

Debt Collection Industry: Machine Learning Approach

Amirhossein Shoghi*

Received: 16 Jan 2020

Approved: 29 Feb 2020

Businesses are increasingly interested in how big data, artificial intelligence, machine learning, and predictive analytics can be used to increase revenue, lower costs, and improve their business processes. In this paper, we describe how we have developed a data-driven machine learning method to optimize the collection process for a debt collection agency. Precisely speaking, we create a framework for the data-driven scheduling of outbound calls made by debt collectors. These phone calls are used to persuade debtors to settle their debt, or to negotiate payment arrangements in case debtors are willing, but unable to repay. We determine daily which debtors should be called to maximize the amount of delinquent debt recovered in the long term, under the constraint that only a limited number of phone calls can be made each day. Our approach is to formulate a Markov decision process and, given its intractability, approximate the value function based on historical data through the use of state-of-the-art machine learning techniques. Precisely, we predict the likelihood with which a debtor in a particular state is going to settle its debt and use this as a proxy for the value function. Based on this value function approximation, we compute for each debtor the marginal value of making a call. This leads to a particularly straightforward optimization procedure, namely, we prioritize the debtors that have the highest marginal value per phone call. We believe that our optimized policy substantially outperforms the current scheduling policy that has been used in business practice for many years. Most importantly, our policy collects more debt in less time, whilst using substantially fewer resources leading to a large increase in the amount of debt collected per phone call.

Keywords: Debt Collection, Artificial Intelligence, Machine Learning, Approximate Dynamic Programming, Prescriptive Analytics.

JEL Classification: G21, G32, H63

1 Introduction

Undoubtedly, one of the most significant challenges that the banking system has faced recently can be considered as the growing unpaid debts, and this issue causes the development of debt collection industry and debts buying in the world (Shoghi, 2019). Indeed, debt collection agencies are the pioneers and specialists of this industry. According to the latest International Debt

* Tamin Andish Pars Co., Iran; amirhossein.shoghi19@gmail.com

Collections Handbook (USA) published in June 2017, the estimated success rate of collections between 01/2014 and 12/2016 in the United States was just at 36.7% (with the same metric for Canada expected to be at 17.7%, UK at 65.8%, Mexico at 38.7%, China at 23.9%, India at 17.3%, Germany at 63.9%, and France at 67.8%). In Figure 1, how the businesses deal with their delinquent debts have been illustrated. Italy, Spain, Germany, France, and Belgium, as well as China, India, Indonesia, Mexico, and Brazil, are found to be more inclined to adopt alternative or additional solutions to traditional debt collection services, compared to the general tendency among the other countries profiled in the report. In general, according to Atradius Collections, companies in Europe appear to be more cautious in using alternative services other than traditional debt collection. Companies in the Americas and APAC show a stronger openness towards service innovation (Mesropyran, 2019). As you see, most of the businesses within the entire world prefer using internal sources to collect their debts instead of other ways, such as using external debt collection agencies. It means that debt collection industries are not intriguing enough to attract more businesses. Thus, using special skills, debt collection industries can provide more banks and large public or private companies the possibility of debt revival to their financial cycle. By the way, considering economic, social, and political changes especially in recent years, lack of new formulas with scientific methods within debt collection companies is felt more and more (Shoghi, 2019).

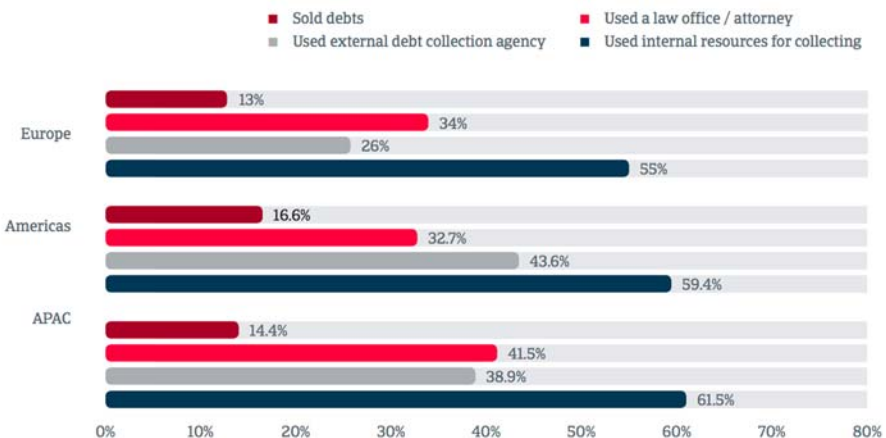


Figure 1. How Companies Dealt with Overdue Invoices in the World. (Mäkikangas et al., 2017)

Artificial Intelligence and its abbreviation AI are the buzz words of today. Companies use them to explain what they do and to make solutions and products sound smarter and more sophisticated than they are. AI is a broad and complex subject. In the broad sense, AI can be described as building computers and robots to do human tasks, which we consider require intelligence such as to build humans like machines, to collect data, and to analyze data. The consequence is that AI can be used to describe almost anything. However, the overuse of the terms makes it near impossible for the layman to separate what is an exact AI-based solution from what is hyperbole. In the world of debt collection, the techniques used to gather and process “big data” to classify debtors and recommend actions regularly use AI and machine learning technologies. These solutions are established and growing. They are back-office solutions used by credit management departments and Debt Collection Agencies (DCAs) to enhance productivity or support business growth (Phillips and Moggridge, 2019). Therefore, using AI solutions, DCAs can boost their methods, thereby, they will elevate their attributions to the debt market.

The collection process usually follows a predefined schedule of letters, SMSs, and phone calls that communicate with increasing urgency the need to repay the debt over time. Ultimately, if the debtor refuses to repay the debt, then legal action can be taken by the collection agency to force repayment. Legal action is expensive and often outside of the collection agency’s control, so it is only viewed as a last resort and avoided as much as possible. In contrary to popular belief, debt collectors generally prefer to cooperate with debtors to repay their debt by offering interest-free extensions, repayment plans, or in some cases waiving parts of the debt if the debtor is genuinely unable to repay. However, this is only possible if the debtor is cooperative and responds to the collectors’ communication attempts (e.g., answers the phone or replies to SMS). Letters and SMSs are mostly automated, but phone calls still require human collectors to physically dial a number and have a conversation with the debtor. It is integral to the collection process because debt collection is highly emotional, and an experienced collector can decipher the needs and problems of the debtor and determine the best course of action to maximize the likelihood of repayment. However, debt collection agencies generally have a large number of open cases and the number of phone calls that it can make is limited by human resources. Under these constraints, it becomes infeasible to call every debtor and a method to select debtors to call becomes necessary. Not calling a debtor who needs human persuasion results in further delinquency and higher risk for non-repayment, but calling a debtor who does

not require additional persuasion results in wasted effort. Our goal is to identify under which conditions phone calls are most effective in eliciting eventual repayment and creating an optimal schedule of calls to each debtor while abiding by the capacity constraints faced by the collector (Wang et al., 2018).

In this paper, we present a framework for the data-driven scheduling of outbound phone calls made by debt collectors. That is, we determine daily which debtors a debt collector should call to maximize the amount of delinquent debt recovered in the long term, under the constraint that only a limited number of phone calls can be made each day. These phone calls are used to persuade debtors to settle their debt, or to negotiate payment arrangements (e.g., a payment plan) in case debtors are willing, but unable to repay their debt. Scheduling these calls is challenging, as it is difficult to assess the value of making a phone call to a debtor. It is because a priori the outcome of making a call is uncertain, and the extent to which a request attributes to repayment is non-trivial. In general, the effect of phone calls on the repayment behavior of debtors depends on numerous interacting features, such as the time since the previous phone call, whether the debtor answered the call before, the amount of debt owed, the time of the month, and the persuasiveness of the agent who is calling. It is unclear what the effect of these (interacting) features is on the outcome of phone calls and, hence, on the effectiveness of a schedule of phone calls. This lack of structure and understanding drives our belief that a flexible non-parametric machine learning method would be most appropriate to leverage data for optimizing actions.

To this end, we show that the problem of scheduling phone calls is naturally formulated as a Markov decision process (MDP), but that prohibitively ample state space is required to capture the dynamics of the collection process appropriately. To alleviate this, we show how state-of-the-art machine learning methods can be used in approximate dynamic programming (ADP) framework that is interpretable, highly scalable, and data-driven.

To the best of our knowledge, the current paper incorporates modern machine learning methods into an ADP framework that is validated through a controlled field experiment in a real-life business setting. We take the problem of dynamically scheduling outbound calls for a debt collector, as naturally described by an MDP, and approximate state values using supervised machine learning. More precisely, we construct a binary classification problem to predict, based on a debtor's state, the likelihood with which a debtor is going

to repay its debt. The debtor's state space is high dimensional and incorporates all static and dynamic information that characterizes a debtor at a given point in time. For value function approximation, we multiply the likelihood with which a debtor settles its debt by the size of the debt, thereby obtaining an estimate for the expected value of a debtor given its current state. In doing so, we overcome the curse of dimensionality inherent to this problem by inferring the value of a debtor's state based on historical data in a highly scalable and flexible manner. Based on our value function approximation, we compute for each debtor the marginal value of a phone call, which is dependent on the change in the value function if we spend another phone call on this debtor. It leads to a particularly straightforward optimization procedure. Namely, we prioritize the debtors that have the highest marginal value per phone call. The result is an interpretable policy (debtors with the highest marginal value on the effort are prioritized), highly scalable, and data-driven. Besides, the optimization procedure allows for straightforward implementation in business practice: arrivals of new debtors are naturally incorporated, and an appropriate number of phone calls can be determined to be made on a given day, depending on the debt collector's capacity (van de Geer et al., 2018).

In summary, this paper contributes to the existing literature on business analytics, data-driven optimization, and that of ADPs in the following ways: i) we add to the debt collection optimization literature by presenting a novel, scalable, and flexible framework for daily data-driven scheduling of outbound calls; ii) we incorporate state-of-the-art machine learning methods to the ADP framework, which takes advantage of higher-order feature interactions and results in superior out-of-sample model fit for value function approximation compared to benchmark models; and iii) we open the proverbial machine learning black box and identify generalizable insights for the improved scheduling of outbound debt collection phone calls.

2 Subject Literature

2.1 Debt Collection Optimization

More than half a century ago, Mitchner and Peterson (1957) considered the problem of optimizing the collection of delinquent debt at Bank of America for various types of loans, such as car loans, personal loans, and real estate loans. They formulated the problem of collecting a debt as an optimal stopping problem, in which the duration with which the collector should pursue the debtor was optimized, taking into account the cost of doing so. Their results show a potential increase in net profit of 33%.

Fifteen years later, Liebman (1972) developed a simple Markov decision process for optimizing credit control policies. They solve an example problem with four delinquency states; two amount owed states, two recent experience states, and three action strategies (Abe et al. (2010), De Almeida Filho et al. (2010), and Miller et al. (2012)). However, the curse of dimensionality quickly becomes a significant challenge and no further progress on this topic was made until more recently.

In Abe et al. (2010) and the accompanying paper, Miller et al. (2012), a framework for debt collection optimization presented that, of the current work, is closest to the approach considered in this paper. In Abe et al. (2010), the collection process is modeled as a constrained MDP, which explicitly takes business, legal, and resource constraints into account. Subsequently, given the intractability of the MDP, a constrained Q-learning algorithm is proposed employing which they can obtain a policy. In Miller et al. (2012), the deployment of this methodology at the New York State Department of Taxation and Finance is described for which they report an increase in collecting delinquent debt by 8 percent over the first year, where they would otherwise have projected a rise of 2-4 percent.

Also, from the operations domain, De Almeida Filho et al. (2010) present a study on the optimization of debt collection in the context of consumer lending. In their work, they offer a dynamic programming approach in which the monthly decision epochs pertain to deciding which action to take in the month to come. The value function corresponds to the future net discounted recovery rate, and the transitions are assumed to be deterministic. Since the model assumes homogeneous debtors, the approach is especially useful to predict collection performance and resource requirements for aggregated portfolios of debtors for which it is reasonable to assume homogeneity. The authors refer to the importance and potential of tailoring the collection process to the individual debtor, but note that the data required for this purpose are hardly ever available in practice.

Van de Geer et al. (2018) used consumer debt collection with machine learning to optimize the debt collection process with regards to phone calls, emails, and letters. In their work, they used LightGBM, which works exceptionally well in practice and is often considered together with XGBoost as the best algorithm for the predictive analytics competitions hosted on Kaggle. Furthermore, it is easy to use and does not require sophisticated feature engineering to achieve excellent performance. So, LightGBM was able to reduce significantly calling effort as a decrease of 21.5% in the number of calls, and it means that it increased the efficiency of the collection process.

2.2 Credit Scoring and Valuation

The field of finance, much research has been done on the credit-granting decision, i.e., whether to grant a loan to a potential new customer. Typically, they make the credit-granting choice for personal loans employing credit scoring, which is a standardized method of assigning a score to potential customers that represent their creditworthiness, see Crook et al. (2007) for a literature review and Lessmann et al. (2015) for a benchmarking study on existing scoring models. On the other hand, the valuation of existing credit, and existing unsecured loans, in particular, is more closely related to our work since the debt that the Collector is trying to collect is mainly outstanding unsecured credit. Although much work has been done on the valuation of corporate lending and secured customer credit, the literature on unsecured consumer credit is sparse. The work of Chehrazi and Weber (2015) on the dynamic assessment of delinquent credit card accounts models the stochastic repayment behavior of individual debtors over time. They derive a self-exciting point process for repayment behavior and estimate the parameters of the process using the generalized method of moments. This model is then used to construct a dynamic collectability score to estimate the probability of collecting from a debit account, thus allowing for the valuation of credit card debt. In a subsequent paper, Chehrazi et al. (2018) formulates a stochastic optimal control problem from the self-exciting point process established in Chehrazi and Weber (2015) and derives a semi-analytic solution. However, this solution was not analyzed empirically nor experimentally validated.

3 Research Methodology

We first provide a high-level overview of the operations of a debt collector. After that, we provide more details on the actual collection process. It is all based on the experiences of our Collector but is illustrative for the debt collection industry in general.

3.1 High-level Overview

In practice, a client that has an overdue debt with a company is placed “in collections”, which means that the debtor is transferred to either a specialized debt collection department within the company or to an external debt collecting agency that works on behalf of the company. In this work, we refer to both as a debt collector, i.e., a debt collector can be either the debt owner itself or a third-party debt collection agency working on behalf of the debt owner. The debtor typically incurs a collection fee that is added to the original debt to cover the additional costs of recovering the debt and is regulated in

many countries. In the problem that we consider, the collection fee is independent of the amount of debt owed and constant across debtors.

Once placed in collections, the debt collector pursues the debtor to settle the debt plus the collection fee by sending out letters and SMSs, and through phone calls made by its agents. Messages and SMSs are mostly automated, but phone calls still require human collectors to physically dial a number and have a conversation with the debtor. Human collectors are integral to the collection process because debt collection is highly emotional, and an experienced collector can decipher the needs and problems of the debtor and determine the best course of action to maximize the likelihood of repayment. Amongst the Collector's clients are utility providers, credit facilitators, and health care providers, which operate in the business-to-consumer market. The process of the Collector comprises two phases. Upon arrival, a debtor first enters the collection phase, in which the Collector pursues the debtor to repay the debt plus the collection fee through letters, SMS, and phone calls. During this phase, the collector acts cooperatively towards the debtor, and can offer payment plans if debtors are willing, but not able to pay on a short-term notice. As such, this phase can take from a few days (in case the debtor pays immediately) up to a few months (in case the debtor does not pay at all or gets involved in a payment plan). In contrary to popular belief, debt collectors generally prefer to cooperate with debtors to repay their debt by offering interest-free extensions, repayment plans, or in some cases waiving parts of the debt if the debtor is genuinely unable to repay. However, this is only possible if the debtor is cooperative and responds to the collectors' communication attempts (e.g., answers the phone or replies to email).

When the Collector is unsuccessful in recovering the debt during the collection phase, it chooses to either write off the debt or invoke a legal procedure. The former happens when, for example, the debtor is deceased or has declared bankruptcy. However, legal action is expensive and often outside of the collection agency's control, so we only view it as a last resort and avoid it as much as possible. The latter means that a bailiff is invoked. He will send out a subpoena and ultimately can confiscate property if necessary. Whether a debtor is escalated to the legal phase or written off is determined case by case and depends, amongst other things, on the amount of outstanding debt and the likelihood of recovering the debt through legal procedures. Since this phase requires legal assessment by an expert, it is costly, and the outcome is highly uncertain. Hence, recovering debt before the legal aspect is deemed beneficial for both the collector and debtor. As such, the legal phase is excluded from the optimization procedures proposed in this work, and our

objective is to maximize recovered debt during the collection phase, which is described in greater detail in the following section.

3.2 Collection Phase

The collection phase is characterized by four sequential letters (sent via both post and SMS panel simultaneously), where each letter has a seven-day payment notice and communicates with increasing urgency the necessity to repay the debt. The messages are sent between seven to ten days of each other. The fourth and final letter communicates the severe (financial) consequences of the legal procedure that the Collector possibly invokes if the debtor does not settle. In between the letters (or after the final letter), the Collector is free to call debtors at its discretion. It is considered a vital tool during the collection process, as the phone calls allow the agents to inform the debtor about the situation along with the consequences of non-payment, and to make an assessment of whether the debtor is willing and/or able to pay. Figure 2 provides a schematic illustration of the collection process.

The optimization problem that the Collector encounters is deciding each day in which debtors should call to maximize recovered debt, given the finite and inflexible capacity of its workforce. In practice, this implies that the Collector has to decide on a prioritization on the debtor portfolio that indicates which debtors should be called first. Currently, the Collector's policy is to schedule a phone call each time a debtor has received a new letter. Besides, if a debtor agreed on a payment plan and failed to comply with its conditions, a call is scheduled as well. In case, capacity is insufficient, the Collector's managerial staff makes an assessment of which debtors should be called first. Given the labor-intensive nature of the phone calls, the gains from optimizing the prioritization of calls are potentially substantial.

Contrary to some belief, machine learning is not a black box, and it's always possible to analyze the predictions made concerning the feature values used to make the predictions. In the following sections, we are going to suggest a method of prioritizing phone calls and find which features we can link to better calling efficiency. We mean our approach to understanding the collection phase better than to make definitive conclusions on the debt collection process.

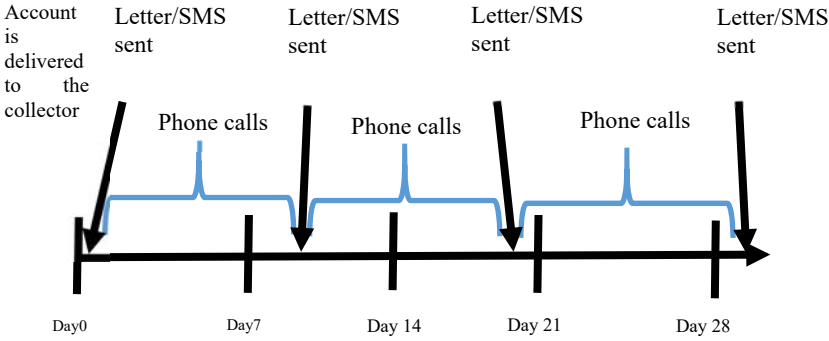


Figure 2. The Standard Operating Timeline of a Debt Collection Agency.

4 Results and Findings

The problem of optimizing debt collection efforts over time in the current context is formulated as an MDP with an infinite time horizon and decision epochs in discrete time. It suits the approach of the Collector since, in principle, the Collector operates indefinitely, and decisions made at discrete points in time (i.e., daily). To formalize the MDP, we assume that at any given point in time, the Collector has at most $N \in \mathbb{N}$ debtors in its portfolio. Practically, this means we set N arbitrarily large such that the Collector never has more than N debtors in its portfolio.

4.1 State Space

We denote the state space of each of the (at most) N debtors by χ and the state space of the portfolio of debtors by $\bar{\chi} := \chi^N$ (i.e., the N -fold Cartesian product of χ). In our formulation, each of the N parts of the state space is utilized by different debtors over time, each part χ of the state space $\bar{\chi}$ functions as a slot for storing the information of one of the debtors. A slot becomes available for new arriving debtors once efforts on the existing debtor are terminated because the debt is recovered or written-off.

The state-space below is chosen to accommodate the data to which we apply our methodology. We divide the debtor state-space χ into debtor-specific features, historical-interaction features and seasonalities as follows. The superscripts B, I, N, and C indicate whether it is a binary, integer, numerical, or categorical variable, respectively.

Debtor specific: 1) initial debt amount ^N, 2) customer tenure ^I, 3) has partially repaid debt ^B, 4) repayment plan in place ^B, 5) phone number is available ^B, 6) mobile number is available ^B, 7) product type ^C, 8) amount repaid already ^N, 9) Collector collected from debtor before ^B, 10) average income in the postal code area of the debtor ^N, 11) share of people under 30 in postal code area of the debtor ^N, 12) current substatus ^C, 13) passed final letter ^B. Here, 2) indicates when the debtor became a customer of the debt owner: the exact time was not provided. Instead we have an integer that represents the inverse order in which the debtor became a customer relative to all customers of the debt owner (the more significant the value means the debtor was a customer of the debt owner for a more extended time); 5) pertains to whether the debt owner provided the Collector with a phone number of the debtor. If this is not the case, the Collector may still be able to call the debtor by searching manually in publicly available resources for potential phone numbers that match to the name and address of the debtor; 7) refers to the product or service that the debtor purchased and led to the debt; 12) refers to a debtor's state description used internally by the Collector to characterize a debtor at a given point in time; 13) refers to whether the debtors have received the final (i.e., fourth) letter.

Table 1

Observable collector-debtor interactions.

Variable	Type	Description
Debtor ID	Integer	-
Date	Date	-
Communication type	Categorical	Letter-SMS or Phone call
Communication direction	Binary	In or outbound (regards to "communication type")
Reached	Binary	In case of outbound phone call
Document type	Categorical	In case outbound Phone call or SMS
Promised to pay	Binary	If debtor promised to pay

Historical interaction: 1) has answered a phone call ^B, 2) promised to repay ^B, 3) number of previous collector-debtor interactions ^I, 4) number of previous phone calls ^I, 5) days since promise to repay ^I, 6) days since last collector-debtor interaction ^I, 7) days since last phone call ^I, 8) days since last answered phone call ^I, 9) days since last incoming contact ^I, 10) days since last incoming SMS ^I, 11) days since last incoming phone call ^I. Here, 2) and 5) refer to the event in which the debtor has (verbally) promised the Collector to

settle the debt; in 3) and 6) the word 'interaction' includes both collector- and debtor-initiated communication efforts (Table 1), and also phone calls that did not get through count as an interaction, thereby using it in a broader sense than usual.

Seasonality: day of week ^C, week of month ^C.

For features where missing values are possible, a unique integer is used as a replacement for missing values. An example of this is the feature days since the last phone call for cases where no phone calls have previously been made to the debtor.

4.2 Action Space and Value Function

Regarding the action space, for a given day let $a \in \{0,1\}^N$ describe which debtors will be called: $a_i = 1$ for $i \in \{1,2, \dots, N\}$ indicates that a call is made to debtor i and $a_i = 0$ means no call is made on a given day. In some cases, it is undesirable to call a debtor (e.g., when the debt is currently being further investigated because the debtor disputed the debt). Therefore, we construct the action space as follows. Let $i \in \{1,2, \dots, N\}$, $x = (x_1, \dots, x_N) \in \bar{\chi}$ with $x_i \in \chi$, and let $\mathcal{A}'(x_i)$ denote the action space of debtor i , so that $\mathcal{A}'(x_i)$ equals $\{0\}$ if no call is allowed and $\{0,1\}$ when a call to debtor i is allowed. Besides, let $c_t \in \mathbb{N}$ denotes the (deterministic) capacity on day t , i.e., the maximum number of phone calls that can be made on day t , where t counts the number of days since the collection process was initiated. Accordingly, we define

$$\mathcal{A}_t(x) := \{(a_1, \dots, a_N) : a_i \in \mathcal{A}'(x_i), i = 1, \dots, N, \sum_{i=1}^N a_i \leq c_t\}$$

as the action space on day t . In case of slot i of the state space is not used, $\mathcal{A}'(x_i) = \{0\}$ for all $x_i \in \chi$. Furthermore, on day t , for $x, y \in \bar{\chi}$ and $a \in \mathcal{A}_t(x)$, let $p(x, a, y)$ denote the probability of moving from state x on day t to state y on day $t+1$, when choosing action a and let $r(x, a, y)$ denote the amount of debt recovered (i.e., repaid and received) when moving from state x to state y choosing action a . The possible arrival of new debtors is implicitly incorporated in $p(x, a, y)$. Then, the optimality equation becomes

$$V_t(x) = \max_{a \in \mathcal{A}_t(x)} \sum_{y \in \bar{\chi}} p(x, a, y) (r(x, a, y) + \gamma V_{t+1}(y))$$

For $t = 0, 1, 2, \dots$, where $V_t(x)$ denotes the total expected discounted reward when being in state $x \in \bar{\chi}$ at day t , and $\gamma \in (0,1)$ denotes an appropriate discount rate. The function $V_t : \bar{\chi} \rightarrow \mathbb{R}$ is often referred to as the value function.

Since the state space $\bar{\chi}$ consists of all debtor information, the formulated MDP has a high-dimensional state space. Moreover, parts of the state space are unbounded (e.g., the number of collector-debtor interactions). It makes it intractable to solve the MDP, even numerically. The MDP, however, has structural properties that facilitate the computation of near-optimal policies. First, the debtors in the portfolio behave independently of each other, i.e., changes to the state and repayment probability of one debtor do not affect the repayment probability of the other debtors. Second, the dependence in the problem formulation is only due to the capacity constraint c_t on day t . Hence, a natural approximation that breaks the dependence arises when the Collector solves a stochastic knapsack problem based on the state of the debtors in the portfolio on that day. The knapsack has size c_t on day t , and the expected value of each item in the knapsack will be given by the expected gain in the value function from calling the debtor. Note that in this formulation, the discount factor naturally disappears since future arrivals do not affect current decisions. In the next section, we elaborate on how to estimate the value function of each debtor.

4.3 Value Function Approximation with Machine Learning

We use value function approximation (VFA) to approximate the value of the states of the MDP described in the previous section. Any function can be used to approximate the value function, including radial basis functions, polynomials, neural networks, and decision trees (Bertsekas and Tsitsiklis, 1995). VFA has been successfully applied in optimization in a variety of problems, such as large-scale resource allocation, (Powell and Topaloglu, 2006) multi-priority patient scheduling, (Patrick et al., 2008) and autonomous inverted helicopter flight (Ng et al., 2006). Recent breakthroughs in machine learning, notably convolutional neural networks, have sparked the field of deep reinforcement learning, which allows for VFA through visual images. For example, AlphaGo was able to exploit this approach by successfully approximating the 10^{170} state-space in the game of Go and defeat the world's best human players (Silver et al., 2016).

In this paper, we use another state-of-the-art machine learning algorithm for VFA; namely, gradient boosted decision trees (GBDT). It is a more suitable algorithm for prediction problems that are arranged in the standard tabular structure and has been the dominant algorithm in winning well over half of all machine learning competitions in 2015, including the KDD Cup. (Chen & Guestrin, 2016). It was also found by Olson et al. (2018) to be the best algorithm when benchmarked against twelve other algorithms for 165

publicly available classification problems. In Section 4.3.3, we provide details on the GBDT algorithm.

We use the GBDT model to construct a mapping $\hat{V} : \bar{\chi} \rightarrow \mathbb{R}$ that approximates the value function, thereby circumventing the problem of having to solve $V_t(x)$. This approximation is used to optimize the actions, i.e., to determine which debtors are to be called on a given day. In the following two sections, we show how we construct the mapping \hat{V} (Section 4.3.1) and optimize actions based on this approximation (Section 4.3.2).

4.3.1 Estimating the Predicted Repayment Probability

To approximate the value of a debtor being in a particular state, we estimate the debtor's predicted repayment probability (PRP), which is defined as the likelihood of recovering the full debt during the collection phase. Partially repaid cases are considered to be unpaid as the Collector only receives credit for fully collected cases. Our approach is to estimate the PRP based on historical data employing a GBDT model as follows. Suppose a certain debtor is $k \in \mathbb{N}$ days into the collection process, and consider all closed cases that once were k days into the collection process as well, i.e., all closed cases that either did not settle their debt within k days or were not written off within k days. We use these closed cases to train a GBDT model that predicts the likelihood of recovering the debt of the debtor currently considered. We formalize this procedure as follows.

Let $n \in \mathbb{N}$ be the total number of closed cases in our dataset, i.e., cases for which the debt was either recovered or written off, and for which the debtor is no longer be contacted. Let $i \in \{1, 2, \dots, n\}$ and define $\tau_i \in \mathbb{N}$ as the total number of days debtor i spent in the collection process. For all $s \in \{1, 2, \dots, \tau_i\}$, let $x_i^{(s)} \in \chi$ be the state of debtor i at s days since arrival. We optimize the phone calls during the first $K \in \mathbb{N}$ days of the collection process of each debtor. Although, theoretically, K is unbounded, in our practical implementation we set K such that virtually all calling efforts take place in the first K days. For all $k \in \{1, \dots, K\}$, denote by

$$\mathcal{J}_k := \{i : k \leq \tau_i, i \in \{1, \dots, n\}\}$$

the index set containing all the closed cases in the dataset that were still in the collection process k days after arrival.

Furthermore, we denote by $y_i \in \{0, 1\}$ the eventual outcome of the collection process: $y_i = 1$ if the debt of debtor i was fully recovered after τ_i days, i.e., during the collection phase, and $y_i = 0$ otherwise, meaning that the

debt was either written off or recovered after legal actions. Hence, $x_i^{(k)}$ and y_i are the state of debtor i after k days and the eventual outcome of the collection process, respectively, for all $k \in \{1, \dots, K\}$ and all $i \in \mathcal{J}_k$.

Our approach is to train one GBDT model for each number of days since arrival $k \in \{1, \dots, K\}$ as follows. Let $k \in \{1, \dots, K\}$. Then, we train model k by using $(x_i^{(k)})_{i \in \mathcal{J}_k}$ as features (or independent variables) and $(y_i)_{i \in \mathcal{J}_k}$ as target (or dependent) variables. We denote the trained GBDT model by $f_k : \mathcal{X} \rightarrow (0,1)$, where f_k maps the state of a debtor after k days to a prediction for the likelihood that the debt is eventually recovered. This likelihood is exactly the PRP that we introduced earlier on, i.e., if debtor $i \in \mathcal{J}_k$ is in state $x \in \mathcal{X}$ after k days, then $f_k(x)$ represents its PRP.

We train a single model for each number of days since arrival because the data is unbalanced in the sense that there are many more observations for debtors that are earlier in the collection process (i.e., $|\mathcal{J}_k| \geq |\mathcal{J}_{k+1}|$ for each $k \in \{1, \dots, K-1\}$). This is because cases are closed as soon as the debt is fully recovered or written off. If we train a single model, this could cause the GBDT model to be biased toward better predicting the early part of the process at the expense of the latter part. To alleviate this, we follow the aforementioned approach in which we split the data by days after arrival into K sets and train K models.

Summarizing, we compute the PRP of a debtor on a given day by considering debtors that once were in a similar situation before, given that the days since arrival is highly correlated with the rest of the collection process.

4.3.2 Approximating the Value Function

To approximate the value of the state of a particular debtor, we multiply the debtor's PRP by its outstanding debt. More precisely, let $x = (x_1, \dots, x_N) \in \bar{\mathcal{X}}$ be the state of the debtor portfolio at a certain point in time and let $k_i \in \mathbb{N}$ denote the number of days debtor $i \in \{1, \dots, N\}$ has been in the collection process. Our approximation for the value of being in state x is

$$\hat{V}(x) := \sum_{i=1}^N f_{k_i}(x_i) \cdot debt_i$$

where $debt_i$ denotes debtor i 's current outstanding debts. When slot $i \in \{1, \dots, N\}$ of the state space is not used, we set $debt_i=0$. Observe that, when the objective is to maximize the number of fully collected cases (irrespective of the amount of debt recovered), we can accommodate for this by setting $debt_i=1$ for all $i \in \{1, \dots, N\}$.

Our proposed approximation in $\hat{V}(x)$ equation implies that we consider the Collector's portfolio on a particular day as an assortment of independent debtors in different states of the collection process, and compute the value of the portfolio as a sum of their individual values. This approximation allows us to evaluate policies by computing the difference in PRP with and without making a phone call to a particular debtor.

To formalize this, let $\psi : \chi \rightarrow \chi$ be the mapping that takes as input a debtor's state and then updates this state as follows: i) increase the feature number of previous collector-debtor interactions by one; ii) increase the feature number of previous phone calls by one; iii) set the feature days since last collector-debtor interaction to zero, and iv) set the feature days since last phone call to zero (see also the state space description in Section 4.1). Our approach is to determine the marginal value of making an additional call to debtor $i \in \{1, \dots, N\}$ by computing

$$[f_{k_i}(\psi(x_i)) - f_{k_i}(x_i)] \cdot \text{debt}_i.$$

Recall that f_{k_i} maps debtor i 's state to the PRP, i.e., to a prediction of the likelihood that debtor i will eventually repay, without needing to explicitly consider potential future states. Hence, the last equation provides us with a measure to compare the added value of calling different debtors. Naturally, the policy on day t is to call the c_t debtors for which the equation is the highest (recall that c_t denotes the capacity of the Collector on day t).

4.3.3 Gradient Boosted Decision Trees

GBDT, also called gradient boosting machines and multiple additive regression trees, falls under the general paradigm of ensemble methods in machine learning (Dietterich, 2000). The algorithm works by constructing multiple decision trees using the classification and regression trees algorithm (CART) (CART, Breiman et al., 1984) and combining these into a so-called committee, in which the predictions of the individual trees are combined to form one prediction (usually via a weighted average). We first describe how CART works and what its drawbacks are. Then, we explain how ensembles of trees overcome these drawbacks. Finally, we describe the GBDT algorithm and discuss why it works for our problem of predicting the repayment of debt.

The CART algorithm works by recursively partitioning the feature space into non-overlapping rectangular subsets and making a prediction for the target variable for each of these subsets. It is done by splitting, in each recursion, the feature that minimizes a certain error metric (e.g., mean squared error or Gini impurity). This procedure is myopic in the sense that the

partitioning decision does not consider future partitionings. As a result, CART does not guarantee a globally optimal partitioning.

A major drawback of CART is its propensity to overfit on training data, which results in a model that generalizes poorly to unseen data. Ensembles of CART models have been successfully used to overcome this. Early ensembling techniques, such as bootstrapped aggregating, commonly referred to as bagging, work by generating multiple versions of a prediction algorithm by using randomly selected subsamples of the training data (Breiman, 1996). The random forest algorithm is an example of a bagging algorithm. Subsampling observations via bootstrapping add variation to the training data, which leads to building significantly different trees, resulting in reductions in error rate by 20-89 percent (Breiman, 2001).

Unlike bagging, where we build the trees independently, GBDT builds trees sequentially. It is called boosting and works as follows. The goal of GBDT is to minimize a loss (or: objective) function that maps the predictions to a score that measures the quality of the predictions. Theoretically, any differentiable function can be used as a loss function. We use the logarithmic loss function, which is the standard choice for binary classification problems and is defined as follows. Suppose we are training model f_k for $k \in \{1, \dots, K\}$ then the logarithmic loss function $L : \mathbb{R}^{|\mathcal{J}_k|} \rightarrow \mathbb{R}$ is

$$L(z) := -\left(\sum_{i \in \mathcal{J}_k} y_i \cdot \log(\sigma(z_i)) + (1 - y_i) \cdot \log(1 - \sigma(z_i))\right)$$

where $\sigma : \mathbb{R} \rightarrow (0,1)$ is defined by $\sigma(u) := (1 + e^{-u})^{-1}$ for $u \in \mathbb{R}$. The GBDT algorithm repeats Step 1-3 below a pre-specified number of times, where $\varepsilon > 0$ is set as a hyper-parameter:

Step 0. Initialize with $z_i \leftarrow \sigma^{-1}\left(\frac{1}{|\mathcal{J}_k|} \sum_{j \in \mathcal{J}_k} y_j\right)$ for all $i \in \mathcal{J}_k$.

Step 1. Compute the gradient of the loss function $\frac{\partial L(z)}{\partial z_i} = \sigma(z_i) - y_i$ for all $i \in \mathcal{J}_k$.

Step 2. Train a regression tree using $-(\sigma(z_i) - y_i)$ for all $i \in \mathcal{J}_k$ as the target variables.

Step 3. Update $z \leftarrow z + \varepsilon z'$, where $z' \in \mathbb{R}^{|\mathcal{J}_k|}$ are the predictions from Step 2. Go to Step 1.

When the algorithm terminates, $\sigma(z_i)$ is GBDT's prediction for y_i for all $i \in \mathcal{J}_k$.

By iteratively building regression trees on the negative gradient in Step 2, newly built trees are optimized for observations that are difficult to predict, thereby improving the overall model fit with each iteration. For a more

detailed discussion on GBDTs we refer the reader to Friedman (2001) and Friedman et al. (2001).

Improving the model fit of the training data does not guarantee generalization to unseen data. Therefore, a cap in the number of iterations is required to prevent overfitting. The cap in the number of iterations, along with other hyper-parameters, such as maximum depth per tree, can be tuned using a training-validation framework. The implementation of GBDT used in this paper is LightGBM, which is a fast and distributed open-source GBDT framework developed by Microsoft (Ke et al., 2017).

CART, and GBDT, in particular, are well suited for our prediction problem for two reasons. First, since CART works by partitioning the data, it is invariant to monotonic transformations of the features. It differs from models such as logistic regression, where substantial efforts in finding the best functional transformations of the features are required to tune the model to achieve better prediction performance. It means that we can directly use the debtor's collection state as features in a CART model without performing any functional transformations. Second, as a consequence of recursive partitioning, CART implicitly takes into account feature interactions that can lead to improved prediction accuracy and better state-value approximations. Again, for other models such as logistic regression, the feature interactions must be defined manually.

5 Conclusion

We apply machine learning and approximate dynamic programming to help a debt collection agency optimize its collection process. Using data recorded from its historical collection interactions and outcomes we develop a method to intelligently select which debtors the collection agency should call for a given day. To be more precise, this paper considers the problem of deciding on a daily basis which debtors a debt collection agency should call, given that only a limited amount of calls can be made by its agents. It is a challenging optimization problem since, at any given time, a debtor portfolio consists of a large collection of heterogeneous debtors that are at different stages in the collection process. Our approach is to formulate an MDP and approximate it through data-driven machine learning methods, thereby circumventing dimensionality issues by relying on historical data. This approach revolves around computing, for each debtor at each state, the predicted repayment probability (PRP), and inferring the marginal increase in PRP when making an additional phone call.

This method implemented at an insurance company in the Netherlands and conducted a controlled field experiment. The results of the experiment show a relative increase of 14% in collected debt and a decrease of 22% in calling effort when using the suggested method as compared to the current collection process. Most notably, this policy leads to an increase in the amount of debt collected per outbound call from 38.53 euros to 56.72 euros, leading to a 47.2% improvement in return on calling effort. The improvement comes mostly from selecting debtors that have been in the collection process longer, have not been contacted recently, and have not previously answered calls nor promised to repay their debt. In general, the proposed policy puts more emphasis on debtors that are harder to collect from, and calls are scheduled later in the collection process.

Collectors can use our proposed technique to prognosticate the best time and the best forum by which to contact your debtors. Then you can prioritize calls by the time of day a debtor is most likely to respond. It is also capable of providing its users with SMS, Live Chat, and Email integration potentially. It can vastly reduce your telephone costs while also increasing connection and collection rates. By and large, using these method companies can optimize the performance of their employees. Besides, it allows them to understand their debtors' needs better, to react faster than ever to a changing landscape, and to vastly improve debt recovery rates.

References

- Abe, N., Melville, P., Pendus, C., Reddy, C. K., Jensen, D. L., Thomas, V. P., Bennett, J. J., Anderson, G. F., Cooley, B. R., Kowalczyk, M., Domick, M., & Gardinier, T. (2010). *Optimizing Debt Collections Using Constrained Reinforcement Learning*. Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 75-84 (ACM).
- Phillips, L., & Moggridge, P. (2019). Artificial Intelligence in Debt Collection. *Credit Control Journal and Asset & Risk Review*, 40(2).
- Bertsekas, D. P., & Tsitsiklis, J. N. (1995). *Neuro-Dynamic Programming: an Overview*. Decision and Control, 1995, Proceedings of the 34th IEEE Conference on, volume 1, 560-564 (IEEE).
- Breiman, L., Friedman, J., Stone, C. J., & Olshen, R. A. (1984). *Classification and Regression Trees (CRC press)*.
- Breiman, L. (1996). Bagging Predictors. *Machine Learning*, 24(2), 123-140.
- Breiman, L. (2001). Random Forests. *Machine Learning*, 45(1), 5-32.
- Chehrizi, N., & Weber, T. A. (2015). Dynamic Valuation of Delinquent Credit-Card Accounts. *Management Science*, 61(12), 3077-3096.

- Chehrizi, N., Glynn, P., & Weber, T. A. (2018). Dynamic credit-collections optimization. *Management Science*, Forthcoming Available at SSRN: <https://ssrn.com/abstract=2593728>.
- Chen, T., & Guestrin, C. (2016). *Xgboost: A Scalable Tree Boosting System*. Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, 785-794 (ACM).
- Crook, J. N., Edelman, D. B., & Thomas, L. C. (2007). Recent Developments in Consumer Credit Risk Assessment. *European Journal of Operational Research* 183(3), 1447-1465.
- De Almeida Filho, A. T., Mues, C., Thomas, L. C. (2010). Optimizing the Collections Process in Consumer Credit. *Production and Operations Management*, 19(6), 698-708.
- Dietterich, T. G. (2000). Ensemble Methods in Machine Learning. *International Workshop on Multiple Classifier Systems*, 1-15 (Springer).
- Friedman, J. H. (2001). Greedy Function Approximation: A Gradient Boosting Machine. *Annals of Statistics*, 1189-1232.
- Friedman, J., Hastie, T., & Tibshirani, R. (2001). *The Elements of Statistical Learning*, volume 1 (Springer series in statistics New York, NY, USA).
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T. Y. (2017). Lightgbm: A Highly Efficient Gradient Boosting Decision Tree. *Advances in Neural Information Processing Systems*, 3149-3157.
- Lessmann, S., Baesens, B., Seow, H. V., Thomas, L. C. (2015). Benchmarking State-Of-The-Art Classification Algorithms for Credit Scoring: An Update of Research. *European Journal of Operational Research*, 247(1), 124-136.
- Liebman, L. H. (1972). A Markov Decision Model for Selecting Optimal Credit Control Policies. *Management Science*, 18(10), B-519.
- Mäkikangas, M., Schaufeli, W., Leskinen, E., Kinnunen, U., Hyvönen, K., & Feldt, T. (2017). Handling Delayed Payments. *Global Collections Review*, 17(9), 6–12. Retrieved from <https://atradiuscollections.com/global/publications/global-collections-review.html>.
- Mesropyan, E. (2019). FinTech Startups Changing the Debt Collection Experience for Businesses & Consumers. Retrieved from <https://gomedici.com/fintech-startups-changing-debt-collection-experience-for-businesses-consumers>.
- Miller, G., Weatherwax, M., Gardinier, T., Abe, N., Melville, P., Pendus, C., Jensen, D., Reddy, C. K., Thomas, V., Bennett, J., Anderson, G., & Cooley B. (2012). Tax Collections Optimization for New York State. *Interfaces*, 42(1), 74-84.
- Mitchner, M., & Peterson, R. P. (1957). An Operations-Research Study of the Collection of Defaulted Loans. *Operations Research*, 5(4), 522-545.
- Ng A. Y., Coates, A., Diel, M., Ganapathi, V., Schulte, J., Tse, B., Berger, E., & Liang, E. (2006). Autonomous Inverted Helicopter Flight via Reinforcement Learning. *Experimental Robotics IX*, 363-372 (Springer).

- Olson, R. S., La Cava, W., Mustahsan, Z., Varik, A., & Moore, J. H. (2018). Data-Driven Advice for Applying Machine Learning to Bioinformatics Problems. *Pacific Symposium on Biocomputing*, Volume 23, 192 (NIH Public Access).
- Patrick, J., Puterman, M. L., & Queyranne, M. (2008). Dynamic Multi Priority Patient Scheduling for a Diagnostic Resource. *Operations Research*, 56(6), 1507-1525.
- Powell, W. B., & Topaloglu, H. (2006). Approximate Dynamic Programming For Large-Scale Resource Allocation Problems. *Models, Methods, and Applications for Innovative Decision Making*, 123-147 (INFORMS).
- Shoghi, A. (2019). *Blockchain and Evolution in the Debt Collection Industry*. Proceedings of the 8th Annual Conference on Electronic Banking and Payment Systems, Code: 1579 (EBPS).
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, B., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016) Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature*, 529(7587), 484-489.
- Van de Geer, R., Wang, Q., & Bhulai, S. (2018). *Data-Driven Consumer Debt Collection via Machine Learning and Approximate Dynamic Programming*. SSRN. doi: 10.2139/ssrn.3250755.
- Wang, Q., van de Geer, R., & Bhulai, S. (2018). Data-driven debt collection using-machine learning and predictive analytics. Retrieved from <https://blogs.oracle.com/datascience/data-driven-debt-collection-using-machine-learning-and-predictive-analytics>.