Harnessing Collective Intelligence in P2P Lending

Henry K. Dambanemuya and Emőke-Ágnes Horvát, Northwestern University

1. INTRODUCTION

Peer-to-Peer (P2P) lending platforms connect individual borrowers to lenders who compete through a bidding process to invest outside of traditional financial institutions [Agrawal et al. 2014; Belleflamme et al. 2014; Mollick 2014]. Although such online financing is a relatively new phenomenon, several studies have already investigated determinants of successful fundraising [Freedman and Jin 2008; Collier and Hampshire 2010; Etter et al. 2013; Marom and Sade 2013; Greenberg et al. 2013; Althoff et al. 2014; Xu et al. 2014; Lu et al. 2014; Mollick 2014; Iyer et al. 2015; Colombo et al. 2015; Agrawal et al. 2015; Ahlers et al. 2015; Horvát et al. 2015; Vismara 2016a; 2016b; Vulkan et al. 2016]. Understanding who will receive funding is a crucial component of the crowd financing model. Typically, lenders consider loan characteristics such as the requested amount, interest rate, and credit score when making lending decisions. From their perspective, the key outcome is whether a borrower pays the loan on time [Serrano-Cinca et al. 2015; Emekter et al. 2015; Iyer et al. 2015]. However, determining whether a borrower will pay their loan on time is extremely hard, even for experienced institutional lenders. Untrained individuals 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. We, therefore, 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., we study the long-term success of projects. To this end, we develop a set of simple features based on the amount and timing of individual contributions. We then aggregate these features to describe the lender crowd for a project. Finally, we train classification models with features that describe the loan, borrower, and lending dynamics to investigate which factors are associated with the loan payment. Our models predict with high accuracy loan payment from features that summarise lenders' actions and are thus determinants of long-term project success. Finally, 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.

2. PREDICTING LONG-TERM SUCCESS

We predict long-term success in 415,157 project listings created between November 2005 and October 2008 on Prosper.com, the oldest P2P lending marketplace in the US. *Each listing contains a set of features that describe attributes of the borrower, terms of the loan, and behaviour of the lenders who contributed to the loan.* A full description of these features is available from [Dambanemuya and Horvát 2019]. 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) [Breiman 2001], Classification and Regression Trees (CART), Adaptive Boosting (ADB) [Freund and Schapire 1997], Logistic Regression (LR), Gaussian Naive Bayes (GNB), and Quadratic Discriminant Analysis (QDA). We rely on Scikit-Learn's Python API for method implementation [Pedregosa et al. 2011]. 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 [Breiman 2017]. 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.

3. RESULTS

Using all features, we were able to predict loan payment 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 I shows the evaluation results for all six machine learning models. Our results, obtained with different sets of features, are comparable to the current state-of-the-art model by Ceyhan et al. [Ceyhan et al. 2011].

Table I. : Predicting payment of P2P loans from characteristics of the proposed project (loan features), the credit history of the borrower (borrower features), and aggregate summaries of the timing and amount of contributions (lender features). Random Forest classifier yielded the 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

Using the best performing model from the previous experiment, random forest, we compare the predictive performance of lending dynamics, borrower, and loan features. Lending dynamics features account for 38.6% of the predictive performance and therefore add significant value to the prediction of loan payment. Using the random forest variable importance (viRF) measure to determine individual feature importance scores, we further observe that lending dynamics features have relatively higher viRF scores compared to most borrower features and other loan features such as credit grade, requested loan amount, and loan age. Average lender age and experience are among the most crucial 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 contributions and the coefficient of variation, has medium importance among our set of features.

We also 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. [Iyer et al. 2015]. 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 gain in estimation accuracy is not negligible given our sample size (25,000 listings). It suggests that lending decisions provide more than a simple summation of loan and borrower information furnished through the platform. Additionally, lenders' collective intelligence varies by loan category. We observe that the wisdom of the lending crowd is most prominent in the auto loan category and is also statistically significant for all other categories except student debt (see Figure 1, left). Finally, an analysis of the importance of lender features by loan category shows that average lender experience is overall the most important. At the same time,

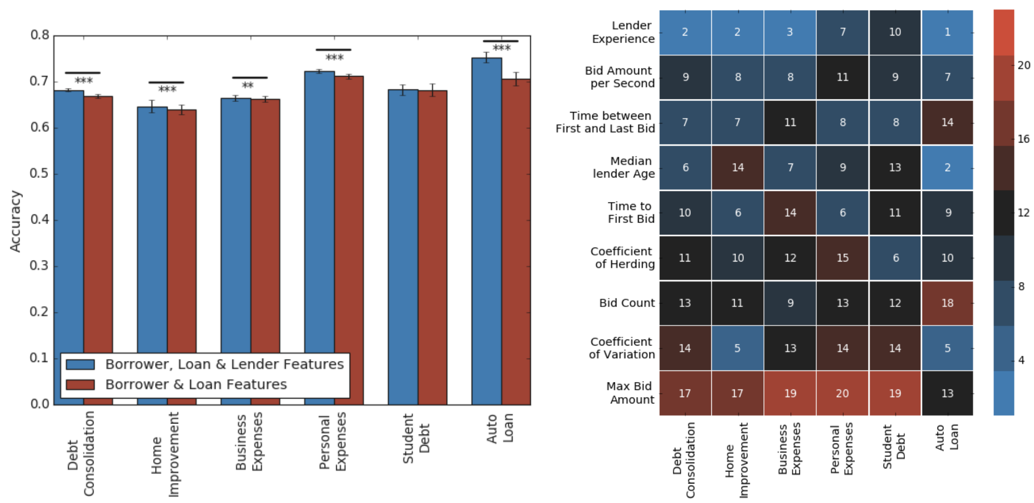

Fig. 1: **Left:** Model accuracy by feature group and project category. Lender features contribute significant gains in estimation accuracy indicative of collective intelligence beyond summarising loan and borrower information in all loan categories but student debt. **Right:** Analysis of feature ranking by project category. Displayed are only the ranks for lending dynamics features. We observe no systematic ranking patterns among lending features which suggests that one can hardly reduce lender characteristics to a single aspect of lender behaviour that is consistent across all loan categories.

the maximum bid amount is the least essential predictor for most loan categories. Note that the analysis in Figure 1 (right) 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. This suggests that one can hardly reduce lender characteristics to a single aspect of lender behaviour that is consistent across all six loan categories.

4. CONCLUSION

The tasks tackled by lenders on the examined platform are 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 is expected to find use also beyond the context of crowd financing. Our contributions were threefold. First, we provided new knowledge about signals deduced from lending behaviour that can contribute to the efficiency of crowd financing in online marketplaces. Second, we added to the growing literature on the wisdom of crowds by providing new insights about different expressions of collective intelligence and potential ways to harvest it. Third, the proposed collective intelligence signals are general and quickly 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. We believe that a better understanding of lender determinants of long-term project success promoted in this work will help improve the efficiency of capital allocation in P2P markets.

ACKNOWLEDGMENTS

The authors would like to thank Brian Uzzi and Jayaram Uparna for providing the data. The work is supported by the U.S. National Science Foundation under Grant No. IIS-1755873.

REFERENCES

Ajay Agrawal, Christian Catalini, and Avi Goldfarb. 2014. Some simple economics of crowdfunding. *Innovation Policy and the Economy* 14, 1 (2014), 63–97.

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.

Gerrit KC Ahlers, Douglas Cumming, Christina Günther, and Denis Schweizer. 2015. Signaling in equity crowdfunding. *Entrepreneurship Theory and Practice* 39, 4 (2015), 955–980.

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*.

Paul Belleflamme, Thomas Lambert, and Armin Schwienbacher. 2014. Crowdfunding: Tapping the right crowd. *Journal of Business Venturing* 29, 5 (2014), 585–609.

Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.

Leo Breiman. 2017. *Classification and regression trees*. Routledge.

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.

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.

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.

Henry K Dambanemuya and Em? ke-Ágnes Horvát. 2019. Harnessing collective intelligence in P2P lending. In *Proceedings of the 10th ACM Conference on Web Science*. 57–64.

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.

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.

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.

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.

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.

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.

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.

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.

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).

Ethan Mollick. 2014. The dynamics of crowdfunding: An exploratory study. *Journal of Business Venturing* 29, 1 (2014), 1–16.

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.

Carlos Serrano-Cinca, Begona Gutierrez-Nieto, and Luz López-Palacios. 2015. Determinants of default in P2P lending. *PloS One* 10, 10 (2015), e0139427.

Silvio Vismara. 2016a. Equity retention and social network theory in equity crowdfunding. *Small Business Economics* 46, 4 (2016), 579–590.

Silvio Vismara. 2016b. Information cascades among investors in equity crowdfunding. *Entrepreneurship Theory and Practice* (2016).

Nir Vulkan, Thomas Åstebro, and Manuel Fernandez Sierra. 2016. Equity crowdfunding: A new phenomena. *Journal of Business Venturing Insights* 5 (2016), 37–49.

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.