Proceedings of Machine Learning Research vol 211:1–13, 2023 5th Annual Conference on Learning for Dynamics and Control

A Learning and Control Perspective for Microfinance

Xiyu Deng XIYUD@ANDREW.CMU.EDU
Christian Kurniawan CHRISTIAN.PARYOTO@GMAIL.COM
Adhiraj Chakraborty ADHIRAJC@ANDREW.CMU.EDU
5000 Forbes Ave, Pittsburgh, PA, USA

Assane Gueye ASSANEG@ANDREW.CMU.EDU
Regional ICT Center of Excellence Bldg Plot No A8, Kigali Special Economic Zone Phase II, Rwanda

Niangjun Chen NIANGJUN_CHEN@SUTD.EDU.SG
8 Somapah Road, Singapore

Yorie Nakahira YNAKAHIR@ANDREW.CMU.EDU
5000 Forbes Ave, Pittsburgh, PA, USA

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

While microfinance has excellent potential for poverty reduction, microfinance institutions (MFIs) are facing sustainability hardships due to high default rates. Existing methods in traditional finance are not directly applicable to microfinance due to the following unique characteristics: (a) insufficient prior loan histories to establish a credit scoring system; (b) applicants may have difficulty providing all the information required by MFIs to predict default probabilities accurately, and (c) many MFIs use group liability (instead of collateral) to secure repayment. In this paper, we present a novel control-theoretic model of microfinance that accounts for these characteristics and an algorithm to optimize the financing decision in real-time. We characterize the convergence conditions to Pareto-optimum. We demonstrate that the proposed method produces fast decisions and is robust against missing information while still accounting for financial inclusion, fairness, social welfare, sustainability, and the complexities induced by group liability. To the best of our knowledge, this paper is the first to connect microfinance and control theory.
Keywords: Microfinance, Control System, Learning
Source Code: https://github.com/xiyudeng/microfinance

1. Introduction

Microfinance is a category of financial services that gives small loans to low-income people who may not have access to or be eligible for conventional finance (Armendáriz and Morduch (2010); Kamanza (2014)). Microfinance has demonstrated potential for poverty reduction, financial inclusion, and economic development (Mersland and Strøm (2010); Hermes and Hudon (2018); Milana and Ashta (2020)). Despite the proven potential, many microfinance institutions (MFIs) have experienced sustainability hardships, primarily due to the increase in loan default rates (Nawai and Shariff (2012); Addae-Korankye (2014)).

Specifically, MFIs face the following challenges. First, existing approaches, such as credit-score-based or binary classification methods, focus mostly on offline learning, assuming that accurate homogeneous data is abundantly available. Thus, the approaches cannot be directly applied to microfinance without prior loan histories or proper financial systems to establish such learning or credit-scoring procedures accurately (Njagi and Njoka (2021); Alamoudi and Othman (2021)). Moreover, due to the lack of proper mechanisms, it is difficult for some applicants to provide sufficient information for estimating their credit scores and default probability accurately (Arandia and Hepp (2021); Maulana and Umam (2017)). Second, there is increasing evidence that loan approval algorithms based on black-box machine learning techniques may be biased and discriminatory against minorities (Zliobaite (2015); Corbett-Davies and Goel (2018)). Such biases are particularly problematic in fulfilling the objectives of microfinance to provide more opportunities for disadvantaged populations and underdeveloped regions (Alexander-Tedeschi and Karlan (2006); Agier and Szafarz (2013)). Third, the presence of multiple populations/regions also poses additional challenges in allocating microfinance resources to balance different fairness/inclusion objectives. Due to the lack of methodologies that can systematically balance the risks, fairness, and multi-faceted objectives of microfinance, MFIs have relied heavily on the judgment of loan officers, resulting in decisions that sometimes let Portfolio at Risk (PAR) exceed a level to sustain microfinance operations (Yimga (2016); Huo and Fu (2017)).

Our contributions: To address these challenges, we establish a novel microfinance model and propose an algorithm for learning loan approval policies that account for the aforementioned challenges. Our method makes decisions in an online manner that is robust against missing information, accounts for group liability and social fairness, and converges to optimal policy parameters. In addition to theoretical guarantees (Section 3.2), we show via simulation (Section 4) that our method degrades gracefully for increasing levels of missing information in the applications and exploits the potential of group liability while avoiding its pitfalls. The proposed method systematically optimizes competing objectives such as risks, socioeconomic impacts, and active and passive fairness among different groups. The prioritization among different objectives can be specified in the utility function, and the policy has an interpretable structure that informs which factors contributed positively/negatively to applicant approvals. To the best of our knowledge, this paper is the first to use control-theoretic techniques to learn microfinance policy parameters without relying on credit scores directly.

Related works: The most common approach for evaluating regular finance and microfinance is based on credit scores (Ala'raj and Abbod (2016); Shi et al. (2019); Ampountolas et al. (2021); Klaff (2004); Puro et al. (2010)). Some studies predict the default probability of the applicants utilizing logistic regression and its extensions (Bolton et al. (2010); Sohn et al. (2016)). Other approaches consider the loan approval process as a binary classification problem to be solved using machine learning methods such as discriminant analysis (Baesens et al. (2003)), logistic regression (Ala'raj and Abbod (2016); Vaidya (2017)), neural networks (Abdou et al. (2008); Chen et al. (2018); Condori-Alejo et al. (2021); Zhao et al. (2015); Correa et al. (2011)), random forests (Wang et al. (2012); Van Sang et al. (2016)), and SVMs (support vector machines) (Huang et al. (2007); Chen and Li (2010)). However, as mentioned, most of these approaches focus on offline learning, assuming that accurate homogeneous data is abundantly available. In addition, regular loans are given to individuals with collateral, whereas microfinance often uses group liability, where all individuals in the group are liable if any borrower defaults, to secure repayment (Lehner (2009); Kodongo and Kendi (2013); Haldar and Stiglitz (2016)). Group liability can improve the repayment rate by incentivizing members to look after each other, but it has the pitfalls of inducing defaults for borrowers who otherwise have the ability to repay. In addition, because the approaches for granting regular loans do not sufficiently account for the complexities of group liability, group loans have tended to result in higher default rates (Nandhi (2012); Allen (2016)).

2. Problem Statement

Notation: We use capital letters for random variables, *e.g.*, S, and lowercase letters for their specific realization, *e.g.*, s. A square bracket is used to represent the entries of a vector, *e.g.*, $s = \big[s[1], s[2], \cdots, s[n]\big]^\intercal$, and a regular bracket is used for the input of a function, *e.g.*, $f(x)$. We use the notion $(x)_+$ as $\max(0, x)$.

Microfinance model: A microfinance application can be modeled by application properties, MFI's decisions, and loan outcomes. We consider the setting that MFI receives applications and makes decisions at each lending period, indexed by $t \in \{1, 2, \cdots, T\}$. An application is parameterized by the group size M, intrinsic features that govern default probability $S \in \mathcal{S}$, and the MFI's accessible information $\hat{S} \in \hat{\mathcal{S}}$. At the beginning of each lending period t, the MFI receives N_t financing applications, indexed by $i \in \mathcal{N}_t = \{1, 2, \cdots, N_t\}$. Application i has group size $m_{i,t} \overset{i.i.d}{\sim} \mathbb{P}(M = m_{i,t})$, unobserved underlying features $s_{i,t} \overset{i.i.d}{\sim} \mathbb{P}(S)$, and accessible information $\hat{s}_{i,t} \sim \mathbb{P}(\hat{S} \mid S = s_{i,t})$. When some information in S is unavailable, it corresponds to the empty value $\emptyset$ entries in $\hat{S}$. The set of the available information in $\hat{S}$ is denoted by $U(\hat{S}) = \{j : \hat{S}[j] \neq \emptyset\}$. The MFI's lending decision is denoted by a random variable A:

$$A = \begin{cases} 1 & \text{for approval,} \\ 0 & \text{for rejection.} \end{cases} \tag{1}$$

The MFI's approval/rejection probability, $\mathbb{P}(A \mid \hat{S}, M)$, for a certain application is based on the lending policy π_Z, *i.e.*,

$$\mathbb{P}(A \mid \hat{S}, M) = \pi_Z(\hat{S}, M, A). \tag{2}$$

Here, π_Z is controlled by policy parameter Z, defined later in (12).

As the amount of loans given out by MFIs is normally small for each individual, we assume the amount of loan and its interest rate are identical among members within the group and, without loss of generality, are set as 1 and r, respectively. Thus, an approved application of group size M receives a loan of size M and must return the principal and interest of $M \cdot (r + 1)$ at the end of the lending period, where the loan liability is imposed on the whole group. The outcome of the loan (the ability of the applicant to return) is given by

$$B = \begin{cases} 1 & \text{for return,} \\ 0 & \text{for default.} \end{cases} \tag{3}$$

We assume S is independently drawn from some underlying population feature distribution $\mathbb{P}(S)$; $\hat{S}$ is determined based on how features are reflected in the accessible information $\mathbb{P}(\hat{S} \mid S)$; and the outcome of the application is governed by $\mathbb{P}(B \mid S, M) = \mathbb{P}(B \mid S, \hat{S}, M)$, which does not depend on $\hat{S}$ given S and M.

Microfinance decision criteria: The MFI uses policy π_{z_t} to decide on MFI's action $a_{i,t} \sim \mathbb{P}(A \mid \hat{S} = \hat{s}_{i,t},\ M = m_{i,t}) = \pi_{z_t}(\hat{s}_{i,t}, m_{i,t}, a_{i,t})$. At the end of the lending period, the MFI observes the loan outcome $b_{i,t} \overset{i.i.d}{\sim} \mathbb{P}(B \mid S = s_{i,t},\ M = m_{i,t})$ and learns (updates) the policy parameter to z_{t+1}, which is to be used in the next lending period. Here, z_t represents the learned policy parameter at time t and is updated to optimize the objectives specified by MFIs. Microfinance has multifaceted (non-mutually exclusive) objectives including financial inclusion, fairness, social and economic impact, and sustainability. These objectives are captured by a utility function $\mathcal{R}(\{\hat{s}_{i,t}, m_{i,t}, a_{i,t}, b_{i,t}\}_{i \in \mathcal{N}_t})$ for all applications $i \in \mathcal{N}_t$. We defined $V(z_t)$ as the expected utility of the lending period t with policy π_{z_t} and control parameter z_t, *i.e.*,

$$V(z_t) = \mathbb{E}\big[\mathcal{R}(\{\hat{s}_{i,t}, m_{i,t}, a_{i,t}, b_{i,t}\}_{i \in \mathcal{N}_t})\big]. \tag{4}$$

In this study, we considered the utility function for each application in the following form,

$$R(\hat{s}, m, a, b) = \begin{cases} m(r+e); & a = 1, b = 1, \\ m(-1+e); & a = 1, b = 0, \\ 0; & a = 0. \end{cases} \tag{5}$$

Here, $e \in \mathbb{R}_+$ is the financial inclusion factor to motivate the MFI toward approving more applications[1]. Then, $\mathcal{R}(\{\hat{s}_{i,t}, m_{i,t}, a_{i,t}, b_{i,t}\}_{i \in \mathcal{N}_t})$ will be a function of $R(\hat{s}, m, a, b)$ and other objectives such as fairness.

Accounting for fairness. Microfinance decisions should be fair and avoid discrimination against certain populations or regions. Besides the explicit discrimination feature that we could remove directly, we also consider the following two types of implicit fairness.

Type 1 (Outcome fairness): Type 1 fairness actively sets a target approval rate $\Pi^*(\xi)$ for applications with attribute ξ. For example, the loan approval policy of an MFI has a target of at least Π^* approval rate for female applicants as a criterion for gender equality (Khaleghi (2020)). For type 1 fairness, we want to have $\mathbb{P}\left(a_{i,t} = 1 \mid i \in \mathcal{N}_{t,\xi}\right) \geq \Pi^*(\xi), \forall \xi \in \Xi$, where $\Xi = \{\xi, \xi', \cdots\}$ is the set of attributes for which we have a target approval rate and $\mathcal{N}_{t,\xi} = \{i \in \mathbb{N} : i \in \mathcal{N}_t, \hat{S}_\xi \in \hat{S}\}$ is the set of applications with attribute ξ at time t. To achieve type 1 fairness, we can design the reward function as

$$\mathcal{R}\left(\{s_{i,t}, m_{i,t}, a_{i,t}, b_{i,t}\}_{i \in \mathcal{N}_t}\right) = \text{other objectives} - \mathcal{F}_1 \cdot \sum_{\xi \in \Xi} \left(\Pi^*(\xi) - \rho_{\xi,t}\right)_+, \tag{6}$$

where $\mathcal{F}_1 \in \mathbb{R}_+$ is a weighing factor that reflects the relative importance placed upon type 1 fairness, and $\rho_{\xi,t}$ is the current approval rate for applications with ξ, *i.e.*,

$$\rho_{\xi,t} = \frac{1}{|\mathcal{N}_{t,\xi}|} \sum_{i \in \mathcal{N}_{t,\xi}} a_{i,t}. \tag{7}$$

Type 2 (Statistical parity): Type 2 fairness enforces fairness among applications with different attributes, for example, male and female. If we would like to have type 2 fairness among applications with attributes ξ and ξ', then we should have $\mathbb{P}\left(a_i = 1 \mid i \in G_\xi\right) \approx \mathbb{P}\left(a_i = 1 \mid i \in G_{\xi'}\right)$. To enforce type 2 fairness, we can adjust our reward as,

$$\mathcal{R}\left(\{s_{i,t}, m_{i,t}, a_{i,t}, b_{i,t}\}_{i \in \mathcal{N}_t}\right) = \text{other objectives} - \mathcal{F}_2 \cdot \|\rho_{\xi,t} - \rho_{\xi',t}\|, \tag{8}$$

1. MFIs often receive subsidies from international development agencies and governments to help offset high risks of lending without collateral.

where a larger value of $\mathcal{F}_2$ indicates more emphasis is put on type 2 fairness.

Policy learning objectives: The MFI updates z_t to converge to the optimal parameters,

$$z^* = \arg\max_z V(z), \tag{9}$$

such that the cumulative utility converges quickly to that of the optimal policy, with a low policy exploration cost, quantified by

$$V(z^*) - \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} V(z_t)\right]. \tag{10}$$

The expectation $V(z)$ is taken over the probability measurement involving group liability and missing information. Thus, by construction, the optimization of (9) and (10) also accounts for the above-mentioned challenges. Additionally, we take interpretability into the design consideration by imposing structures in policy π_z so that the parameter z informs how much each entry of available information $\hat{S}$ and the group size M have contributed toward approvals or denials.

3. Methodology

3.1. Proposed Algorithm

We propose a lending policy π_z and consider the decision policy in the form

$$\pi_z(\hat{s}, m, a) = L(q), \tag{11}$$

where z is parameterized by

$$z = [\phi^\intercal, \epsilon^\intercal, \gamma^\intercal]^\intercal \in \mathcal{Z} \subset \mathbb{R}^{2n+|\Xi|}. \tag{12}$$

Here, n is the number of features, $\phi \in R^n$ gives weight to the relative importance of each feature to the decision, $\epsilon \in R^n$ accounts for missing information, and $\gamma \in R^{|\Xi|}$ handles the fairness considerations. We consider q as a simple linear combination of the features given by

$$q = \frac{1}{n}\sum_{j\in U(\hat{s})} \phi[j]\hat{s}[j] + \epsilon[j] + \sum_{\xi} \gamma_\xi \, \mathbb{1}\{\hat{s} \in \hat{S}_\xi\}, \tag{13}$$

$\mathbb{1}\{\hat{s} \in \hat{S}_\xi\}$ is an indicator function for $(\hat{s} \in \hat{S}_\xi)$. $\hat{S}_\xi$ refers to feature values that may introduce discrimination in microfinance lending decisions, such as gender, race, ethnicity, etc. and $L(q)$ can be any continuously-differentiable monotonically-increasing function of q that map the domain of q to $[0, 1]$. For example, we consider the following choice of $L(q)$:

$$L(q) = \frac{2\exp(q)}{1+\exp(q)} - 1. \tag{14}$$

We update the policy parameter z_t according to:

$$\hat{z}_{t+1} = z_t + \alpha_t F_{z_t}, \tag{15}$$

$$z_{t+1} = \operatorname*{proj}_{\mathcal{Z}}(\hat{z}_{t+1}), \tag{16}$$

where $\alpha_t > 0$ is the step size to update the parameters at lending period t. The choice of α_t is crucial for the convergence and learning speed of the algorithm and is studied theoretically in Corollary 2. Readers familiar with learning theory would anticipate F_{z_t} as the gradient of the objective $V(z_t)$ as shown in Lemma 3. The explicit form of F_{z_t} is given by

$$F_{z_t}[k] = \left(\mathcal{R}(\{s_{i,t}, m_{i,t}, a_{i,t}, b_{i,t}\}_{i\in\mathcal{N}_t}) - \bar{R}_t\right) w_t[k], \tag{17}$$

$$w_t[k] = \sum_{i=1}^{N_t} w_{i,t}[k], \tag{18}$$

$$w_{i,t}[k] = \frac{1}{\pi_{z_t}(\hat{s}_{i,t}, m_{i,t}, a_{i,t})} \frac{\partial \pi_z(\hat{s}_{i,t}, m_{i,t}, a_{i,t})}{\partial z[k]}, \tag{19}$$

$$\bar{R}_t = \frac{1}{t-1} \sum_{\tau=1}^{t-1} R(\{s_{i,t}, m_{i,t}, a_{i,t}, b_{i,t}\}_{i\in\mathcal{N}_t}). \tag{20}$$

Here, $\dfrac{\partial \pi_{z_t}(\hat{s}_{i,t}, m_{i,t}, a_{i,t})}{\partial z[k]} = g(\hat{s}, k)\dfrac{dL(q)}{dq}$ is the partial derivative of $\pi_z(\hat{s}_{i,t}, m_{i,t}, a_{i,t})$ with respect to the k-th entry of z evaluated at z_t, where g is defined to be

$$g(\hat{s}, k) = \begin{cases} \hat{s}[k]; & k \le n, k \in U(\hat{s}), \\ 1; & n+1 \le k \le 2n, k \in U(\hat{s}), \\ \gamma_{k-2n}; & 2n+1 \le k \le 2n + |\Xi|, \\ 0; & \text{otherwise.} \end{cases} \tag{21}$$

The value of $R(\hat{s}_{i,t}, m_{i,t}, a_{i,t}, b_{i,t}) - \bar{R}_t$ and the update size of z_t will become small when $\bar{R}_t$ is sufficiently close to the sample average. To make sure that the updated parameters stay in the allowable domain $\mathcal{Z}$, step (16) projects $\hat{z}_{t+1}$ onto domain $\mathcal{Z}$. The above procedures are summarized in Algorithm 1.

Algorithm 1 Policy Update

Initialized z_1
for each lending period t **do**
  for each application i **do**
    Generate the decision of application i with:

$$a_{i,t} = \begin{cases} 1; & \text{with probability } \pi_{z_t}(\hat{s}_{i,t}, 1), \\ 0; & \text{with probability } \pi_{z_t}(\hat{s}_{i,t}, 0). \end{cases}$$

    Observe outcome $b_{i,t} \in \{0, 1\}$.
    Gain utility $R(\hat{s}_{i,t}, a_{i,t}, b_{i,t})$.
  end for
  Compute F_{z_t} from (17).
  Update z_{t+1} based on (15) and (16).
end for

3.2. Optimality and Convergence Analysis

Here, we provide conditions for ensuring the proposed algorithm converges to optimal parameters (Theorem 1). Along the way, we explain the ideas behind the updating rules (15) in Lemma 3, and find an appropriate choice for the step size α_t in Corollary 2, based on the results of Theorem 1.

Convergence condition. Algorithm 1 converges to the optimal parameters when the following conditions are fulfilled. **(1)** $L(q)$ is concave over the domain of z; **(2)** the set of admissible policy parameters $\mathcal{Z}$ satisfies

$$\mathbb{E}\left[\|z_{t_1} - z_{t_2}\|^2\right] \leq D^2,\ \forall z_{t_1}, z_{t_2} \in \mathcal{Z}; \tag{22}$$

(3) the second moment of the stochastic gradient is bounded, *i.e.*,

$$\mathbb{E}\left[\|F_{z_t}\|^2 \mid z_t\right] \leq G^2. \tag{23}$$

Here, D restricts the step size in z to avoid overshooting, while G limits the gradient to prevent skipping possible optimum. The convergence conditions are formally stated in the theorem below.

Theorem 1 *Assuming conditions 1, 2, and 3 above hold, let $C(T)$ be defined by*

$$C(T) = \sum_{t=1}^{T} \alpha_t, \tag{24}$$

and $z^* = (\phi^*, \epsilon^*, \gamma^*)$ be defined in (9). Then, Algorithm 1 gives the following performance:

$$\mathbb{E}\left[\sum_{t=1}^{T} \big(V(z^*) - V(z_t)\big)\right] \leq \frac{1}{2}\left(\frac{1}{\alpha_T} D^2 + G^2 C(T)\right). \tag{25}$$

Theorem 1 gives the following relation between step size and convergence speed.

Corollary 2 *When step size is chosen to be $\alpha_t = \dfrac{D}{G\sqrt{t}}$, we have*

$$\mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} (V(z^*) - V(z_t))\right] \leq \frac{3DG}{2\sqrt{T}}. \tag{26}$$

Theorem 1 relies on the property that

$$\mathbb{E}\left[F_{z_t} \mid z_t\right] = \nabla_z V(z_t). \tag{27}$$

Lemma 3 *The updating rules* (15) *and* (16) *given* (17) - (21) *will also satisfy* (27).

Here, from an online learning perspective, one can interpret that the algorithm performs a form of stochastic gradient on utility function V in the presence of missing information where, from a control perspective, $V(z)$ can be interpreted as a Lyapunov function. From Theorem 1, optimization problem (9) converges to the optimal parameters as $T \to \infty$ for the concave utility function. When the utility function of interest is not concave, further care might be needed. Proposition 4 gives conditions under which the utility function is concave.

Proposition 4 *If the approval probability $L(q)$ is a concave function of q, then the objective function $V(z)$ is concave in z.*

The decision rule (14) in Section 2 is an example of a concave function in q over the positive domain. Proposition 4 and condition 3.2 imply that the expected rewards derived from concave decision rules such as (14) are also concave. The proofs of the above theorem, corollary, lemma, and proposition are derived in Kurniawan et al. (2022).

4. Experiment and Results

4.1. Experimental Settings

We considered the utility function (5). In our study, we set the interest $r = 0.35$ (Kneiding and Rosenberg (2008)). We performed simulations for lending periods of $t = 1$ to $t = 500$. At each lending period, we chose $N_t = 10$ applicants, taken from synthetic and artificial pools of applicants containing 10^6 data samples each, generated in advance. The synthetic data pools were generated based on a real microfinance loan dataset from the Kiva platform (Hartley (2010); Kiva Data) with the synthetic data vault algorithm (Patki et al. (2016); Zhang et al. (2019)). The artificial data pools were generated from 30 different artificial distributions. We considered 26 types of distribution with bounded values of feature information in the range of $[0, 4]$ and 4 distribution types with unbounded values of feature information. From the 26 bounded distributions, 8 of them have negative weight for the features. Each loan application i comes with group size $m \in \mathbb{Z}$, generated from the uniform distribution $U(1, 100)$. We then compared our proposed algorithm against the benchmark algorithms, *i.e.*, the credit-score-based method, perceptron, random forests, support vector machines (SVMs), and logistic regression algorithms. A complete description of the data generation procedures and the benchmark algorithms can be found in Kurniawan et al. (2022).

4.2. Performance Comparisons

Robustness against missing data. In this study, we set $e = 0$, and performed the simulation 50 times for each case. We simulated missing information by choosing a missing probability $\mathbb{P}_{missing}$. Here, the feature information $s_{i,t}$ still had all information, but the corresponding observed entry in $\hat{s}_{i,t}$ was empty with probability $\mathbb{P}_{missing}$. We varied the missing probability equal to 0, 10%, 25%, and 50%.[2] To ease the comparison, we normalized the utilities as follows,

$$\{\tilde{\mathcal{V}}(z^*)\} = 2\left(\frac{\{\mathcal{V}(z^*)\} - \min\{\mathcal{V}(z^*)\}}{V_{perfect}}\right) - 1, \tag{28}$$

where $V_{perfect}$ is the utility when perfect knowledge of repayment probability is available, $\{\mathcal{V}(z^*)\}$ is a set of converged utilities of compared algorithms, and $\{\tilde{\mathcal{V}}(z^*)\}$ is a set of normalized converged utilities of compared algorithms.

Figure 1(*a*) shows the performance degradation of the algorithms to a varying level of missing feature information. While the performance of all algorithms decreased with the ratio of missing entries, the proposed algorithm degraded more gracefully than others. This is because, unlike the other algorithms, the proposed algorithm is designed to differentiate the empty entry so that the

2. We did not assume any underlying distribution with respect to the different categories within S that will be accessible.

missing information does not affect the decision policy as much (see Section 2), resulting in an approach that is robust against missing information.

Ensuring fairness among different groups. To investigate the fairness of our algorithm, we introduced a supposedly discriminative feature to the feature vector. This discriminative feature had three discrete values $\{0, 2, 4\}$ and did not affect the actual repayment probability. The algorithm then ran for some target ratio of the discriminative feature with a value of 0 to investigate the outcome fairness (type 1). To investigate statistical parity (type 2) fairness, we considered the discriminative feature with values 0 and 4. Figure 1(*b*) shows the box plot for the acceptance rate for the minority group over the course of the simulation. For type 1 fairness, we can see that by setting the target ratio $\Pi(\xi) = 0.4$, the mean of the acceptance rate is above 0.4. For type 2 fairness, the acceptance rate gap for both groups stays relatively small throughout.

Improved tradeoffs between default risk vs. financial inclusion. There is a tradeoff between default risk vs. financial inclusion because a higher approval rate comes at the expense of higher default risk. To investigate this property, we varied the loan subsidy level e from 0 to 1 with a 0.05 interval. We then recorded the final approval and default rates of each value of e. Figure 1(*c*) shows such tradeoffs for the algorithms tested. The proposed algorithm allows us to systematically tradeoff default risk and financial inclusion by varying the loan subsidy level e. The perceptron, random forests, SVMs, and logistic regression algorithms do not have the flexibility to do so because they cannot be optimized for a utility function. The credit-score-based method is not visible in the current plot range due to large performance degradation in the presence of missing data (10% missing information). The proposed algorithm achieves a reduced default rate for an identical approval rate.

Ability to deal with diverse microfinance distributions. We examined the performance of the compared algorithms using the dataset generated from 30 different distributions and the Kiva dataset as mentioned previously. We considered the cases with 10% and 20% missing information on the artificial and synthetic datasets, respectively. Figures 1(*d*) and 1(*e*) capture the statistic of the normalized steady-state utilities. On average, our algorithm converges to a higher utility than the other algorithms, suggesting that our algorithm can learn a more optimal policy even with incomplete information, achieving our first design goal. This result also shows that the proposed approach can better seize and learn the influence caused by different group sizes. Although it can be seen that the random forest and proposed algorithms have similar performance, the proposed algorithm can adapt better to the dynamics of social and economic conditions.

Adaptation to changes. We studied the adaptability of the proposed algorithm to the dynamics in social and economic conditions by changing the distribution of the dataset in the middle of the simulation. To highlight adaptation to the changes, we chose the distributions that would have similar repayment probability distributions but have opposite feature weights. Figure 1(*f*) shows the performance when application distribution changes at the 150*th* lending period, without the knowledge of if and when the distribution has changed. As can be seen, the proposed algorithm recovers faster than the other algorithms as it uses immediate feedback from the latest samples to perform quick adaptation. In contrast, the other algorithms, which are primarily designed for offline use, end up putting the same emphasis on data from both before and after the changes.

5. Conclusion

We presented a novel control-theoretic model for microfinance lending strategy. The model solves (a) the insufficient past data problem, (b) the missing applicants' information problem, and (c) the

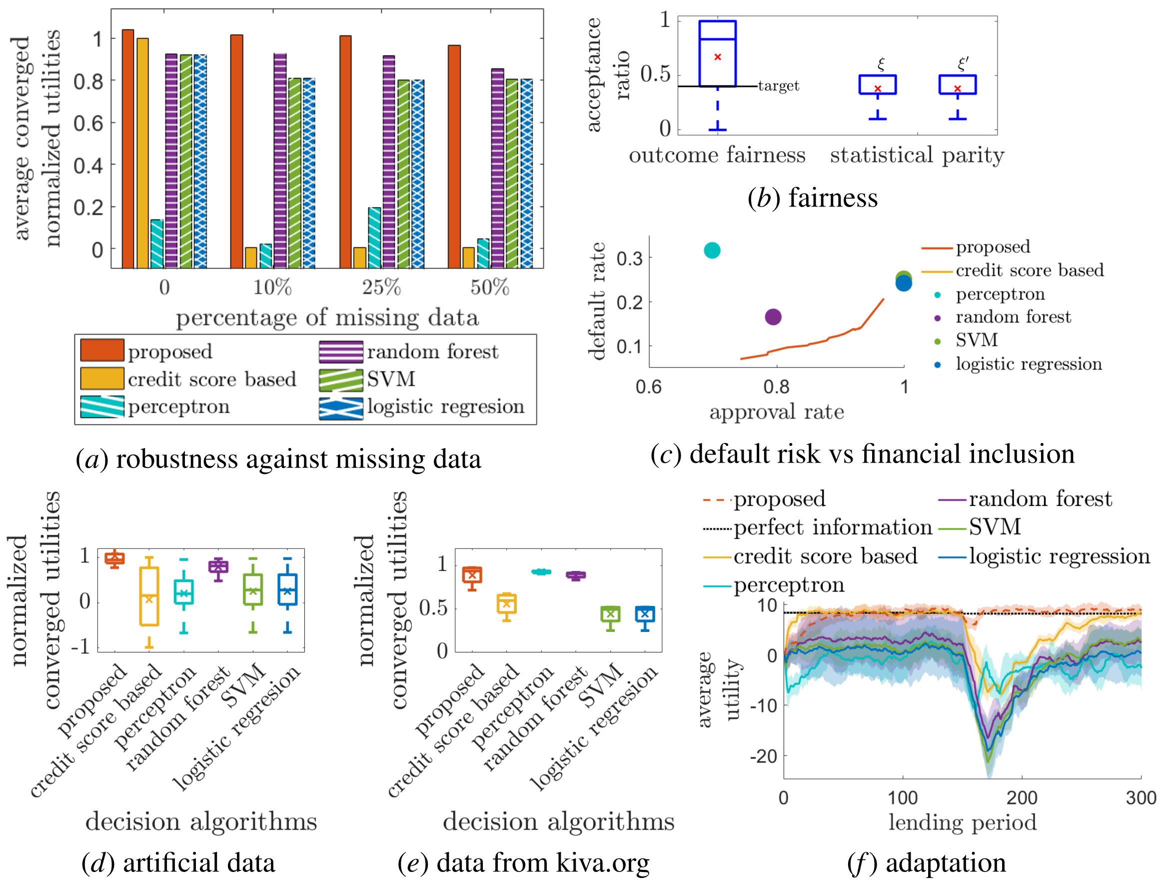

Figure 1: (a) Average converged cumulative normalized utilities for varying rates of missing data. (b) Acceptance ratio statistics of the discriminative features to show fairness types 1 and 2. (c) Tradeoffs between default probability (risk) vs approval rate (financial inclusion). The statistics of the normalized steady-state utilities for (d) diverse applicant distributions and (e) dataset from kiva. (f) The average utility when the distribution changes at the 150th lending period. In subfigures (b), (d), and (e) the mean values are shown by '$\times$', and the first, second, and third quartiles, as well as the maximum and minimum values, are shown as the boxplot. Except for subfigures (a) and (d), the results were generated for 10% of missing information. 20% missing information was considered for the Kiva dataset (d). Except for (c), we set $e = 0$. We consider group liability for all the subfigures.

group liability structure. Extensive empirical results from numerous artificial and synthetic datasets showed several notable performances upon benchmark models, such as robustness against missing data, adaption to changes, and financial inclusion tradeoffs. In addition, we proposed penalty methods for two different fairness scenarios to avoid discrimination in decisions. We hope our model will be useful for achieving the United Nation's Sustainable Development Goals and can help more people in underdeveloped regions have a better life.

References

Hussein Abdou, John Pointon, and Ahmed El-Masry. Neural nets versus conventional techniques in credit scoring in egyptian banking. *Expert Systems with Applications*, 35(3):1275–1292, 2008.

Alex Addae-Korankye. Causes and control of loan default/delinquency in microfinance institutions in ghana. *American International Journal of Contemporary Research*, 4(12):36–45, 2014.

Isabelle Agier and Ariane Szafarz. Microfinance and gender: Is there a glass ceiling on loan size? *World development*, 42:165–181, 2013.

Arwa Abubaker Alamoudi and Anwar Hasan Abdullah Othman. Available financing resources for islamic microfinance institutions to alleviate poverty-cash waqf approach. *Journal of Islamic Finance*, 10:076–084, 2021.

Maher Ala'raj and Maysam F Abbod. Classifiers consensus system approach for credit scoring. *Knowledge-Based Systems*, 104:89–105, 2016.

Gwendolyn Alexander-Tedeschi and Dean Karlan. Microfinance impact: Bias from dropouts. *Unpublished manuscript, Department of Economics, Yale University*, 2006.

Treb Allen. Optimal (partial) group liability in microfinance lending. *Journal of Development Economics*, 121:201–216, 2016.

Apostolos Ampountolas, Titus Nyarko Nde, Paresh Date, and Corina Constantinescu. A machine learning approach for micro-credit scoring. *Risks*, 9(3), 2021. doi: 10.3390/risks9030050.

Osmar Arandia and Saskia Hepp. Poverty alleviation through financial practices: The importance of microfinance: The importance of microfonance. *Handbook on Ethics in Finance*, pages 397–421, 2021.

Beatriz Armendáriz and Jonathan Morduch. *The Economics of Microfinance*. The MIT Press, 2nd edition, 2010. ISBN 978-0262513982.

Bart Baesens, Tony Van Gestel, Stijn Viaene, Maria Stepanova, Johan Suykens, and Jan Vanthienen. Benchmarking state-of-the-art classification algorithms for credit scoring. *Journal of the operational research society*, 54(6):627–635, 2003.

Christine Bolton et al. *Logistic regression and its application in credit scoring*. PhD thesis, University of Pretoria, 2010.

Fei-Long Chen and Feng-Chia Li. Combination of feature selection approaches with svm in credit scoring. *Expert systems with applications*, 37(7):4902–4909, 2010.

Ya-Qi Chen, Jianjun Zhang, and Wing WY Ng. Loan default prediction using diversified sensitivity undersampling. In *2018 International Conference on Machine Learning and Cybernetics (ICMLC)*, volume 1, pages 240–245. IEEE, 2018.

Henry Ivan Condori-Alejo, Miguel Romilio Aceituno-Rojo, and Guina Sotomayor Alzamora. Rural micro credit assessment using machine learning in a peruvian microfinance institution. *Procedia Computer Science*, 187:408–413, 2021.

Sam Corbett-Davies and Sharad Goel. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*, 2018.

Alehandro Correa, Andres Gonzalez, and Camilo Ladino. Genetic algorithm optimization for selecting the best architecture of a multi-layer perceptron neural network: a credit scoring case. In *SAS Global Forum*. Citeseer, 2011.

Antara Haldar and Joseph E Stiglitz. Group lending, joint liability, and social capital: Insights from the indian microfinance crisis. *Politics & Society*, 44(4):459–497, 2016.

Scott E Hartley. Kiva. org: Crowd-sourced microfinance and cooperation in group lending. *Available at SSRN 1572182*, 2010.

Niels Hermes and Marek Hudon. Determinants of the performance of microfinance institutions: A systematic review. *Journal of economic surveys*, 32(5):1483–1513, 2018.

Cheng-Lung Huang, Mu-Chen Chen, and Chieh-Jen Wang. Credit scoring with a data mining approach based on support vector machines. *Expert systems with applications*, 33(4):847–856, 2007.

Xiaoguang Huo and Feng Fu. Risk-aware multi-armed bandit problem with application to portfolio selection. *Royal Society open science*, 4(11):171377, 2017.

Rashid Mwatemo Kamanza. Causes of default on micro - credit among women micro - entrepreneurs in kenya. a case study of women enterprise development fund (wedf) msambweni constituency. *IOSR Journal of Economics and Finance*, 3:32–47, 2014.

Amirali Khaleghi. *Towards Achieving Gender Equality in Automated Loan Approval Processes*. Master thesis, Delft University of Technology, 2020.

Kiva Data. Kiva Loans, 2011. URL `https://stat.duke.edu/datasets/kiva-loans`.

Leslie Gross Klaff. New internal hiring systems reduce cost and boost morale. *Workforce*, 2004.

Christoph Kneiding and Richard Rosenberg. Variations in microcredit interest rates. *Consultative Group to Assist the Poor (CGAP)*, 2008.

Odongo Kodongo and Lilian G Kendi. Individual lending versus group lending: An evaluation with kenya's microfinance data. *Review of Development Finance*, 3(2):99–108, 2013.

Christian Kurniawan, Xiyu Deng, Adhiraj Chakraborty, Assane Gueye, Niangjun Chen, and Yorie Nakahira. A Learning and Control Perspective for Microfinance. *arXiv preprint arXiv:2207.12631*, 2022.

Maria Lehner. Group lending versus individual lending in microfinance. Technical report, SFB/TR 15 Discussion Paper, 2009.

Hartomi Maulana and Khoirul Umam. Identifying financial exclusion and islamic microfinance as an alternative to enhance financial inclusion. *International Journal of Islamic Business and Economics (IJIBEC)*, pages 99–106, 2017.

Roy Mersland and R Øystein Strøm. Microfinance mission drift? *World development*, 38(1):28–36, 2010.

Carlo Milana and Arvind Ashta. Microfinance and financial inclusion: Challenges and opportunities. *Strategic Change*, 29(3):257–266, 2020.

Mani A Nandhi. Incidence of loan default in group lending programme. *CMFWorking Paper*, 2012.

Norhaziah Nawai and Mohd Noor Mohd Shariff. Factors affecting repayment performance in microfinance programs in malaysia. *Procedia-Social and Behavioral Sciences*, 62:806–811, 2012.

Joram Nyaga Njagi and Charity Njoka. Microfinance reforms and financial inclusion in kenya. *International Journal of Current Aspects in Finance, Banking and Accounting*, 3(1):54–72, 2021.

N. Patki, R. Wedge, and K. Veeramachaneni. The synthetic data vault. In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 399–410, Oct 2016. doi: 10.1109/DSAA.2016.49.

Lauri Puro, Jeffrey E Teich, Hannele Wallenius, and Jyrki Wallenius. Borrower decision aid for people-to-people lending. *Decision Support Systems*, 49(1):52–60, 2010.

Baofeng Shi, Xue Zhao, Bi Wu, and Yizhe Dong. Credit rating and microfinance lending decisions based on loss given default (lgd). *Finance Research Letters*, 30:124–129, 2019.

So Young Sohn, Dong Ha Kim, and Jin Hee Yoon. Technology credit scoring model with fuzzy logistic regression. *Applied Soft Computing*, 43:150–158, 2016.

Ashlesha Vaidya. Predictive and probabilistic approach using logistic regression: application to prediction of loan approval. In *2017 8th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, pages 1–6. IEEE, 2017.

Ha Van Sang, Nguyen Ha Nam, and Nguyen Duc Nhan. A novel credit scoring prediction model based on feature selection approach and parallel random forest. *Indian Journal of Science and Technology*, 9(20):1–6, 2016.

Gang Wang, Jian Ma, Lihua Huang, and Kaiquan Xu. Two credit scoring models based on dual strategy ensemble trees. *Knowledge-Based Systems*, 26:61–68, 2012.

Jules Yimga. The impact of high microfinance growth on loan portfolio. *Journal of International Development*, 28(5):697–714, 2016.

Kevin Zhang, Felipe Alex Hofmann, and Plamen Valentinov Kolev. The synthetic data vault., 2019. URL https://sdv.dev/.

Zongyuan Zhao, Shuxiang Xu, Byeong Ho Kang, Mir Md Jahangir Kabir, Yunling Liu, and Rainer Wasinger. Investigation and improvement of multi-layer perceptron neural networks for credit scoring. *Expert Systems with Applications*, 42(7):3508–3516, 2015.

Indre Zliobaite. A survey on measuring indirect discrimination in machine learning. *arXiv preprint arXiv:1511.00148*, 2015.