Submitted to *Information Systems Research*
manuscript ISR-2022-638.R2

Gaining a Seat at the Table: Enhancing the Attractiveness of Online Lending for Institutional Investors

Ram D. Gopal
Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom, ram.gopal@wbs.ac.uk, https://orcid.org/0000-0003-4241-9355

Xiao Qiao
School of Data Science, City University of Hong Kong, and Hong Kong Institute for Data Science, Hong Kong, xiaoqiao@cityu.edu.hk, https://orcid.org/0000-0002-1274-175X

Moris S. Strub
Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom, moris.strub@wbs.ac.uk, https://orcid.org/0000-0002-6303-6700

Zonghao Yang
School of Data Science, City University of Hong Kong, Hong Kong, zonghao.y@my.cityu.edu.hk, https://orcid.org/0000-0002-5376-9505

Although online lending enjoyed explosive growth in the past decade, its market size remains small compared to other financial assets. The risk of losing money, stringent government regulations, and low awareness of the benefits have hampered the realization of the full potential of the online lending market. Because online loans are an emerging asset class, investors may not be aware of the investment performance of online loans compared to other assets, and it remains an open question whether online loans offer sufficiently attractive returns to warrant inclusion in an asset allocation decision. To attract lenders, platforms must provide an appealing investment opportunity which entails construction of portfolios of loans that investors find attractive. We propose general characteristics-based portfolio policies (GCPP), a novel framework to overcome the difficulties associated with portfolio construction of loans. GCPP directly models the portfolio weight of a loan as a flexible function of its characteristics and does not require direct estimation of the distributional properties of loans. Using an extensive dataset spanning over one million loans from 2013 to 2020 from LendingClub, we show that GCPP portfolios can achieve an average annualized internal rate of return (IRR) of 8.86% to 13.08%, significantly outperforming an equal-weight portfolio of loans. We then address the question of whether online loans can earn competitive rates of return compared to traditional investment vehicles through six market indices covering stocks, bonds, and real estate. The results demonstrate that a portfolio of online loans earns competitive or higher rates of return compared to traditional asset classes. Furthermore, the IRRs of the loan portfolios have small correlations with the benchmark index IRRs, pointing toward significant diversification benefits. Together, we demonstrate that GCPP is an approach that can help platforms better serve both borrowers and lenders en route to growing their business.

Key words: online lending; portfolio optimization; neural networks; fintech

1. Introduction

Online lending[1] is one of the earliest and most successful manifestations of the financial technology (fintech) revolution (Goldstein et al. 2019), based on the core principle of democratization of finance and powered by technology. Participants can directly borrow and lend from one another through an intermediary platform. The loan approval process, credit risk assessment, interest rate determination, and the match of borrowers and lenders are entirely operated online. Platforms' innovative approaches in using digital technologies to enhance loan underwriting and improve consumer experience have served as important driving forces for the burgeoning industry.

Online lending increases financial inclusion, helps to reduce discrimination, and achieves a more egalitarian allocation of resources (Bartlett et al. 2022, Katsamakas and Sánchez-Cartas 2022, Agarwal et al. 2020). By assigning appropriate interest rates and providing desirable return streams, platforms can entice investor participation in a novel asset class. Due to lower overhead costs compared to traditional banks, online lending platforms can offer lower interest rates to borrowers and higher returns for lenders (U.S. Treasury 2016), thereby disintermediating banks. Platforms also promote greater participation and fairness by enabling access to credit for those individuals who otherwise cannot borrow. Market segments that have been under-served by traditional lenders are allowed access to funds that empower economic growth. Online lending also played an important role in the recovery from the 2008 financial crisis (Gopal and Schnabl 2022). The rise, in turn, prompts banks to expand credit access for consumers who obtain online loans (Balyuk 2022).

In recent years, the online lending market has experienced significant growth around the globe. In the U.S., the new lending of online consumer loans is estimated to be $38 billion in 2020 with an annual growth rate of 11.9% (Berg et al. 2022). The U.K. market experienced a 20.7% annualized growth rate from 2017 to 2022, reaching $335.8 million.[2] In China, the market has undergone explosive growth since 2012. In 2014, the trading volume in online lending increased at an average rate of around 10% per month. At its peak in 2015, the accumulative trading volume reached over $150 billion (Huang and Wang 2021). Online lending has significantly changed the landscape of the credit market. Online lenders are major suppliers of credit to small businesses (Gopal and Schnabl 2022) and their market share of U.S. mortgage lending increased from 2% to 8% from 2010 to 2016 (Fuster et al. 2019).

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). The risk of losing money, stringent government regulations, and low awareness of the benefits of online loans have hampered the realization of the full potential of the online lending market. Online loans are unsecured debt not backed up by collateral or protected by a guarantor. If a borrower stops paying, the lender of that loan takes a loss. Platforms need to meet the lender's needs by providing appropriate interest rates to compensate the risk of default. However, some evidence suggests a contradiction between the priorities of the lenders and those of the platforms (Klein et al. 2021), suggesting that platforms have not properly catered to their investors. The rapid growth in online lending has been accompanied by regulatory changes intended to improve individual investor protection, leading to the ongoing shift in the investor base from individual lenders to institutional lenders (U.S. Treasury 2016, Nemoto et al. 2019). Because online loans are an emerging asset class, investors may not be aware of the investment performance of online loans compared to other assets, and it remains an open question whether online loans offer sufficiently attractive returns to warrant inclusion in an asset allocation decision.

To attract lenders, platforms must provide an appealing investment opportunity. An investor faced with an investment opportunity in online loans needs to determine which loans to invest in and how much capital to allocate to each loan. There are two key issues associated with building a portfolio of online loans. The first issue is that loans may be charged off, and the set of loans that eventually charged off makes up a different distribution compared to those that do not.

Beyond binary classification, the two distributions contain additional information about online loans. Much of the literature on online loans focuses on assessing the credit risk of loans, but these papers seldom explore other potentially useful distributional information. A binary classification of loans is likely not sufficient for the construction of an attractive portfolio of loans.

The second issue pertains to the inputs needed to construct a portfolio. 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). The sheer size of the number of loans makes it prohibitive to estimate the covariance matrix - LendingClub 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.

We propose general characteristics-based portfolio policies (GCPP), a novel framework to overcome the difficulties associated with portfolio construction of loans. GCPP directly models the portfolio weight of a loan as a flexible function of its characteristics and does not require direct estimation of the distributional properties of loans. The portfolio is optimized through maximizing the average utility of a representative investor. The empirical performance of GCPP hinges on the form of dependence of weights on the model inputs. We explore two particular specifications of the weight function - a linear form similar to the approach in Brandt et al. (2009) and a nonlinear form based on neural networks. Compared with linear functions, neural networks consider rich nonlinear interactions of loan characteristics and search across a broader portfolio weight space, potentially leading to better portfolio performance.

Our loan sample is collected from LendingClub, the world's largest online lending platform.[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 from 2013 to 2020. In comparison, existing studies of online loans tend to contain tens or hundreds of thousands of loans. 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 equal-weight portfolio is also offered as part of LendingClub's automated investing tool.

The models are updated using an expanding window. We initially train the models using data from 2013 - a random split of 90-10 is used to divide all loans originating in 2013. 90% is used to estimate model parameters, and 10% is reserved as the validation set used to tune the hyperparameters. We then test the model performance on loans that originate in 2014. Next, we expand the training set by one year, using loans that originate in 2013 and 2014 to train, and 2015 to test. We repeat this procedure until the end of our sample. Model performance on the test set can be considered as an out-of-sample test. We obtain a separate loan portfolio each month, using the most recent model parameter estimates, which can be compared to benchmark portfolios over the same investment horizon.

Loan portfolios formed using GCPP display significantly higher performance compared to the equal-weight portfolios. The equal-weight portfolios offer an average out-of-sample annual IRR of 6.55%. A linear GCPP has an average IRR of 8.86%, whereas a nonlinear GCPP leads to an average IRR of 13.08%. These results indicate that general characteristics-based portfolio policies constitute a valuable framework for constructing a portfolio of online loans that can expand and improve the investment opportunity set of lenders.

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 propose a binary search algorithm that transforms an unconstrained portfolio into constrained portfolios. 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]

The stochastic nature of loan arrivals places an additional constraint on investors. Because new loans become available intermittently, an investor who wishes to form a portfolio would be faced with loans that are almost fully funded, loans that are partially funded, and loans without any funding. The upper bound the investor can allocate to a loan is then not the total loan amount, but the total available loan amount at the time of portfolio formation. Our method readily handles these restrictions, by treating the available investable amount of a loan as the upper limit in our binary search algorithm. Therefore, our proposed loan portfolios are implementable in real time through the following steps: 1) forming the unconstrained portfolio, 2) taking stock of the investable amount in each available loans and the intended investment amount of the investor, and 3) transforming the unconstrained portfolio into an implementable portfolio through the binary search algorithm.

Having established a sophisticated investment strategy for online loans, we turn to the question of whether online loans can provide competitive rates of return compared to traditional investment vehicles. Although online lending enjoyed explosive growth in the past decade, its market size remains small compared to other financial assets. Online lending makes up approximately 0.9% of the total non-mortgage consumer credit market in the U.S. (Berg et al. 2022). The global stock market capitalization is nearly $100 trillion, and the aggregate bond market is estimated to be $119 trillion worldwide.[5] In order to attract investors to online loans, platforms have to ensure that loan portfolios offer adequate returns compared to common investment options such as stocks, bonds, and real estate. The majority of investable assets in the U.S. is invested passively in products that track major indices, diversified portfolios that track the broad market movements in different assets. We compare our loan portfolios with six indices covering a range of markets: The 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 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. 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. The IRRs calculated from the these cash-flow matched 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 online loans offer a worthwhile expansion of the investment opportunity set.

The ratio of the IRR of a loan portfolio to a benchmark portfolio, called public market equivalent (PME), provides a measure that compares the investment performance of two sets of cash flows (Kaplan and Schoar 2005). A PME greater than one indicates that the loan portfolio is more profitable than the benchmark portfolio. The PMEs of the nonlinear GCPP portfolio compared to the six indices are 1.93, 4.84, 17.99, 6.00, 2.81, and 2.82 on the out-of-sample period. These results demonstrate that a portfolio of online loans earn competitive or higher rates of return compared to traditional asset classes. In comparison, the PME of the equal-weight portfolio of loans relative to the S&P 500 is 0.93, suggesting that a naive portfolio of loans cannot outperform the stock market. The difference in average PME between the GCPP and equal-weight loan portfolios indicates the crucial role that sophisticated portfolio optimization plays in determining the attractiveness of online loans as an investment. Furthermore, the IRRs of the loan portfolios have small correlations with the benchmark index IRRs, pointing toward significant diversification benefits.

Our findings demonstrate that online loans are an attractive novel asset class that possesses high rates of return and low correlation to traditional investments. Online loans can therefore expand the investor's opportunity set. Our approach achieves superior performance because the GCPP framework is designed to capture the unique properties of online loans. An investor with heavy exposure to traditional assets can diversify her holdings by investing in online loans to enhance overall portfolio return while limiting her risk. Desirable returns and diversification benefits of online loans would encourage greater investor participation, thereby promoting the continued growth of online lending markets.

GCPP is related to parametric portfolio policies of Brandt et al. (2009) with some important differences. Methodologically, GCPP generalizes Brandt et al. (2009) along two dimensions. First, GCPP incorporates a binary search algorithm which allows the portfolio to comply with investment constraints of online loans. The minimum investment amount, total requested amount, and the currently available amount can all impose restrictions on how much an investor can allocate to a loan. These constraints are of practical relevance, as the unconstrained optimal portfolio differs substantially from the constrained optimal portfolio. Second, GCPP considers general goodness functions which can be linear or nonlinear, subsuming the linear specification in Brandt et al. (2009), thereby advancing beyond the linear factor model tradition that underpins much of the empirical asset pricing and portfolio optimization literature.

The GCPP framework identifies mismatches in the risk of a loan and its interest rate set by LendingClub, and it tends to assign higher portfolio weights to loans with lower grades, those deemed riskier by LendingClub's proprietary credit assessment. Compared to the overall composition of loan grades available on LendingClub, the GCPP portfolios invest a greater fraction in loans with grades D, E, F, and G, and a lower fraction in loans with grades A, B, and C. To the extent that high-risk borrowers find it more difficult to obtain loans with favorable terms, our framework advocates expanded access to credit for these borrowers. If GCPP were viewed from the platform's perspective, credit access should be expanded to more high-risk borrowers at lower interest rates.

Therefore, GCPP is an approach that not only offers investors more attractive opportunities, but also has the potential to improve financial inclusion for high-risk borrowers.

The remainder of the paper is organized as follows. Section 2 covers the dataset. Section 3 introduces our portfolio framework and applies it to online loans. Section 4 investigates practical constraints faced by loan investors. Section 5 examines whether online loans expand the investor's opportunity set through a comparison with other asset classes. Section 6 concludes.

1.1. Literature Review

The extant literature has focused on two main approaches to construct a portfolio of online loans. The first approach attempts to predict the credit risk of online loans using loan and borrower characteristics, and construct portfolios based on credit risk prediction.[6] Researchers have explored a variety of methods for credit risk assessment, including single 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). 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 the 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.

The second approach to loan portfolio construction applies portfolio optimization methods to online loans. Guo et al. (2016) approximate the return and risk of a current loan as weighted averages of past loans, where the weights are assigned based on the charged-off probabilities, and then construct a mean-variance (MV) optimal portfolio. Following Guo et al. (2016), researchers have developed other similarity measures (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. Indeed, solutions to MV portfolio problems can be highly unstable given the estimation error of the input parameters, see for example Michaud (1989) or DeMiguel et al. (2009). This general problem is accentuated for investment in online loans because the heterogeneity of individual loans makes it difficult to precisely estimate their expected returns and risk. 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. Data

We collect all the loans listed on LendingClub between 2007Q1 and 2020Q3.[7] It is shown in the data that all loans listed on LendingClub are fully funded by lenders.[8] During this period, LendingClub expanded rapidly, both in the number of loans and the total dollar amount funded. 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.

LendingClub offers loans that mature either 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 *interest rate* and *loan grade*, and its related borrower characteristics including *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 data collection, most loans originating after May 2017 have not yet been completed. We therefore do not include loans that originated after May 2017. Loans originating in May 2017 mature in June 2020. Hence, our sample ends in June 2020. Our dataset contains 1,158,476 loans with 83 loan and borrower characteristics. The details of the characteristics can be found in Online Appendix A.1. 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 (Balyuk 2022, Fu et al. 2021, Guo et al. 2021).

2.1. Performance Measures

LendingClub follows a posted pricing mechanism. The platform assigns a credit grade (A1, ... 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. We only consider loans in their terminal status, either "fully-paid" or "charged-off."

The cash flow of a loan consists of five possible parts: Principal, interest, late fees, sales proceeds, and recovery amount.[9] From the lender's perspective, the typical set of cash flows for a loan may look like the following: $\{-P_0, P_1, \ldots, P_T\}$, consisting of an initial cash outflow P_0 and monthly cash inflows $\{P_t, t \in \{1, \ldots, T\}\}$ 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 net of collecting fees divided by the number of months between origination and the last payment. The cash flows for a portfolio of loans can be generated by summing up the cash flows of constituent loans.

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,

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

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

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

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, $r_\omega = \sum_{i=1}^{N} \omega_i r_i$, is the weighted sum of individual loan ROIs, which allows us to compute derivatives of the objective with respect to the model parameters. In contrast, the portfolio IRR is some unknown function of which the derivatives with respect to the parameters are not readily available.

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 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 differences in the summary statistics between the completed and charged-off loans, indicating different distributions depending on the loan status.

3. Portfolio Construction for Online Loans

A potential investor in online loans is faced with a portfolio choice problem: Which loans are worth investing in, and how much to allocate to each loan?[11] Existing literature mainly rely on two

Table 1 Summary Statistics[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	
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
Length of credit history, months	195.37	197.07	185.92
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.

portfolio construction methods, credit risk filtering (Fu et al. 2021) and mean-variance optimization (Guo et al. 2016). Credit risk filtering overlooks additional useful information from the loan return distribution, and ignores commensurate compensation for taking on additional risk (See Online Appendix B for more details). Mean-variance optimization suffers from imprecise input estimates - the expected returns vector and covariance matrix, while central to sensible portfolio weights, are particularly difficult to estimate for online loans due to a large cross section and a short history. 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. This is a rather strong assumption with limited theoretical justification.

We propose general characteristics-based portfolio policies (GCPP), a novel framework to overcome the difficulties associated with portfolio construction of loans. GCPP does not require estimated expected returns, covariances, or other distributional properties as inputs but instead directly models portfolio weights through a flexible function of the underlying asset's characteristics. While GCPP shares some common features to the parametric portfolio policies (PPP) of Brandt et al. (2009), there are important conceptual and methodological differences. Brandt et al. (2009) consider investments in stocks while we consider investments in online loans. Stocks offer repeated investment opportunities whereas a given online loan is a one-off opportunity. While historical returns are readily available for stocks, historical returns are not available for online loans. Furthermore, investment amount constraints due to the stochastic arrival of loans, the minimum investment amount, and the maximum investable amount combine to impose restrictions on the investor, a consideration important to online loans but absent for stocks. While the motivation for PPP is dimensionality reduction from a large number of stocks to a small number of characteristics, we are primarily interested in developing a tailored approach to investment in online loans. For this purpose, we must consider portfolios that are characteristics-based which readily handle investment constraints.

Methodologically, GCPP generalizes PPP along two key dimensions. First, PPP considers a linear specification which presumably was motivated by the prevalence and popularity of linear factor models in the finance literature. In contrast, we consider general goodness functions. This choice is driven by the consideration that in the case of online loans, we may expect lots of nonlinearities. For example, the appropriate interest rate for a loan does not have to be a linear function of its charged-off probability. We view GCPP as our contribution to advance beyond the linear factor model tradition that underlines Brandt et al. (2009) in the portfolio optimization literature. Second, we add binary search as a component of GCPP that allows us to comply with investment constraints of central importance to online loans. Investments in online loans must consider restrictions placed on the investable dollar amount. Absent such considerations, the resulting portfolio can differ meaningfully from the target portfolio. Investment constraints do not typically present a problem for stock investments, and are not a part of PPP.

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 a vector of K loan and borrower characteristics, x_i. The values of characteristics may drift over time, and standardization ensures that the features have stationary first and second moments. 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

$$\hat{x}_i(k) = \frac{x_i(k) - \frac{1}{N}\sum_{j=1}^N x_j(k)}{\sqrt{\frac{1}{N-1}\sum_{j=1}^N \left(x_j(k) - \frac{1}{N}\sum_{\ell=1}^N x_\ell(k)\right)^2}}, \quad k = 1, \ldots, K. \tag{4}$$

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 $\kappa \in \mathbb{N}$. $\kappa = K$ in a linear setting, whereas κ can be much larger than K in a nonlinear setting. We build a portfolio of loans by assigning loan i the weight ω_i corresponding to the relative goodness of the loan:

$$\omega_i = \frac{g(\hat{x}_i;\theta)}{\sum_{j=1}^N g(\hat{x}_j;\theta)}, \quad i = 1, \ldots, N. \tag{5}$$

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

$$\sup_{\theta\in\Theta} \mathbb{E}\left[u(r_\omega)\right] = \sup_{\theta\in\Theta} \mathbb{E}\left[u\left(\sum_{i=1}^N \omega_i \cdot r_i\right)\right] = \sup_{\theta\in\Theta} \mathbb{E}\left[u\left(\sum_{i=1}^N \frac{g(\hat{x}_i;\theta)}{\sum_{j=1}^N g(\hat{x}_j;\theta)} \cdot r_i\right)\right]. \tag{6}$$

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

$$u(r_\omega) = \frac{(1 + r_\omega)^{1-\gamma}}{1-\gamma}, \tag{7}$$

where γ is the risk aversion coefficient[12].

To solve problem (6), 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 (6) to

$$\sup_{\theta \in \Theta} \sum_{t=1}^{\tau} u \left(\sum_{i=1}^{N_t} \frac{g(\hat{x}_{ti}; \theta)}{\sum_{j=1}^{N_t} g(\hat{x}_{tj}; \theta)} \cdot r_{ti} \right). \tag{8}$$

An implicit assumption of our formulation is that loan and borrower 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.

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, different investor risk preferences may be reflected through choices of the utility function and the risk aversion parameter. Second, GCPP offers the flexibility to include further variables. We illustrate the case of incorporating the probability of being charged-off as a case study in Online Appendix I.1.

3.2. The Goodness Function

3.2.1. Truncated Linear Portfolio Policy. To operationalize GCPP, we must choose a functional form for our goodness measure g. We first 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 policy is given by

$$g_L(\hat{x}_i; \theta) = \max(0, \bar{\omega}_i + \frac{1}{N} \theta^\top \hat{x}_i), \;\; i \in \{1, \ldots, N\}, \tag{9}$$

where $\bar{\omega}_i$ is the portfolio weight for loan i in a benchmark portfolio. For the loan portfolios, we use an equal-weight portfolio as the benchmark.[13] Under the truncated linear policy, the number of parameters κ is equal to the number of loan characteristics K. The kth parameter θ_k corresponds to the weight of the kth characteristic. A positive θ_k translates into a larger portfolio weight for a loan that scores highly on the standardized kth characteristic.

Two observations are in order. First, the truncated linear portfolio policy embeds significant dimensionality reduction compared to the mean-variance framework. The average number of loans each month in our dataset is roughly 18,000. Forming an MV portfolio requires estimating 36,000 parameters even if one ignores the dependence among loans. In contrast, the optimization problem for truncated linear portfolio policy has just 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 in the vector θ^*. Standard errors can be computed via a bootstrap experiment, discussed in more detail in Appendix E.1. We can then perform statistical inference on the estimated feature weight vector to identify the most critical characteristics for portfolio construction.

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 (see Online Appendix C for a schematic diagram). 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 goodness function defined on the above neural network $g_{NN}: \mathbb{R}^K \times \Theta \to [0, \infty)$ can be written as

$$g_{NN}(\hat{x}_i; \theta) = \rho\left(W_{(1)}^\top P(W_{(0)}^\top \hat{x}_i + b_{(0)}) + b_{(1)}\right), \quad i \in \{1, \ldots, N\}, \tag{10}$$

where $W_{(0)}$, $W_{(1)}$, $b_{(0)}$, and $b_{(1)}$ are $\mathbb{R}^{144\times 64}$, $\mathbb{R}^{64\times 1}$, $\mathbb{R}^{64\times 1}$, and $\mathbb{R}^{1\times 1}$ matrices of parameters, and $\theta = \{W_{(0)}, b_{(0)}, W_{(1)}, b_{(1)}\}$. As the truncated linear policy, $\hat{x}_i$ are standardized characteristics to ensure stationarity of the first and second moments of the features. The activation function ρ governs the type of nonlinear transformation from one layer to the next. We use the sigmoid function as our activation function.[14]

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:

$$\omega_i = \frac{g_{NN}(\hat{x}_i;\theta)}{\sum_{j=1}^{N} g_{NN}(\hat{x}_j;\theta)} = \frac{\rho\left(W_{(1)}^{\top} P(W_{(0)}^{\top}\hat{x}_i + b_{(0)}) + b_{(1)}\right)}{\sum_{j=1}^{N} \rho\left(W_{(1)}^{\top} P(W_{(0)}^{\top}\hat{x}_j + b_{(0)}) + b_{(1)}\right)} \tag{11}$$

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

$$\sup_{\theta\in\mathbb{R}^{\kappa}} \sum_{t=1}^{\tau} u\left(\sum_{i=1}^{N_t} \frac{\rho\left(W_{(1)}^{\top} P(W_{(0)}^{\top}\hat{x}_{ti} + b_{(0)}) + b_{(1)}\right)}{\sum_{j=1}^{N_t} \rho\left(W_{(1)}^{\top} P(W_{(0)}^{\top}\hat{x}_{tj} + b_{(0)}) + b_{(1)}\right)} \cdot r_{ti}\right). \tag{12}$$

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.[15] These alternative specifications tend to overfit the training sample and exhibit volatile learning curves on the validation sample. Our experiment suggests that shallower neural networks may be more suitable in a financial setting, 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 for financial data.

3.3. Portfolio Optimization

We pre-process LendingClub data before proceeding with the empirical analysis. Missing values are filled with zero or the maximum entry as applicable. We also create binary variables to capture any additional information missing values may imply. For example, missing entries in *number of months since last delinquency* imply that the borrower does not have any past delinquency. 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. One widely used variable in the literature of credit risk assessment is *length of credit history* (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 equals one if the borrower reports a job title and zero otherwise. *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.

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. *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 Online Appendix C. We train our model using an expanding window. The initial training sample includes all loans that originate in 2013, 90% of which is randomly selected to estimate model parameters and the remaining 10% is reserved 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, random small subsets of data rather than the full training set. We set the batch size to $N = 256$.[16] 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 fast and how stable the objective function changes when evaluated on the validation set, and we set the other hyperparameters associated with *Adam* following the suggestions from Kingma and Ba (2015).[17] 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

We evaluate the out-of-sample portfolio performance of our GCPP framework, comparing it against various benchmarks, including equal-weight, mean-variance, credit-risk filtering, and regression-based portfolios.

An equal-weight strategy uniformly invests in all available loans. Although the equal-weight strategy appears simple and naive, there are two reasons for choosing it as a benchmark. First, an equal-weight portfolio tends to outperform more sophisticated portfolio methods on out-of-sample tests because it does not contain estimation error (DeMiguel et al. 2009). This makes it a tough benchmark to overcome. Second, the automated investing tool provided by LendingClub asks investors to set a fixed dollar amount to invest in each loan, thereby offering each investor an equal-weight portfolio.

We follow the literature (Guo et al. 2016, Fu et al. 2021) to construct mean-variance and credit risk filtering strategies. We also construct regression-based portfolios based on machine learning predictions of expected returns. Specifically, we consider linear regression, gradient-boosted regression, and neural network regression with the same architecture used for the nonlinear GCPP. We use these methods to predict the loan ROI. For each method, a portfolio is constructed by investing uniformly in the top 10% of the loans based on their ROI predictions. We provide further details for the construction of these benchmark portfolios in Online Appendix H.

Figure 1 Comparison of Portfolio Performance

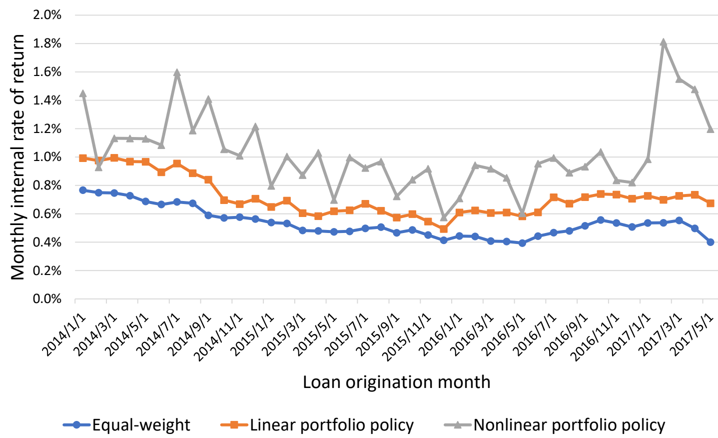

3.4.1. Out-of-Sample Portfolio Performance. We apply the GCPP framework to LendingClub loans. Each month, we form portfolios using loans originating in that month. Since the models used to construct loan portfolios are trained using information prior to the evaluation year, all of the performance evaluation presented is out-of-sample.

Figure 1 plots the internal rates of return of GCPP and the equal-weight portfolio 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%. The nonlinear portfolio further improves performance, showing IRRs between 0.6% and 1.8%. The volatility of the nonlinear portfolio policy is higher than the other methods as a result of the neural network's ability to explore a much broader portfolio weight space. 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.

Table 2 provides a comparison between GCPP and the benchmark portfolios. Our out-of-sample period includes an average of 21,384 loans available each month. By construction, the equal-weight and nonlinear GCPP portfolios invest in all available loans. The linear GCPP and the risk-filtering portfolios invest in just over half of the available loans. The linear portfolio is associated with improved portfolio performance compared to the equal-weight and risk-filtering portfolios, with higher levels of variability in portfolio ROI and monthly IRR as captured by standard deviation.

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. The investor must allocate equal amounts to each loan, resulting in an equal-weight portfolio. Panel C contains a comparison with the equal-weight portfolios of loans within specific credit grades to mimic the investor's choice set when choosing the automated investing tool. These portfolios are all outperformed by the GCPP portfolios. We perform one-sided paired sample t-tests to compare the portfolio returns of the nonlinear portfolio against alternative portfolios. For ROI, utility, and IRR, the outperformance of the nonlinear portfolio is statistically significant at the 1% level.

The gradient-boosted regression (GB) portfolio outperforms the other regression-based portfolios and has the best forecast on loan ROIs with an out-of-sample R-squared of 4.9%. Notably, the regression-based portfolios, including GB regression and the neural network, can outperform the linear portfolio policy while capturing a similar level of risk based on the standard deviations and charged-off rates. The nonlinear portfolio policy shows consistent outperformance in terms of utility, portfolio ROI, and monthly IRR. This result emphasizes the profitability of the nonlinear policy in online loans and highlights the importance of exploring a general goodness function our framework.

We implement the mean-variance portfolio construction method described by Guo et al. (2016) by estimating the expected returns and risks using past loans with similar attributes. However, we find this approach to be extremely computationally intensive for online loans. To address this issue, we compare the mean-variance portfolio to our GCPP framework using a smaller dataset about the same size as Guo et al. (2016) (more details in Appendix H.2). The mean-variance approach yields unstable performance with a standard deviation in monthly IRR of 1.25% compared to 0.35% and 0.27% for the linear and nonlinear policies, respectively. The average utility of the mean-variance portfolio is -0.9587, while the linear and nonlinear policies achieve average utilities of -0.9424 and -0.9260.

The above results indicates that the linear and nonlinear GCPP approaches can add value for investors by improving performance. The nonlinear portfolio policy is able to explore a broader

Table 2 Summary Statistics of Loan Portfolios

	Utility	ROI	SD - ROI	IRR	SD - IRR	# Funded	$\max \omega_i$	Charged-off rate
A: General Characteristics-Based Portfolio Policy								
Linear	-0.9324	7.28%	1.86%	0.71%	0.14%	10810	0.357%	14.61%
Nonlinear	-0.8938	12.01%	3.90%	1.03%	0.27%	21384	5.635%	13.13%
B: Benchmark Strategies								
Equal weight	-0.9546	4.78%	1.40%	0.53%	0.10%	21384	0.005%	14.78%
Risk filtering	-0.9443	5.91%	1.06%	0.62%	0.08%	11723	0.010%	8.73%
Linear regression	-0.9539	4.86%	1.51%	0.54%	0.11%	2139	0.054%	14.88%
GB regression	-0.9173	9.04%	1.91%	0.85%	0.13%	2139	0.054%	14.07%
NN regression	-0.9239	8.26%	1.92%	0.78%	0.13%	2139	0.054%	13.74%
C: Automated Investing								
EW - A	-0.9582	4.37%	0.61%	0.49%	0.04%	4762	0.025%	5.42%
EW - B	-0.9490	5.38%	1.25%	0.57%	0.09%	7255	0.017%	11.68%
EW - C	-0.9542	4.83%	1.59%	0.54%	0.12%	5972	0.020%	18.83%
EW - D	-0.9605	4.16%	2.29%	0.50%	0.17%	2483	0.047%	25.36%
EW - E	-0.9710	3.11%	3.71%	0.42%	0.29%	730	0.155%	31.62%
EW - F	-0.9899	1.40%	6.42%	0.24%	0.54%	154	0.819%	38.04%
EW - G	-1.0306	-1.58%	12.20%	-0.07%	1.13%	28	7.182%	43.31%

weight space compared to the linear policy. To maximize expected utility, the nonlinear policy takes more extreme positions compared to the linear policy or the equal-weight portfolio, leading to more concentrated positions. The maximum portfolio weight, $\max \omega_i$, shows that the nonlinear portfolio can take weights several times larger than the linear portfolio. The results show that the nonlinear policy rightly identifies 'good' investment opportunities, and produces superior excess returns. As a result of a lack of diversification, the nonlinear portfolio yields higher variability from month to month (0.27% versus 0.14% for the linear policy). The last column of the table 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. The risk-filtering approach and an equal weight on loans with grade A have the lowest charged-off risk.

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 a given time, loans with the same credit grade are assigned the same interest rate. 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

Figure 2 **Average Composition of Loan Portfolios**

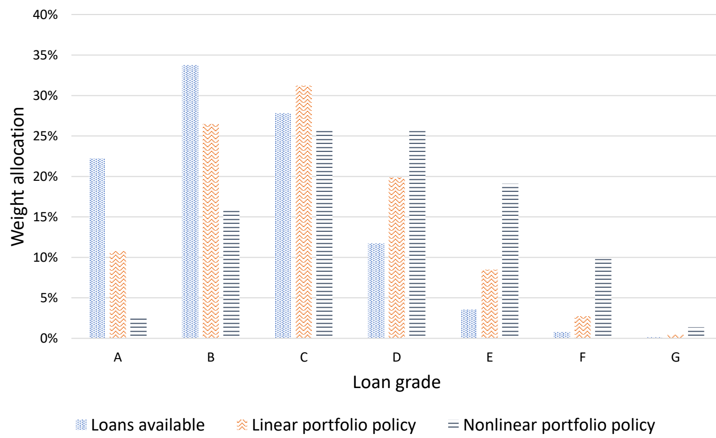

composition differences suggest differential risk-reward opportunities across credit grades. Figure 2 displays the average loan grade distribution (see Online Appendix D for time-varying portfolio composition). In order to maximize the lender's expected utility, both linear and nonlinear policies assign higher weights to loans with lower grades, those that LendingClub views as having higher credit risks and ascribes higher interest rates. Note that the linear and nonlinear portfolio policies are implementable even when the weight assigned to a loan grade is larger than the percentage of available loans because the total portfolio size of an investor is typically much smaller than the totality of all available loans.

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 risk and expected return by allocating larger weights to loans with grades D, E, F, and G. The GCPP framework applied at the platform level, rather than at the individual investor level, would advocate expanding credit access to high-risk borrowers at lower interest rates.

3.4.3. Bias in Interest Rates. Fu et al. (2021) show that machine credit risk predictions can contain bias even when gender and race are not included as inputs. Because our GCPP framework allows us to identify loans that have high interest rates for the identified level of risk, it provides a test for whether LendingClub loans contain bias. We follow Fu et al. (2021) to obtain the gender and race labels via occupation and location (see Online Appendix F), and we find that the average interest rate for African Americans and the female borrowers to be higher than the average interest rate of all loans (12.2% and 12.6%, versus 12.0%). The disparity in interest rate is statistically significant at a 1% level. That is, LendingClub identifies these two groups as riskier than the platform's average loans. If the risk calibration of these borrower groups is appropriate, then the portfolio weight assigned by our portfolio strategy to these loans should not significantly deviate from their share in the overall set of available loans.

Figure 3 illustrates that both the linear and nonlinear GCPP assign higher weights for the African American group.[18] Our results are consistent with the interpretation that African Americans are assigned somewhat higher interest rates than non-African Americans with similar loan and borrower characteristics. However, we do not observe a weight difference for the female group, which suggests that women are given interest rates commensurate with their risk profiles.[19] Our finding is consistent with Shen et al. (2022) and related works that online lending helps close the gender gap in access to credit. Our GCPP framework can help identify potential bias against minority groups, thereby fostering fairer credit market access to all potential borrowers.

Figure 3 **Portfolio Allocation to Borrower Groups**

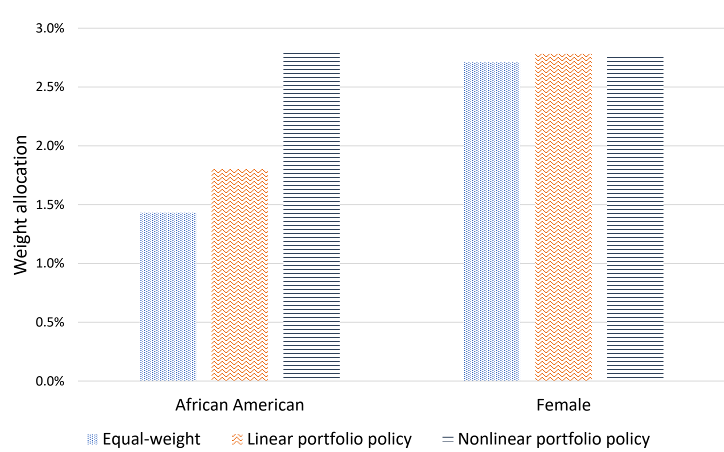

4. Practical Investment Constraints

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, and investors can allocate arbitrarily large amounts to a loan without limit. In reality, the minimum amount one can invest in a loan is \$25. Lenders also cannot lend more than the total loan amount (or the currently available amount), which places an upper limit on the investable capital. We propose a binary search method to handle investment amount constraints.

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. For each loan i, the truncated weight is

$$\tilde{\omega}(M) = \begin{cases} 0 & \text{if } \omega_i * M < 25 \\ U_i/M & \text{if } \omega_i * M > U_i \\ \omega_i & \text{otherwise} \end{cases} \tag{13}$$

where $\$U_i$ is the available loan amount for loan i and ω_i is the unconstrained portfolio weight. The actual investment amount after truncation is given by a function f that maps the preliminary portfolio to the desired portfolio:

$$f(M) = \left(\sum_{i=1}^{N} \tilde{\omega}_i(M) \right) \times M. \tag{14}$$

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 Online Appendix G.1.

We investigate the portfolio performance of the constrained nonlinear portfolio with an investable amount ranging from \$1,000 (1k) to \$100 million (100m). The smaller amounts serve as proxies for retail investors, whereas the larger amounts simulate the problem institutional investors face. 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. Online Appendix G contains additional details. Figure 4 illustrates that 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. Investors can select the investment amount that fits their return-risk expectation.

Figure 4 Average and Standard Deviation of IRRs for Investment-Amount Constrained Portfolios

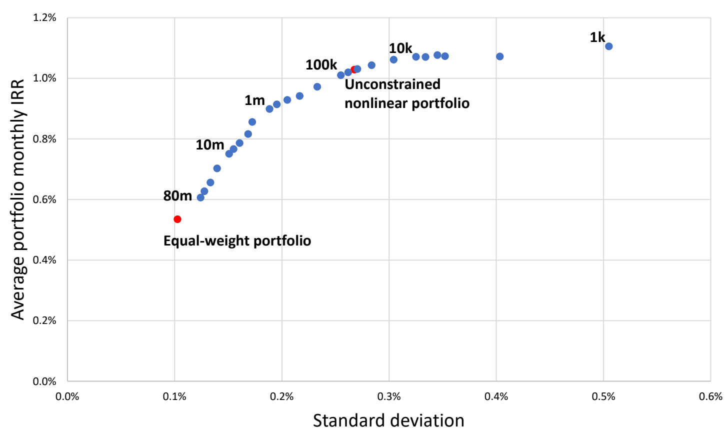

To better understand the make-up of the constrained portfolios, we calculate the sum of weight deviation between a constrained portfolio (ω_i) and the benchmark unconstrained portfolio ($\omega_{i,b}$):

$$\text{Sum of Weight Deviation} = \frac{1}{2}\sum_{i=1}^{N_t} |\omega_i - \omega_{i,b}| \tag{15}$$

Two completely different portfolios have a sum of weight deviation of 100%, while two highly similar portfolios have a sum close to zero (Cremers and Petajisto 2009). The 100k portfolio has weights closest to the unconstrained portfolio - just a 17% difference in portfolio weights. Portfolios of larger or smaller sizes tend to deviate further from the unconstrained baseline; small portfolios come up against the 25-dollar minimum investment limit, and 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 91%, which illustrates the significant differences between these two strategies; the vast majority of the positions in the two portfolios do not match.

5. Performance Comparison to Other Asset Classes

Before LendingClub shifted its investor base from retail to institutional investors, the number of retail users on the platform steadily increased. This suggests that a growing fraction of investors consider online loans to be beneficial for their portfolios. However, the total number of investors in online loans remains small compared to those in other asset classes. For example, in the United States, a half of all U.S. households invest in the stock market, and nearly all institutional investors have equity exposure (Bhutta et al. 2020). Let us step back to answer a broader question: Is it worthwhile to invest in online loans, especially when compared to more traditional assets such as stocks and bonds?

We consider three benchmark assets: Stocks, bonds, and real estate. The S&P 500 Index is one of the most representative gauges of public equity investments, as much of stock market investments are in the form of index funds or exchange-traded funds tracking the S&P 500. Therefore, we use the S&P 500 Index as the stock market benchmark. We consider equally prominent indices in the bond and real estate markets. Bond market indices include the Bloomberg U.S. Aggregate Bond Index, which tracks all investable bonds in the U.S.,[20] and S&P U.S. Treasury Bond Indices for 1-3, 3-5, and 10-20 years, which track U.S. government debt at different maturities. The MSCI U.S. REIT Index is used as the benchmark for real estate investments.[21]

Portfolio IRR is not directly comparable to index returns on stock, bonds, or real estate because online loans do not have an actively-traded secondary market; investors must hold loans to maturity or the charged-off date, and they cannot realize the calculated IRR figures before then. While it is difficult to construct a return measure for online loans similar to liquid markets, we 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 any future cash flows also in the S&P (or liquidate some positions to satisfy cash outflows). 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. Public market equivalent (PME) provides a direct comparison metric (Kaplan and Schoar 2005):[22]

$$PME = \frac{IRR_{\text{loan}}}{IRR_{\text{benchmark}}}. \tag{16}$$

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

We compare \$100,000 investments in Table 3, since the \$100,000 loan portfolio best resembles the optimal unconstrained portfolio while satisfying practical constraints. 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. 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[23], which indicates that stocks offer a higher 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. In comparison, an equal-weight portfolio of online loans provides a higher IRR than that of 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. The nonlinear policy also significantly reduces the correlation between online loans and other assets with a fixed-income component.

Existing literature on asset allocation emphasizes the importance of a portfolio perspective in addition to performance analysis of individual assets Campbell and Viceira (2001), Judd et al.

Table 3 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 Treasuries		
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 Treasuries		
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 Treasuries			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%

(2011). We compare two asset allocation schemes, one including online loans, and one without. A portfolio of traditional assets excluding online loans consists of 60% invested in the S&P 500 Index, 15% Bloomberg U.S. Aggregate Bond Index, 15% S&P U.S. Treasury Bond 10-20 Year Index, and 10% MSCI U.S. REIT Index. A portfolio including online loans consists of 55% S&P 500 Index, 10% Bloomberg U.S. Aggregate Bond Index, 10% S&P U.S. Treasury Bond 10-20 Year Index, 10% MSCI U.S. REIT Index, and 15% in online loans based on the nonlinear portfolio policy. The portfolio including online loans consistently outperforms the portfolio excluding loans. The monthly IRR increases by 21.4% from 0.53% to 0.64%, while the standard deviation decreases by 13.9% from 0.26% to 0.22%. Please refer to Online Appendix J for more details.

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, offering 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 an 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

There is an ongoing shift in the investor base of online loans from retail to institutional. Whether this shift can be successful and sustainable crucially depends on whether online loans can offer attractive returns to a sophisticated investor. Our answer is affirmative: Online loans constitute an attractive new asset class when sophisticated techniques for portfolio optimization are applied. We thus predict that the shift in investor base will be successful and recommend institutional investors to gain a seat at the table and consider online lending as part of their asset allocation decision.

Our main theoretical contribution lies in the introduction of general characteristics-based portfolio policies as a novel framework for portfolio construction that is suitable for investment in online loans. This framework bypasses the difficulty of estimating distributional properties of loans by modelling portfolio weights as flexible functions of loan characteristics. While our framework is general, we put particular emphasis on linear portfolio policies and a nonlinear policy based on neural networks. The flexibility of GCPP allows it to be readily extended to include other input variables such as the predictions of charged-off probability or loan returns generated from machine learning algorithms. The GCPP portfolios outperform an equal-weight portfolio of loans out-of-sample. In particular, the nonlinear portfolio policy performs the best. These results are robust when considering important practical constraints on loan portfolios such as the minimum and maximum investment amounts. We also find that online loans offer competitive rates of return compared to stocks, bonds, and real estate, while offering diversification benefits. Crucially, comparing an equal-weight portfolio of loans with stocks would lead to the conclusion that online loans do not constitute an attractive asset class. Thus, sophisticated portfolio optimization is important when deciding whether to invest in online loans.

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, while enabling greater access to credit for high-risk borrowers. Furthermore, our results indicate that there is room for improvement in how platforms set interest rates. Interest rates on LendingClub may not always be reflective of the underlying risk of loans, and our model uncovers such mismatches. Indeed, the mismatch between risk and interest rate is the main driver behind the superior performance of the GCPP framework. A more efficient mechanism to determine interest rates has the potential to resolve much of the investors' hesitation and make investment in online loans more attractive for investors who do not have the abilities to construct sophisticated portfolios of loans.

For practitioners interested in applying the GCPP framework, two comments are in order. First, the outperformance of the nonlinear GCPP is driven by the amount of available training data and mismatches between interest rate and risk. A greater degree of mismatches provides more opportunities to profit, and more data allow GCPP to identify these mismatches more precisely. Second, since the nonlinear GCPP is a complex model, the possibility of overfitting is considerably higher compared to simpler models. In fact, the primary factor that could lead to underperformance of GCPP lies in overfitting. Although overfitting does not appear to be the case in our setting (all comparisons are out-of-sample), researchers and practitioners must be keenly aware of this possibility.

While the GCPP framework is particularly suitable for online loans, its range of applications is potentially much broader. Our framework can be applied to 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. Potential bias against certain borrower groups is also worthy of a deeper investigation. We leave these possibilities to future research.

Acknowledgments

The work described in this paper was partially supported by InnoHK initiative, The Government of the HKSAR, and Laboratory for AI-Powered Financial Technologies. The authors are grateful to the senior editor, the associate editor, and three anonymous reviewers for their detailed comments that have greatly improved the paper. The authors also thank Jinguang Yang for valuable discussions.

Endnotes

[1]Online lending is known initially as peer-to-peer (P2P) lending, which connects individual borrowers with individual lenders. As the business evolves to include institutional lenders, people now describe it as online lending or marketplace lending (Treasury, 2016).

[2]The estimated market size in the U.K. is from IBIS World.

[3]The GCPP framework is generalizable to other major online lending platforms that adopt the fixed-rate mechanism. Prosper and Funding Circle started off using auction-based mechanisms but moved to the fixed-rate mechanism in 2011 and 2015, respectively. The current specification of GCPP cannot be directly applied to the auction-based setting, and the specification would have to be adjusted (See more details in Online Appendix I.2).

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

[5]The total market capitalization of global stock markets is obtained from the World Bank. The estimated bond market size is from the Securities Industry and Financial Markets Association (SIFMA).

[6]Notable papers include Iyer et al. (2016); Herzenstein et al. (2008); Emekter et al. (2015); Serrano-Cinca et al. (2015). Some studies find that incorporating alternative information, such as appearance (Pope and Sydnor, 2011; Duarte et al., 2012), social network on the online lending platforms (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.

[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.

[8]LendingClub lists requested loans on its marketplace after platform screening, and investors decide whether to fund them. In principle, some listed loans may be subsequently rejected if there is insufficient investor interest. Kleinberg et al. (2018) illustrates how a selective labels problem arises if the availability of the outcome data is not random. In our setting, there would be a potential bias if we were to only focus on the accepted loans. We download all listed loans on LendingClub, and we find that none of the listed loans were rejected. Therefore, our results are not driven by the selective labels problem. Please refer to Online Appendix K for additional discussion.

[9]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.

[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.

[11]In this paper, we study a single-period portfolio problem. In practice, investors may face a multi-period problem where they need to decide on the optimal capital allocation across time as well as across loans. Under such setting, investors can consider a laddering strategy (Cheung et al. 2010) that uniformly invests in the optimal portfolio to diversify their risk exposure across loan origination time.

[12]We fix $\gamma = 2$ for all calculations. We also tested for different values for γ, including $\gamma \in \{0, 1, 2, 4, 6, 8, 10\}$, and found that our results remained qualitatively robust. More details is available in Online Appendix H.4

[13]The benchmark choice depends on specific context. Other benchmark choices can be easily accommodated by Equation (9). 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.

[14]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 $y \in \mathbb{R}^{64}$.

[15]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.

[16]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.

[17]$\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-8}$

[18]African American and the female borrowers consist 1.4% and 2.7% of all loans, respectively. Note that these numbers only represent the borrowers we can identify with high confidence, not all borrowers that belong to these groups.

[19]The difference in portfolio weight allocation to African Americans and all loans is statistically significant at a 1% level using a one-sided paired two-sample t-test. However, we cannot reject the null hypothesis that the portfolio weight is the same for the female group and all loans.

[20]All bonds include corporate debt, government debt, mortgage-backed securities, and asset-backed securities.

[21]All prices are adjusted for splits and dividend and capital gain distributions.

[22]PME was proposed by Kaplan and Schoar (2005) as a performance metric for private equity. Like loans, private equity investments also cannot be easily sold prior to an exit date.

[23]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.

References

Agarwal S, Alok S, Ghosh P, Gupta S (2020) Financial inclusion and alternate credit scoring for the millennials: Role of big data and machine learning in fintech. Working paper, Business School, National University of Singapore, Singapore.

Agarwal S, Ambrose BW, Liu C (2006) Credit lines and credit utilization. *J. Money Credit Bank.* 38(1):1–22.

Balyuk T (2022) FinTech lending and bank credit access for consumers. *Management Sci.* 0(0).

Bartlett R, Morse A, Stanton R, Wallace N (2022) Consumer-lending discrimination in the FinTech era. *J. Financial Econom.* 143(1):30–56.

Berg T, Fuster A, Puri M (2022) Fintech lending. *Annu. Rev. Financial Econ.* 14(1):187–207.

Bhutta N, Bricker J, Chang AC, Dettling LJ, Goodman S, Hsu JW, Moore KB, Reber S, Henriques Volz A, Windle R (2020) Changes in us family finances from 2016 to 2019: Evidence from the survey of consumer finances. *Fed. Reserve Bulletin* 106(5).

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

Byanjankar A, Mezei J, Heikkilä M (2021) Data-driven optimization of Peer-to-Peer lending portfolios based on the expected value framework. *Intell. Syst. Account. Finance Manag.* 28(2):119–129.

Campbell JY, Viceira LM (2001) Who should buy long-term bonds? *Am. Econ. Rev.* 91(1):99–127.

Cheung CS, Kwan CC, Sarkar S (2010) Bond portfolio laddering: a mean-variance perspective. *J. Applied Finance* 20(1).

Chi G, Ding S, Peng X (2019) Data-driven robust credit portfolio optimization for investment decisions in P2P lending. *Math. Prob. Eng.* 2019.

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

Cremers KM, Petajisto A (2009) How active is your fund manager? A new measure that predicts performance. *Rev. Financial Stud.* 22(9):3329–3365.

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

Fu R, Huang Y, Singh PV (2021) Crowds, lending, machine, and bias. *Info. Sys. Res.* 32(1):72–92.

Fuster A, Plosser M, Schnabl P, Vickery J (2019) The role of technology in mortgage lending. *Rev. Financial Stud.* 32(5):1854–1899.

Goldstein I, Jiang W, Karolyi GA (2019) To fintech and beyond. *Rev. Financial Stud.* 32(5):1647–1661.

Gopal M, Schnabl P (2022) The rise of finance companies and fintech lenders in small business lending. *Rev. Financial Stud.* 35(11):4859–4901.

Gu S, Kelly B, Xiu D (2020) Empirical asset pricing via machine learning. *Rev. Financial Stud.* 33(5):2223–2273.

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

Guo Y, Zhou W, Luo C, Liu C, Xiong H (2016) Instance-based credit risk assessment for investment decisions in P2P lending. *Eur. J. Oper. Res.* 249(2):417–426.

Huang H, Wang ML (2021) The fall of online P2P lending in China: A critique of the central-local co-regulatory regime. *Bank. & Fin. L. Rev.* 36(3):481–515.

Judd KL, Kubler F, Schmedders K (2011) Bond ladders and optimal portfolios. *Rev. Financial Stud.* 24(12):4123–4166.

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

Kate B, Nicholas M (2019) Peer-to-Peer pressure: Do the risks outweigh the rewards? *Financial Times* URL `https://www.ft.com/content/6bf2c806-4a6b-11e9-8b7f-d49067e0f50d`.

Katsamakas E, Sánchez-Cartas JM (2022) Network formation and financial inclusion in P2P lending: A computational model. *Systems* 10(5):155.

Kingma DP, Ba JL (2015) Adam: A method for stochastic optimization. Bengio Y, LeCun Y, eds., *Proc. Int. Conf. Learn. Representations* (San Diego, CA).

Klein G, Shtudiner Z, Zwilling M (2021) Why do peer-to-peer (P2P) lending platforms fail? The gap between P2P lenders' preferences and the platforms' intentions. *J. Electron. Commer. Res.* 1–30.

Kleinberg J, Lakkaraju H, Leskovec J, Ludwig J, Mullainathan S (2018) Human decisions and machine predictions. *Q. J. Econ.* 133(1):237–293.

Michaud RO (1989) The Markowitz optimization enigma: Is "optimized" optimal? *Financial Anal. J.* 45(1):31–42.

Nemoto N, Storey DJ, Huang B (2019) Optimal regulation of P2P lending for small and medium-sized enterprises. Working paper, Asian Development Bank, Japan.

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

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

Shen X, Tan T, Phan TQ, Keppo J (2022) Unwarranted gender disparity in online P2P lending: Evidence of affirmative action. Working paper, arXiv:2210.07864.

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