Portfolio Choice for Online Loans and Implications for Platforms

Ram D. Gopal Xiao Qiao Moris S. Strub Zonghao Yang *

November 5, 2022

Abstract

Using over one million LendingClub loans from 2013 to 2020, we investigate the suitability of online loans as an investment through the lens of a portfolio optimization framework. We introduce general characteristics-based portfolio policies, a framework which overcomes unique challenges associated with building a portfolio of online loans. Under this framework, we propose a nonlinear portfolio policy based on a shallow neural network. Whereas an equal-weight portfolio achieves an average annual internal rate of return of 6.55%, a nonlinear portfolio leads to an improved annual IRR of 13.08%. The nonlinear portfolio also enables more access to credit by investing more in loans with lower credit grades. To assess the attractiveness of online loans, we compare the performance of the nonlinear portfolio to other benchmark assets, including stocks, bonds, and real estate. We find that in our sample, online loans earn competitive rates of return to the other assets while showing limited comovement. Our results indicate that online loans are an attractive novel asset class for investors, and investors can diversify their holdings by investing in online loans with increased expected returns. Platforms may consider embedding a GCPP framework in a robo-advising system, which would expand the access of sophisticated loan portfolios to a broad set of investors.

Keywords: Online Lending, Portfolio Optimization, Neural Networks, Fintech

---

*Gopal: University of Warwick; Qiao: City University of Hong Kong and Hong Kong Institute for Data Science; Strub: Southern University of Science and Technology; Yang: City University of Hong Kong. Please send correspondence to Qiao (xiaoqiao@cityu.edu.hk) and Yang (zonghao.y@my.cityu.edu.hk). We thank Carrie Wang, Jinguang Yang, conference participants at 2022 Greater China Area Finance Conference, 2022 Conference on Information Systems and Technology, 2022 INFORMS Annual Meeting, seminar participants at HEC Paris and Korea Advanced Institute of Science and Technology, for valuable comments and discussions. We thank Yudong Liang for excellent research assistance when collecting loan data.

1 Introduction

Online lending platforms enable greater participation in credit markets by allowing individuals who otherwise cannot borrow to do so. They also tend to offer lower interest rates than traditional sources of funding such as credit unions and banks (Treasury, 2016). Online lending is an example of how financial technology (fintech) is driving innovation in the financial services sector, changing the nature of business and consumer expectations. In the past decade, the rise of online marketplaces for loans had a significant impact on traditional credit markets and drove the disintermediation of banks (Drummer et al., 2017).

Despite the growing popularity of online lending, major online lending platforms such as LendingClub, OnDeck, and Funding Circle have struggled to turn a profit (Kate and Nicholas, 2019). An additional challenge faced by online lending platforms is the ongoing shift in investor base from mostly individual lenders to institutional lenders (Treasury, 2016). In the U.S., U.K, and China, this shift in the investor base is driven by regulatory changes intended to improve individual investor protection (Nemoto et al., 2019). Whether the transition to institutional investors is viable for platforms crucially depends on whether online loans constitute an attractive investment opportunity for a sophisticated investor.

We study two crucial questions facing prospective investors of online loans: First, is it worthwhile to invest in this new asset class, given her existing portfolio? If the answer is yes, then, second, what is the appropriate investment strategy? The investor needs to compare the attractiveness of online loans with other investable assets, and she needs to solve a portfolio choice problem should she choose to invest in online loans. In this paper, we address these two questions by developing a novel portfolio optimization framework tailored for the specific setting of investment in online loans.

We start with the second major research question. If the investor decides to invest in online loans, she needs to determine which loans to invest in and how much capital to allocate to each loan. That is, her decision becomes a portfolio optimization problem. Online loans fundamentally differ from stocks, where people typically discuss portfolio optimization. Loans have limited investment amounts available, so the lenders cannot invest as much as they want in a single loan. Investors

Figure 1: Histogram of Loan ROIs

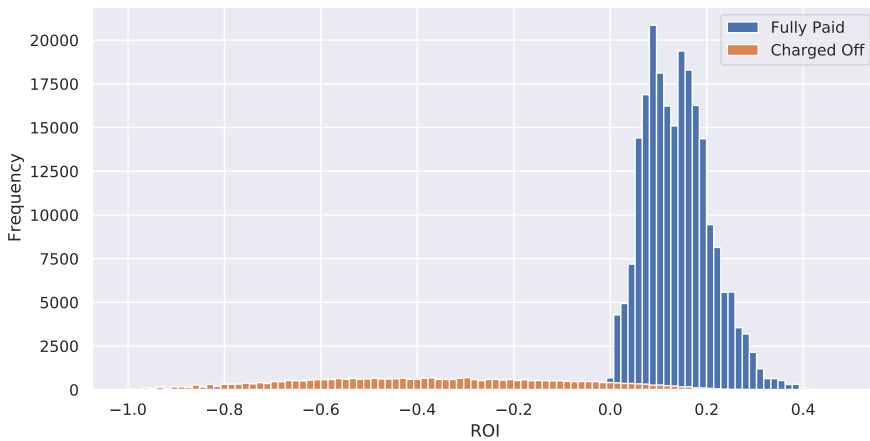

cannot foresee if or when new loans will become available. More importantly, loans do not have a liquid secondary market. Therefore, once a loan is fully funded, the investment opportunity is no longer available.

In particular, there are two key challenges associated with building a portfolio of online loans. The first challenge is that loans may be charged off[1], and the set of loans that eventually charged off makes up a different distribution compared to those that do not. Figure 1 illustrates that the return on investment (ROI) of loans is made up of two distinct distributions, one for loans that charged off and one for fully paid loans. Much of the literature (see Section 1.1) on online loans focuses on assessing the credit risk of loans, but these papers seldomly explore other potentially useful information. Figure 1 shows that beyond binary classification, the two distributions contain additional information about online loans. Therefore, a binary classification of loans is likely not sufficient for the construction of an attractive portfolio of loans.

The second challenge pertains to the inputs needed for portfolio optimization. In a typical portfolio optimization approach, we need to model the joint distribution of returns and then solve for optimal portfolio weights. These steps are difficult to implement for a large number of assets, and they are especially susceptible to unstable results when the inputs are noisy (Michaud, 1989;

[1]If a loan defaults, the borrower is charged a fee for each missed payment. If a loan is delinquent by more than 120 days, the loan status changes from "defaulted" to "charged off" and LendingClub may sell this loan to a third party for collection. We discuss the loan mechanics in more detail in Section 3.1.

DeMiguel et al., 2009). The sheer size of the number of loans makes it prohibitive to estimate the covariance matrix - LendingClub alone lists over 500 new loans daily. The heterogeneity of loans further complicates expected return and covariance matrix estimation. A loan may have unique characteristics due to its borrower using the platform for the first time, or the repayment behavior of a borrower may change under different personal or macroeconomic conditions. These issues make traditional portfolio optimization methods particularly unsuitable for online loans.[2]

We introduce general characteristics-based portfolio policies (GCPP), a framework to overcome the key challenges associated with the portfolio construction of loans. GCPP directly models the portfolio weight in each loan as a function of its characteristics, thereby addressing the challenge of estimating the distributional properties of loans. GCPP extends the linear parametric portfolio policy of Brandt et al. (2009) to a nonlinear portfolio policy based on neural networks. The neural network considers the rich nonlinear interactions of characteristics and allows the portfolio to search in a broader portfolio weight space.

We collect loan samples from LendingClub, the world's largest provider of online loans.[3] To the best of our knowledge, our dataset is the most extensive in the online lending literature to date, covering more than one million loans over a period from 2013 to 2020. Each month, we form an equal-weight portfolio of all available loans, and we calculate its monthly internal rate of return (IRR). The equal-weight portfolio constitutes a natural benchmark against which the performance of optimized portfolios can be evaluated (see, *inter alia*, DeMiguel et al., 2009). Furthermore, the equal-weight portfolio is indicative of the performance of an investor who is interested in online loans but lacks the tools to construct a more sophisticated portfolio, or of an institutional investor seeking broad diversification in online loans.

The portfolio performance is trained and tested following a yearly expanding window. We first train the models using loans from 2013 and test the portfolio performance on loans from 2014. Then, we keep expanding the training set a year ahead and test the performance on the following year until we exhaust all loan samples. In the out-of-sample period, we obtain a separate loan portfolio each month, with a maturity of 36 months. These portfolios can then be compared to

[2]Despite these issues, some researchers have applied a mean-variance framework to construct loan portfolios (Guo et al., 2016; Chi et al., 2019; Byanjankar et al., 2021; Guo et al., 2021). The mean-variance framework works well with symmetric, bell-shaped returns, but loan ROIs are not symmetrically distributed (Figure 1).

[3]Reference available in the Annual Reports of LendingClub at the U.S. Securities and Exchange Commission (SEC).

the equal-weight portfolios over the same investment intervals. The equal-weight portfolios have an average out-of-sample annual IRR of 6.55%, whereas the GCPP approach leads to significant improvements. A linear GCPP - equivalent to the linear parametric portfolio policy of (Brandt et al., 2009) - has an IRR of 8.86%. A nonlinear GCPP leads to superior portfolio performance of an IRR of 13.08%. We find that GCPP assigns higher weights to loans with lower grades, implying our framework is able to enable greater participation in the credit market. These results indicate that general characteristics-based portfolio policies constitute a suitable framework for constructing a portfolio of online loans and, more generally, there can be significant benefits for both borrowers and investors from employing sophisticated techniques for investing in online loans.

As a practical challenge, investors are faced with constraints on the amount that can be invested in an individual loan. On LendingClub, the minimum investment per loan is $25. Investors also cannot lend more than the requested loan amount, which places an upper limit on the investable amount. We investigate the effect of these constraints on the optimal portfolio, with a total investable amount ranging from $1,000 to $100 million representing investors of various sizes. We find that the investment performance remains relatively stable over this range of portfolio sizes, but the binding constraints on the portfolios are not identical. For the smallest investment amounts, the minimum investment per loan becomes binding. Consequently, such a portfolio only invest in a small set of loans. For the largest investment amounts, the upper limit on the investable amount restricts the maximum weight that can be placed on certain loans. In our sample, the $100,000 portfolio deviates the least from the optimal unconstrained portfolio, striking a balance between the minimum and maximum investable amounts.[4]

Our other main research question focuses on whether it is worthwhile to invest in online loans at all given other investment opportunities. For an investor, the decision to invest in online loans ultimately depends on whether the loans offer a worthwhile expansion of her opportunity set. To make this assessment, she needs to compare online loans with other asset classes. An index tracks the performance of a group of assets in a standardized way, thus, provides a comparison benchmark for our loan portfolios. We consider six indexes, i.e., S&P 500 Index, Bloomberg U.S. Aggregate Bond Index, S&P 1-3/3-5/10-20 Year U.S. Treasury Bond Indexes, and MSCI U.S. REIT Index. The

[4]On average, $100,000 is about 0.03% of the total funded loan amount each month in LendingClub.

performance of a loan is typically measured by the return on investment or internal rate of return calculated based on cash flows. However, these measures are not directly comparable to index returns. Loans do not have a liquid secondary market where investors can freely buy and sell, so investors cannot earn the calculated ROI or IRR until the loan matures. Loan portfolios also have different cash flows compared to index portfolios: Loan investors typically lend all the money upfront and receive monthly payments until each loan is either fully paid or charged off. In order to make a fair comparison between a portfolio of online loans and indexes, we construct portfolios with identical cash flows as the loans, but with the cash flows invested in the index funds. The IRRs of these portfolios are then directly comparable to the IRRs of the loan portfolios. The levels and comovement of loan and index IRRs shed light on whether it is advisable to invest in online loans.

To quantify and evaluate the attractiveness of online loans relative to index investments, we construct portfolios whose cash flows are identical to the $100,000 nonlinear loan portfolios but instead invest in the respective index funds. We then compare the IRRs of these portfolios with those of the investment in online loans. The average public market equivalent (PME)[5], the ratio of the respective IRRs of the loan portfolio to the six benchmark index portfolios, is 1.93, 4.84, 17.99, 6.00, 2.81, and 2.82 in our out-of-sample period. The result indicates that, with our nonlinear policy, the loan loans earn much higher rates of return compared to stocks, bonds, and real estate. Furthermore, the monthly IRRs of the loan portfolios have a correlation less than 0.22 with those of the indexes using the same cash flows, suggesting significant diversification benefits. In contrast to the performance of the nonlinear portfolio, the average PME of an equal-weight portfolio of loans against the benchmarks is just 0.93 (< 1), 2.41, 8.79, 3.38, 1.51, and 2.18, respectively. The difference in average PME between the nonlinear and equal-weight loan portfolios indicates a crucial role that sophisticated portfolio optimization plays in expanding the investor's opportunity set. Indeed, online loans are much less attractive if one were to compare the performance of a naïve equal-weight portfolio of loans with the S&P 500 Index.

Our findings shed light on the investor's consideration of whether and how to invest in online loans. The nonlinear portfolios have high average IRRs that are not correlated with invest-

[5]Public market equivalent is a common metric used to measure the performance of private equity investment compared to a stock market index. This metric is suitable for the context of online lending because the liquidity and a lack of secondary market for online loans (Treasury, 2016) are similar to those characteristics of private equity deals.

ments in the index funds, indicating an attractive novel asset class for investors. An investor with heavy stock market exposure can diversify her holdings by allocating some capital to online loans to increase the overall portfolio return with reduced risk. Our framework achieves superior performance because the GCPP framework is designed to capture the unique properties of online loans.

While our study was performed under the light of the ongoing shift in investor base, we emphasize that our results are by no means limited to institutional investors. Instead, we provide recommendations to help platforms attract and retain their customer base in more general terms. Our findings suggest that investors in online loans can improve their investment performance through more sophisticated portfolio construction. Platforms may consult our findings to help investors better manage their loan investments, or platforms may consider embedding the GCPP framework in a robo-advising system that can offer personalized portfolio recommendations. Such a system would be straightforward to implement and can be quickly retrained to incorporate additional loan characteristics. Our results also indicate that investors can benefit from reduced minimum investment limits on individual loans. As the minimum investment limit increases, the average IRR for the nonlinear portfolio does not vary much, but the portfolio standard deviation rises due to decreased diversification across loans. Investors' opportunity set can improve if platforms set the minimum investment limit as low as it is operationally feasible.

The remainder of the paper is organized as follows. Section 2 introduces the dataset and the mechanics of online loans. Sections 3 and 4 discuss the portfolio choice problem in online lending. Section 5 investigates whether online loans expand the investor's opportunity set. Section 6 concludes.

1.1 Literature Review

A large literature attempts to predict the credit risk of online loans using loan and borrower characteristics[6] (Iyer et al., 2016; Herzenstein et al., 2008; Emekter et al., 2015; Serrano-Cinca et al., 2015). Researchers have explored a variety of methods for credit risk assessment, including sin-

[6]Some studies find that incorporating alternative information, such as appearance (Pope and Sydnor, 2011 and Duarte et al., 2012), social network on the online lending platforms (see Lin et al., 2013, Everett, 2015, Liu et al., 2015), and text description in the loan application (Du et al., 2020, Xu and Chau, 2018, Wang et al., 2020, Jiang et al., 2018), can improve the accuracy in assessing the credit risk of borrowers.

gle learner algorithms such as logistic regression and support vector machine (Serrano-Cinca et al., 2015; Cho et al., 2019) to more complex ensemble algorithms such as random forest and eXtreme Gradient Boosting (XGBoost) (Fu et al., 2021; Xia et al., 2017; Malekipirbazari and Aksakalli, 2015). Wang et al. (2019) propose a deep learning model named "NeuCredit" which takes consumer purchase history and payment records as inputs. This model significantly outperforms other ML algorithms (also see Yang et al., 2018 and Manzo and Qiao, 2021).

While credit risk assessment is an important aspect of online loans, it does not fully characterize loans. Serrano-Cinca and Gutiérrez-Nieto (2016) argue that because lenders may not lose the full amount in a loan that is charged off and loans with higher charged-off probability may have higher interest rates, assessing the profit of online loans rather than their credit risk can better reflect their investment potential. In the same spirit, we also argue that investors should not only consider the probability of a loan to be charged-off, but also additional information in the return distribution when making investments in online loans.

A separate strand of literature attempts to solve a portfolio optimization problem facing investors of online loans. Guo et al. (2016) propose an instance-based credit risk assessment framework to form loan portfolios. The authors approximate the return and risk of a current loan as weighted averages of past loans, where the weights are assigned based on a similarity measure. Given the estimates for expected returns and risk, Guo et al. (2016) construct a mean-variance (MV) optimal portfolio. Following Guo et al. (2016), researchers have developed new ways to measure the similarity of loans (Guo et al., 2021; Byanjankar et al., 2021) and extended the optimization problem to a robust MV optimization problem based on a relative entropy method (Chi et al., 2019).

Mean-variance optimization heavily relies on the quality of model inputs; solutions to MV optimization problems can be highly unstable given the estimation error of the input parameters (Michaud, 1989; DeMiguel et al., 2009). The heterogeneity of individual loans makes it difficult to precisely estimate their expected returns and risk, and the estimation errors are carried through in the portfolio construction process. Furthermore, Guo et al. (2016) and related work ignore covariances in their MV setup because they are infeasible to estimate for such a large cross section of loans. Our work contributes to the literature by introducing a generalized portfolio framework suitable for online loans, overcoming the challenges faced by mean-variance optimization. The GCPP approach bypasses the need to estimate average returns and the covariance matrix through a direct parameterization of portfolio weights in terms of loan characteristics.

2 Dataset

We collect loan data from LendingClub between 2007Q1 and 2020Q3.[7] During this period, LendingClub expanded rapidly, both in the number of loans and the total dollar amount funded (see Figure 2). The number of loans grew by 76% per year, whereas the dollar amount grew by 87% per year on average. In 2019 alone, lenders funded 500,000 loans worth a total of $8.5 billion.

Figure 2: LendingClub Funded Loans per Year

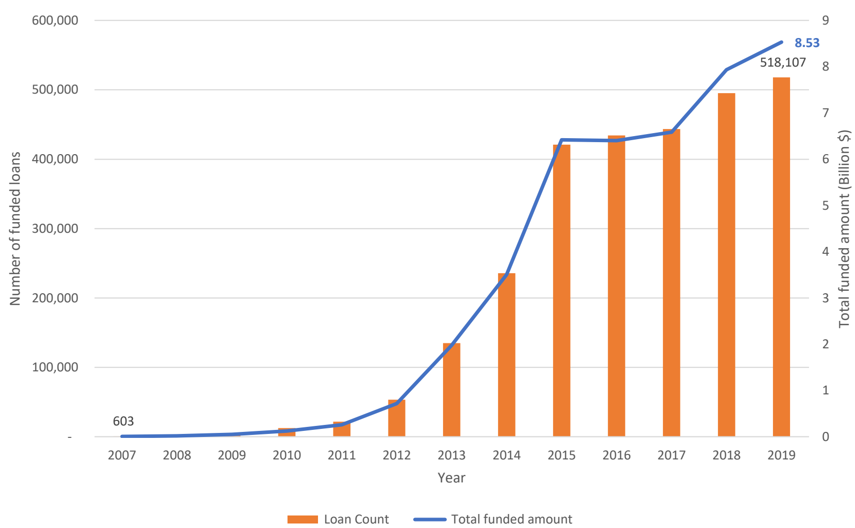

LendingClub loans either mature after 36 months or 60 months. Our analysis centers around 36-month loans, which allows us to work with more recent data. A number of characteristics are associated with each loan such as the *amount of loan borrowed*, *interest rate*, *loan grade*, and its related borrower characteristics including the *number of delinquent loans in the account*, *gross income*, and *debt-to-income ratio*. Over time, LendingClub has required more detailed information in its loan application, resulting in an increase in loan and borrower characteristics from 46 in 2007 to 83 in 2013. In order to balance having a broad set of characteristics and a long history of data, our sample starts in January 2013. At the time of collecting the loan data, most loans originating after May

[7]Historical loan listings with detailed loan characteristics are publicly available up to 2020. LendingClub stopped providing listing data after the platform stopped offering loans to retail investors at the end of 2020.

2017 have not yet been completed. We therefore do not include loans that originated after May 2017. After the above preprocessing, our dataset contains 1,158,476 loans with 83 loan and borrower characteristics. The details of the characteristics can be found in Appendix A.

To the best of our knowledge, our dataset is the most extensive sample of online lending platforms to date. Our data cover over one million loans, whereas existing studies contain tens or hundreds of thousands of loans (Fu et al., 2021; Guo et al., 2021; Yang et al., 2018).

2.1 Performance Measures

LendingClub follows a posted pricing mechanism. The platform assigns a credit grade (A1, A2, ..., A5, ..., G1, ..., G5) to each loan, and, at any given point in time, loans with the same credit grade are all assigned the same interest rate. Once a loan is issued, the borrower pays a fixed monthly installment based on the loan principal and interest rate. If the borrower misses a monthly payment, the loan is considered to be in default. The borrower is charged late fees for each missed payment and experiences a negative impact on her credit profile. If a loan is delinquent by more than 120 days, the loan status changes from "defaulted" to "charged off." LendingClub may proceed to sell this loan to a third party collection agency, and the lenders will receive a pro-rata share of the sale proceeds and any recovery amount. For its services, LendingClub charges an origination fee to the borrowers and a service fee to the lenders.

The cash flow of a loan consists of five possible parts: Principal, interest, late fees, sales proceeds, and recovery amount.[8] From the lender's perspective, the typical set of cash flows for a loan may look like the following: $\{-P_0, P_1, \dots, P_T\}$, consisting of an initial cash outflow P_0 and monthly cash inflows $\{P_t, t \in \{1, \dots, T\}\}$ until the loan is fully paid or charged off. We can approximate the monthly payments of a loan by the average monthly payment, i.e., the total payments received divided by the number of months between origination and the last payment.[9] The cash flows for a portfolio of loans can be generated by summing up the cash flows of constituent loans.

[8] Lenders are charged a 1% service fee for each payment they receive. If borrowers pay back early, LendingClub calculates service fees in a manner protecting lenders' interest. Given that LendingClub does not disclose the details of this calculation, we do not consider this service fee in calculating loan returns. Since this fee can potentially apply to all loans, it does not affect our cross-sectional comparison across loans to identify the best investment opportunity.

[9] The monthly payments information is not available in our dataset. We approximate the monthly payment based on *loan origination date*, *last payment date*, *total payments*, and *collection recovery fee*.

Internal rate of return (IRR) is a common metric used to evaluate the attractiveness of a stream of cash flows. The IRR is the discount rate that makes the net present value (NPV) of all cash flows equal to zero:

$$0 = NPV = \sum_{t=1}^{T} \frac{P_t}{(1+IRR)^t} - P_0. \tag{1}$$

Return on investment is an alternative metric to measure loan performance. ROI is calculated as the ratio of an investment's net profit (or loss) to its initial cash outlay. We compute the ROI of individual loans based on the cumulative discounted payment (CDP) received by the lender,

$$CDP = \frac{P_1}{(1+\frac{d}{12})^1} + \frac{P_2}{(1+\frac{d}{12})^2} + \cdots + \frac{P_T}{(1+\frac{d}{12})^T}, \tag{2}$$

where d is a discount rate that reflects the time value of money. In this study, we fix $d = 2\%$ per year.[10] ROI is then defined as:

$$ROI = \frac{CDP - P_0}{P_0}. \tag{3}$$

Neither metric is a perfect measure of loan performance. IRR is determined implicitly and cannot be expressed in analytical terms, while ROI can be expressed in analytical terms. Assume a loan portfolio of N loans, denoted as $\omega \in \Delta^{N-1}$, where $\Delta^{N-1} = \{\omega \in \mathbb{R}^N : \sum_{i=1}^{N} \omega_i = 1\}$. The portfolio ROI r_ω is the weighted sum of individual loan ROIs (r_i for loan i, $i = 1, \ldots, N$):

$$r_\omega = \sum_{i=1}^{N} \omega_i r_i. \tag{4}$$

In this way, we can derive the derivatives of the objective with respect to the model parameters. The portfolio IRR cannot - it is some unknown function of which the derivatives with respect to the parameters are not readily available. Using ROI in the objective function enables us to use common optimization algorithms to solve the optimization problem.

However, the computation of ROI requires the specification of an additional parameter, the discount rate in Equation (2). Since the calculation of IRR does not need such a discount rate, it requires fewer assumptions. Considering their advantages and disadvantages, we leverage the

[10]Because all loans share the same value for d, a different value for d does not affect the proposed portfolio optimization framework, and the empirical results remain qualitatively unchanged.

tractability of ROI when solving the portfolio optimization problem, and we primarily present IRR comparisons for portfolio performance evaluation.

Table 1 shows the summary statistics for our dataset. The three panels describe variables related to loan outcome, loan characteristics, and borrower characteristics. The overall charged-off rate is 15.3% and the average ROI is 3.13%. We observe clear discrepancies in the summary statistics between the completed and charged-off loans, indicating different distributions depending on the loan status.

3 Portfolio Choice for Online Loans

Lenders of online loans are faced with a portfolio choice problem: Which loans are worth investing in, and how much to allocate to each loan? As discussed in Section 1.1, researchers mainly rely on two existing methods for investment in online loans, credit risk filtering (Fu et al., 2021; Cho et al., 2019) and the mean-variance framework (Guo et al., 2016, 2021). These two methods both have their justifications and limitations.

The credit risk filtering approach uses a binary classification algorithm to predict the outcomes of individual loans and seeks to filter out the loans with the highest charged-off probability. Confident determination of loan status proves to be a difficult task. Even sophisticated machine learning algorithms such as gradient boosting and neural networks perform poorly in distinguishing loans that are eventually charged off from those that are not. Xia et al. (2017) summarize the performance of a wide range of algorithms on credit risk assessment using LendingClub dataset including logistic regression, decision tree, XGBoost, random forest, and neural networks. Let loans that are charged off be assigned to a positive label. XGBoost, the best-performing algorithm, has a false negative rate of 29.8%, and a false positive rate of 36.1%.[11] The former value indicates that the investor will invest in 29.8% of the charged-off loans, and the latter value means she will not invest in 36.1% of the fully-paid loans. Thus, if an investor solely relies on binary classification to make her investment decisions, her portfolio will contain a substantial fraction of charged-off loans and she will miss many investment opportunities by filtering out fully-paid loans.

[11] $FNR = \frac{FN}{FN+TP}$, $FPR = \frac{FP}{TN+FP}$ where *FN* is the number of false negative predictions, *FP* is the number of false positives, *TP* is the number of total positive observations, and *TN* is the total number of negative observations.

Table 1: Summary Statistics of the Dataset[a]

	Funded	Completed	Charged-off
		Panel A: Outcome	
Sample size	1,158,476	981,416	177,060
ROI, %	3.13	11.69	-44.29
IRR, %	-1.16	1.22	-14.33
		Panel B: Loan characteristics	
Loan amount, $	12,641	12,612	12,802
Installment, $	419	416	438
Interest rate average[b], %	12.04	11.66	14.11
Loan Grade, %			
A	23.63	26.21	9.33
B	33.53	34.71	27.00
C	26.85	25.46	34.56
D	12.02	10.49	20.55
E	3.24	2.61	6.77
F	0.61	0.46	1.48
G	0.11	0.07	0.32
		Panel C: Borrower characteristics	
Stated annual income, $	76,240	77,568	68,880
Income verifiable, %	66.74	65.53	73.47
Debt-to-income ratio, %	17.92	17.64	19.46
Is homeowner, %	11.64	11.52	12.32
FICO score range, %			
660-700	60.95	58.80	72.89
700-750	30.76	32.05	23.61
750-800	7.04	7.75	3.12
800-850	1.24	1.40	0.38
Inquiries within 6 months[c]	0.60	0.58	0.75
Length of credit history, month	195.37	197.07	185.92
Current balance to credit limit ratio, %	57.16	56.56	60.47
Accounts opened in past 12 months	2.14	2.07	2.50
Percentage of bankcard accounts over 75% limit, %	42.01	41.24	46.25
Total credit revolving balance, $	15,562	15,847	13,980
Total credit balance excluding mortgage, $	47,819	48,306	45,120
Current balance of all accounts, $	131,689	136,596	104,492
Average current balance of all accounts, $	12,670	13,185	9,817
Total revolving credit limit, $	32,573	33,469	27,605

[a] The dataset consists of loans originating from January 2013 to May 2017.
[b] The minimum interest rate is 5.31% and the maximum interest rate is 30.99% in our sample.
[c] Inquiry is a pull of the credit report when the borrower applies for a loan, whether the loan application is successful or not. The number of inquiries indicates how many times that the borrower applied for loans across all lending sources.

Moreover, we observe that the return on investment of loans consists of a mixture of two distinct distributions (Figure 1). While binary classification divides the overall ROI distribution into two parts, it ignores other distributional properties. The algorithm does not tell lenders which loans are attractive because not all loans in the predicted ″charged-off″ (or ″fully-paid″) group are equally important for portfolio construction. An investor who solely relies on binary classification for portfolio construction overlooks additional useful information from the loan return distribution.

Riskier loans are assigned higher interest rates by LendingClub. As such, it is prudent for an investor to consider a trade-off between higher expected payments and additional risk. Existing work that considers such a risk-return trade-off is typically based on the mean-variance framework (Markowitz, 1952), a constrained optimization problem that seeks to minimize portfolio variance under the condition that the portfolio expected return exceeds an acceptable baseline value. Under this framework, expected returns and the covariance matrix are required inputs into the portfolio optimization problem. Precise estimates of these quantities are central to obtaining sensible portfolio weights. However, due to the heterogeneity and short history of loans, precise estimation becomes a grueling challenge. As mean-variance optimization suffers from the well-known issue that the portfolio weights are sensitive to the quality of the inputs (Michaud, 1989; DeMiguel et al., 2009), loan portfolios based on poorly-estimated expected returns and covariances may not be expected to expand the investor's opportunity set.

The number of unique loans also makes it prohibitive to estimate the covariance matrix. LendingClub typically lists more than 500 new loans daily. At any given point in time, there are tens of thousands of available loans to choose from. Forming a portfolio of 10,000 loans requires the estimation of more than 50 million parameters. Existing works that apply mean-variance optimization to online loans sidestep the above issues by omitting covariance estimation under the assumption that loans are mutually independent (Guo et al., 2016, 2021; Byanjankar et al., 2021; Chi et al., 2019). This is a rather strong assumption with limited theoretical justification.[12]

We propose a novel portfolio optimization framework called general characteristics-based portfolio policies, which does not require direct estimation of expected returns or the covariance matrix. Instead, GCPP directly models portfolio weights as a flexible function of the underlying

[12] For interested readers, we compare our proposed portfolio strategy with the mean-variance approach in Appendix D.1.

asset's characteristics. The portfolio is optimized through maximizing the average utility of a representative investor. GCPP shares a similar motivation with parametric portfolio policies (PPP) of Brandt et al. (2009), while generalizing the linearity-based PPP to allow portfolio weights to have arbitrarily complex nonlinear dependence on asset characteristics.

3.1 General Characteristics-Based Portfolio Policies

Assume there are N available loans at a given instance. Each loan has an ROI r_i and is associated with a vector of K loan and borrower characteristics, $x_i \in \mathbb{R}^K$, $i = 1, \ldots N$. The values of characteristics may drift over time, and standardization ensures that the features have stationary first and second moments. For example, the average loan *interest rate* varies between 11% and 14% in our sample, whereas the standard deviation hovers around 4% with some jumps to 5%. We denote by $\hat{x}_i$ the component-wise standardized loan characteristics with mean zero and standard deviation one, i.e., if $x_i(k)$ and $\hat{x}_i(k)$ are the kth components of the vectors x_i and $\hat{x}_i$, then

$$\hat{x}_i(k) = \frac{x_i(k) - \frac{1}{N}\sum_{j=1}^{N} x_j(k)}{\sqrt{\frac{1}{N-1}\sum_{j=1}^{N}\left(x_j(k) - \frac{1}{N}\sum_{\ell=1}^{N} x_\ell(k)\right)^2}}, \quad k = 1, \ldots, K. \tag{5}$$

We seek to to map the loan characteristics to the goodness of a loan. To this end, we define a function $g : \mathbb{R}^K \times \Theta \to [0, \infty)$, $(\hat{x}, \theta) \mapsto g(\hat{x}, \theta)$, where $\hat{x}$ is the vector of standardized loan characteristics and $\theta \in \Theta \subseteq \mathbb{R}^\kappa$ a vector of parameters for some $\kappa \in \mathbb{N}$.[13] We discuss later in Section 3.3.1 that, after data pre-processing, each loan is associated with $K = 144$ loan characteristics in our set of data. As of now, we do not limit g to belong to a particular family of parameterized functions to illustrate the generality of our approach. Two specific choices will later be discussed in greater detail.

We build a portfolio of loans by assigning loan i the weight ω_i corresponding to the relative goodness of the loan, i.e.,

$$\omega_i = \frac{g(\hat{x}_i; \theta)}{\sum_{j=1}^{N} g(\hat{x}_j; \theta)}, \quad i = 1, \ldots, N. \tag{6}$$

Note that while the goodness of a loan depends on the characteristics of the loan alone, the portfolio

[13] $\kappa = K$ in a linear setting, whereas κ can be much larger than K in a nonlinear setting.

weight given to that loan depends on the characteristics of all available loans.

The lender's problem is to choose the value of parameters $\theta^* \in \Theta$ that maximize the expected utility of the portfolio return r_ω:

$$\sup_{\theta \in \Theta} \mathbb{E}\left[u(r_\omega)\right] = \sup_{\theta \in \Theta} \mathbb{E}\left[u\left(\sum_{i=1}^{N} \omega_i \cdot r_i\right)\right] = \sup_{\theta \in \Theta} \mathbb{E}\left[u\left(\sum_{i=1}^{N} \frac{g(\hat{x}_i;\theta)}{\sum_{j=1}^{N} g(\hat{x}_j;\theta)} \cdot r_i\right)\right]. \tag{7}$$

We consider the utility function $u : [0, \infty) \rightarrow \mathbb{R}$ as a constant relative risk aversion (CRRA) utility function:

$$u(r_\omega) = \frac{(1 + r_\omega)^{1-\gamma}}{1 - \gamma}, \tag{8}$$

where γ is the risk aversion coefficient[14].

To solve problem (7), we approximate the unknown distribution of loan returns with their empirical distribution. Suppose our data set consists of τ temporally separate investment opportunities. At each investment opportunity, there are N_t available loans with characteristics $x_{ti} \in \mathbb{R}^K$ and historical returns r_{ti}, $t = 1, \ldots, \tau$. We standardize loan characteristics at each time instance to obtain $\hat{x}_{ti}$. While the number of available loans may vary over time, the number of loan characteristics remains constant across time. Approximating the unknown distribution of loans through their empirical distribution allows us to transform (7) to

$$\sup_{\theta \in \Theta} \sum_{t=1}^{\tau} u\left(\sum_{i=1}^{N_t} \frac{g(\hat{x}_{ti};\theta)}{\sum_{j=1}^{N_t} g(\hat{x}_{tj};\theta)} \cdot r_{ti}\right). \tag{9}$$

An implicit assumption of our formulation is that loan characteristics capture all relevant aspects of the joint return distribution necessary for forming the optimal loan portfolio. This assumption allows us to avoid the difficult problem of estimating the distribution of loan returns. Instead, we obtain portfolio weights by directly maximizing a parameterized function of loan characteristics. The dimensionality of the problem is reduced from the vast number of loans to a manageable number of characteristics. As such, our framework is able to handle an immense number of investable assets - even if the number of loans were to grow by several additional orders of magnitude.

[14]We fix $\gamma = 2$ for all calculations. We also tested for different values for γ, including $\gamma \in \{1, 3, 4\}$, and found that our results remained qualitatively robust.

The general characteristics-based portfolio policies framework allows for straightforward extensions, at least two of which are of interest for the online loan setting. First, investors have different risk preferences, which can be reflected through choosing the appropriate utility function. We use the CRRA utility function, in which one can set the risk aversion parameter γ to a suitable value. Other utility functions may be used in place of CRRA utility to reflect alternative investor preferences.

Second, GCPP allows for the integration of additional characteristics, as well as provides statistical tests of whether an additional characteristic adds significant value in the presence of all other characteristics. For instance, our framework can easily accommodate the credit risk assessment of loans by including an estimated charged-off probability as an additional loan characteristic. While lenders may give up potentially attractive investment opportunities by removing all loans that appear risky, the predicted probability of charge off may nevertheless provide useful information for portfolio construction. For the brevity of the paper, we illustrate the case of incorporating the probability of being charged-off as a case study in Appendix C.

3.2 The Goodness Function

To utilize GCPP, we must choose a functional form for our goodness measure g. We explore two important cases: A linear specification in Section 3.2.1 and a nonlinear specification based on neural networks in Section 3.2.2.

3.2.1 Truncated Linear Portfolio Policy.

We explore a specification of the goodness function g as a linear function of loan characteristics, with a non-negativity constraint on portfolio weights since investors cannot take a short position in a loan. This specification corresponds to a modified version of parametric portfolio policies for stock investments of Brandt et al. (2009).

The goodness function corresponding to a truncated linear portfolio weight function $g_L : \mathbb{R}^K \times \mathbb{R}^K \to [0, \infty)$ is defined by

$$g_L(\hat{x}_i; \theta) = \max(0, \bar{\omega}_i + \frac{1}{N}\theta^\top \hat{x}_i), \;\; i \in \{1, \ldots, N\}, \tag{10}$$

where $\bar{\omega}_i$ is the portfolio weight for loan i in a benchmark portfolio. We consider using the equal-weight portfolio, $\bar{\omega}_i = 1/N$, as the benchmark.[15] The optimization objective of the truncated linear policy can be obtained by substituting (10) into (9).

Under the truncated linear policy, the number of parameters is equal to the number of loan characteristics, i.e., $\kappa = K$, and there are no further constraints on the parameters such that $\Theta = \mathbb{R}^K$. Thus, there are 144 model parameters to be estimated. The kth parameter θ_k corresponds to the weight of the kth characteristic. A large positive θ_k translates into a larger portfolio weight for a loan that scores highly on the standardized kth characteristic.

We make two observations on the truncated linear portfolio policy specification. First, the truncated linear portfolio policy embeds significant dimensionality reduction compared to the mean-variance portfolio optimization framework. The average number of loans each month in our dataset is 18,437. For the MV framework, forming a monthly portfolio of 18,000 loans requires estimating 36,000 parameters (18,000 means and 18,000 variances) even if one ignores the dependence among loans.[16] In contrast, the optimization problem for truncated linear portfolio policy merely has 144 parameters, which grows only with the number of characteristics. As the entire portfolio is optimized by choosing a comparatively small number of parameters, the linear portfolio reduces the risk of overfitting, resulting in robust performance in-sample and out-of-sample.

Second, the truncated linear portfolio policy provides a simple framework for testing feature importance. We can test the importance of individual features or a group of features by selecting elements of the feature importance vector θ^*. Like the least-squares estimator or the maximum-likelihood estimator, parameter estimates of θ^* are obtained by maximizing the objective function. Standard errors can be computed via a bootstrap experiment, discussed in more detail in Section 3.5.1. We can then perform statistical inference on the estimated feature weight vector to identify the most critical characteristics for the portfolio formation.

[15]The benchmark choice depends on specific context. Other benchmark choices can be easily accommodated by (10). For a portfolio of stocks, the benchmark is usually an equal-weight or value-weight portfolio (DeMiguel et al., 2009). A value-weight portfolio is unsuitable for online loans because there is no available "value" proxy for online loans.

[16]When dependence between loans is taken into account, one needs to estimate 162,027,000 parameters.

3.2.2 Nonlinear Portfolio Policy.

Online loans have over 100 loan and borrower characteristics, some of which exhibit high correlation with one another. The high dimensionality of the feature space and the presence of multicollinearity can cause unstable estimation of coefficients θ in a linear weight function. A nonlinear policy can explore richer interactions and more complex transformations of features, which potentially leads to better portfolio performance.

We specify a nonlinear portfolio policy as a neural network with a single hidden layer, as illustrated in Figure 3b. This nonlinear portfolio policy uses the same set of loan and borrower characteristics, so the input layer consists of $K = 144$ neurons. The hidden layer is made up of 64 hidden neurons, and the output layer has only one element, a scalar representing the goodness of the loan.

The activation function (ρ and P) in a neural network governs the type of nonlinear transformation from one layer to the next. The goodness function defined on the above neural network $g_{NN} : \mathbb{R}^K \times \Theta \to [0, \infty)$ can be written as

$$g_{NN}(\hat{x}_i; \theta) = \rho\left(W_{(1)}^\top P(W_{(0)}^\top \hat{x}_i + b_{(0)}) + b_{(1)}\right), \quad i \in \{1, \ldots, N\}, \tag{11}$$

where $W_{(0)}$, $W_{(1)}$, $b_{(0)}$, and $b_{(1)}$ are $\mathbb{R}^{144\times64}$, $\mathbb{R}^{64\times1}$, $\mathbb{R}^{64\times1}$, and $\mathbb{R}^{1\times1}$ matrices of parameters, respectively. The parameter θ of the weight function is a κ-vector obtained by flattening and concatenating the set of the neural network parameters $\{W_{(0)}, b_{(0)}, W_{(1)}, b_{(1)}\}$. In the shallow neural network under our setup, $\kappa = 9,345$ and $\Theta = \mathbb{R}^\kappa$. Recall that, as the truncated linear policy, $\hat{x}_i$ are standardized characteristics to ensure stationarity of the first and second moments of the features.

We use the sigmoid function as our activation function.[17] By the properties of the sigmoid function, the nonlinear policy allocates an arbitrarily small but nonzero weight on a loan with unattractive traits. In comparison, the truncated linear policy would place a zero weight on such a loan. In other words, the nonlinear portfolio policy invests in all available loans whereas the truncated linear policy invests in a subset of loans. In practice, minimum investment constraints apply so it

[17]The sigmoid function, $\rho : \mathbb{R} \to (0, 1)$, $\rho(y) = \frac{1}{1+e^{-y}}$, is applied to a single neuron. Let $P : \mathbb{R}^{64} \to \mathbb{R}^{64}$ denote the activation function applied on the vector of hidden neurons, then $P(y) = (\rho(y_1), \ldots, \rho(y_{64}))$ for $y \in \mathbb{R}^{64}$.

Figure 3: Neural Network Representation

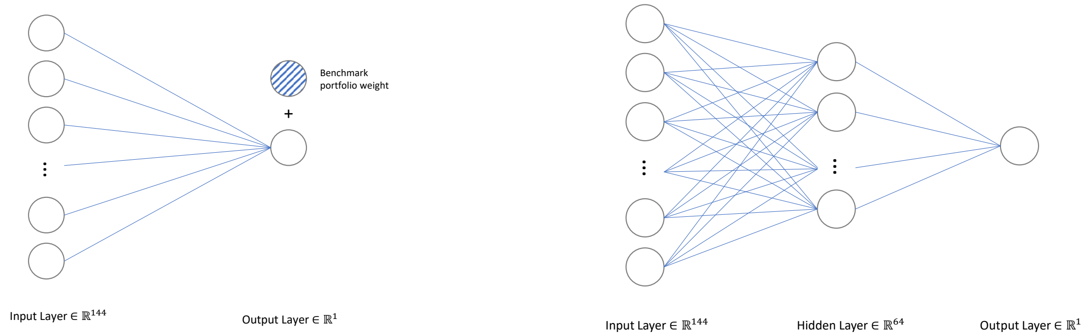

(a) Truncated Linear Portfolio Policy

(b) Nonlinear Portfolio Policy

is not feasible to allocate an arbitrarily small amount to a loan. These constraints lead to the removal of undesirable loans from the nonlinear portfolio. We offer a deeper discussion on practical constraints in Section 4.

The goodness function provides the attractiveness of each loan in relative terms, and the corresponding portfolio weights are given by a normalization of loan goodness:

$$\omega_i = \frac{g_{NN}(\hat{x}_i;\theta)}{\sum_{j=1}^{N} g_{NN}(\hat{x}_j;\theta)} = \frac{\rho\left(W_{(1)}^{\top} P(W_{(0)}^{\top}\hat{x}_i + b_{(0)}) + b_{(1)}\right)}{\sum_{j=1}^{N} \rho\left(W_{(1)}^{\top} P(W_{(0)}^{\top}\hat{x}_j + b_{(0)}) + b_{(1)}\right)} \tag{12}$$

and the optimal parameters $\theta^* = (W_{(0)}, W_{(1)}, b_{(0)}, b_{(1)}) \in \mathbb{R}^{\kappa}$ can be obtained by solving

$$\sup_{\theta\in\mathbb{R}^{\kappa}} \sum_{t=1}^{\tau} u\left(\sum_{i=1}^{N_t} \frac{\rho\left(W_{(1)}^{\top} P(W_{(0)}^{\top}\hat{x}_{ti} + b_{(0)}) + b_{(1)}\right)}{\sum_{j=1}^{N_t} \rho\left(W_{(1)}^{\top} P(W_{(0)}^{\top}\hat{x}_{tj} + b_{(0)}) + b_{(1)}\right)} \cdot r_{ti}\right). \tag{13}$$

The proliferation and success of deep learning in a wide range of applications may suggest that additional hidden layers in our neural network architecture can further improve model performance. We investigate alternative architectures including a larger number of hidden neurons and deeper network structure.[18] These alternative specifications tend to overfit the training sample and exhibit volatile learning curves on the validation sample. Our experiment suggests that shallower

[18] In particular, we tried four alternative neural network architectures: 1) neural network with a single layer of 256 hidden neurons; 2) neural network with a single layer of 128 hidden neurons; 3) neural network with two hidden layers of 128 and 64 hidden neurons; 4) neural network with two hidden layers of 64 and 32 hidden neurons.

neural networks may be more suitable in a financial setting, possibly because financial data are considerably smaller compared to other settings (e.g., computer vision), and they typically exhibit low signal-to-noise ratio. Our observation is consistent with the findings in Gu et al. (2020) that "shallow" learning outperforms "deep" learning in a stock market application.

The nonlinear portfolio policy based on neural networks generalizes the truncated linear policy. The latter is a special case of the former which directly connects model inputs with the output layer without any intervening hidden layers (Figure 3a). A fixed benchmark portfolio weight (the shaded neuron) is added to the output. The short-sale transformation in the truncated policy has the same functional form as a Rectified Linear Unit (ReLU, $\rho(y) = \max(0, y)$) activation function applied to the output neuron. Thus, we have a common optimization framework for both linear and nonlinear portfolios.

3.3 Portfolio Optimization

3.3.1 Data Preparation.

We pre-process LendingClub data before proceeding with the empirical analysis. Some features contain missing values, and we replace these missing values with zero or the maximum entry as it is applicable. Some missing entries carry significant information. For example, missing entries in *number of months since last delinquency* imply that the borrower does not have any past delinquency. In that scenario, we create a binary variable to capture the embedded information. For categorical variables, we apply one-hot encoding to transform categorical features to numerical features to allow straightforward interpretation.

Aside from the existing characteristics available from LendingClub, we define several additional features which may be related to loan outcomes. One widely used variable in the literature of credit risk assessment is *length of credit history*, which indicates how long the borrower has been engaged with credit lending (Fu et al., 2021). Salary income is an importance source of cash flow borrower use to repay loans. The monthly loan payment imparts varying degrees of pressure on borrowers of different income level. To capture this dynamic, we take the ratio of the loan payment and the monthly income as *payback pressure*. *Job title* is an elective manual input in loan applications which contains an enormous number of unique entries. We summarize this information in a binary feature, *job title report*, which indicates whether the borrower reports her job title. *Credit line utilization* shows how deeply in debt a borrower is (Agarwal et al., 2006), reflecting personal finance practices and ability to repay. After pre-processing, there are a total of 144 features associated with each loan.

3.3.2 Optimization.

We solve the optimization problem using the *Adam* algorithm (Kingma and Ba, 2015), a stochastic optimization method widely used in the field of deep learning. Intuitively, *Adam* is a combination of gradient descent with momentum and root mean square propagation. As *Adam* only requires first-order gradients, it is computationally efficient. We provide more details on the optimization algorithm in Appendix B.

We train our model using an expanding window. The initial training sample includes all loans that originate in 2013, from which a randomly-selected 10% of loans are used as a validation set. We test the performance of this model on loans originating in 2014. Then, we expand the training set to include all loans from both 2013 and 2014, and we test this updated model on loans originating in 2015. We repeat this process until the end of our sample.

Model parameters are evaluated and updated using mini-batches, i.e., random small subsets of data rather than the full training set. We set the batch size to $N = 256$.[19] Figure 4 presents our neural network architecture for the nonlinear portfolio policy with a batch of loans. Batches enable *Adam* to perform more updates in each iteration through the full training sample (i.e., one *epoch*), leading to faster convergence of the loss function. In our empirical analysis, we find that mini-batch optimization greatly improves the resulting portfolio. Compared to solving the optimization problem using the full sample, the randomness of the mini-batch helps the optimizer escape from saddle points of the objective function and lead to better portfolio performance.

After specifying the network architecture, there are two hyperparameters that require tuning: The *learning rate* and the number of training *epochs*. We tune the *learning rate* based on how

[19]We also tried different values for *batch size*, including 64, 128, and 512. The loss function all converges well, and the choice of *batch size* does not make a significant difference in convergence rate.

Figure 4: Neural Network with Batching Training

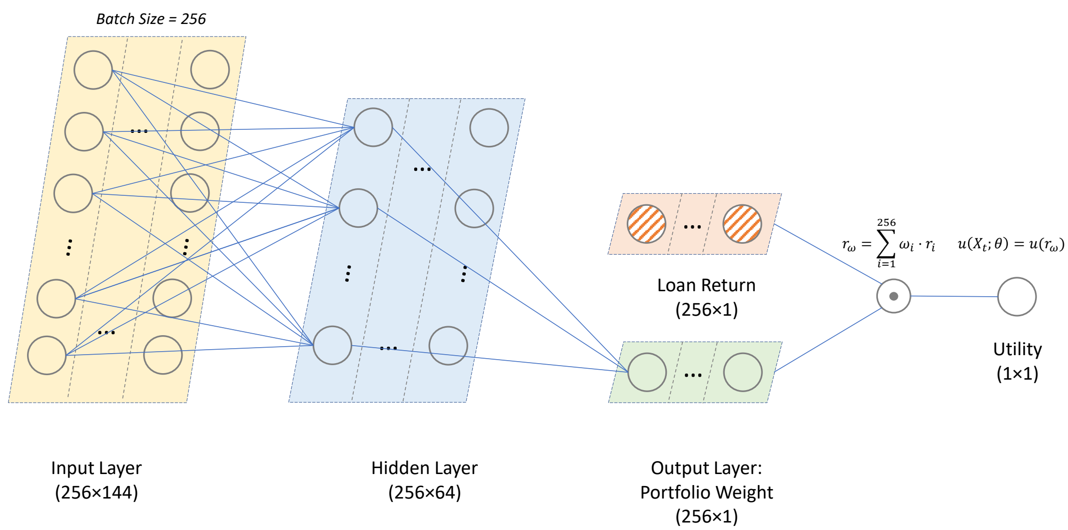

fast and how stable the objective function on the validation set changes, and we set the other hyperparameters associated with *Adam* following the suggestions from Kingma and Ba (2015), i.e., $\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-8}$. The number of *epochs* is chosen at the point which the curvature of validation learning curve starts to decrease significantly (the "elbow" of the learning curve) to avoid overfitting. Both the linear and nonlinear policies are implemented using *PyTorch*, an open source machine learning framework developed by Meta AI.

3.4 Portfolio Performance

3.4.1 Out-of-sample Portfolio Performance.

We apply the GCPP framework to LendingClub loans. Each month, we form portfolios using the linear and nonlinear portfolio policies, and compare the portfolio performance to an equal-weight strategy that uniformly invests in all available loans. Since the models used to construct loan portfolios are trained using information prior to the evaluation year, we present out-of-sample performance evaluation for the period from January 2014 to May 2017.

Figure 5 plots the internal rates of return of the loan portfolios in each month. The equal-weight portfolio exhibits monthly IRR values between 0.4% and 0.8%. The linear portfolio yields higher IRRs compared to the equal-weight portfolio every month, ranging between 0.5% and 1.0%.

Figure 5: Comparison of Portfolio Performance

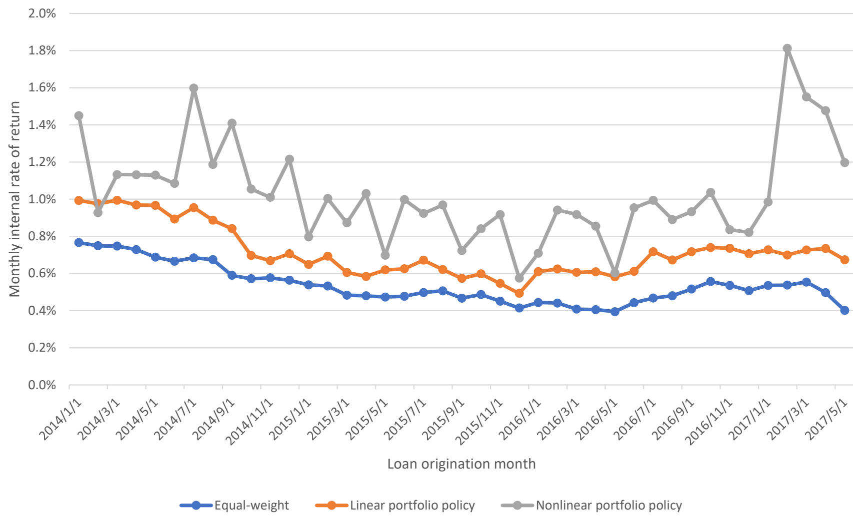

The nonlinear portfolio results in still better performance, between 0.6% and 1.8%. Note that the equal-weight portfolio of all loans has a positive IRR every month, indicating that the average loan interest rates are set sufficiently high to more than offset the average charged-off risk.

Panel A of Table 2 presents summary statistics for competing loan portfolios. Our out-of-sample period includes an average of 21,384 loans available each month. By construction, the equal-weight and nonlinear portfolios invest in all available loans. The linear portfolio invests in just more than half of the available loans given the short-sale constraint (Equation 10). The linear portfolio is associated with improved portfolio performance compared to the equal-weight portfolio, with similar levels of variation in monthly IRR as captured by standard deviation.

The maximum portfolio weight ($\max \omega_i$) shows the nonlinear portfolio taking more concentrated positions compared to the linear portfolio, implying that the nonlinear neural network is able to explore a broader portfolio weight space. As a result of this flexibility, the nonlinear portfolio yields substantially higher IRRs and utility values with higher variation from month to month (0.27% versus 0.10% and 0.14%). The last column shows the charged-off rate for each portfolio. The two general characteristics-based portfolio policies have less exposure to default risk than the equal-weight portfolio.

Table 2: Summary Statistics of Loan Portfolios

	Utility	ROI	SD - ROI	IRR	SD - IRR	# Funded	$\max \omega_i$	Charged-off rate
				A: Portfolio Strategies				
Equal weight	-0.9546	4.78%	1.40%	0.53%	0.10%	21384	0.01%	14.78%
Linear	-0.9324	7.28%	1.86%	0.71%	0.14%	10810	0.36%	14.61%
Nonlinear	-0.8938	12.01%	3.90%	1.03%	0.27%	21384	5.63%	13.13%
				B: Automated Investing				
EW - A	-0.9582	4.37%	0.61%	0.49%	0.04%	4762	0.03%	5.42%
EW - B	-0.9490	5.38%	1.25%	0.57%	0.09%	7255	0.02%	11.68%
EW - C	-0.9542	4.83%	1.59%	0.54%	0.12%	5972	0.02%	18.83%
EW - D	-0.9605	4.16%	2.29%	0.50%	0.17%	2483	0.05%	25.36%
EW - E	-0.9710	3.11%	3.71%	0.42%	0.29%	730	0.16%	31.62%
EW - F	-0.9899	1.40%	6.42%	0.24%	0.54%	154	0.82%	38.04%
EW - G	-1.0306	-1.58%	12.20%	-0.07%	1.13%	28	7.18%	43.31%

3.4.2 Portfolio Composition.

LendingClub assigns a credit grade to each loan, A1 being the most credit-worthy category and G5 being the least credit-worthy. At any point in time, loans with the same credit grade are assigned the same interest rate. Loan investors on LendingClub have the option to use its automated investing tool, where an investor can specify her investment amount and allocation across the credit grades. LendingClub will then automatically invest appropriate amounts among the selected credit grades. We explore whether optimized portfolios outperform specific credit grade segments in Figure 6, pitting the general characteristics-based portfolio policies against equal-weight loan portfolios within each credit grade. The general characteristics-based portfolio policies consistently outperform loans across credit grades. As the credit grade raises from A to G, the portfolio IRR for the equal-weight portfolio becomes more volatile.[20] More details are presented in Panel B of Table 2.

To the extent that interest rates on loans compensate lenders commensurately for the credit risk they take, one may expect a similar portfolio composition for an optimized portfolio compared to the universe of all available loans. Significant composition differences suggest differential risk-reward opportunities across credit grades. Figure 7 displays the unconditional loan grade distribu-

[20]We do not display the equal-weight portfolio for loans from grades F and G in Figure 6 because those portfolios are too volatile. The maximum (minimum) portfolio IRR can be as high (low) as 2.6% (-2.4%).

Figure 6: Equal-weight Loan Portfolio

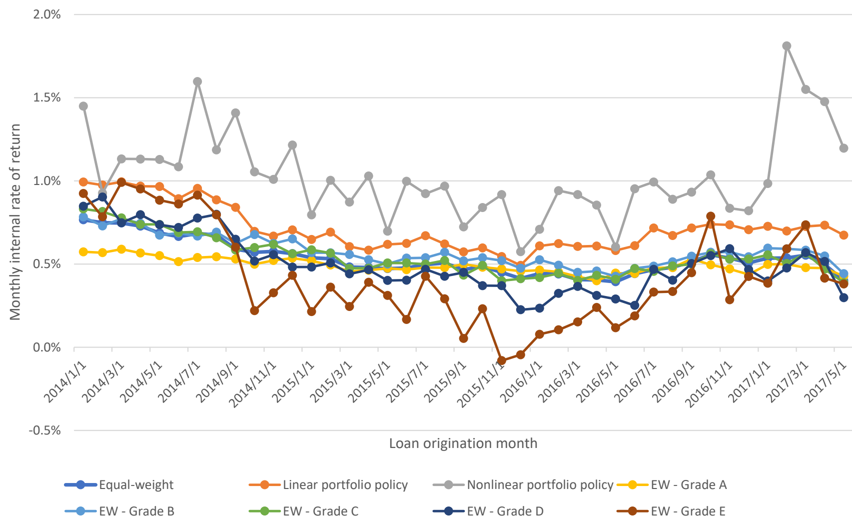

tion, and Figure 8 plots the time-varying portfolio composition. In order to maximize the lender's expected utility, both linear and nonlinear policies assign higher weights to loans with lower grades, those that LendingClub views as having higher credit risks and ascribes higher interest rates. It is possible that a relatively safe loan is designated by LendingClub as risky and receives a high interest rate, or a relatively risky loan is designated as safe and receives a low interest rate. The general characteristics-based portfolio policies exploit these mismatches to improve the investment opportunity in the online loan market. The nonlinear policy more aggressively exploits this mismatch between perceived risk and expected return by allocating larger weights to loans with grades D, E, F, and G. The superior portfolio performance of the nonlinear policy demonstrates the merit in investing in lower-grade loans. By allocating more investment to loans with lower grades, our GCPP framework promotes a more egalitarian approach to lending, allowing borrowers of limited means or less-than-stellar past records better access to credit markets.

3.5 Feature Importance

GCPP implemented in its linear or nonlinear form can improve the investment opportunity set of online loan lenders. The framework relies on capturing the joint relation among loan and borrower characteristics and portfolio weights. In this section, we examine the importance of the input variables. The identity of the most important features may reveal the underlying economic

Figure 7: Unconditional Composition of Loan Portfolios

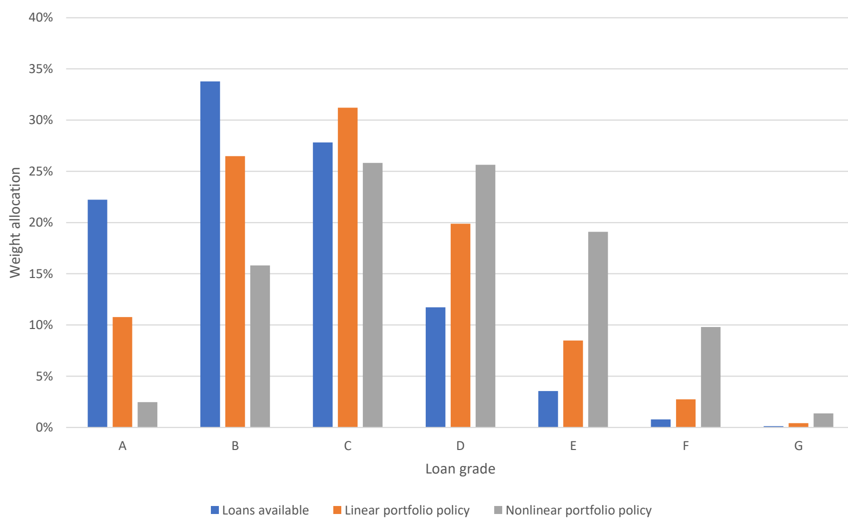

Figure 8: Time Series of Weight Allocation to Loan Grades

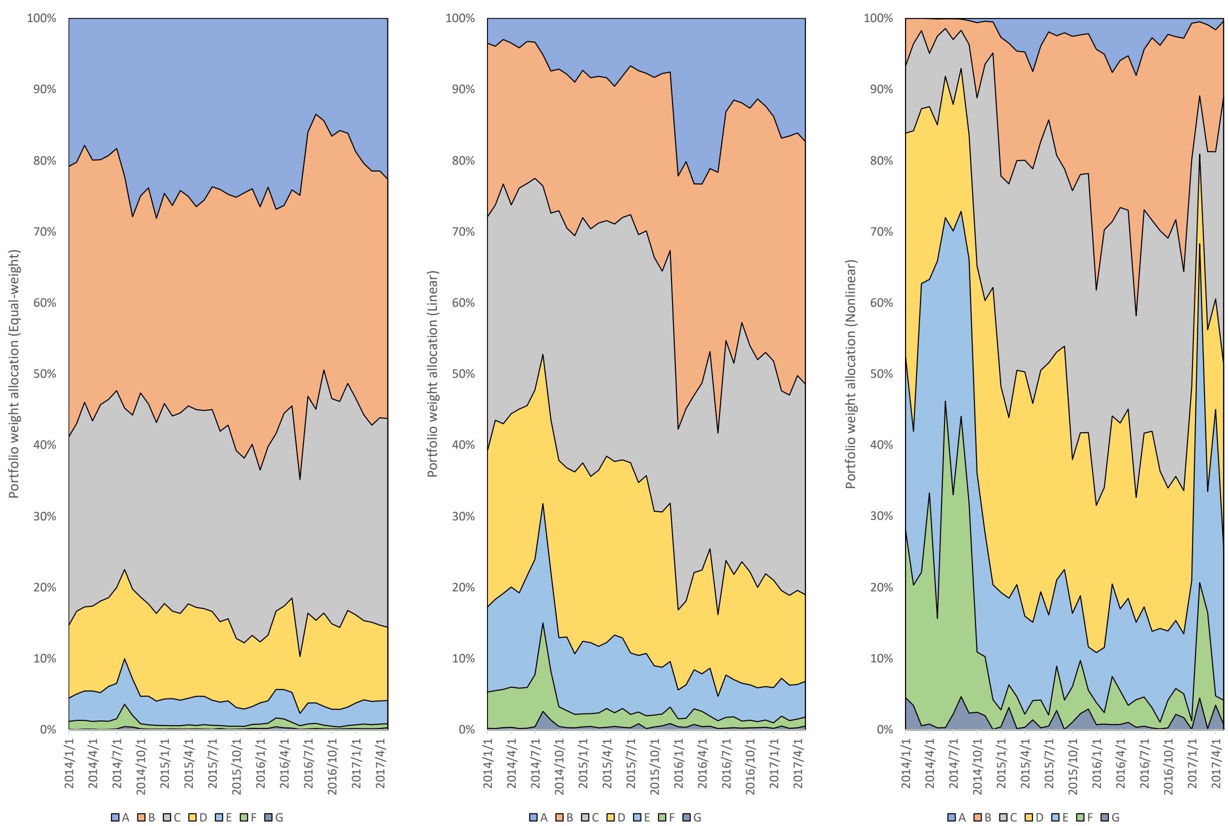

(a) Equal-Weight (b) Linear Portfolio Policy (c) Nonlinear Portfolio Policy mechanisms that determine the risk-reward trade-off in online loan markets.

The weight function of the linear policy is a linear combination of the features, making statistical inference and model interpretation particularly easy. We use a bootstrap method to compute statistical significance. While the nonlinear portfolio policy yields better portfolio performance, its interpretability is more opaque compared to the linear policy. In fact, there is active research to improve the interpretability of neural networks. We utilize a recently introduced method, permutation feature importance (Fisher et al., 2019), to study the contribution of features.

3.5.1 Bootstrap.

We evaluate the statistical significance of each feature through a bootstrap experiment. We draw with replacement from loans in the training set to construct a new sample of the same size as the original training set, and we estimate the feature weight vector θ^* on this new sample. We repeat this step 1,000 times to obtain a bootstrapped distribution of θ^*.

A large feature weight, whether positive or negative, indicates that a loan with a feature value deviating from the average would receive either an above average or below average tilt in its portfolio weight. A feature weight close to zero implies little shift in portfolio weight based on this feature. Therefore, we are interested in whether the coefficient of a feature equals zero.

Given the bootstrap distribution of the feature importance vector, we check whether zero falls in the 95% confidence interval. 60 of 144 features have coefficients significantly different from zero. Figure 9 plots the significant coefficients (red triangle) and the corresponding 95% confidence intervals (blue line), omitting significant coefficients related to geographic regions.[21]. We discuss coefficient interpretation in Section 3.5.3.

3.5.2 Permutation Feature Importance.

Permutation feature importance (Fisher et al., 2019) is a useful tool to quantify the contributions of input variables for a neural network. The calculation entails computing the change in the model's objective function after permuting a specific feature while holding other features fixed. This

[21] Of the 60 features with significant coefficients, 27 out of 51 binary variables presenting the borrowers' states of residence are significant, which suggests that LendingClub may use geographic location as a determinant in setting interest rates. We omit those variables in Figure 9.

Figure 9: Features with Significant Coefficients in Linear Portfolio Policy

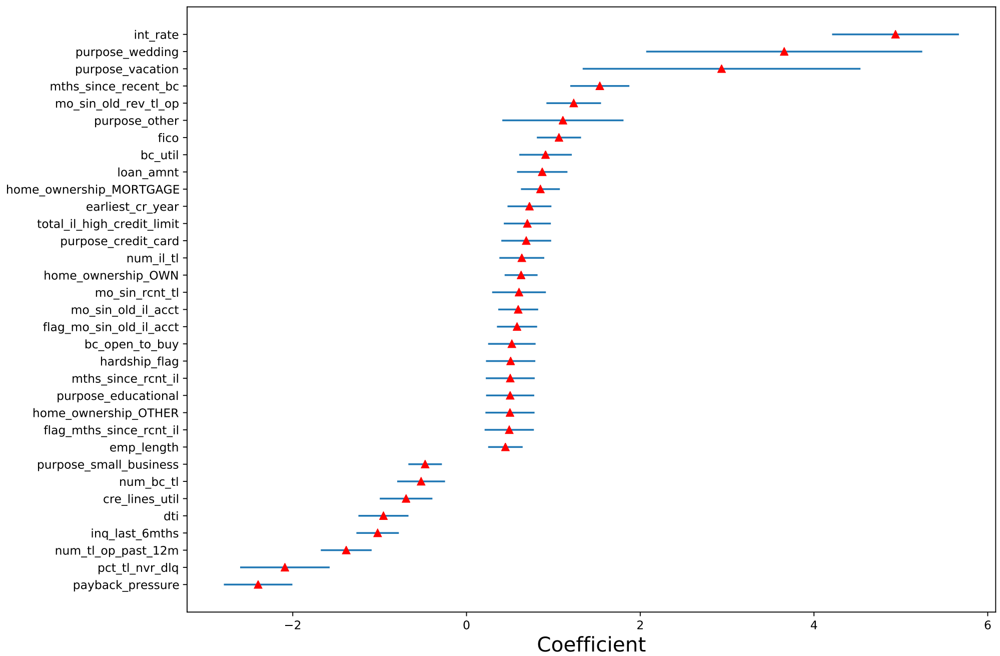

step isolates the contribution to model performance for a single feature. If model performance were similar under the original data and the permuted data, the permuted feature carries little marginal predictive power and is unlike to contribute to model prediction. Significantly different model performance under the two datasets provides evidence that original distribution of the permuted feature carries important information for improving model performance. After permutation, the scrambled distribution no longer contains useful information.

We calculate permutation feature importance as the ratio of utilities based on the permuted data $u(\hat{X}_{perm};\theta^*)$ and the original data $u(\hat{X}_{orig};\theta^*)$:

$$FI = \frac{u(\hat{X}_{perm};\theta^*)}{u(\hat{X}_{orig};\theta^*)}. \tag{14}$$

Note utilities are negative values. A feature importance greater than 1 indicates that the utility based on the original data is higher than that that of the permuted data, so the feature contributes towards model prediction. The higher the feature importance, the stronger the model relies on this feature for prediction.

We train the nonlinear portfolio policy using loans originating in 2013, and test the permuta-

Figure 10: Feature Importance for Nonlinear Portfolio Policy

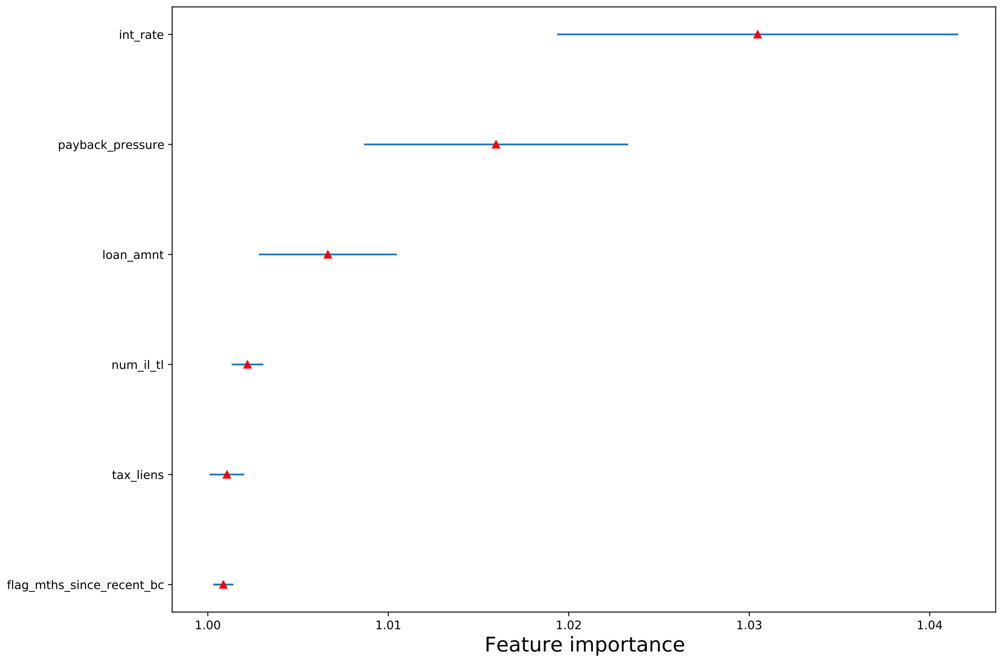

tion feature importance on loans originating in 2017.[22] We permute each of the 144 features 1000 times to get an empirical distribution of its feature importance. A one-tailed tests shows there are eight features whose scores are statistically greater than 1 at a 95% confidence level. Figure 10 plots the feature importance score and the 95% confidence interval for these features.

Unlike the linear portfolio policy case in which the sign of a coefficient implies a change in loan goodness, the permutation feature importance only informs us the importance of features but not the sign. To this end, we examine whether a feature contributes to a marginal increase or decrease in loan goodness, which will further translate to its portfolio weight. Figure 11 plots the corresponding loan goodness predicted by the neural network (red line) as the feature of interest changes from its minimum to maximum, while holding all input values constant. An increasing curve indicates the increase of the feature value results in greater goodness and, thus, higher portfolio weight assigned by the nonlinear portfolio policy. Each figure also contains the histogram of the feature of interest. We demonstrate the six significant features excluding state indicator variables, and the panels are arranged by permutation feature importance.

[22] We choose loans originating in 2017 as our test set to best mimic the practical investment scenario. Because loans usually take 36 months to maturate, it is until 2017 that lenders can observe the outcomes for loans originating in 2013.

Figure 11: Feature Effect on Portfolio Weight

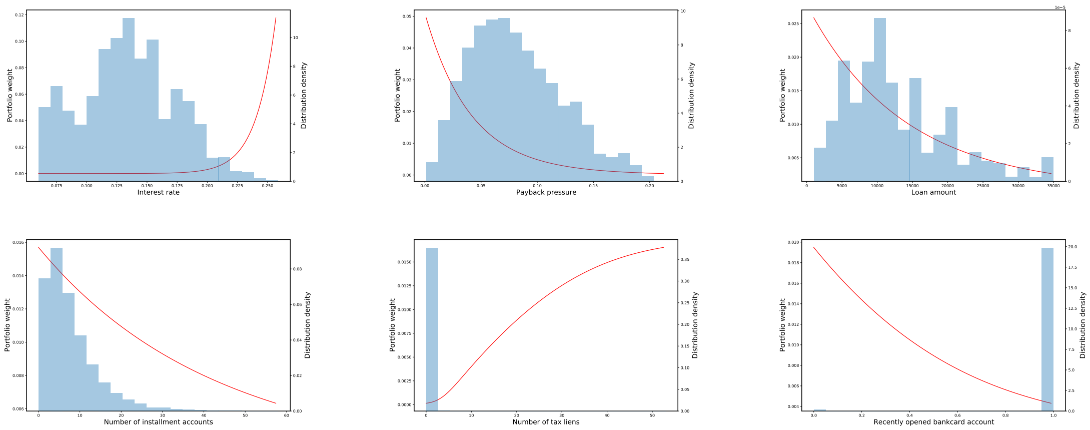

3.5.3 Discussion.

Linear and nonlinear portfolio policies share the common goal of maximizing the expected utility associated with a portfolio, but contain significant differences in the formulation of the goodness function. Given their common goal, it is perhaps not surprising to observe some agreement in the most significant features. *Interest rate* and *payback pressure* are identified as the most important features in both the linear and nonlinear setting. Other significant features shared by both specifications include *loan amount*, *number of installment accounts*, and *months since most recent bankcard account opened*. However, there exists some disagreement on feature effect. For example, linear policy assigns higher weight to loans with larger loan amount, whereas nonlinear policy assigns lower weight. Because the nonlinear portfolio policy offers more attractive opportunities for lenders, we focus our discussion on the nonlinear model.

Interest rate and *payback pressure* are the two most important loan and borrower characteristics. *Interest rate* sets a benchmark for the cash flows of a loan. If the borrower of the loan does not stop payment, the lender of the loan is expected to receive constant interest payments until the maturity date.

The credit grade assigned by LendingClub determines the interest rate of a loan: The lower the credit grade, the higher the interest rate. To maximize the lender's utility, the nonlinear portfolio policy chooses to allocate more weights to loans with higher interest rates but do not appear to be correspondingly risky given their characteristics. This approach to portfolio construction is fundamentally different from the one taken by the credit risk assessment approach. Rather than throwing out all loans that appear risky, the nonlinear policy considers the reward of these loans via the rich interactions among features, expected risk, and expected payoffs. Significantly enhanced portfolio performance of the nonlinear portfolio policy adds corroborating evidence that loans assigned to a low credit grade in fact contain promising investment opportunities if some are less risky relative to their interest rate, and loans assigned to a high credit grade may not be good investments if they carry more risk than rewarded by their interest rate.

Payback pressure is the ratio between loan installment and borrower income, which measures the pressure exerted by the loan on the borrower. The nonlinear portfolio policy prefers borrowers with low payback pressure. *Loan amount* is another important feature closely related to *payback pressure*. The nonlinear policy also prefers a low value for *loan amount*, perhaps for similar reasons as *payback pressure*. These variables reveal the ability of the borrower to repay their loans. All else equal, a borrower more able to repay her loan is less likely to default on her payments. The nonlinear portfolio assigns lower weights to loans with characteristics that indicate lower willingness to pay. Loans associated with borrowers who have poor credit records, such as having more installment accounts or tax liens, or borrowers who recently had frequent financial activities such as opening a new credit line, receive reduced weight in the portfolio.

Homeownership and *loan purposes* are two categorical variables which further reveal the willingness and ability of borrowers to service loans. Figure 12a shows that the nonlinear portfolio policy assigns higher portfolio weights to loans associated with borrowers who own their primary residence, and lower weights to those loans whose borrowers rent. Figure 12b illustrates how different categories of *loan purposes* are treated by the nonlinear policy. Values that suggest a poor record of personal finances are assigned lower weights, such as *debt consolidation*. Spending related to families are assigned higher weights, including *house*, *wedding*, *home improvement*, *medical* and *vacation*. It is also interesting to note that the nonlinear portfolio policy advocates investing more in *renewable energy*, *small business*, and *other*[23], apparently showing prioritization for loans with stated purposes related to family, small business, and renewable energy.

[23]The "other" purpose may include: paying back family and friends loan, and emergency expenses such as car breaks down, pet gets sick, or water heater needs to be replaced immediately.

Figure 12: Effects of Home Ownership and Loan Purpose

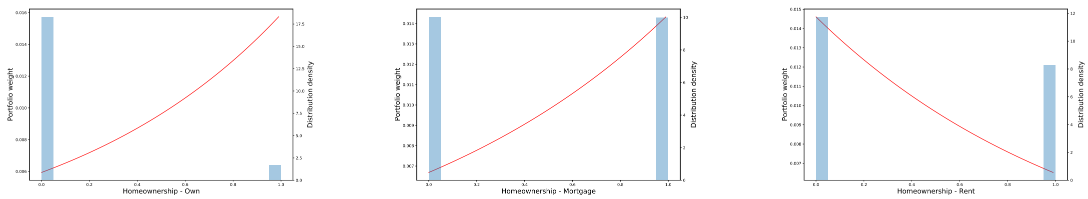

(a) Home Ownership

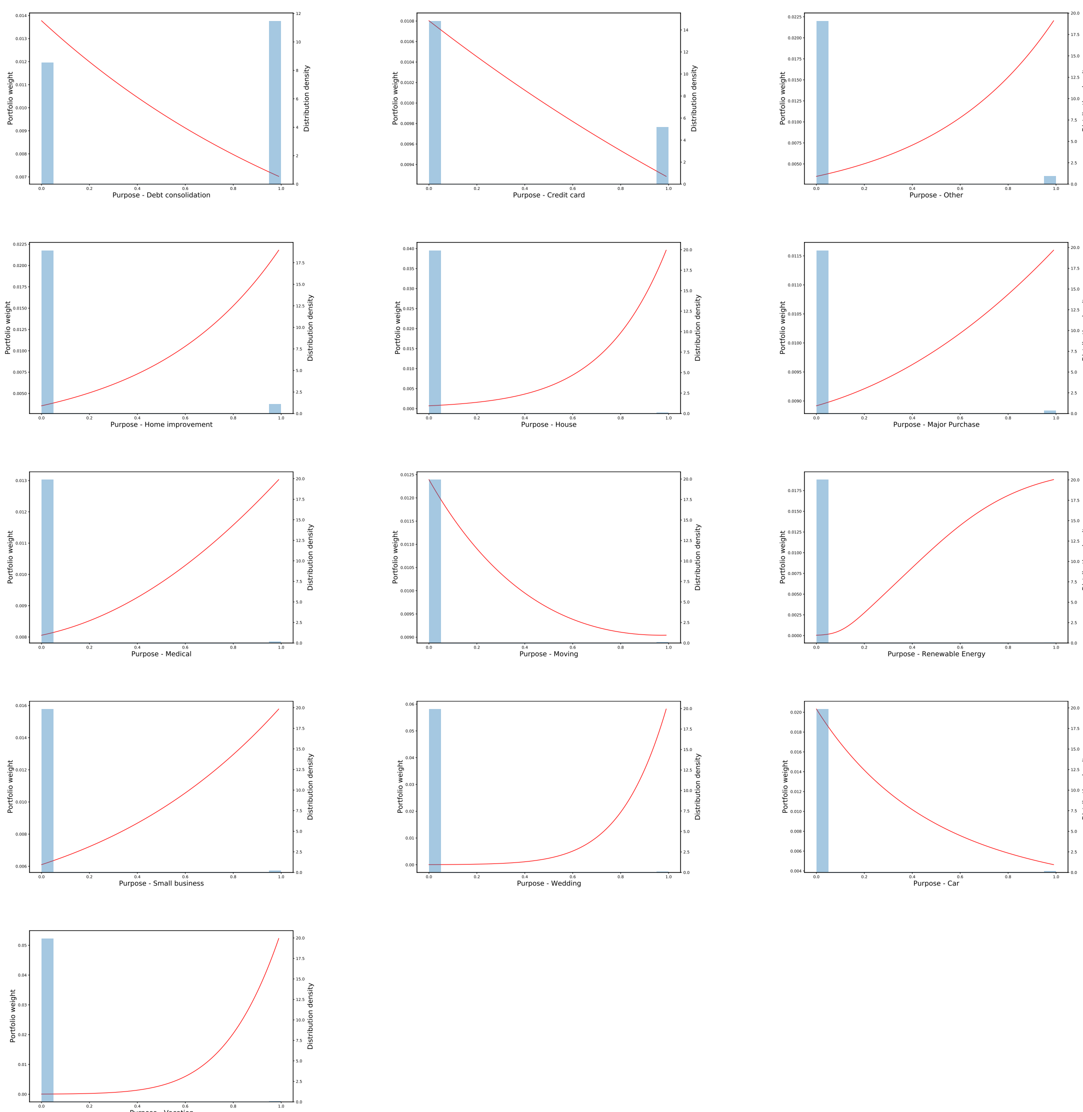

(b) Loan Purpose

4 Practical Investment Constraints

Our analysis so far has focused on idealized loan portfolios, abstracting from practical considerations faced by LendingClub users. In particular, we assumed each loan is infinitely divisible such that a lender can allocate an arbitrarily small amount to it. In reality, the minimum amount one can invest in a loan is \$25. Lenders also cannot lend more than the total loan amount, which places an upper limit on the investable capital. How can we handle such practical concerns? We propose a binary search method to construct loan portfolios under the investment-amount constraint. We then examine the performance of the constrained portfolios in relation to the unconstrained nonlinear portfolio. We then explore the impact of alternative minimum investment restrictions for lenders.

4.1 Constrained Portfolio Construction

The procedure for constructing an investment-amount constrained loan portfolio is as follows. We first train the unconstrained nonlinear portfolio to obtain the parameters θ^*. Then, we restrict the portfolio weights to accommodate the constraints placed on each loan.

Suppose an investor wants to invest $\$F$ in a portfolio of loans. We cannot simply form a nonlinear portfolio with $\$F$, since some weights may be truncated by the investment amount constraints. Instead, we form a preliminary portfolio with an amount $\$M$ with the property that after weight truncation results in the desired portfolio. Let $f : \mathbb{R} \to \mathbb{R}$ denote the function that maps the preliminary portfolio to the desired portfolio. For each loan i, the truncated weight is

$$\tilde{\omega}(M) = \begin{cases} 0 & \text{if } \omega_i * M < 25 \\ U_i/M & \text{if } \omega_i * M > U_i \\ \omega_i & \text{otherwise} \end{cases} \tag{15}$$

where $\$U_i$ is the available loan amount for loan i and ω_i is the portfolio weight given by the weight function (Equation 12). The actual investment amount after truncation is given by

$$f(M) = \left(\sum_{i=1}^{N} \tilde{\omega}_i(M) \right) \times M. \tag{16}$$

Since each weight i is truncated, the actual investment amount $f(M)$ is less than \$$M$. The function f is a monotonically increasing in M. Thus, we can use a binary search algorithm to search for the amount M such that $|f(M) - F| < 25$. We provide pseudocode for our binary search method in Algorithm 1.

Algorithm 1 Binary Search Method

Require: Total investment amount: F; Loan amount of N_t available loans: *loan amount* (a N_t-vector)
  Initialize UP = 1e20; LOW = 0; *invested amount* = 0; ▷ UP should be sufficiently high
  while $|$*invested amount* $- F| > 25$ **do**
    MID = (UP + LOW) / 2
    Get the truncated portfolio weights given the investment amount is MID: *truncated portfolio weight*; a N_t-vector
    Update *invested amount* = sum (*truncated portfolio weight*) $\cdot$ MID
    if $F >$ *invested amount* **then**
      Update UP = MID
    else if $F <$ *invested amount* **then**
      Update LOW = MID
    end if
  end while ▷ After the while loop, MID is the M we search for
  return MID

4.2 Portfolio Performance

We test the portfolio performance of the constrained nonlinear portfolio with an investable amount ranging from \$1,000 to \$100 million to investigate the effect of the investment-amount constraint on investors of various sizes. The smaller amounts serve as proxies for retail investors, whereas the larger amounts simulate the problem institutional investors face. All results presented are out-of-sample.

Figure 13 compares the portfolio IRRs for the investment-amount constrained nonlinear portfolios. For ease of exposition, the figure only includes four constrained portfolios and the unconstrained portfolio. Investment-amount constrained portfolios have similar average IRRs compared to the unconstrained nonlinear portfolio. The 1k portfolio has noticeably larger month-to-month fluctuations, whereas the 100k and 1m portfolios closely track the unconstrained portfolio. The 10m portfolio has somewhat lower IRRs. A comparison of all constrained portfolios of various sizes is shown in Table 3.

Figure 13: Investment-Amount Constrained Portfolio Performance

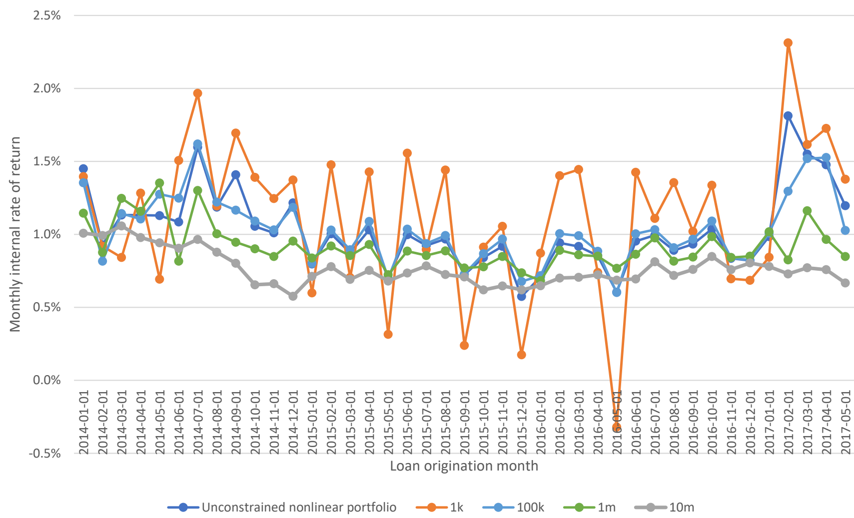

Figure 14 illustrates the relation between investment amount and portfolio IRR. Portfolios with the largest investable amounts have markedly smaller fluctuation in IRRs, and relatively lower average IRRs. Portfolios with smaller investable amounts experience larger changes in their IRRs, up to five times more volatile compared to larger portfolios. Larger variation in IRRs likely results from the lack of diversification for the smaller portfolios.

To better understand the make-up of the constrained portfolios, we follow Cremers and Petajisto (2009) and define the "sum of weight deviation" to quantify the difference in portfolio weights between a constrained portfolio (ω_i) and the benchmark unconstrained portfolio ($\omega_{i,b}$):

$$\text{Sum of Weight Deviation} = \frac{1}{2}\sum_{i=1}^{N_t} |\omega_i - \omega_{i,b}| \tag{17}$$

The sum of the weight differences is divided by two to avoid double counting. Two completely different portfolios have a sum of weight deviation of 100%, while two highly similar portfolios have a sum close to zero. Figure 15 plots the sum of weight deviation of the investment-amount constrained portfolios. The $100,000 portfolio has weights closest to the unconstrained portfolio, with a sum of weight deviation of 17%. Portfolios of larger or smaller sizes tend to deviate further from the unconstrained baseline, up to 60% on one end and more than 80% on the other end. Portfolios with small investment amounts come up against the 25-dollar minimum investment limit, and

Figure 14: Average and Standard Deviation of IRRs for Investment-Amount Constrained Portfolios

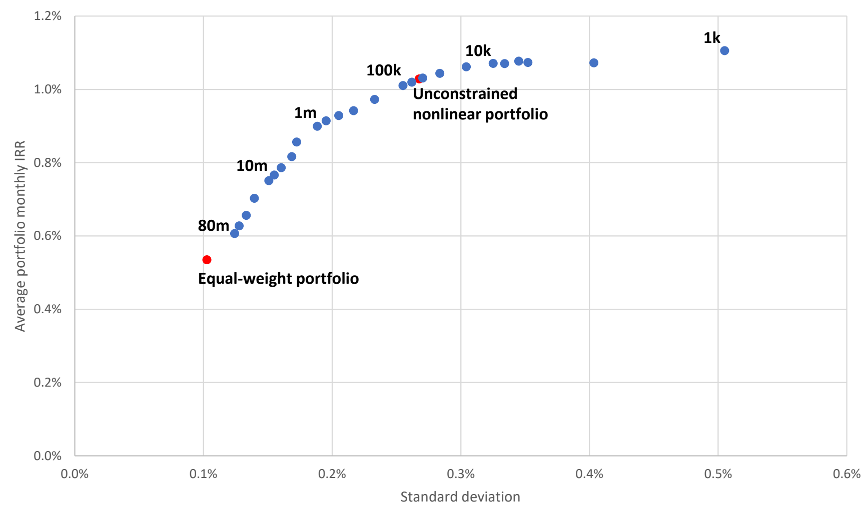

portfolios with large investment amounts are more likely to be restricted by the total loan amounts. The sum of weight deviation for the equal-weight portfolio relative to the unconstrained nonlinear portfolio is 90.89%, which illustrates the significant differences between these two strategies; the vast majority of the portfolio positions of the two portfolios do not match.

Table 3 presents summary statistics for the investment-amount constrained portfolios and the benchmark unconstrained nonlinear portfolio, as well as the equal-weight portfolio. The first few columns in Table 3 characterize the performance of the investment-amount constrained portfolios compared to the unconstrained portfolio. The average monthly IRR for the unconstrained nonlinear portfolio is 1.03% in the test period, compared to 0.53% for the equal-weight portfolio. The average portfolio IRRs for the investment-amount constrained portfolios are comparable to the benchmark, ranging between 0.60% to 1.11%. The variability of IRR, a decreasing function of portfolio size, shows larger disparity across portfolios that goes from 0.51% for the 1k portfolio to 0.10% for the 100m portfolio. Note that in the presence of a minimum investment requirement of $25, the equal-weight portfolio is no longer feasible for smaller portfolios because its allocation in some loans is too small. The equal-weight portfolio may also be constrained by the total requested loan amount. Sufficiently small loans may received a weight from the equal-weight strategy exceeding the total requested amount. In contrast, the results presented for the investment-amount constrained portfolios are feasible IRRs that investors can obtain.

Figure 15: Sum of Weight Deviation for Investment-Amount Constrained Portfolios

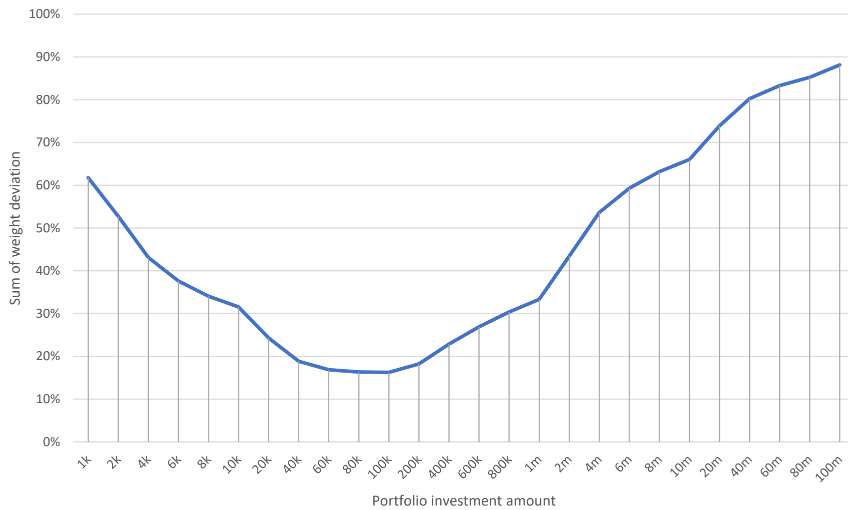

Note. The benchmark portfolio is the unconstrained nonlinear portfolio.

Table 3: Summary Statistics for Investment-Amount Constrained Portfolios

	Utility	ROI	SD - ROI	IRR	SD - IRR	# funded	$\max \omega_i$	Charged-off rate
Equal weight	-0.9546	4.78%	1.40%	0.53%	0.10%	21384	0.01%	14.78%
Benchmark[a]	-0.8938	12.01%	3.90%	1.03%	0.27%	21384	5.63%	13.13%
1k	-0.8833	13.66%	7.17%	1.12%	0.51%	19	13.86%	11.94%
10k	-0.8867	12.94%	4.46%	1.09%	0.31%	109	7.52%	12.70%
20k	-0.8875	12.82%	4.13%	1.08%	0.29%	174	6.85%	12.69%
40k	-0.8896	12.53%	3.78%	1.06%	0.26%	275	6.01%	12.95%
60k	-0.8912	12.32%	3.57%	1.05%	0.25%	358	5.38%	13.13%
80k	-0.8923	12.17%	3.45%	1.04%	0.24%	427	4.90%	13.27%
100k	-0.8933	12.04%	3.37%	1.03%	0.23%	493	4.59%	13.36%
200k	-0.8977	11.48%	3.04%	0.99%	0.20%	748	3.67%	13.83%
400k	-0.9012	11.03%	2.75%	0.96%	0.19%	1125	3.05%	14.14%
600k	-0.9029	10.81%	2.53%	0.95%	0.17%	1426	2.51%	14.07%
800k	-0.9047	10.58%	2.34%	0.93%	0.16%	1688	2.11%	14.19%
1m	-0.9065	10.35%	2.20%	0.92%	0.15%	1923	1.91%	14.37%
10m	-0.9252	8.10%	1.64%	0.77%	0.12%	7200	0.35%	15.48%
100m	-0.9455	5.78%	1.35%	0.60%	0.10%	20487	0.04%	15.47%

[a] The benchmark is the unconstrained nonlinear portfolio.

The third last column shows the number of funded loans. Although the nonlinear portfolios invest in all available loans in theory, investment-amount constraints limit the number of funded loans. The second last column shows the maximum portfolio weight for a single loan. We observe that the nonlinear portfolio takes larger positions than the equal-weight portfolio, but the positions are not extreme. The maximum portfolio weight decreases in the monthly investment amount as more loans are invested. The final column displays the charged-off rate of each portfolio. The average charged-off rate for the unconstrained nonlinear portfolio is 13.13%. The investment-amount constraint will filter out the loans with relatively small weights and truncate the loans with relatively large weights. To the extent riskier loans tend to receive smaller weights and safer loans larger weights, smaller portfolios filter out more risky loans and larger portfolios restrict investments in safer loans. Therefore, the charged-off rate is lower for smaller portfolios and higher for larger portfolios.

4.3 Minimum Investment Amount

On LendingClub, the minimum investment amount for a loan is $25. Is LendingClub optimizing the lender's investment opportunity set in setting the $25 minimum, or would lenders be better off if this minimum were set to a different amount? To answer these questions, we examine the influence of the minimum investment amount on portfolio performance. In contrast to the previous section in which we fixed the minimum investment amount to $25 and varied the portfolio size, we now fix the investable amount to be $5,000 each month, and we explore portfolio IRRs under different minimum investment amounts ranging from $0 to $200.

Figure 16 plots the average IRRs and their standard deviations for portfolios subject to different minimum investment constraints. The average portfolio IRRs do not vary much as the minimum investment amount increases from $0 to $200. However, the corresponding IRR standard deviations increase two-fold from the low end of the limit to the high end, which indicates twice as much variability in monthly IRRs for larger investment minimums. A larger minimum investment amount restricts the number of loans lenders can allocate capital to, so their portfolios become less diversified and experience more significant fluctuations from one month to the next. The increase variability in IRRs can be more easily appreciated by inspecting the ratio of expected IRR and its

Figure 16: Portfolio Performance under Varying Minimum Investment Amount

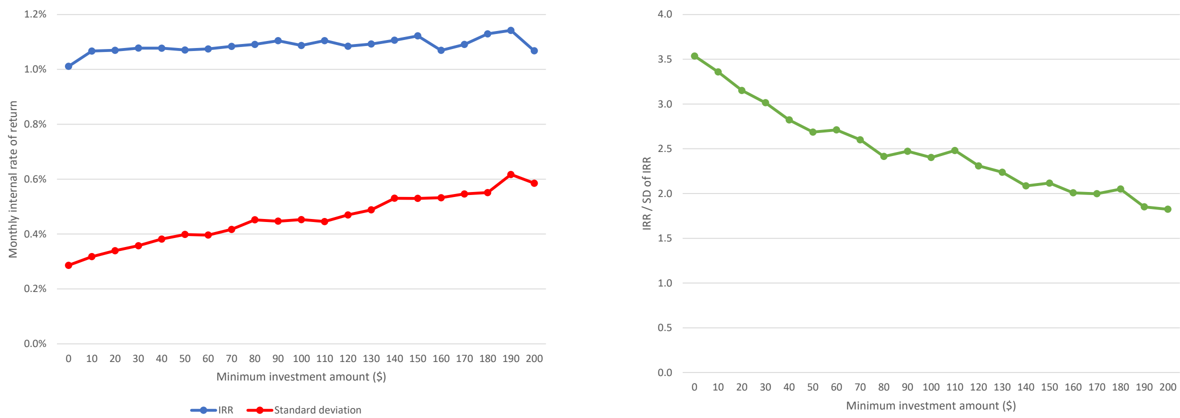

(a) IRR & Standard Deviation

(b) IRR / Standard Deviation

standard deviation (Figure 16b), which decreases from 3.5 to 2.0. A minimum investment amount of 0$ appears desirable for lenders because it allows the construction of a high-performing and reliable set of portfolios. Purely considering the lender's perspective, the most attractive minimum investment amount is $0.

Of course, the lending platform may have additional considerations in setting the minimum investment amount beyond maximizing the investment opportunity set for lenders, such as operational costs and infrastructure development costs. Without additional insight into how LendingClub sets such an amount, we hesitate to speculate on the welfare implications for LendingClub stakeholders. Our results do suggest that the minimum investment amount can have consequences for the performance of a portfolio of loans. With lower minimum investment amounts, a lender can construct more diversified portfolios that maximizes her chance of achieving good investment returns. To the extent it is possible to reduce the minimum investment amount, the platform offers improved investment opportunities and can potentially attract a greater number of participants.

5 Investment Opportunity Set

Let us step back to answer a broader question: Is it worthwhile to invest in online loans? This is an important problem investors face when trying to determine whether to add a new asset class to their existing asset allocation. The answer to this question is also related to how the recent shift in investor base from retail to institutional will impact the business model of online lending platforms.

We consider three benchmark assets: Stocks, bonds, and real estate. In the United States, stocks are the most popular asset class. Half of all U.S. households invest in the stock market, and nearly all institutional investors have equity exposure (Bhutta et al., 2020). The S&P 500 Index is perhaps the most representative of all public equity investments, and much of stock market investments are in the form of index funds or exchange-traded-funds tracking the S&P 500. Hence, we use the S&P 500 Index as the stock market benchmark. Bonds and real estate investments have cash flows reminiscent of online loans. We use their corresponding indexes as comparison benchmarks, including Bloomberg U.S. Aggregate Bond Index[24], S&P 1-3/3-5/10-20 Year U.S. Treasury Bond Indexes, and MSCI U.S. REIT (Real-Estate Investment Trust) Index. We adopt the monthly closing price for each index, and all prices are adjusted for splits and dividend and capital gain distributions.

5.1 Evaluation Metric

Portfolio IRR is not directly comparable to index returns in stock, bonds, or real estate because online loans do not have an actively-traded secondary market. Without an active secondary market, investors must hold loans to their maturity or charged-off date, and they cannot realize the calculated IRR figures before then. Online loans resemble private equity investments as private equity also cannot be easily sold prior to an exit date. Investors usually have to hold their private equity investments for several years before they can realize a return, and these investments also often impose an investment amount restriction. Given these similarities, we borrow a commonly used evaluation metric for private equity (Kaplan and Schoar, 2005), public market equivalent (PME), in our assessment of online loan portfolios. Public market equivalent is calculated as ratio of the IRR of the loan portfolio to the IRR of a benchmark portfolio:

$$PME = \frac{IRR_{\text{loan}}}{IRR_{\text{benchmark}}}. \tag{18}$$

A PME greater than 1 indicates the loan portfolio outperforms the benchmark index portfolio.

While it is difficult to construct a return measure for online loans similar to liquid markets, we

[24] Bloomberg U.S. Aggregate Bond Index tracks all investable bonds in the U.S., including corporate debt, government debt, mortgage-backed securities, and asset-backed securities.

Table 4: Online Loans Versus Stocks, Bonds, and Real Estate

Loan strategy	PME	Corr.	Pr(win)	PME	Corr.	Pr(win)
	A: S&P 500 Index			D: 3-5 Year T-Bond Index		
Equal-weight	0.93	-0.32	44%	3.38	0.49	100%
Linear	1.49	-0.12	49%	4.64	0.26	100%
Nonlinear	1.93	-0.21	66%	6.00	0.13	100%
	B: Aggregate Bond Index			E: 10-20 Year T-Bond Index		
Equal-weight	2.41	0.44	100%	1.51	0.46	100%
Linear	3.56	0.31	100%	2.14	0.58	100%
Nonlinear	4.84	0.14	100%	2.81	0.24	100%
	C: 1-3 Year T-Bond Index			F: REIT Index		
Equal-weight	8.79	0.44	100%	2.18	0.48	76%
Linear	13.47	0.26	100%	2.49	0.46	88%
Nonlinear	17.99	0.08	100%	2.82	0.18	95%

are able to compute IRRs for stocks, bonds, and real estate. We consider portfolios with identical cash flows as our loan portfolios that instead invest in other assets. Suppose an investor wants to form a $100,000 portfolio of online loans, but also has the opportunity to invest that amount in the stock market. We compare the performance of such a portfolio with one that invests the amount in the S&P 500 Index, and reinvests or sell any future cash flows also in the S&P. The loan and stock portfolios will receive the same cash flows over time but possibly have different terminal values, allowing us to calculate two directly comparable internal rates of return.

5.2 Online Loans Versus Stocks, Bonds, and Real Estate

The $100,000 loan portfolio best resembles the optimal unconstrained portfolio while satisfying practical constraints. We compare the $100,000 loan portfolio to investments in benchmarks in Table 4. In the table, winning probability, Pr(win), refers to the fraction of investments where the loan portfolio achieves a higher IRR than the benchmark portfolio. Due to the distinct nature of bonds with different duration, we include more than one benchmark for bonds.[25]

The passive, equal-weight loan strategy provides a starting point for comparison. The PME of this portfolio with respect to the S&P 500 is 0.93 (< 1), which indicates that stocks offer a higher in-

[25] We trim the outliers in calculating the average PME, because a small IRR of the benchmark portfolio in the denominator can lead to an enormously high PME. These small IRRs are due to buying high and selling low of the benchmark portfolio. PME values more than one standard deviation away from the mean are identified as outliers. We do include these outliers in the calculation of correlation and winning probability.

ternal rate of return than online loans. The portfolio IRR of loans does not tend to comove with that of the S&P 500, as indicated by a correlation of -0.32, leading to significant diversification benefits. In comparison, an equal-weight portfolio of online loans provides a higher IRR than that of bonds, treasury bonds, and real estate. Perhaps due to shared fixed-income traits, online loans are more correlated with bonds and real-estate investment, with correlations exceeding 0.40.

Our active strategies, the linear and nonlinear policies, show significant improvements in performance. The linear loan portfolio proves to have a somewhat higher IRR compared to that of the S&P 500 portfolio (PME = 1.49). The nonlinear loan portfolio displays an even higher PME of 1.93, nearly doubling the IRR of a diversified portfolio of stocks. The nonlinear policy also significantly reduces the correlation between online loans and other assets with a fixed-income component.

The above results suggest that online loans can expand the opportunity set of investors who tend to build their portfolios around traditional asset classes. Online loans offer investors high rates of return and low correlations with traditional asset classes, carrying the potential to improve overall portfolio performance with limited risk. The improved PME and reduced correlation of the nonlinear portfolio indicate the crucial role that sophisticated portfolio optimization plays in addressing whether or not a typical investor should include online loans in her investment opportunity set. The combined advantages of high returns and diversification suggest online loans offer an attractive novel asset class.

6 Conclusion

In this paper, we introduce a novel framework for portfolio construction suitable for online loans. General characteristics-based portfolio policies model portfolio weights as flexible functions of loan characteristics, bypassing the difficulty in estimating expected returns and the covariance matrix necessary for traditional portfolio optimization. Linear portfolio policies explored in previous literature is subsumed by a nonlinear policy based on neural networks. The GCPP framework can be readily extended to include other loan and borrower characteristics, such as the predictions of charged-off probability and loan return generated from machine learning algorithms.

Our results indicate that the nonlinear portfolio policy leads to superior performance compared to an equal-weight portfolio of loans, while enabling greater access to credit. The results hold if we consider practical constraints on loan portfolios imposed by minimum and maximum investment amounts. We also find that online loans earn competitive rates of return to stocks, bonds, and real estate, while offering diversification benefits.

Our findings have practical implications for online lending platforms. With the help of the GCPP framework, online loans can be made more attractive to potential lenders. Platforms may consider embedding such a framework in a robo-advising system, which would expand the access of sophisticated loan portfolios to a broad set of investors. Platforms should also be aware that the minimum investment amount in loans can constrain lenders and potentially influence how they make their investment decisions. Lowering this limit as much as operationally feasible would be in the interest of lenders.

While the GCPP framework is particularly suitable for online loans, its range of applications is potentially much broader. Investors can readily apply our framework to investments in other asset classes such as stocks, bonds, or options, all of which are associated with numerous useful attributes that help characterize the risk-return trade-off in those markets. GCPP can also be applied to asset allocation decisions across asset classes or geographic regions. We leave these possibilities for future research.

Acknowledgments

Ram Gopal gratefully acknowledges support by the Gillmore Centre for Financial Technology at the Warwick Business School. Xiao Qiao gratefully acknowledges financial support from City University of Hong Kong Project No. 7005690. Moris Strub gratefully acknowledges funding through the National Natural Science Foundation of China under Grant No. 72050410356. Zonghao Yang gratefully acknowledges financial support by InnoHK initiative, The Government of the HKSAR, and Laboratory for AI-Powered Financial Technologies.

A Loan and Borrower Characteristics

We collect 2,925,493 loans on LendingClub that originating between 2007Q1 and 2020Q3. In our dataset, each loan has 20 loan characteristics and 63 borrower characteristics. The details of the loan and borrower characteristics are documented in Table A.1 and A.2.

Table A.1: Loan Characteristics

Feature Name	Description
id	A unique LC assigned ID for the loan listing
title	The loan title provided by the borrower
grade	LC assigned loan grade
installment	The monthly payment owed by the borrower if the loan originates
int_rate	Interest Rate on the loan
issue_d	The month which the loan was funded
loan_amnt	The listed amount of the loan applied for by the borrower. If at some point in time, the credit department reduces the loan amount, then it will be reflected in this value
funded_amnt	The total amount committed to that loan at that point in time
loan_status	Current status of the loan
purpose	A category provided by the borrower for the loan request.
collection_recovery_fee	Post charge off collection fee
sub_grade	LC assigned loan subgrade
term	The number of payments on the loan. Values are in months and can be either 36 or 60
total_pymnt	Payments received to date for total amount funded
total_rec_int	Interest received to date
total_rec_late_fee	Late fees received to date
total_rec_prncp	Principal received to date
recoveries	post charge off gross recovery
last_pymnt_d	Last month payment was received
verification_status	Indicates if income was verified by LC, not verified, or if the income source was verified

B Optimization Algorithm: Adaptive Moment Estimation

We use the Adaptive moment estimation (*Adam*; Kingma and Ba, 2015) as our optimization algorithm. The objective is to maximize the expected utility function (Equation 9), which we denote as $u(\hat{X};\theta)$. In each update t, *Adam* first evaluate the gradient $g_t = \nabla_\theta u(\hat{X}_t;\theta)$, i.e., the vector of partial derivatives of $u(\hat{X}_t)$ w.r.t. θ, for a subsample X_t through backpropogation of the network. The algorithm then updates exponential moving averages of the gradient (m_t) and the squared

Table A.2: Borrower Characteristics

Feature Name	Description
emp_title	The job title supplied by the borrower when applying for the loan
emp_length	Employment length in years
home_ownership	The home ownership status provided by the borrower during registration or obtained from the credit report
annual_inc	The self-reported annual income provided by the borrower during registration
zip_code	The first 3 numbers of the zip code provided by the borrower in the loan application
addr_state	The state provided by the borrower in the loan application
dti	A ratio calculated using the borrower's total monthly debt payments on the total debt obligations, excluding mortgage and the requested LC loan, divided by the borrower's self-reported monthly income
delinq_2yrs	The number of 30+ days past-due incidences of delinquency in the borrower's credit file for the past 2 years
earliest_cr_line	The month the borrower's earliest reported credit line was opened
fico_range_low	The lower boundary range the borrower's FICO at loan origination belongs to
fico_range_high	The upper boundary range the borrower's FICO at loan origination belongs to
inq_last_6mths	The number of inquiries in past 6 months (excluding auto and mortgage inquiries)
mths_since_last_delinq	The number of months since the borrower's last delinquency
mths_since_last_record	The number of months since the last public record
open_acc	The number of open credit lines in the borrower's credit file
pub_rec	Number of derogatory public records
revol_bal	Total credit revolving balance
revol_util	Revolving line utilization rate, or the amount of credit the borrower is using relative to all available revolving credit
total_acc	The total number of credit lines currently in the borrower's credit file
collections_12_mths_ex_med	Number of collections in 12 months excluding medical collections
mths_since_last_major_derog	Months since most recent 90-day or worse rating
acc_now_delinq	The number of accounts on which the borrower is now delinquent.
tot_coll_amt	Total collection amounts ever owed
tot_cur_bal	Total current balance of all accounts
mths_since_rcnt_il	Months since most recent installment accounts opened

The table continues in the next page.

Feature Name	Description
total_rev_hi_lim	Total revolving high credit/credit limit
avg_cur_bal	Average current balance of all accounts
bc_open_to_buy	Total open to buy on revolving bankcards
bc_util	Ratio of total current balance to high credit/credit limit for all bankcard accounts
chargeoff_within_12_mths	Number of charge-offs within 12 months
delinq_amnt	The past-due amount owed for the accounts on which the borrower is now delinquent
mo_sin_old_il_acct	Months since oldest bank installment account opened
mo_sin_old_rev_tl_op	Months since oldest revolving account opened
mo_sin_rcnt_rev_tl_op	Months since most recent revolving account opened
mo_sin_rcnt_tl	Months since most recent account opened
mort_acc	Number of mortgage accounts
mths_since_recent_bc	Months since most recent bankcard account opened
mths_since_recent_bc_dlq	Months since most recent bankcard delinquency
mths_since_recent_inq	Months since most recent inquiry
mths_since_recent_revol_delinq	Months since most recent revolving delinquency
num_accts_ever_120_pd	Number of accounts ever 120 or more days past due
num_actv_bc_tl	Number of currently active bankcard accounts
num_actv_rev_tl	Number of currently active revolving trades
num_bc_sats	Number of satisfactory bankcard accounts
num_bc_tl	Number of bankcard accounts
num_il_tl	Number of installment accounts
num_op_rev_tl	Number of open revolving accounts
num_rev_accts	Number of revolving accounts
num_rev_tl_bal_gt_0	Number of revolving trades with balance greater than 0
num_sats	Number of satisfactory accounts
num_tl_120dpd_2m	Number of accounts currently 120 days past due (updated in past 2 months)
num_tl_30dpd	Number of accounts currently 30 days past due (updated in past 2 months)
num_tl_90g_dpd_24m	Number of accounts 90 or more days past due in last 24 months
num_tl_op_past_12m	Number of accounts opened in past 12 months
pct_tl_nvr_dlq	Percent of trades never delinquent
percent_bc_gt_75	Percentage of all bankcard accounts greater than 75% of limit
pub_rec_bankruptcies	Number of public record bankruptcies
tax_liens	Number of tax liens
tot_hi_cred_lim	Total high credit/credit limit
total_bal_ex_mort	Total credit balance excluding mortgage
total_bc_limit	Total bankcard high credit/credit limit
total_il_high_credit_limit	Total installment high credit/credit limit
hardship_flag	Flags whether or not the borrower is on a hardship plan

gradient (v_t):

$$m_t = \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot g_t, \tag{B.1}$$

$$v_t = \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot g_t^2, \tag{B.2}$$

where hyperparameters $\beta_1, \beta_2 \in [0, 1)$ control the exponential increase rates of these moving averages. Lastly, *Adam* updates θ by

$$\theta_t = \theta_{t-1} + \alpha \cdot \frac{m_t}{\sqrt{v_t} + \epsilon}, \tag{B.3}$$

where α is the *learning rate* and ϵ is usually small number to prevent any division by zero.

C Extension on GCPP Framework

The GCPP framework can be extended to include other loan and borrower characteristics. This section uses the predicted charged-off risk as a case study to demonstrate how the extension works, whether it improves portfolio performance, and whether the added feature is statistically significant.

C.1 Credit Risk Assessment Model

We train a credit risk assessment model with our dataset using XGBoost (Chen and Guestrin, 2016). The XGBoost classifier predicts the likelihood of individual loans (y) being charged off based on the loan and borrower characteristics (x). XGBoost is a tree ensemble model which uses multiple additive functions from the space of regression trees to make a prediction. For each loan i,

$$\hat{y}_i = \mathbb{P}(y_i = 1 | x_i). \tag{C.1}$$

The loan outcome is a binary variable, with 1 representing charged off and 0 fully paid.

We conduct data pre-processing as in Section 3.3.1. In tuning the XGBoost model, we use a grid search method to find the best combination of hyperparameters of the classifier based on the validation sample separated from the training sample. The loan dataset is highly imbalanced in the two classes; only 18.8% of loans ever charged off. Thus, we oversample the minority group until we have 50-50 proportions for the two classes in the training set so that the model has enough samples in the minority class to learn the decision boundary effectively. We apply the SMOTE method (Chawla et al., 2002), the most widely used and effective approach for oversampling. The XGBoost model predicts the charged-off probabilities of loans and requires a threshold value to transform the probabilities into a binary variable. The threshold value is another hyperparameter in the model and is optimized on the validation sample.

We use f1 score as the evaluation metric for the binary classification model, defined as a harmonic mean of the precision and the recall:

$$f1\ score = 2 \cdot \frac{precision \cdot recall}{precision + recall} \tag{C.2}$$

where precision is computed as the number of true positives divided by the total number of predictions and recall is the number of true positives divided by the total number of positive instances.

We need to point out that, by comparing the predictive performance documented in the literature, charge offs are more difficult to predict for loans from LendingClub than those from Prosper during the period it was using an auction mechanism. Our XGBoost model achieves an f1 score of 0.344 on the test set, with a threshold value of 0.15 as the cutoff for loans that are likely to be charged off. The precision and recall are 0.26 and 0.53, respectively. The corresponding area under the receiver operating characteristic curve (AUC) is 0.685. Our prediction is more accurate than similar studies which also use LendingClub loans: Guo et al. (2021) achieve an f1 score of 0.275, whereas Cho et al. (2019) obtain 0.594 for the area under the AUC curve. Although our model indicates some predictive power for charged-off rates, much of the variation in loan charged-offs remains unexplained.

We follow Fu et al. (2021) and apply an equal-weight strategy on loans whose charged-off risk is lower than 0.15 based on the prediction by XGBoost. Figure C.1 compares the out-of-sample monthly IRR of the equal-weight, equal-weight (filtered), and parametric portfolios. The equal-weight portfolios of filtered loans consistently outperforms the equal-weight portfolio of all loans. The two parametric portfolios still have dominant performance compared to the equal-weight port-

Figure C.1: Comparison of Portfolio Performance

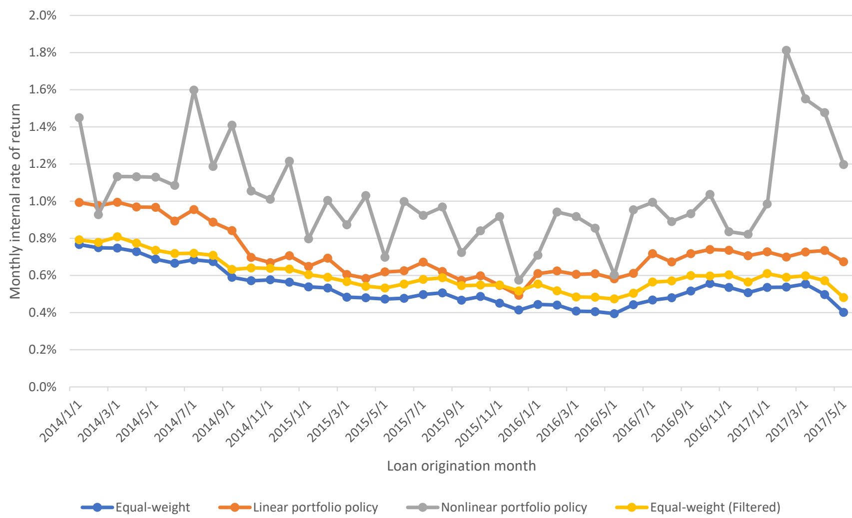

folio after filtering the risky loans.

The observation that equal-weight portfolio of filtered loans consistently outperforms the equal-weight portfolio of all loans implies that the online lending market, at least LendingClub, is not efficient. This finding indicates that adjusting for charged-off probabilities, the loans that are more likely to charge off have lower average returns compared to loans that are less likely to charge off. Consistent with Emekter et al. (2015), interest rates on riskier loans, as identified by XGBoost, do not contain a sufficiently large risk premium to compensate investors for the additional charged-off risk. In an efficient market, investors are expected to be compensated for taking on additional risk. Loans with higher likelihood of charge off may be expected to carry higher levels of interest rates compared to loans with lower likelihood. Absent any financial frictions, one may expect that observed ROI, net of charge offs, are higher for risky loans: Riskier loans have higher interest rates that just offset the greater number of charged-off events. What we instead observe is lower ROIs for risky loans and higher ROIs for safer loans.

C.2 GCPP with Charged-off Prediction

Our GCPP framework can be easily extended to incorporate the charged-off prediction by including it as an additional characteristic. Figure C.2 demonstrates the portfolio performance of the portfolios with the charged-off prediction by the trained XGBoost. The formed portfolios are all

Figure C.2: GCPP with Charged-off Prediction

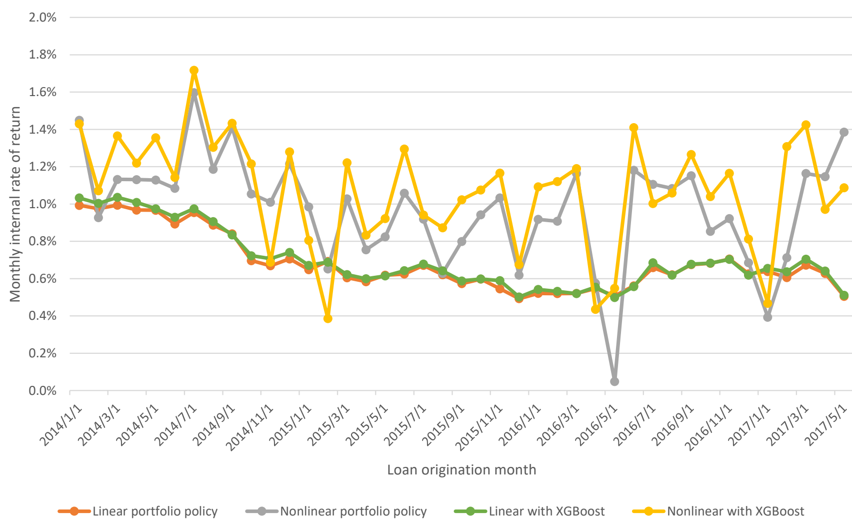

based on trained models using loans in 2013.

In both linear and nonlinear portfolio policies, the charged-off prediction lifts the portfolio performance of the origination portfolios. Following Section 3.5, we conduct bootstrap experiments to test the significance of the added new feature. The results show that the charged-off probability is statistically significant and among the top two most important features under linear and nonlinear policies.

The result implies that the XGBoost can capture information not learned by the shallow neural network. As the main strand of research in online lending, we show that credit risk assessment is beneficial under our framework in portfolio construction. The promising result leaves room for future studies to integrate other machine learning algorithms into our general characteristics-based portfolio policies framework.

D Robustness Check on Portfolio Performance

This section conducts several robustness check on the portfolio performance of our GCPP framework.

D.1 Mean-Variance Approach

The mean-variance approach is a classic portfolio optimization technique and is one of the main strands of research in portfolio construction for online loans. It is, thus, an advanced benchmark portfolio. We discussed the limitations of the mean-variance approach for online loans in Section 1.1. In this subsection, we compare the portfolio performance of our parametric portfolio policies with the mean-variance approach.

D.1.1 Portfolio Formulation.

The mean-variance approach needs to estimate the input parameters first and then solve the portfolio optimization problem. We briefly introduce the instance-based credit risk assessment framework proposed in Guo et al. (2016) and refer the interested readers to the original paper for the details.

One can estimate the expected returns and risks using the past loans with similar attributes. For any given loan, its expected return (ROI) and risk can be written as a weighted average of the relevant quantities of past loans:

$$\hat{\mu} = \sum_{j=1}^{N} w_j r_j \qquad \hat{\sigma}^2 = \sum_{j=1}^{N} w_j (r_j - \hat{\mu})^2, \tag{D.1}$$

where w_j is the weight of loan j in predicting the outcome of the candidate loan. The weight w_j is calculated based on the similarity of the two loans, measured by the difference in default probabilities:

$$s_j = |p - p_j|, \tag{D.2}$$

where p denotes the default probability obtained from the trained XGBoost as in Appendix C.

To determine the function that maps the similarity into optimal weight, Guo et al. (2016) apply the kernel regression:

$$w_j = \frac{K(\frac{p-p_j}{h})}{\sum_{l=1}^{N} K(\frac{p-p_l}{h})}, \tag{D.3}$$

where $K(\cdot)$ is a Gaussian kernel function and h is the bandwidth which determines the proportion of local versus remote information used in the summation. h is a hyperparameter - We follow Guo et al. (2016) and select h^* using the leave-one-out cross validation.

Once we estimated the expected returns and risk for N investable loans, we obtain the minimum variance portfolio (ω_{mv}) by solving the following optimization problem:

$$\begin{aligned} \min_{\omega_{mv}} \quad & \sum_{i=1}^{N} \omega_{mv,i}^2 * \hat{\sigma}_i^2 \\ \text{subject to} \quad & \sum_{i=1}^{N} \omega_{mv,i} * \hat{\mu}_i > r^* \\ & \sum_{i=1}^{N} \omega_{mv,i} = 1 \\ & 25 \le \omega_{mv,i} * M \le U_i \quad \text{or} \quad \omega_{mv,i} = 0; \quad i \in \{1, \ldots, N\} \end{aligned} \tag{D.4}$$

In the above optimization, r^* is the expected portfolio return and M is the total investment amount. The last constraint considers the minimum investment limit of \$25 and the available loan amount (U_i) for every loan i.

In our implementation of the mean-variance approach for online loans[26], we find this approach computationally prohibitive and practically infeasible. The estimation of expected returns and risk is non-parametric, which takes a great deal of time and demands a large memory. Since investors cannot foresee loans' arrival, they can only perform this procedure after new loans are publicly available. The investment opportunities may already be missing after the estimation. [27] In sharp contrast, our GCPP framework requires much less computing power - A laptop computer can train the shallow neural network with millions of loans for less than one hour. More importantly, GCPP preempts the estimation time because parameters are estimated using historical loans. With the estimated parameters, forming a portfolio is instantaneous.

D.1.2 Portfolio Performance.

Given the limitation in computation, we compare the mean-variance portfolio with our generalized parametric portfolios using a smaller dataset sampled from our original data. In particular,

[26] We thank Yanhong Guo for the kind help in reexamining our replication code of Guo et al. (2016).

[27] We have a Linux server with 8 CPU-cores and a total memory of 32 gigabytes. This machine is not able to complete the estimation if we include all our loan samples.

we randomly sample 10,000 loans originating in 2013 as the training set, and 1,000 loans each month from January 2014 to May 2017 as the test set. Then, each month of the out-of-sample period, we form loan portfolios with 1,000 sampled loans using the mean-variance approach and our GCPP framework. Even with a smaller sample, the mean-variance still takes more than two hours to estimate and solve the optimization problem, which is impractical for investment in online loans. For a fair comparison, we retrained the neural networks for the linear and nonlinear policies using the same training set as the mean-variance approach. Because the mean-variance approach uses the default probability predicted by XGBoost, we also include it as a feature for the generalized parametric portfolios as in Appendix C.

Figure D.1: Comparison of Portfolio Performance

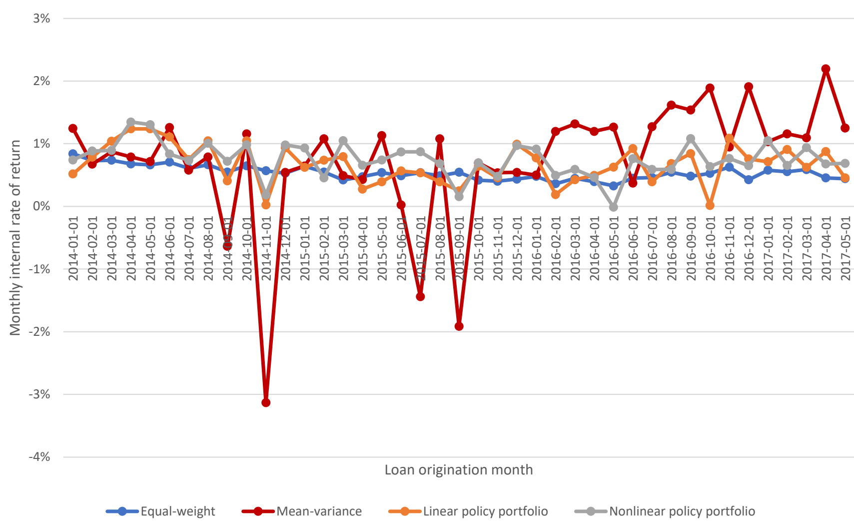

Figure D.1 plots the out-of-sample monthly IRRs of the compared portfolios with a monthly investment of $10,000. The out-of-sample R-squared of the expected returns predicted by Guo et al. (2016) is 1.56%. Consistent with Michaud (1989), given the noisy input parameters, the solution of the minimum-variance optimization problem is unstable, leading to volatile out-of-sample performance for the mean-variance portfolio. The small sample size also increases volatility for the linear and nonlinear parametric portfolios. Table D.1 demonstrates the summary statistics for the loan portfolios. The nonlinear policy portfolio has the highest utility and portfolio IRR. The average IRR of the mean-variance portfolio is close to that of the nonlinear policy portfolio. However, the large

Table D.1: Summary Statistics of Loan Portfolios

	Utility	IRR	SD - IRR	Information ratio
Equal-weight	-0.9543	0.54%	0.11%	0.00
Mean-variance	-0.9316	0.73%	0.99%	0.19
Linear	-0.9389	0.67%	0.31%	0.48
Nonlinear	-0.9263	0.75%	0.28%	0.86

volatility of its portfolio performance leads to a significantly worsened information ratio.[28]

D.2 Relative Portfolio Performance.

Figure 5 shows a secular decline in performance for the equal-weight portfolio of all loans. Several factors may contribute to this trend. First, the charged-off rate of loans grew steadily after 2013, which may erode the realized rates of return of loan portfolios. Second, LendingClub went public in 2014 and subsequently received increased media attention. As a result, the number of loan applications rose dramatically, possibly making it more difficult for the platform to process loan issues. Third, the average interest rate set by LendingClub shows a similar decline with the equal-weight portfolio IRRs, which marks lower *ex ante* expected returns for loans.

How do general characteristics-based portfolio policies handle shifts in loan performance under changing market conditions? If GCPP portfolios dependably allocates loan weights in an optimal manner, we may expect a consistent improvement in performance of these portfolios compared to equal-weight portfolios. We define the scaled excess portfolio IRR $\hat{r}_\omega$ to measure the relative performance of a GCPP portfolio against the equal-weight portfolio of all loans:

$$\hat{r}_\omega = \frac{r_\omega - r_{\text{ew}}}{r_{\text{ew}}}, \tag{D.5}$$

where r_ω and r_{ew} are the monthly IRRs for the GCPP and the equal-weight portfolios, respectively.

Figure D.2 presents the scaled excess portfolio ROI for the nonlinear portfolio policy. This relative measure remains positive throughout our sample period, exhibiting economically large performance improvement in the 50% to 150% range. There is no obvious time trend, suggesting that

[28] The information ratio compares the performance of one portfolio with a benchmark portfolio, i.e., an equal-weight portfolio. It is defined as the ratio of average excess IRR and the corresponding standard deviations.

Figure D.2: Relative Performance of Nonlinear Portfolios Against Equal-weight Portfolios

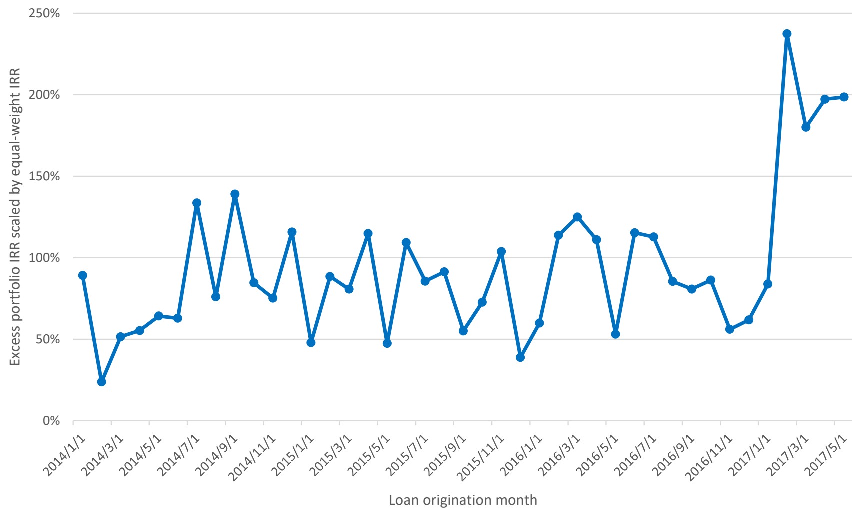

the relative outperformance of the nonlinear portfolio to the equal-weight portfolio does not deteriorate out-of-sample. We also observe steady outperformance of the linear portfolio policy relative to the equal-weight portfolio, albeit to a smaller degree. Although the average loan on LendingClub shows a decline in performance over time, the ability of the GCPP portfolios to outperform the equal-weight portfolio remains consistent.

D.3 Weekly Loan Portfolio

Throughout this paper, we form loan portfolios every month. However, lenders do not have the luxury of choosing their portfolios from all listed loans in a month. New loans will be quickly filled up and not available for lenders. Thus, we test the portfolio performance on a shorter time interval. Loans are listed on the market for two weeks until they are fully funded. Considering lenders usually fund the loans earlier than the deadline, we form loan portfolios every week.

Note that our original dataset does not have the specific issue dates for loans but only the origination month. To get the daily timestamp, we web-crawled LendingClub's public sales reports and merged the crawled data with our original dataset. The sample size for the merged dataset reduced from 1,158,476 loans to 513,307 loans. The average number of loans each week is 2,421.

In terms of the utility of portfolio ROI, which we optimize for, the nonlinear portfolio continues to demonstrate a superior performance. The average out-of-sample utility is -0.956, -0.932, and -

Figure D.3: Portfolio Performance on Weekly Loan Portfolios

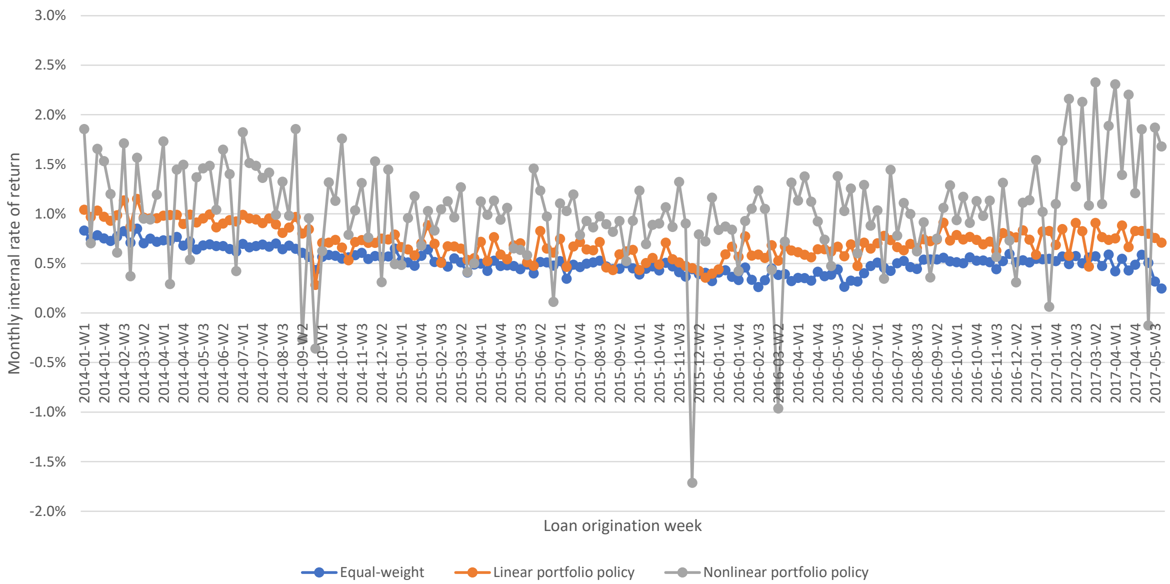

0.896, for the equal-weight, linear, and nonlinear portfolios, respectively. Figure D.3 demonstrates the out-of-sample monthly IRRs for portfolios generated weekly. The nonlinear portfolio leads to a higher expected IRR with higher volatility. It outperforms the linear portfolio 79.3% of the time. However, the nonlinear portfolio occasionally presents large drawdowns. For example, a negative monthly IRR appears five times during the test sample period. The large drawdowns are caused by aggressive weight allocations. The average maximum weight on a single loan for the nonlinear portfolio is 20%. The five weeks with a negative IRR are the extreme cases when the portfolio allocates more than 30% of the weight to a single loan, and the loan is charged off. On the other hand, a linear portfolio is much more stable and outperforms the equal-weight portfolio 95.7% of the time.

References

Agarwal, S., B. W. Ambrose, and C. Liu (2006). Credit lines and credit utilization. *Journal of Money, Credit and Banking 38*(1), 1–22.

Bhutta, N., J. Bricker, A. C. Chang, L. J. Dettling, S. Goodman, J. W. Hsu, K. B. Moore, S. Reber, A. Henriques Volz, and R. Windle (2020). Changes in us family finances from 2016 to 2019: Evidence from the survey of consumer finances. *Federal Reserve Bulletin 106*(5).

Brandt, M. W., P. Santa-Clara, and R. Valkanov (2009). Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *Review of Financial Studies 22*(9), 3411–3447.

Byanjankar, A., J. Mezei, and M. Heikkilä (2021). Data-driven optimization of Peer-to-Peer lending portfolios based on the expected value framework. *Intelligent Systems in Accounting, Finance and Management 28*(2), 119–129.

Chawla, N. V., K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research 16*, 321–357.

Chen, T. and C. Guestrin (2016). XGBoost: A scalable tree boosting system. In B. Krishnapuram, M. Shah, A. Smola, C. Aggarwal, D. Shen, and R. Rastogi (Eds.), *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794. ACM, New York, NY.

Chi, G., S. Ding, and X. Peng (2019). Data-driven robust credit portfolio optimization for investment decisions in P2P lending. *Mathematical Problems in Engineering 2019*.

Cho, P., W. Chang, and J. W. Song (2019). Application of instance-based entropy fuzzy support vector machine in Peer-to-Peer lending investment decision. *IEEE Access 7*, 16925–16939.

Cremers, K. M. and A. Petajisto (2009). How active is your fund manager? A new measure that predicts performance. *Review of Financial Studies 22*(9), 3329–3365.

DeMiguel, V., L. Garlappi, and R. Uppal (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies 22*(5), 1915–1953.

Drummer, D., S. Feuerriegel, and D. Neumann (2017). Crossing the next frontier: The role of ICT in driving the financialization of credit. *Journal of Information Technology 32*(3), 218–233.

Du, N., L. Li, T. Lu, and X. Lu (2020). Prosocial compliance in P2P lending: A natural field experiment. *Management Science 66*(1), 315–333.

Duarte, J., S. Siegel, and L. Young (2012). Trust and credit: The role of appearance in Peer-to-Peer lending. *Review of Financial Studies 25*(8), 2455–2484.

Emekter, R., Y. Tu, B. Jirasakuldech, and M. Lu (2015). Evaluating credit risk and loan performance in online Peer-to-Peer (P2P) lending. *Applied Economics 47*(1), 54–70.

Everett, C. R. (2015). Group membership, relationship banking and loan default risk: The case of online social lending. *Banking and Finance Review 7*(2), 15–54.

Fisher, A., C. Rudin, and F. Dominici (2019). All models are wrong, but many are useful: Learning a variable's importance by studying an entire class of prediction models simultaneously. *Journal of Machine Learning Research 20*(177), 1–81.

Fu, R., Y. Huang, and P. V. Singh (2021). Crowds, lending, machine, and bias. *Information Systems Research 32*(1), 72–92.

Gu, S., B. Kelly, and D. Xiu (2020). Empirical asset pricing via machine learning. *Review of Financial Studies 33*(5), 2223–2273.

Guo, Y., S. Jiang, H. Qiao, F. Chen, and Y. Li (2021). A new integrated similarity measure for enhancing instance-based credit assessment in P2P lending. *Expert Systems with Applications 175*, 114798.

Guo, Y., W. Zhou, C. Luo, C. Liu, and H. Xiong (2016). Instance-based credit risk assessment for investment decisions in P2P lending. *European Journal of Operational Research 249*(2), 417–426.

Herzenstein, M., R. L. Andrews, U. M. Dholakia, and E. Lyandres (2008). The democratization of personal consumer loans? Determinants of success in online Peer-to-Peer loan auctions. *Bulletin of the University of Delaware 15*(3), 274–277.

Iyer, R., A. I. Khwaja, E. F. Luttmer, and K. Shue (2016). Screening peers softly: Inferring the quality of small borrowers. *Management Science 62*(6), 1554–1577.

Jiang, C., Z. Wang, R. Wang, and Y. Ding (2018). Loan default prediction by combining soft information extracted from descriptive text in online Peer-to-Peer lending. *Annals of Operations Research 266*(1), 511–529.

Kaplan, S. N. and A. Schoar (2005). Private equity performance: Returns, persistence, and capital flows. *Journal of Finance 60*(4), 1791–1823.

Kate, B. and M. Nicholas (2019). Peer-to-Peer pressure: do the risks outweigh the rewards? *FT.com*.

Kingma, D. P. and J. L. Ba (2015). Adam: A method for stochastic optimization. In Y. Bengio and Y. LeCun (Eds.), *International Conference on Learning Representations (ICLR)*. San Diego, CA.

Lin, M., N. R. Prabhala, and S. Viswanathan (2013). Judging borrowers by the company they keep: Friendship networks and information asymmetry in online Peer-to-Peer lending. *Management Science 59*(1), 17–35.

Liu, D., D. J. Brass, Y. Lu, and D. Chen (2015). Friendships in online Peer-to-Peer lending. *MIS Quarterly 39*(3), 729–742.

Malekipirbazari, M. and V. Aksakalli (2015). Risk assessment in social lending via random forests. *Expert Systems with Applications 42*(10), 4621–4631.

Manzo, G. and X. Qiao (2021). Deep learning credit risk modeling. *Journal of Fixed Income 31*(2), 101–127.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance 7*(1), 77–91.

Michaud, R. O. (1989). The Markowitz optimization enigma: Is "optimized" optimal? *Financial Analysts Journal 45*(1), 31–42.

Nemoto, N., D. J. Storey, and B. Huang (2019). Optimal regulation of P2P lending for small and medium-sized enterprises.

Pope, D. G. and J. R. Sydnor (2011). What's in a picture? Evidence of discrimination from prosper.com. *Journal of Human Resources 46*(1), 53–92.

Serrano-Cinca, C. and B. Gutiérrez-Nieto (2016). The use of profit scoring as an alternative to credit scoring systems in Peer-to-Peer (P2P) lending. *Decision Support Systems 89*, 113–122.

Serrano-Cinca, C., B. Gutiérrez-Nieto, and L. López-Palacios (2015). Determinants of default in P2P lending. *PLoS One 10*(10), e0139427.

Treasury, U. (2016). Opportunities and challenges in online marketplace lending. *US Treasury, Washington, DC*.

Wang, D., Q. Wu, and W. Zhang (2019). Neural learning of online consumer credit risk. *Available at SSRN 3398981*.

Wang, Z., C. Jiang, H. Zhao, and Y. Ding (2020). Mining semantic soft factors for credit risk evaluation in Peer-to-Peer lending. *Journal of Management Information Systems 37*(1), 282–308.

Xia, Y., C. Liu, Y. Li, and N. Liu (2017). A boosted decision tree approach using bayesian hyper-parameter optimization for credit scoring. *Expert Systems with Applications 78*, 225–241.

Xu, J. J. and M. Chau (2018). Cheap talk? The impact of lender-borrower communication on Peer-to-Peer lending outcomes. *Journal of Management Information Systems 35*(1), 53–85.

Yang, Z., Y. Zhang, B. Guo, B. Y. Zhao, and Y. Dai (2018). Deepcredit: Exploiting user cickstream for loan risk prediction in P2P lending. In H. Jeff, S. Kate, and W. Ingmar (Eds.), *Proceedings of the International AAAI Conference on Web and Social Media*, Volume 12. AAAI, California, CA.