Communications in Statistics–Theory and Methods 2007;36(13):2395–403.