
Bandits in Matching Markets: Ideas and Proposals for Peer Lending

Soumajyoti Sarkar*

Arizona State University, USA
soumajyoti.sarkar@asu.edu

Abstract

Motivated by recent applications of sequential decision making in matching markets, in this paper we attempt at formulating and abstracting market designs in peer lending. In the rest of this paper, what will follow is a paradigm to set the stage for how peer lending can be conceived from a matching market perspective with sequential design making embedded in it. We attempt at laying the stepping stones toward understanding how sequential decision making can be made more flexible in peer lending platforms and as a way to devise more fair and equitable outcomes for both borrowers and lenders. The goal of this paper is to provide some ideas on how and why lending platforms conceived from the perspective of matching markets can allow for incorporating fairness and equitable outcomes when we design lending platforms.

1 Introduction

Sequential decision making in two sided markets like consumers and producers has been part of bidding in e-commerce platforms like eBay, eBid for a very long time. Not only that, P2P platforms like Prosper for a very long time allowed lenders to bid on projects for peer microlending until they switched to posted price mechanism Ceyhan et al. [2011]. However, for most peer microlending platforms like Kiva, LendingClub among others, sequential decision making is either not available or limited in its functionality while investor funding cycles generally have a monopoly on who they fund². In this paper, we therefore attempt at abstracting the concept of sequential decision making in two sided markets through a centralized platform for peer lending where both sides have a chance to log their own preferences which are options generally not available in peer lending platforms. What follows then, is the added capability of such centralized platforms to ensure mechanisms to mitigate bias that can be implicit in such platforms Sarkar and Alvari [2020].

As mentioned, in such P2P platforms, there are mainly two sides to the market: the borrowers who want to borrow money from others for their projects or startups and the lenders who lend money to borrowers. The traditional rule has been that these sides involve in two sided trading following the Dutch Auction Mechanism Kumar and Feldman [1998], Wei and Lin [2017]. In platforms like Kiva and Prosper, for borrowing money, a borrower is required to create a listing to solicit bids from lenders by describing a few details about itself - the reason of lending (or the category of the project), the required amount, the minimum interest rate it is willing to borrow money at. If the listing receives more money than its soliciting duration, then the bids with lower rates succeed (win) and the bids with lower rates fail (outbid). So, in case the listing does not receive the full money of its soliciting duration, it would be expired and all its bid from lenders so far will have failed. From the lender's perspective, the main signals of interest then constitute the probability of winning the bid, the probability of the loan being fully funded, as well as the returns from the investment. In keeping

*Work done while author was at Arizona State University

²<https://www.inc.com/christine-lagorio/sam-altman-yc-monopoly.html>

with these expectations, often the borrower’s interests are sidelined as it is assumed that its only expectation from the platform and from investors is to get its project or startup funded. It is also tied to the opinion and expectation that the lenders have no other contribution to the borrowers save the money in the execution of the project. But when it comes to fairness in lending, instead of allowing only one funding cycle to match lenders to borrowers without allowing borrowers to explicitly state their preferences over available lenders can be detrimental to how we perceive what is fair. This is similar to what has been typically the college admissions problem Roth and Sotomayor [1989] where students get to apply to multiple colleges and colleges can select the students they want based on a quota. We perceive the lending mechanism in this paper from a similar mechanism.

So in general, as has been with most microlending internet platforms, supply side agents, in this case the lenders, are faced with multiple choices and the agents must make a choice based on their preferences. However, these preferences often tend to encapsulate implicit bias thus leading to unfair playing fields for the borrower Sarkar and Alviri [2020], Burtch et al. [2014]. With recent advances in abstracting such platforms for matching agents to arms, we formulate this problem of matching lenders to borrowers with certain constraints using a centralized platform but which allows users to exploit their preferences over multiple rounds. For the rest of the paper, we lay the foundation for future work that address sequential decision making in such peer lending platforms with ideas from matching markets. The rest of the paper discusses some choices that could be made towards formulating the utilities or rewards on both sides, the mechanism design for sequential decision making over rounds, and finally incorporating welfare constraints to ensure every borrower is included.

2 Framework for Peer Lending

We model the market platform as a market with 2 sides - the lenders denoted by the set of agents $\mathcal{L} = \{l_1, l_2, \dots, l_N\}$ and the borrowers denoted by the set of agents $\mathcal{B} = \{b_1, b_2, \dots, b_K\}$ and we assume that $N \leq K$ (think of startups applying for VC funding, the number of startups who identify themselves as borrowers far exceeds the number of VC investors). We align our model with traditional market designs previously researched Liu et al. [2020], where the authors study two paradigms - centralized and decentralized. In what follows, we will assume our settings are aligned with the centralized mechanism where the platform decides the choice of matching lenders to borrowers at each time step of a sequential decision making game.

We now have a two-sided matching market where the agents on the borrowing side each have their own funding request proposals and their corresponding requested amount which we denote by c_i , where $b_i \in \mathcal{B}$ (the agents identified by the indices). Similarly, the lenders each have an overall budget q_j , where $l_j \in \mathcal{L}$. We will assume that all these capacities are integers by expressing them in the smallest unit of the currency. In addition, each set of agents on one side of the market have the opportunity to submit their preferred rankings of the agents on the other side of the market to the platform based on the rewards they expect to receive from their matching partner (at the end of the game in retrospect) in this market. But the preference of all these agents may be conflicting and many lenders might prefer to lend to the same borrower, while multiple borrowers may prefer to tie up with the same lenders having specific portfolio and interests. These preferences which were earlier the privilege of the lenders exclusively is now assumed to be available to the borrowers as well.

When borrowers rank the lenders, it can be implicitly based on several factors: the interest rate that a lender assigns to a borrower, the portfolio and interests history of the lenders. Note that the same lender can offer different interest rates to different borrowers and the same borrower can be offered different interest rates by the same lender. In our work, we assume that a lender offers the same interest rates to all the borrowers. Similarly, the set of lenders that the borrowers choose to select and rank depends on the overlap in their project sector and the investors’ portfolio topics. As a first step in simplifying our work, we assume that the borrowers only get to submit their preference of the available lenders once at the beginning of the game based on a deterministic function that we simulate to capture behavioral and interest aspects. We denote the utility that a borrower b_i gets from being funded in part or wholly by lender l_j as $u_i(l_j)$. On the other hand, the lender utilities or “rewards” comprise a deterministic component similar to the borrowers which captures the similarity between the borrowers and its own interests, and a stochastic component which captures the uncertainty associated with the returns from the project (and this is also based off

the interest rate that the lender sets for the borrower). Accordingly, we denote the “true” utility or reward that the lender gets at the end of the completion of the project by $u_j(b_i)$. Note that this true utility cannot be determined over the period of matching since typically, the returns that a lender could get would depend on how well the project is executed, which cannot be determined prior to matching. Note that importantly, in centralized matching platform as ours, the ranking preferences of the lenders and borrowers are not available by design to any of the agents participating in the game.

Additionally, we lay out the following desiderata for our market model that allows us to understand how to design mechanisms for getting the best matches for both the borrowers and the lenders. Each lender l_j can be matched to at most one borrower while each borrower b_i can be matched to multiple lenders based on the capacity c_i requested. That is, we consider the case of many-to-one matching markets Bodine-Baron et al. [2011]. Additionally, in the final assignment list of a borrower to multiple lenders, a correct assignment entails that the sum of the donated amounts of the lenders be at least as much as c of the borrower. Such mechanisms are currently followed in platforms like GoFundMe or Prosper Full Coverage lending model where the borrower only gets the project funded when the sum of amounts lent, match or exceed the requested amount. Also, each borrower and each lender have strict preference rankings of agents on the other side in our work and we assume that there cannot be any ties in the ranking list submitted by the agents on either side of the market.

Each funding cycle consists of a set of borrowers with particular interests and the set of lenders available for the funding. The matching occurs over multiple rounds (or steps) for one funding cycle. At the start of the funding cycle, the list of borrowers and lenders are made available in the platform. The preference orders of the borrowers and the lenders can be captured in the following way:

$$\begin{aligned} b \succ_j b' &\iff u_j(b) > u_j(b') \\ l \succ_i l' &\iff u_i(l) > u_i(l') \end{aligned}$$

The final objective in peer lending is to successfully match a borrower to a set of lenders which happens only when the borrower’s requested amount c_i is matched.

3 Desiderata for matching with welfare constraints

In the rest of the paper, we propose some steps to achieve the goals of meeting both borrower and lender preferences using matching markets and bandit settings.

3.1 Matching objective

In each round of one funding cycle, the objective of the platform is to match borrowers and lenders after obtaining the following: (1) the ranking preferences of the borrowers and the lenders, (2) the fixed utility functions of the borrowers and the revised utility functions of the lenders after each round. So at each round of a funding cycle, the platform solves a multi-objective optimization problem that aims at matching the borrowers and the lenders. The number of rounds can be decided in two ways after this: (1) keeping the number of rounds fixed and the solution obtained at the last round decides the end result for all borrowers and lenders and (2) preemptive removal of borrowers and lenders from further rounds when they get matched successfully.

To decide a matching between \mathcal{B} and \mathcal{L} , we introduce the binary decision variable $\mathbf{x} := (x_{bl})_{(b,l) \in \mathcal{B} \times \mathcal{L}}$ such that $x_{bl} = 1$ if the loan from lender l is assigned and accepted by borrower b and 0 otherwise. It is supposed that when the borrower b and the lender l are matched, the borrower and the lender gain utilities of $u_b(l)$ and $u_l(b)$ respectively (we propose ideas of what these utilities could be in the next section). Accordingly, $u_{bl} = u_b(l) + u_l(b)$. The total utility of a matching $\mathbf{x} \in \{0, 1\}^{|\mathcal{B}| \times |\mathcal{L}|}$ is given by $\sum_{b \in \mathcal{B}} \sum_{l \in \mathcal{L}} u_{bl} x_{bl}$. In our work, we consider a many-one matching where each borrower is matched to multiple lenders and each lender matched to a single borrower. Defining a binary decision variable $\mathbf{w} := (w_{b,l}) \in \{0, 1\}^{|\mathcal{B}| \times |\mathcal{L}|}$, the matching objective can be formulated as:

$$\text{maximize} \quad \lambda_1 \sum_{b \in \mathcal{B}} \sum_{l \in \mathcal{L}} u_{bl} x_{bl} - \lambda_2 \sum_{b \in \mathcal{B}} \sum_{l \in \mathcal{L}} w_{bl} \quad (1)$$

$$\begin{aligned} \text{subject to} \quad & \sum_{l \in \mathcal{L}} x_{bl} q_l \geq c_b \quad (\forall b \in \mathcal{B}) \\ & x_{bl} \leq 1 \quad (\forall l \in \mathcal{L}) \\ & c_b x_{bl} + c_b \sum_{l' >_b l} x_{bl'} + \sum_{b' >_l b} x_{b'l} \\ & \geq c_b (1 - w_{bl}) \quad (2) \\ & x_{bl} \in \{0, 1\} \quad (\forall b \in \mathcal{B}, \forall l \in \mathcal{L}) \\ & w_{bl} \in \{0, 1\} \quad (\forall b \in \mathcal{B}, \forall l \in \mathcal{L}) \end{aligned}$$

The reader can refer to such formulations of matching Abeledo and Blum [1996] for further reading. Briefly these constraints satisfy the following: (1) the borrower's requested amount must exceed the sum of investments from matched lenders (2) the lenders can only be matched to one borrower and (3) the number of blocking pairs should be minimized in accordance with the original stable matching constraints Roth et al. [1993]. However under our additional constraints, stability is not guaranteed.

3.2 Welfare constraints

The concept of what constitutes a fair outcome in peer lending under conditions in our design has been closely studied previously Chen and Ghosh [2011], Bodine-Baron et al. [2011]. In our framework, we can design efficient utility functions u_b and u_l for the borrower and the lender respectively that allow for matching respecting desired outcomes of fairness. As an example, when the borrowers and lenders' interests overlap significantly as well as the expertise of the lenders collides with the borrower's projects, we can assign higher utilities for these pairings even though their preference rankings might be different.

3.3 Bandit Setting

Finally, we end this paper describing a couple of possible ways to incorporate sequential decision making that allows for lenders to exploit their preferences in uncertain environments. We want to reiterate that the matching occurs repeatedly for multiple rounds - in each round the lender has multiple borrowers to rank, however the rewards from the borrower (which depends on how much the borrower expects to make out of the project) is not available to the lenders. This is bandit setting Das and Kamenica [2005] where at each round, the platform provides a pseudo-reward to the lender, that comes from the borrower it is matched to in that round. This pseudo reward can be a machine learning model outcome based on borrower history and the attributes of the project and is one of the caveats of a centralized platform that allows for more fair outcomes.

The two ways to design a sequential decision framework are as follows:

- At each round, the lender ranks the borrowers based on a Upper Confidence Bound Liu et al. [2020] and as the lenders play more rounds, their regrets are attuned accordingly. So while the borrower rankings of the lenders are fixed at every round, the lenders get to decide their preferences at every round until the matching ends. This ensures sufficient opportunities for exploitation of the arms (borrowers).
- At each, the lender ranks the borrowers based on a Upper Confidence Bound, however once a lender is matched to its top ranked borrower and the borrower is also matched to its top ranked lender, the lender is removed from the game. This preemptive setup allows other lenders more chances to revise their rankings in a more certain manner.

3.4 Final Thoughts

We consider a case of a centralized matching platform which requests proposals from borrowers about their preferences over the lenders or the agents through their personal rankings of the lenders. Then the platform decides the matching by allowing the lenders to interact with the platform over multiple time steps by either accepting or rejecting the assigned borrower at a time step. From a lender perspective, this schema thus allows them to get matched without having the information of the actual returns while allowing them certain flexibility to exploit their options. Bandits in

crowdsourcing platforms have been used before Biswas et al. [2015] albeit not in a lending scenario and in recent years, researchers have focused on mechanism design in bandit settings as well Kandasamy et al. [2020]. The goal of this paper has been to lay out some ideas in which centralized peer lending platforms can be abstracted from a matching market perspective and how bandits could play a role in mechanism design. These matching markets allow for more privacy as well as ensuring equitable outcomes and in our situation can be achieved by designing proper utility functions of the agents. Similarly in future one could design decision making in which lenders can elicit information about their peer choices as well as networks have been known to aid funding situations Horvát et al. [2015].

References

- H. Abeledo and Y. Blum. Stable matchings and linear programming. *Linear algebra and its applications*, 245: 321–333, 1996.
- A. Biswas, S. Jain, D. Mandal, and Y. Narahari. A truthful budget feasible multi-armed bandit mechanism for crowdsourcing time critical tasks. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1101–1109, 2015.
- E. Bodine-Baron, C. Lee, A. Chong, B. Hassibi, and A. Wierman. Peer effects and stability in matching markets. In *International Symposium on Algorithmic Game Theory*, pages 117–129. Springer, 2011.
- G. Burtch, A. Ghose, and S. Wattal. Cultural differences and geography as determinants of online prosocial lending. *Mis Quarterly*, 38(3):773–794, 2014.
- S. Ceyhan, X. Shi, and J. Leskovec. 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*, pages 547–556, 2011.
- N. Chen and A. Ghosh. A market clearing solution for social lending. In *IJCAI*, volume 11, pages 152–157, 2011.
- S. Das and E. Kamenica. Two-sided bandits and the dating market. In *IJCAI*, volume 5, page 19, 2005.
- E.-Á. Horvát, J. Uparna, and B. Uzzi. 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*, pages 226–233, 2015.
- K. Kandasamy, J. E. Gonzalez, M. I. Jordan, and I. Stoica. Mechanism design with bandit feedback. *arXiv preprint arXiv:2004.08924*, 2020.
- M. Kumar and S. I. Feldman. Internet auctions. In *USENIX Workshop on Electronic Commerce*, volume 3, pages 49–60, 1998.
- L. T. Liu, H. Mania, and M. Jordan. Competing bandits in matching markets. In *International Conference on Artificial Intelligence and Statistics*, pages 1618–1628. PMLR, 2020.
- A. E. Roth and M. Sotomayor. The college admissions problem revisited. *Econometrica: Journal of the Econometric Society*, pages 559–570, 1989.
- A. E. Roth, U. G. Rothblum, and J. H. Vande Vate. Stable matchings, optimal assignments, and linear programming. *Mathematics of operations research*, 18(4):803–828, 1993.
- S. Sarkar and H. Alvari. Mitigating bias in online microfinance platforms: A case study on kiva. org. *ECML PKDD SoGood*, 2020.
- Z. Wei and M. Lin. Market mechanisms in online peer-to-peer lending. *Management Science*, 63(12):4236–4257, 2017.