

Harnessing Collective Intelligence in P2P Lending

Henry K. Dambanemuya
Northwestern University
Evanston, Illinois
hdambane@u.northwestern.edu

Emőke-Ágnes Horvát*
Northwestern University
Evanston, Illinois
a-horvat@northwestern.edu

Abstract

Crowd financing is a burgeoning phenomenon that promises to improve access to capital by enabling borrowers with limited financial opportunities to receive small contributions from individual lenders towards unsecured loan requests. Faced with information asymmetry about borrowers' credibility, individual lenders bear the entire loss in case of loan default. Predicting loan payment is therefore crucial for lenders and for the sustainability of these platforms. To this end, we examine whether the "wisdom" of the lending crowd can provide reliable decision support with respect to projects' long-term success. Using data from Prosper.com, we investigate the association between the dynamics of lending behaviour and successful loan payment through interpretable classification models. We find evidence for collective intelligence signals in lending behaviour and observe variability in crowd wisdom across loan categories. We find that the wisdom of the lending crowd is most prominent in the auto loan category, but it is statistically significant for all other categories except student debt. Our study contributes new insights on how signals deduced from lending behaviour can improve the efficiency of crowd financing thereby contributing to economic growth and societal development.

Keywords

Crowdsourcing, Crowdfunding, Herding, Signalling, Wisdom of Crowds, Decision-Making, Dynamics

ACM Reference Format:

Henry K. Dambanemuya and Emőke-Ágnes Horvát. 2019. Harnessing Collective Intelligence in P2P Lending. In *11th ACM Conference on Web Science (WebSci '19)*, June 30–July 3, 2019, Boston, MA, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3292522.3326040>

1 Introduction

Peer-to-Peer (P2P) lending services are social platforms that connect individual borrowers to lenders who compete through a bidding process to invest outside of traditional financial institutions [1, 5, 26]. Although such online financing is a relatively new phenomenon, several studies have already investigated determinants of successful fundraising [2–4, 11, 12, 15, 16, 19–21, 24–26, 31–34]. Understanding who will receive funding is a crucial component

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebSci '19, June 30–July 3, 2019, Boston, MA, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6202-3/19/06...\$15.00

<https://doi.org/10.1145/3292522.3326040>

of the crowd financing model. However, crowd financing platforms also need to recognise how to attract and keep lenders. From the lenders' perspective, the key outcome is whether a borrower pays the loan on time [14, 21, 30]. Making this determination is extremely hard, even for experienced institutional lenders, let alone for untrained individuals who face additional information asymmetry compared to offline lenders, because they have less access to factual information about borrowers' credibility such as credit history, income, or employment. Lenders bear the entire loss in case of loan default and might quit using crowd financing, which could impact the viability of several online platforms. Typically, lenders consider loan characteristics such as amount, interest rate, and credit score when making lending decisions. However, lenders may also have additional criteria and since information about bids is made public on the platform, their behaviour can influence others. Such social learning is prevalent when individuals find it difficult to evaluate an object's true quality [28, 29] and can lead to information cascades and herding [6, 13], which might result in sub-optimal outcomes. Research on P2P lending suggests, however, that herding can help novice lenders identify which borrowers are the best risk [35]. Motivated by this work, we propose to scrutinise the collective intelligence of the lending crowd and identify a set of indicators that are associated with accurate predictions of loan success. We argue that lenders can capitalise on each other's evaluation of the borrowers' creditworthiness to improve their decision-making as a collective.

Specifically, *we investigate whether the wisdom of the lending crowd can help estimate the likelihood that a borrower will pay the loan years down the line*, i.e., the long-term success of borrowers' projects. To this end, we develop a set of features that characterise lending behaviour and account for potential learning effects and changes in the lender population. We then aggregate these features to describe the lender crowd contributing to individual projects. Finally, we train classification models with loan, borrower, and lending dynamics features to investigate which factors are associated with loan payment. By improving on a random-estimator baseline, we establish a basis for non-trivial prediction of loan payment from features which summarise lenders' actions and are determinants of long-term project success. We conclude by investigating the lending dynamics features that make prediction possible and further investigate whether the lending crowd is wiser in predicting the long-term success of certain project categories than others. The premise of our work is not to debate whether lender dynamics are more important than borrower or loan features in evaluating loan payment. Rather, *we argue that dynamics in lending behaviour add significant value to the estimation accuracy of machine learning models for predicting loan payment*. Since it is not known how these dynamics manifest themselves, we are compelled to investigate a range of potential collective intelligence indicators and test which

signals contribute to the reliable inference of projects' long-term success.

First and foremost, our findings provide new insights on how lender behaviour can improve crowd financing efficiency and thereby contribute to economic growth and societal development. Specifically, we complement existing research on who receives funding [2–4, 11, 12, 15, 16, 19–21, 24–26, 31–34] with an investigation of the lesser understood determinants of loan payment or default. Additionally, our approach also informs the study of broadly relevant problems of collective decision-making under high uncertainty. Crowd financing is essentially a prototypical case of decision-making with substantial associated ambiguity: It is mediated by an online marketplace and presents challenges in determining, for instance, the veracity of identities, reliability of claims, merit of ideas, and their expected appeal to an indefinite crowd. The tasks tackled by lenders are therefore representative of the problems that arise in a suite of other P2P platforms such as consumer-to-consumer e-commerce websites (eBay), as well as task-service (TaskRabbit) and hospitality (AirBnB) platforms. By creating a framework that separates domain-specific predictors of long-term success from generalisable indicators of collective intelligence, the latter are expected to find use also beyond the context of crowd financing.

The rest of the paper is organised as follows: Section 2 summarises related research on crowd financing. Section 3 introduces data from Prosper.com. Section 4 describes measures and models used for predicting long-term success, while Section 5 provides our results. We conclude the paper in Section 6.

2 Related Work

Crowdfunding is an emerging topic of research in fields ranging from social computing [11, 15, 23, 24, 34] to economics [1, 5, 26]. A key strand of literature examines which projects receive funding. Factors such as description quality [25], specific language use [4], project updates [34], the ability to attract funders early in the campaign [12, 15], promotional activities on social media [24], and the project creator's reputation [11] as well as their social capital [3, 16, 20, 31] have been found to be associated with higher rates of fundraising. Some of this research highlights the role of herding and information cascades among investors in successfully raising capital [2, 32], leaving the question open whether the crowd's herding was rational or not. In the low-information P2P lending context, Zhang and Liu argued that herding can be beneficial for novice lenders [35]. They found that, in general, lenders are savvy enough to know when to follow the herd, but that they are also more inclined to herd on well-funded listings.

Due to these potentially adverse herding effects, when determining whether the crowd produced actual collective intelligence, one needs to look beyond who received funding and tackle questions like who paid the loan back, set up a business successfully, or produced the promised project deliverables? Collecting data on such long-term outcomes is difficult and hence literature examining determinants of long-term project success is scarce. Few exceptions have studied default in P2P lending marketplaces: Serrano-Cinca et al. demonstrated via a logistic regression model that borrowers' credit information is the most predictive factor of loan default [30]. The authors also found evidence for the relevance of borrowers'

annual income, homeownership status, credit history, and indebtedness to loan payment. Emekter et al. found similar loan and borrower features to be important predictors of default [14]. To extend the problem's relevance beyond lenders' profitability, Klafft advocated for a thorough analysis from a lender's perspective as a path to increase the future business potential of P2P lending platforms. Acknowledging the severity of the information asymmetry problem, the author suggested that lenders following simple, portfolio-based investment rules can attain acceptable returns for all credit rating categories, with the exception of high-risk loans. Finally, in another study examining the performance of P2P lending platforms, Iyer et al. found that non-expert lenders predict the likelihood of loan default with 45% greater accuracy than the borrowers' exact credit score and achieve 87% of the predictive power of the best possible default predictor that an econometrician could have constructed using all available standard financial information about borrowers.

In a different spirit, centred on the temporal progression of lenders' bidding behaviour, the study by Ceyhan et al. used Prosper.com data to predict both short- and long-term project success [10]. Instead of making inferences based on a large number of traditional creditworthiness indicators, the authors modelled the temporal dynamics of bidding to grasp "how the market feels" and to gain information that was not present in borrower characteristics like credit grade and debt-to-income-ratio. Due to similarities in terms of data and focus, this study is most closely related to our research. Ceyhan et al. fit the progression of bids with a sigmoid function and leave it unclear how their model relates to collective intelligence theories at large.

In sum, our literature review indicates that while much is understood about short-term funding success (i.e., who receives funding), little is known about long-term payment success (i.e., who delivers). Moreover, aside from the Ceyhan et al. study, previous work on long-term success has concentrated on loan and borrower determinants of default. We see this as a knowledge gap in understanding the wisdom of lending crowds in P2P marketplaces. To fill this gap, our work concentrates on determinants of long-term project success deduced from lending dynamics that could improve crowd financing efficiency.

3 Data: Prosper Marketplace

Prosper.com is a peer-to-peer lending platform that allows borrowers to receive funding from members of a large online marketplace. Borrowers request loans by creating *listings* for a specific amount between \$1,000 and \$25,000 and specifying the maximum interest rate they are willing to pay if the listing turns into a loan. Lenders then *bid* to fund a fraction of the amount at a chosen interest rate. When a listing reaches at least 100% of its requested amount, bids with the lowest interest rates are pooled into a single *loan* awarded to the borrower at a final *interest rate* determined by Prosper.

Lenders make their decisions based on information displayed on the platform. Included on the listings are borrowers' financial information such as *debt-to-income ratio*, *credit grade*, and *Prosper score*, as well as the loan *amount*, project *description* and *category*. These details are available for all the 415,157 listings created between November 2005 and October 2008, belonging to six loan

categories ranging from student loans to auto loans (see Section 4.2 for definitions). In 2009, Prosper changed its privacy policy and stopped showing borrowers' credit score information. To account for this change in platform features, we restrict our study to data prior to this event.

To predict loan payment, we further limit our data to listings that were successfully completed by November 2008 and were either fully paid or defaulted due to, e.g., delinquency, bankruptcy, or a borrower being deceased. This sample excludes loan payments that were still ongoing, late, cancelled, repurchased, or charged off. By design, this sample only considers the activity of borrowers who requested at least one successful loan and lenders with at least one participating bid in a funded listing. The final data contains 3,948,777 bids on 28,935 successfully funded listings. These listings belong to 26,404 borrowers whose projects were funded by 50,264 lenders. Of the 28,935 loans, 18,279 were successfully repaid, while 10,656 defaulted. The average requested loan size for all completed listings was \$6,171 (standard deviation = \$5,585) with a mean of 21.2% in maximum borrower rate. The total amount lent to borrowers through Prosper in this three year period was \$178,556,824.

4 Predicting Long-Term Success

In this section, we describe the experimental setup of our study. We begin by defining the problem of predicting long-term success as a binary classification task, then list the learning methods used for prediction as well as the performance metrics used to evaluate the classifiers. Finally, we discuss our feature selection methodology, providing detailed explanations about how we computed new features from Prosper.com data.

4.1 Experimental Setup

For prediction, we represent listings by a set of features that describe attributes of the borrower, terms of the loan, and features that summarise the behaviour of the lenders who contributed to the loan. We use standard scaling to normalise the feature set by subtracting the mean and scaling to unit variance. Each listing's long-term success is indicated by whether the borrower successfully paid the loan (1) or not (0). We tackle this binary classification problem with a variety of learning methods: Random Forests (RF) [7], Classification and Regression Trees (CART), Adaptive Boosting (ADB) [17], Logistic Regression (LR), Gaussian Naive Bayes (GNB), and Quadratic Discriminant Analysis (QDA). We rely on Scikit-Learn's Python API for method implementation [27].

To perform out of sample tests, in all learning setups, we do 5-fold cross validation and report the classification accuracy, precision, recall, F1 score, and area under the receiver operating characteristic curve (AUC). We also use 5-fold cross validation for parameterisation. To compare the predictive performance of each feature, we use the Random Forest Variable Importance (viRF) computed via Gini Importance [8]. Furthermore, to understand the factors associated with payment success for different loan categories, we run separate classification models and feature importance evaluations within each individual category represented in the data.

4.2 Determinants of Long-Term Success

We proceed by describing factors that contribute to loan payment. We divide these factors into three groups: loan, borrower, and lending dynamics determinants. When designing composite

features within each group, we weigh components of the composite feature using their viRF scores as described below.

4.2.1 Loan Determinants We examine each project's loan features provided by Prosper.com. These features describe standard characteristics of the loan such as the *amount requested*, *monthly loan payment*, *loan term*, estimated annualised loss rate on the loan (*estimated loss*), and the rate the borrower pays on the loan (*interest rate*). We also include the *length of the project description* measured by the number of words in the project description as a proxy of the effort borrowers invested in their listings [10]. For each listing, Prosper also provides a *credit grade* that indicates the loan's estimated average annualised loss rate range. This is calculated based on Prosper's proprietary system and allows the platform to maintain consistency when evaluating individual loan requests.

4.2.2 Borrower Determinants We also make use of standard borrower features. Consistent with existing literature, we consider borrowers' *homeownership* status. Furthermore, on Prosper.com, a member can be simultaneously a borrower, lender, and group leader, which leads to a simple feature that summarises their engagement with the platform by counting their roles (*role count*). For each listing, we additionally measure *borrower age* as the number of days between a borrower became a Prosper member and posted their listing. To evaluate *borrower experience* during this time, we sum up the ratio of the borrower's funded listings to their total listings and the ratio of their paid loans to their total loans. The two components are weighted by their viRF scores obtained from a random forest model that contains both ratios. Borrower experience therefore measures borrowers' *past* success on the platform, prior to their current listing and is zero for each borrower's first project.

To investigate the effect of credit information on the predictive performance of borrower features and to emulate realistic decision-making scenarios where credit information might not be available, we group borrower features into sensitive and non-sensitive features. The borrower features introduced so far were non-sensitive. Sensitive borrower features include: *i.) credit volatility* defined as the difference between a borrower's highest and lowest credit score rating as obtained from reporting agencies; *ii.) Prosper score*, which is a custom risk score built using historical Prosper.com data; and *iii.) debt-to-income ratio*.

4.2.3 Determinants of Lending Dynamics Intuitively, the more people a listing attracts, the more likely it is that a project meets its funding goals and the more compelled the borrower will be to pay back the lenders. As the most straightforward lender feature, we thus count the total *number of bids* for each listing and expect this to correlate with success (c.f. Ceyhan et al. [10]).

Since we also anticipate that lenders' investment outcomes will improve with the amount of time they spend on the platform and the number of funding attempts they make, we compute lender features that correspond to borrower age and experience. In this case, however, we aggregate age and experience over every lender who contributed to a specific loan. The age of individual lenders is defined as the number of days between they became Prosper members and made a bid on the current listing. Lenders' age for a given listing is therefore characterised by the *median lender age*. Lenders' experience is defined as the sum of the fraction of lenders'

winning bids (i.e., the ones that have not been outbid by others), their contributions to funded as opposed to not funded listings, and participation in successfully paid loans as opposed to defaulted ones. Components of the sum are weighted according to viRF scores obtained from a random forest model that contains the three fractions individually. *Average lender experience* is computed from the *past* successes of the entire group of lenders, prior to the current listing.

The temporal aspects of the crowd’s bidding dynamics might also indicate confidence in a listing’s merit, hence we develop a set of features that describe the speed at which funds are accumulating, i.e., how fast lenders make their determination of the projects’ potential. For each listing, we compute the *time to first bid* as the difference between the first bid and the time the listing was posted, as well as the *time between the first and last bid* when the listing completed. Additionally, we measure the *coefficient of variation* in inter-bid event times as the ratio between the mean and standard deviation of the time intervals between consecutive bids. This measure quantifies the difference in bidding times between subsequent lenders and indicates potential plateaus and surges in the accumulation of bids.

Finally, we include a set of features that are tied to the accumulated bid amounts. First, based on Vulkan et al.’s observation that the largest amount in a successfully funded project accounted for about 30% of the total capital sought, as opposed to 5.4% of the not funded project’s target [33], we measure the *maximum lender bid amount* for each listing. This amount might not only be important for its contribution towards a project’s fundraising success, but it can also indirectly signal loan quality. Second, we calculate the *mean bid amount per second*, i.e., dollars per unit time. This feature essentially indicates a contribution frequency that captures both the bid amount and the timing of lender activity, and has been shown to correlate with fundraising success [9]. Third, we explicitly quantify herding based on correlations between the contributed amounts. We use a measure that is similar to Goh and Barabási’s memory coefficient [18]. Given the consecutive bid amounts $\{B_1, B_2, \dots, B_N\}$, the *coefficient of herding* is defined as the ratio between the co-variation of the sequences $S1 = \{B_1, B_2, \dots, B_{N-1}\}$ and $S2 = \{B_2, B_3, \dots, B_N\}$ and the product of the standard deviation of these sequences:

$$\frac{1}{N} \sum_{i=1}^N \frac{(B_i - \mu_{S1})(B_{i+1} - \mu_{S2})}{\sigma_{S1}\sigma_{S2}},$$

where N is the number of bids in the loan, μ_{S1}, μ_{S2} and σ_{S1}, σ_{S2} are the mean and standard deviation of the first ($S1$) and second ($S2$) sequence, respectively. While we also considered longer range correlations, pair-wise correlations were the most predictive of long-term success.

The proposed *lending features evolve with respect to each bid* that a lender makes and with each listing that a lender contributes to. This is important, because through the repeated re-computation of lending dynamics features, we account for potential learning effects and changes in lender population without making any assumptions about how learning and attrition affect our diverse lender population. Our features are thus general and may describe collective intelligence in a variety of settings that involve decision-making online. If shown that minimally constraining features deduced from

lenders’ actions are predictive of loan payments, one can argue that the lending crowd’s behaviour can help improve crowd financing efficiency through better predicting which loans might be successful.

5 Results

Using all the 23 features introduced above, we were able to predict project long-term success with a random forest accuracy of 0.7147, 95% confidence interval: (0.7144, 0.7149), and AUC score of 0.707 (a random estimator would achieve an AUC of 0.5). Table 1 shows the evaluation results for all six machine learning models. Our results, obtained with different sets of features, are comparable to the currently state-of-the-art model by Ceyhan et al. [10].

Table 1: Long-term success prediction results. Random Forest classifier yielded best estimation results with an AUC score of 0.707.

Model	Accuracy	Precision	Recall	F-Score	AUC
QDA	0.704	0.766	0.765	0.766	0.682
CART	0.645	0.721	0.714	0.718	0.620
GNB	0.694	0.749	0.775	0.762	0.665
RF	0.715	0.797	0.736	0.765	0.707
LR	0.626	0.865	0.483	0.620	0.677
ADB	0.725	0.754	0.837	0.793	0.685

Correlations between features. Before comparing the predictive power of loan, borrower, and lending features, we investigate the correlations between features belonging to individual groups (see Table 2). Among loan features, there are unsurprising positive correlations between *i.*) the amount requested and the monthly loan payment, *ii.*) the amount requested and the credit grade, *iii.*) the estimated loss and interest rate, as well as *iv.*) credit grade and monthly loan payment. Equally expected is the negative correlation between credit grade and estimated loss. The highest correlation among borrower features is between borrower age and experience. This is to be expected given that we built both of these features to describe borrowers’ familiarity with the platform. We also find that the Prosper score correlates somewhat with borrower age and credit volatility, giving us an intuition for the criteria Prosper’s algorithm uses when assigning this score. We learn more about the components of Prosper score by inspecting its correlations with borrower features. Accordingly, Prosper score has a negative correlation with interest rate, estimated loss, and loan age, as well a positive correlation with credit grade.

The correlations among lending dynamics features are more interesting. There is a strong positive correlation between the maximum amount bid by a single lender and the coefficient of variation for inter-bid times, meaning that in the presence of a high maximal bid, inter-bid times are spread more widely. We find further, albeit smaller positive correlations between the coefficient of herding and average lender experience, the coefficient of herding and time between first and last bid, as well as bid count and time between first and last bid. The remaining correlations are rather small, indicating that the developed lending features are different across the considered sample of projects. Furthermore, most of the correlations between lending dynamics features and loan or borrower features are weak, demonstrating that lending decisions are based

Table 2: Correlations between loan, borrower, and lending features. Several weak correlations between lending dynamics features and loan or borrower features indicate that lending decisions are based on additional subtle information that is directly observed from the platform. * $p < 0.05$; ** $p < 0.01$; * $p < 0.001$.**

	Loan Features							Lender Features								Borrower Features								
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	
1. Interest Rate	-																							
2. Monthly Loan Payment	-0.074***	-																						
3. Estimated Loss	0.480***	0.105***	-																					
4. Description Length	0.160***	0.073***	0.185***	-																				
5. Credit Grade	-0.700***	0.369***	-0.409***	-0.228***	-																			
6. Loan Age	-0.013*	-0.039***	-0.045***	0.081***	-0.206***	-																		
7. Amount Requested	-0.155***	0.985***	0.061***	0.059***	0.424***	-0.043***	-																	
8. Average Lender Experience	-0.0794***	-0.048***	-0.128***	-0.003	0.097***	-0.046***	-0.043***	-																
9. Bid Amount Per Second	-0.159***	0.058***	-0.113***	-0.105***	0.186***	-0.090***	0.070***	0.040***	-															
10. Time Between First and Last Bid	-0.274***	0.173***	-0.119***	0.054***	0.256***	-0.157***	0.196***	0.085***	-0.128***	-														
11. Median Lender Age	-0.027***	-0.004	0.008	-0.098***	0.207***	-0.688***	-0.003	0.091***	0.082***	0.053***	-													
12. Time to First Bid	0.129***	-0.114***	0.096***	0.114***	-0.198***	-0.047***	-0.119***	-0.051***	-0.102***	-0.186***	0.021***	-												
13. Coefficient of Herding	-0.004	0.083***	-0.068***	0.030***	0.139***	-0.033***	0.081***	0.269***	-0.028***	0.246***	-0.026***	-0.097***	-											
14. Bid Count	-0.312***	0.728***	-0.067***	0.068***	0.450***	-0.174***	0.768***	-0.003	0.037***	0.417***	0.055***	-0.102***	0.153***	-										
15. Coefficient of Variation	0.106***	0.312***	0.101***	0.110***	0.054***	0.034***	0.298***	0.020***	-0.048***	0.072***	0.011	0.022***	0.133***	0.118***	-									
16. Max Lender Bid Amount	-0.003	0.386***	0.089***	0.064***	0.068***	0.016**	0.382***	-0.123***	0.039***	-0.019**	0.047***	0.028***	-0.206***	0.132***	0.709***	-								
17. Prosper Score	-0.345***	-0.070***	-0.418***	-0.133***	0.396***	-0.678***	-0.041***	0.121***	0.153***	0.217***	0.480***	-0.009	0.080***	0.147***	-0.093***	-0.089***	-							
18. Debt to Income Ratio	0.026***	0.079***	0.058***	0.066***	0.015*	0.027***	0.072***	0.004	-0.027***	0.005	-0.038***	0.013*	0.009	0.028***	0.058***	0.056***	-0.018**	-						
19. Borrower Age	-0.008	-0.024***	-0.019**	0.029***	0.060***	-0.291***	-0.022***	0.021***	0.007	0.053***	0.263***	0.037***	-0.000	0.031***	0.018**	0.018**	0.214***	-0.010	-					
20. Credit Volatility	0.069***	0.063***	0.004	-0.012*	0.057***	-0.470***	0.061***	0.020***	0.020***	0.066***	0.306***	-0.200***	-0.005	0.096***	0.008	0.031***	0.271***	-0.011	0.123***	-				
21. Role Count	-0.216***	-0.009	-0.141***	0.027***	0.191***	0.038***	0.003	0.032***	-0.008	0.092***	0.016**	-0.002	0.035***	0.062***	-0.001	-0.011	0.088***	-0.005	0.190***	-0.134***	-			
22. Borrower Experience	-0.045***	0.008	-0.020***	-0.005	0.062***	-0.177***	0.012*	0.021***	0.031***	0.034***	0.183***	-0.006	0.010	0.041***	0.026***	0.028***	0.139***	-0.009	0.575***	0.070***	0.120***	-		
23. Homeownership	-0.166***	0.231***	-0.021***	-0.054***	0.325***	-0.186***	0.243***	0.096***	0.058	0.114***	0.171***	-0.056***	0.056***	0.252***	0.058***	0.051***	0.133***	0.029***	0.051***	0.106***	0.070***	0.049***	-	

also on more subtle information than what is directly provided by the platform. The only high positive correlations we find between lending and loan features are with bid count. As expected, more lenders contributing to a listing is associated with a higher requested amount, higher credit grade, and more monthly loan payment. Finally, the correlations between median lender age and borrower features, indicate that experienced lenders tend to prefer borrowers who have used the platform for longer, have a high Prosper score, and even have more volatile credit. The latter could be attractive to lenders, because it gives them the opportunity to charge higher interest rates.

We also inspect correlations between characteristics of borrowers and loan payment outcome. For example, people with low credit grade are considered risky borrowers and likely to have a higher estimated loss on return (Pearson coefficient: $r = -0.409, p < 0.001$) and therefore high interest rate ($r = -0.700, p < 0.001$, c.f. Klafft [22]). Consequently, as shown in Figure 1, having a high interest rate significantly decreases the odds of successful loan payment (Odds ratio: $OR = 0.460, p < 0.001$). While people with high Prosper scores and credit grades may be good borrowers, they are also likely

to enjoy favourable loan terms such as low interest rates, which increases the odds of successful payment. While we found no significant differences in borrowers' odds of successful payment based on their homeownership status, a further investigation based on borrowers' level of indebtedness (i.e., debt-to-income ratio) yields interesting results. After grouping borrowers into three levels of indebtedness based on the lower, inter, and upper-quartile ranges in debt-to-income ratio, respectively, we observe that whether or not a borrower is a homeowner significantly increases the odds of successful payment if the borrower has low debt-to-income ratio ($OR = 1.172, p < 0.001$), but it decreases the odds of successful payment otherwise ($OR = 0.869, p < 0.001$). Our data shows that 72.8% of homeowners with low indebtedness successfully pay their loans compared to 63.8% of their non-homeowner peers (see Table 3). On the contrary, only 53% of homeowners with high indebtedness successfully pay their loans compare to 60.7% of non-homeowners with similar levels of debt. These findings suggest that loan payment is affected by a combination of features. Hence, in this study

we employ a set of multivariate models that are capable of investigating the effects of higher-order dependencies between borrower, loan, and lender features on long-term project success.

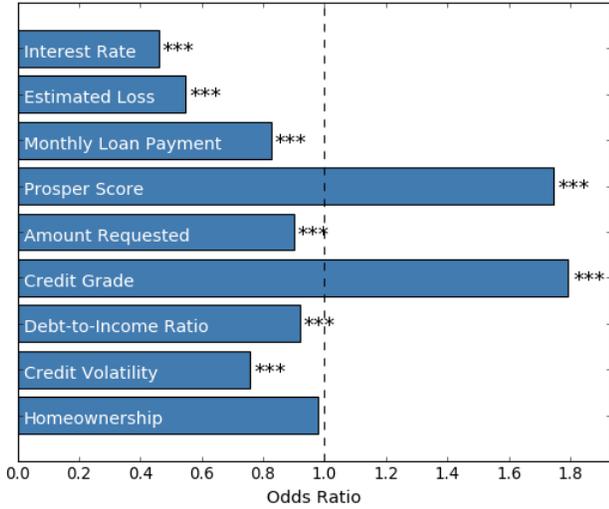


Figure 1: Odds of successful loan payment for observable platform features characteristic of loans and borrowers. High credit grade and Prosper score increase the odds of successful loan payment.

Table 3: Percentage of successfully paid loans by homeownership and level of indebtedness.

		Debt-to-Income Ratio		
		Low	Medium	High
Homeowner	True	72.8%	63.8%	53.8%
	False	63.0%	64.4%	60.7%

Lending dynamics features summarise other features. We further examine how different feature categories perform with respect to predicting long-term success. Using the best performing model from the previous experiment, random forest, we compare the predictive performance of lending dynamics, sensitive and non-sensitive borrower, as well as loan determinants. Unsurprisingly, the model with all feature categories performs best with an accuracy of 0.715. As expected, loan features have higher estimation accuracy (0.658) compared to lending dynamics (0.600) and borrower (0.627) features. We find this observation in agreement with previous studies that demonstrate the relevance of loan characteristics to borrowers' likelihood of fully paying their loans [14, 30]. To investigate the impact of borrowers' credit information, we use borrower features with and without sensitive information. When we train and evaluate the random forest classifier on the two sets of borrower features, we observe that excluding sensitive credit information from the model considerably lowers the estimation accuracy of borrower features by a margin of 0.214. In other words, non-sensitive borrower features attain a significantly lower accuracy than lending dynamics features. However, when we combine sensitive and non-sensitive borrower features with loan and lending dynamics features (i.e., use

all features), we observe that eliminating sensitive credit information has little effect on the model's accuracy (-0.293). This indicates that lending dynamics features must have picked up nearly all of the valuable signal from sensitive borrower information. Our finding suggests thus that despite the differences in lenders' experience and investment expertise, *as a collective, lenders successfully recovered and used credit information available in loan specifications and borrower profiles.* Moreover, the gleaned information is successfully captured in the proposed simple lending dynamics features.

Table 4: Feature importance ranking of predictive features for long-term success categorised into groups and ranked by feature importance score. Lender features have relatively higher viRF scores compared to most loan and borrower features, accounting for 38.6% of the predictive performance. * denotes sensitive credit information.

Group	Feature	Rank	viRF
Loan	Interest Rate	1	0.124
	Monthly Loan Payment	2	0.069
	Estimated Loss	3	0.055
	Description Length	8	0.045
	Credit Grade	15	0.039
	Loan Age	17	0.035
	Amount Requested	18	0.034
	Lender	Average Lender Experience	5
Bid Amount Per Second		6	0.045
Time between First and Last Bid		7	0.045
Median Lender Age		9	0.045
Time to First Bid		10	0.044
Coefficient of Herding		11	0.043
Bid Count		12	0.043
Coefficient of Variation		13	0.042
Max Lender Bid Amount		20	0.027
Borrower		Prosper Score*	4
	Debt-to-Income Ratio*	14	0.041
	Borrower Age	16	0.038
	Credit Volatility*	19	0.028
	Role Count	21	0.025
	Borrower Experience	22	0.019
	Homeownership	23	0.010

Lending dynamics features rank higher than most features. Next, we investigate the relative importance of individual features for prediction. Table 4 shows all features ranked using random forest variable importance (viRF) scores. Loan features are most predictive of long-term success: loan interest rate, estimated loss, and monthly payment consistently rank highest in viRF scores, while loan term and credit grade have poor viRF scores.

Although the top-3 features with the highest viRF scores are loan features, we observe that lending dynamics features have relatively higher viRF scores compared to most borrower features as well as other loan features such as credit grade, requested loan amount, and loan age. Of the three feature groups, lending dynamics features account for 38.6% of the predictive performance and therefore add significant value to the prediction of loan payment. This indicates that lenders do not only evaluate borrowers' *displayed* information when making lending decisions, but also rely on, e.g., potential informal information about borrowers. Furthermore, average lender

age and experience are among the most important lending dynamics features, which means that the more time lenders spend participating on the platform and the more successful they are, the more they contribute to the group’s collective intelligence. Lenders’ bid amount per second and the time between first and last bid are also indicative of successful loan payment. Indirectly linked to this speed of action is herding, which alongside the number of bids and the coefficient of variation, has medium importance among our set of features.

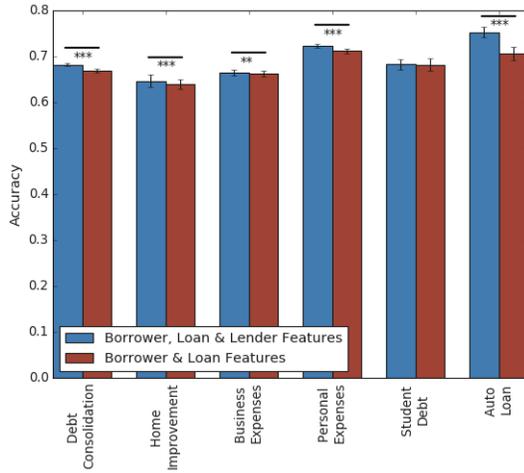


Figure 2: Model accuracy by feature group and project category. Lender features contribute significant marginal gains in estimation accuracy indicative of collective intelligence beyond summarising loan and borrower information.

Lenders demonstrate collective intelligence and their accuracy varies by loan category. We find that lenders demonstrate collective intelligence as lending dynamics features achieve 83.7% of the predictive power of all the indicators of long-term project success combined, which is consistent with the results of Iyer et al. [21]. While a random forest estimator trained only on loan and borrower features yielded an estimation accuracy of 0.7026, 95% confidence interval: (0.7024, 0.7029), adding lending dynamics features to the model slightly improved the prediction accuracy to 0.7147, 95% confidence interval: (0.7144, 0.7149). This marginal gain in estimation accuracy suggests that *lending decisions provide more than a simple summation of loan and borrower information* furnished through the platform. As a point of further inquiry, we investigate how lender wisdom varies by loan category and how the features describing lending dynamics change in importance ranking across those categories.

Adding lending dynamics features to loan and borrower features consistently improves estimation accuracy across loan categories. The wisdom of the lending crowd is most prominent in the auto loan category, but it is also statistically significant for all other categories except student debt (see Figure 2). This means that lenders are going beyond just summarising loan and borrower features. Instead, they augment available direct information with their own perceptions and interpretations of indirect signals, which leads to them being able to select creditworthy borrowers. Lenders are most accurate

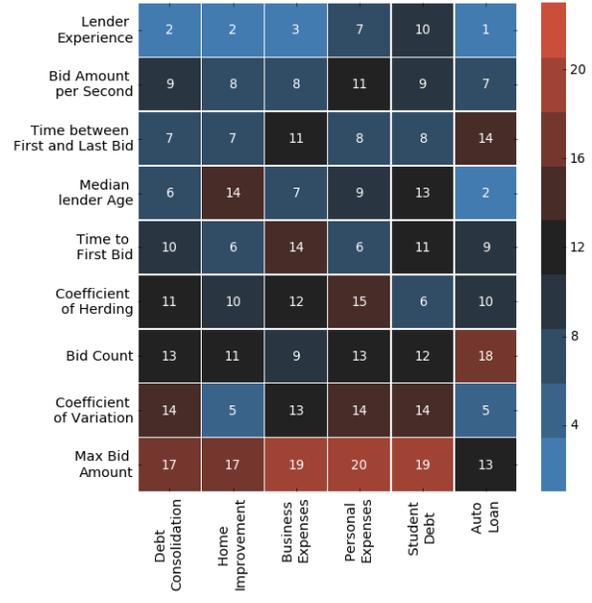


Figure 3: Analysis of feature ranking by project category. Displayed are only the ranks for lending dynamics features. We observe no systematic ranking patterns among lending features.

in the case of auto loans most probably because of their own experiences with this category of personal investment. They are least reliable in assessing the likelihood of successful payment for student loans. Although our data set does not enable more detailed investigations in this direction, we assume that predicting student debt payment is compounded by uncertainties in post-education outcomes that may affect borrowers’ opportunities to pay.

Finally, an analysis of the importance of lender features by loan categories shows that average lender experience is overall the most important, while maximum bid amount is the least important predictor for most loan categories. Note that the analysis underlying Figure 3 contained all three groups of features, but we display only the ranks of lending dynamics features relative to loan and borrower features. Aside of the consistent trends for average lender experience and maximum bid amount, there are no systematic ranking patterns among the other lending features, which suggests that one can hardly reduce lender characteristics to a single aspect of lender behaviour that is consistent across all loan categories.

6 Discussion & Conclusions

The wisdom of crowds idea suggests that untrained lenders working in low-information environments could possess a collective intelligence that allows superior decision-making comparable to or even exceeding that of expert institutional lenders. In P2P lending, a par excellence low-information decision-making context, determining whether borrowers will fully pay their loans is especially hard due to significant information asymmetry about their credibility. Notwithstanding this challenge, we found that lenders demonstrate collective intelligence that can be harnessed to predict future loan payment and help improve efficiencies in P2P markets’ capital allocation.

Our contributions are threefold. First, we provide new knowledge about signals deduced from lending behaviour that can contribute to the efficiency of crowd financing. Characteristics of lending dynamics such as lenders' previous successes, the bid amount per second, as well as the coefficients of variation and herding provide reliable estimates of long-term project success. These findings suggest that lenders' collective intelligence can be harnessed to increase the effectiveness of crowd financing. Second, we contribute to the growing literature on the wisdom of crowds by providing new insights about novel expressions of collective intelligence and potential ways to harvest it. We expect that our results will inform further research into crowd-aware system design on crowd financing platforms and beyond. Third, the proposed collective intelligence signals are general and easily transfer to various other crowd-sourcing settings. They are thus valuable in exploring contexts, where it is less straightforward to establish whether and how individuals delivered on their tasks.

There are several exciting directions for future research. For instance, further investigation into collective intelligence signals that enhance crowd financing outcomes could rely on more detailed observations of social signalling to understand its subtle effects on lenders' collective decision-making. Our results indicate that learning effects in crowds can help reveal innovative strategies to improve lending outcomes, especially for novice lenders. Identifying the conditions that optimise individual learning in this context is another fruitful avenue for future research. Finally, empirically testing the forecasting power of lending dynamics while the fundraising activity is still in progress has considerable practical appeal.

We believe that a better understanding of lender determinants of long-term project success promoted in this paper will help improve the efficiency of capital allocation in P2P markets, ultimately helping individuals get out of debt, improving household economic welfare, and enhancing entrepreneurship.

Acknowledgments

The authors would like to thank Brian Uzzi and Jayaram Uparna for providing the data. The authors would also like to thank the anonymous referees for their valuable comments and helpful suggestions. The work is supported by the U.S. National Science Foundation under Grant No. IIS-1755873.

References

- [1] Ajay Agrawal, Christian Catalini, and Avi Goldfarb. 2014. Some simple economics of crowdfunding. *Innovation Policy and the Economy* 14, 1 (2014), 63–97.
- [2] Ajay Agrawal, Christian Catalini, and Avi Goldfarb. 2015. Crowdfunding: Geography, social networks, and the timing of investment decisions. *Journal of Economics & Management Strategy* 24, 2 (2015), 253–274.
- [3] Gerrit KC Ahlers, Douglas Cumming, Christina Günther, and Denis Schweizer. 2015. Signaling in equity crowdfunding. *Entrepreneurship Theory and Practice* 39, 4 (2015), 955–980.
- [4] Tim Althoff, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2014. How to ask for a favor: A case study on the success of altruistic requests. In *ICWSM*.
- [5] Paul Belleflamme, Thomas Lambert, and Armin Schwienbacher. 2014. Crowdfunding: Tapping the right crowd. *Journal of Business Venturing* 29, 5 (2014), 585–609.
- [6] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. 1998. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives* 12, 3 (1998), 151–170.
- [7] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [8] Leo Breiman. 2017. *Classification and regression trees*. Routledge.
- [9] Gordon Burtch, Anindya Ghose, and Sunil Wattal. 2013. An empirical examination of the antecedents and consequences of contribution patterns in crowd-funded markets. *Information Systems Research* 24, 3 (2013), 499–519.
- [10] Simla Ceyhan, Xiaolin Shi, and Jure Leskovec. 2011. Dynamics of bidding in a P2P lending service: Effects of herding and predicting loan success. In *Proceedings of the 20th International Conference on World Wide Web*. ACM, 547–556.
- [11] Benjamin C Collier and Robert Hampshire. 2010. Sending mixed signals: Multilevel reputation effects in peer-to-peer lending markets. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*. ACM, 197–206.
- [12] Massimo G Colombo, Chiara Franzoni, and Cristina Rossi-Lamastra. 2015. Internal social capital and the attraction of early contributions in crowdfunding. *Entrepreneurship Theory and Practice* 39, 1 (2015), 75–100.
- [13] David Easley and Jon Kleinberg. 2010. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press.
- [14] Riza Emekter, Yanbin Tu, Benjamas Jirasakuldech, and Min Lu. 2015. Evaluating credit risk and loan performance in online Peer-to-Peer (P2P) lending. *Applied Economics* 47, 1 (2015), 54–70.
- [15] Vincent Etter, Matthias Grossglauser, and Patrick Thiran. 2013. Launch hard or go home!: Predicting the success of Kickstarter campaigns. In *Proceedings of the first ACM Conference on Online Social Networks*. ACM, 177–182.
- [16] Seth Freedman and Ginger Zhe Jin. 2008. *Do social networks solve information problems for peer-to-peer lending? Evidence from Prosper.com*. Technical Report 08-43. Indiana University, Bloomington: School of Public & Environmental Affairs, Bloomington, IN.
- [17] Yoav Freund and Robert E Schapire. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. System Sci.* 55, 1 (1997), 119–139.
- [18] K-I Goh and A-L Barabási. 2008. Burstiness and memory in complex systems. *EPL (Europhysics Letters)* 81, 4 (2008), 48002.
- [19] Michael D Greenberg, Bryan Pardo, Karthic Hariharan, and Elizabeth Gerber. 2013. Crowdfunding support tools: Predicting success & failure. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1815–1820.
- [20] Emőke-Ágnes Horvát, Jayaram Uparna, and Brian Uzzi. 2015. Network vs market relations: The effect of friends in crowdfunding. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*. ACM, 226–233.
- [21] Rajkamal Iyer, Asim Ijaz Khwaja, Erzo FP Luttmer, and Kelly Shue. 2015. Screening peers softly: Inferring the quality of small borrowers. *Management Science* 62, 6 (2015), 1554–1577.
- [22] Michael Klafft. 2008. Online peer-to-peer lending: A lenders' perspective. In *Proceedings of the International Conference on E-Learning, E-Business, Enterprise Information Systems, and E-Government, EEE*.
- [23] Yang Liu, Roy Chen, Yan Chen, Qiaozhu Mei, and Suzy Salib. 2012. I loan because...: Understanding motivations for pro-social lending. In *Proceedings of the fifth ACM International Conference on Web Search and Data Mining*. ACM, 503–512.
- [24] Chun-Ta Lu, Sihong Xie, Xiangnan Kong, and Philip S Yu. 2014. Inferring the impacts of social media on crowdfunding. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*. ACM, 573–582.
- [25] Dan Marom and Orly Sade. 2013. Are the life and death of an early stage venture indeed in the power of the tongue? Lessons from online crowdfunding pitches. *Unpublished. Working Paper Hebrew University* (2013).
- [26] Ethan Mollick. 2014. The dynamics of crowdfunding: An exploratory study. *Journal of Business Venturing* 29, 1 (2014), 1–16.
- [27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [28] Matthew J Salganik, Peter Sheridan Dodds, and Duncan J Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 5762 (2006), 854–856.
- [29] Matthew J Salganik and Duncan J Watts. 2008. Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Social Psychology Quarterly* 71, 4 (2008), 338–355.
- [30] Carlos Serrano-Cinca, Begona Gutierrez-Nieto, and Luz López-Palacios. 2015. Determinants of default in P2P lending. *PLoS One* 10, 10 (2015), e0139427.
- [31] Silvio Vismara. 2016. Equity retention and social network theory in equity crowdfunding. *Small Business Economics* 46, 4 (2016), 579–590.
- [32] Silvio Vismara. 2016. Information cascades among investors in equity crowdfunding. *Entrepreneurship Theory and Practice* (2016).
- [33] Nir Vulkan, Thomas Åstebro, and Manuel Fernandez Sierra. 2016. Equity crowdfunding: A new phenomena. *Journal of Business Venturing Insights* 5 (2016), 37–49.
- [34] Anbang Xu, Xiao Yang, Huaming Rao, Wai-Tat Fu, Shih-Wen Huang, and Brian P Bailey. 2014. Show me the money!: An analysis of project updates during crowdfunding campaigns. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 591–600.
- [35] Juanjuan Zhang and Peng Liu. 2012. Rational herding in microloan markets. *Management Science* 58, 5 (2012), 892–912.