Games and Economic Behavior 86 (2014) 367–391

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Auctions for social lending: A theoretical analysis ☆

Ning Chen[a], Arpita Ghosh[b,*], Nicolas S. Lambert[c]

[a] *Division of Mathematical Sciences, Nanyang Technological University, Singapore*
[b] *Department of Information Science, Cornell University, NY, USA*
[c] *Graduate School of Business, Stanford University, 655 Knight Way, Stanford, CA 94305, USA*

ARTICLE INFO

Article history:
Received 27 June 2011
Available online 23 May 2013

JEL classification:
D44
D82

Keywords:
Social lending
Peer-to-peer lending
Auctions
Mechanism design

ABSTRACT

Prosper, today the second largest social lending marketplace with nearly 1.5 million members and $380 million in funded loans, employed an auction mechanism amongst lenders to finance each borrower's loan until 2010. Given that a basic premise of social lending is cheap loans for borrowers, how does the Prosper auction do in terms of the borrower's payment, when lenders are *strategic agents* with private true interest rates? We first analyze the Prosper auction as a game of complete information and fully characterize its Nash equilibria, and show that the uniform-price Prosper mechanism, while simple, can lead to much larger payments for the borrower than the VCG mechanism. We next compare the Prosper mechanism against the borrower-optimal auction in an incomplete information setting, and conclude by examining the Prosper mechanism when modeled as a dynamic auction, and provide tight bounds on the price for a general class of bidding strategies.

1. Introduction

Social lending, or peer-to-peer lending, is an emerging alternative to banks and personal loans, allowing individuals to lend or borrow money to each other directly without the participation of traditional financial intermediaries. Social lending offers borrowers the opportunity to obtain loans at lower interest rates and costs, and lenders with an opportunity for investments with higher rates of return than from banks or other common alternatives. Since its beginning in 2005, social lending has grown to become a major business on the Internet: The total amount of money borrowed using such peer-to-peer loans is projected to reach over $5 billion in 2013, and the total amount funded by the two leaders Prosper and Lending Club increased from $200 million in 2008 to over $1 billion in 2012 in the US market alone.[1] Given the large volume of trade, evidenced by the large number of users and the vast sums of money being lent and borrowed, social lending is a significant component of electronic commerce. This paper studies the auction mechanism used to finance loans by Prosper, one of the largest such marketplaces on the Web.

☆ Part of this work was done while the authors were employed by Yahoo! Research. Preliminary results were presented at the Tenth ACM Conference on Electronic Commerce ["Social Lending", by N. Chen, A. Ghosh, and N.S. Lambert, in: Proc. of the 10th ACM Conference on Electronic Commerce, pp. 335–344]. We are grateful to Preston McAfee, John Morgan, Michael Ostrovsky, David Parkes, Tuomas Sandholm, and to the reviewers and participants of the Tenth ACM Conference on Electronic Commerce for helpful comments and discussions. We are also very thankful to two editors and two anonymous referees for their detailed suggestions which helped to greatly improve the paper. Finally we thank Prosper Marketplace, Inc. for providing auction data. Chen's research is supported by the AcRF Tier 2 grant of Singapore (No. MOE2012-T2-2-071), and Lambert gratefully acknowledges the support of Google Research and the National Science Foundation under grant No. CCF-1101209. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

* Corresponding author.
E-mail addresses: ningc@ntu.edu.sg (N. Chen), arpitaghosh@cornell.edu (A. Ghosh), nlambert@stanford.edu (N.S. Lambert).

[1] Sources: http://www.gartner.com/it/page.jsp?id=1272313 and http://techcrunch.com/2012/05/29/.

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2013.05.004

Prosper is the first social lending site launched in the United States, and is today the second largest site (behind Lending Club) with nearly 1.5 million members and $380 million of funded loans.[2] From its creation in 2005 to 2010, Prosper, which described itself as an "eBay for loans", auctioned off loans amongst interested lenders, using competition amongst lenders to bring down the final interest rate for the borrower.[3] In the Prosper mechanism, borrowers create loan listings, specifying the amount of money they are willing to borrow and a reserve interest rate, which is the maximum rate they are willing to accept. Lenders choose individual listings to bid on, specifying the amount they are willing to lend and their desired interest rate for each loan. In addition to standard criteria such as credit categories and histories, lenders can also consider a borrower's personal story, endorsements from friends, and group affiliations. The bidding starts at the reserve rate, and lenders can then bid down the interest rate in an auction. When the auction ends, Prosper combines the bids with the lowest interest rates into a single loan to the borrower and handles all loan administration tasks including loan repayment and collections on behalf of the matched borrower and lenders. Each time a loan request is successfully completed, Prosper charges a transaction fee to the borrower and a servicing fee to the lenders, in equal proportion.

A basic premise of social lending is cheap loans for borrowers. To what extent does the mechanism used by Prosper to fund loans lead to small borrower payments, *given that lenders act strategically, as selfish and rational agents*? Each lender has a *private interest rate*, which is the minimum interest rate at which she is willing to invest in a particular loan—the rate equivalent to her best outside option. While the Prosper mechanism certainly selects the lenders with the lowest interest rates to finance the loan, this does not necessarily lead to a cheap loan—the rate reported by a lender need not be her true private interest rate, since bidding a higher rate might lead to a better return. (That lenders are strategic is clearly evidenced by their behavior in Prosper: Lenders are allowed to, and indeed do, decrease their rates to increase their allocation through the course of the auction.) Given that lenders behave strategically, how does the choice of mechanism—consisting of an allocation and payment rule—affect the total payment of the borrower?

1.1. Overview of results

We first provide a complete analysis of the Nash equilibria of the Prosper mechanism modeled as a one-shot auction game of complete information (Section 3), which turns out to be a fairly curious mechanism. The Prosper auction is VCG-like, but not quite VCG—it is, in fact, a uniform-price mechanism obtained by applying VCG to a *modified instance* of the problem, as described below. Suppose a borrower wants to borrow an amount D, and each lender i specifies her budget a_i and her offered interest rate b_i. Replace every lender i by a_i dummy lenders with budget 1 and interest rate b_i each. Now run VCG on this new instance to determine the winners and their payments. Recall that the VCG mechanism to buy k identical items from competing sellers (here, $k = D$ and each item is a unit of money) buys from the k cheapest sellers and pays each of them the same price, which is the bid of the $(k+1)$-th lowest bid. Thus, applying VCG to the modified instance yields a solution where all winning lenders receive the same interest rate, which is either the bid of the first loser or the last winner, depending on whether or not the last winner exhausts her budget. Since the Prosper mechanism is a uniform-price mechanism (*i.e.*, every winner receives the same interest rate, also called the *price*), and we are interested in the borrower's payment, we focus on the set of possible prices that can arise in a Nash equilibrium. We show how to completely characterize the equilibria of the Prosper mechanism, by characterizing the set of possible prices and last winners (Section 3.1) when losers are restricted to bid their true interest rate.[4] The characterization of the equilibria can be sharpened further if we restrict ourselves to equilibria where winners do not bid less than their true interest rates (Section 3.2).

Next we use this characterization to compare the Prosper mechanism against the VCG mechanism from the perspective of the borrower's payment (Section 4). Since the Prosper mechanism is not incentive-compatible, we compare the payment in the *best* and *worst* Nash equilibria of the Prosper mechanism against that of the VCG outcome. While neither mechanism dominates the other, we show that the VCG mechanism leads to a payment that is always within a factor of $O(\log D)$ of the cheapest Nash equilibrium of the Prosper mechanism, whereas even the cheapest Nash equilibrium of the Prosper mechanism can be as large as a factor D of the VCG payment (both factors are tight). A similar result holds for the worst Nash equilibrium of the Prosper mechanism.

We next analyze the Prosper auction under alternative information structures. We first investigate an incomplete information setting. Under some assumptions on the joint distribution of budgets and rates, we address the question of deriving the Bayesian-optimal auction for the borrower, and discuss how such an optimal auction differs from the Prosper auction (Section 5). Finally, we examine the Prosper mechanism when modeled as a dynamic auction, and provide tight bounds on the price for a general class of bidding strategies (Section 6).

[2] Data as of August 2012. Source: http://www.prosper.com/.

[3] In 2011 Prosper moved from the auction mechanism to a pre-set rate mechanism in which loans are allocated among lenders on a first-come, first-serve basis. This rate is set by Prosper for each listing as a function of personal information available from the borrower, such as credit scores, debt-to-income ratio, and transaction history.

[4] It turns out that without this restriction, computing the Nash equilibria with the smallest and largest prices is, in general, NP-hard, and hard to approximate within any reasonable factor.

The appendices are organized as follows. (i) Appendix A investigates the complexity of equilibrium computation when bid profiles are not restricted; (ii) Appendix B compares the Prosper auction to other uniform-price mechanisms, and (iii) Appendix C contains the omitted proofs.

1.2. Related work

The literature on social lending thus far has largely focused on empirical studies. Freedman and Jin (2011) study the loan-level data of Prosper auctions over a 2-year period, establishing relations between interest rates, actual returns, default rates, and credit grades. They show that the informational limitation due to the non-disclosure of credit scores is partially compensated by reputation effects generated via the Prosper social network. In a similar fashion, Iyer et al. (2010) examine the extent to which the lenders that participate in Prosper can infer borrowers' creditworthiness by comparing actual credit scores to the interest rates of funded loans. This literature does not attempt to model or analyze the mechanisms used by social lending platforms from a theoretical standpoint. Chen and Ghosh (2011) model the Zopa online lending market, which allows lenders to only specify acceptable borrower *categories* rather than bid for individual borrower's loans, as a two-sided matching market, and investigate the question of what is a computationally feasible, good allocation to clear this market if lenders are additionally allowed to express preferences across borrower categories. This work differs significantly in the market model (due in large part to the differences between Prosper and Zopa) and consequently the nature of the questions addressed—among other differences, Chen and Ghosh (2011) investigate allocations among multiple borrowers and lenders assuming *non-strategic* agents, and address the problems of what constitutes a fair allocation and the algorithmic problem of computing such a fair allocation, whereas our work analyzes strategic behavior in the auction actually deployed by Prosper for a single borrower's loan.

There is an extensive body of work in the economics literature on auctions in general and multi-unit auctions in particular (for a review see, for example, Part II of Krishna, 2010); we discuss the relevant work from this literature next. The auction used by Prosper is a variant of multi-unit auction with uniform price and identical items—and, since different lenders offer to lend different amounts to a borrower, we must consider a multi-unit demand (or rather, supply) auction. In multi-unit auctions, the precise format of the auction matters, even in a private-value setting. While it would be tempting to use the well-known properties of, say, the multi-unit Vickrey auction, and use revenue equivalence to derive the borrower's payment in the Prosper auction, this approach has the difficulty that while the Revenue Equivalence Theorem continues to hold in multi-unit auctions, the revenues of two auctions are only guaranteed to be the same (up to a constant) when the items are allocated to the same bidders. Unfortunately, as opposed their single-unit analogs, multi-unit auctions typically do not allocate units in the same way, even when looking at the common formats such as Vickrey, discriminatory, or uniform-price auctions. For this reason, the specificities of the Prosper auction matter, and the analysis does not reduce to that of a standard auction model.

In unit-demand auctions where multiple items are sold but bidders are each interested in only one item, the mechanics of uniform-price auctions are generally well understood. However the case of multi-unit demand turns out to be a great deal more difficult, and equilibrium analysis of uniform-price auctions with multi-unit demand has only been carried out in special cases. Noussair (1995) and Engelbrecht-Wiggans and Kahn (1998) consider the case where bidders demand at most two units and derive equilibrium behaviors in the sealed-bid private-value setting. In their version of the auction the clearing price is set to be the first-rejected bid. In the same setting, Draaisma and Noussair (1997) derive equilibrium conditions for a different version of the auction, where the price is the last-accepted bid. Still with two-unit demand, Bresky (2009) analyzes the impact of a reserve price on the efficiency and revenue of a uniform-price auction. For general multi-unit demand, Nautz (1995) derives optimal bidding strategies under the assumption that bidders act as price takers. Swinkels (2001), Chakraborty and Engelbrecht-Wiggans (2005) and Katzman (2009) examine properties of the auction outcome in asymptotic cases. Ausubel and Cramton (2002) investigate efficiency properties of the sealed-bid uniform-price multi-unit auction, and describe the general phenomenon of "demand reduction": To pay less for the items won, bidders may prefer to shade their bids in order to lower the market-clearing price, thereby generating inefficiency and reducing the seller's revenue, an important and common phenomenon in multi-unit auctions. Weber (1997) and List and Lucking-Reiley (2000) discuss demand reduction from an empirical and experimental standpoint in sealed-bid uniform-price auctions.

The Prosper auction is also related to multi-item *ascending-price* auctions, the open-format equivalent of the uniform-price auctions, since Prosper releases partial information regarding bids during the auction process. In these ascending-price auctions, the seller starts by announcing a low per-unit price for the items. Every buyer indicates how many units she wants to purchase at that price. The seller then raises the price until supply equals demand, that is, until the number of items demanded equates the number of items auctioned off. There is a single price for these items, set at the market-clearing price (Ausubel, 2003; Krishna, 2010). Ausubel and Schwartz (1999) perform an equilibrium analysis of a model of ascending-price auctions in which two bidders place bids alternatively until the market clears, in a complete information setting. Using a backward induction argument, they show that bidders are best off reducing their demand to the market-clearing condition, so that the market clears at the first round of trading, and the outcome is efficient but generates low revenue. Under similar assumptions, Grimm et al. (2003) and Riedel and Wolfstetter (2006) characterize the unique equilibrium in the open ascending-price dynamic auction game, with multiple rounds of bids but assuming bidders bid simultaneously in every round. They reach a similar conclusion: The market clears at the first round of bids, where every bidder will have reduced demand so as to attain the market-clearing condition. Jun and Wolfstetter (2004) study an instance of the auction under incomplete information. They show that the result of Riedel and Wolfstetter continues to hold with two bidders, two types, and two units. Kagel and Levin (2001) present experimental evidence of demand reduction in both sealed-bid and open formats. Brusco and Lopomo (2008) examine the impact of budget constraints on the incentives for demand reduction and its consequences on social welfare. Brusco and Lopomo (2002) investigate collusion in general multi-unit ascending-price auctions, and show existence of collusive equilibria where bidders share the objects for sale and keep a low price. Aside from the literature on standard ascending-price auctions, Ausubel (2004) introduces a variant that is outcome-equivalent to the Vickrey auction, and characterizes equilibria under complete and incomplete information. Bikhchandani et al. (2011) consider a matroid setting, and show that a dynamic ascending auction can converge to the Vickrey outcome when a greedy algorithm can solve the winner determination problem. However, Prosper gives all winning lenders the same rate, *i.e.*, is restricted to a uniform-price auction, which is not guaranteed by the Vickrey outcome; a major focus of our study is therefore on the characterization of Nash equilibrium.

There is also a growing body of literature devoted to core-selecting auctions, a family of combinatorial auctions designed for its appealing properties in heterogeneous-item environments with complementarities (Milgrom, 2006; Day and Raghavan, 2007; Cramton and Day, 2008; Day and Milgrom, 2008). Day and Milgrom (2008) analyze core-selecting auctions in terms of "truncation strategies". They show that the strategy profiles in which bidders misreport their true valuations by shading their bids uniformly (*i.e.*, where a bidder lowers her bid on all packages by some uniform number) form Nash equilibria under complete information, a result that applies to all core-selecting auctions. Sano (2013) extends these results to the Vickrey-reserve auctions, a larger class of efficient, individually rational auctions in which bidders never pay less than what they would pay in a Vickrey auction. It is easily seen that the Prosper auction discussed in this paper is core-selecting. However, truncation strategies are generally not possible in social lending auctions given the restricted bidding format, which only allows the specification of a budget and interest rate. We therefore describe analogous Nash equilibria where those existing results cannot be applied. In both cases, the results indicate situations in which a specific bidder can, at equilibrium, shade her bid down to the Vickrey-price-setting bid of a (truthful) loser.

In spite of its apparent similarities, the Prosper auction is quite distinct from these existing uniform-price auction formats. In consequence our analysis differs from the above works along several dimensions. First, in a large part of our paper, we study a one-shot auction model under complete information. We choose this approach because an equilibrium analysis of the dynamic Prosper auction appears difficult, even under a complete information setting. In contrast, the previous studies on uniform-price auctions do not combine these two features. Another difference concerns the form of bids that can be submitted—in the aforementioned sealed-bid auction models, bidders bid a full demand curve. In our model, however, a bid is a demand of (actually, an offer to supply) some quantity at some price, and each bidder submits a single bid to prevent them from announcing a demand curve by an appropriate bid combination (we do so for tractability reasons, since equilibrium analysis where bidders announce a demand curve is known to be intractable). Finally, we are also particularly interested in *worst-case* comparisons between the borrower's payment in different mechanisms, an approach that, until recently, has seldom been used in economic analysis.

Finally, our work is also related to the work on *frugal* mechanism design in the computer science literature, which seeks mechanisms with small payments for the buyer in a reverse auction. In "*hiring a team*" problems (Archer and Tardos, 2002; Talwar, 2003; Elkind et al., 2004; Immorlica et al., 2005; Karlin et al., 2005; Chen and Karlin, 2007), a principal wants to hire a team of selfish agents at a low cost to perform a task, where each agent has a private cost for performing her sub-task. Only feasible teams are able to complete the task. In the context of social lending, one can consider lenders as agents, and a feasible team is simply one whose total budget is greater than or equal to the borrower's demand. However, our work differs significantly from these *hiring a team* problems in that the system of feasible sets in our setting is quite different from those considered in the frugality literature, and the existing results do not apply to the feasible sets in our social lending setting. Also, rather than attempt to derive the optimal incentive-compatible frugal mechanism, we examine the most commonly used social lending mechanism, and compare it with other natural alternatives. To the best of our knowledge, this is the first paper to propose a theoretical analysis of the auction mechanisms used in social lending.

2. Preliminaries

2.1. The Prosper marketplace

To motivate the model of social lending auctions used throughout the paper, we start with a brief description of the actual mechanism employed by Prosper until 2010.[5] Auctions are initiated by members who want to borrow money (hereafter borrowers), which can be any amount from \$1000 to \$25,000. The term of a loan can be one, three or five years. Borrowers also specify a reserve interest rate, which is the maximum interest rate above which they are not willing to borrow. This rate could be, for example, the interest rate that would apply if the borrower were to obtain the funds from a bank. By convention all rates are annualized and net of fees.

The normal duration of the bidding process is fourteen days, unless the borrower closes the auction early. Members who wish to lend money (hereafter lenders) can participate in the auction by bidding the amount they wish to lend to this

[5] The Prosper interface is described on the company's Web site and the prospectuses filed at the Securities and Exchange Commission.

particular borrower (which can be any amount greater equal $25), together with an interest rate, which is the rate they seek to receive from the loan.

Prosper auctions are dynamic, semi-open auctions, in the sense that bids are submitted sequentially and Prosper releases bidding information through the course of the auction. Specifically, every time a bid is received, Prosper recomputes the auction outcome—the rates and amounts allocated to lenders—*as if* the auction were to close at that time. Given a set of bids, Prosper allocates the loan amount requested by the borrower to the bidding lender with the lowest interest rate. If the amount bid does not cover the entire loan, Prosper allocates what remains of the loan to the bidding lender with the second-lowest rate, and so on until the entire amount has been allocated, or until no bidding lender remains. In case of ties, bids placed earliest in time takes precedence over later bids. Once allocations are determined, Prosper sets a single loan rate that applies to all winning lenders (*i.e.*, lenders who receive a positive allocation). This loan rate is set to be the interest rate demanded by the winning lender whose bid amount is only partially allocated (*i.e.*, the lender who bid the highest interest rate amongst the winning bids but did not lend the full amount he bid), or, if there is no winning lender with such a partial allocation, to the next highest rate which is the lowest interest rate among the losing bids (or, if neither applies, to the reserve rate set by the borrower). The auction outcome is announced publicly at all times. In addition, all information about losing bids is public, while for winning bids, Prosper displays only the amount bid and lender identity (the interest rate requested by the winning bidders remains undisclosed). Lenders who are outbid can compete by placing new bids, as long as the rate bid is at least 5 basis points (0.05%) below the current loan rate. Note that while a lender can submit any number of bids in the course of the auction, these cannot be withdrawn: once placed, bids are irrevocable.[6]

Once the auction closes, the loan is made provided the bids received cover the full amount requested by the borrower. Lenders transfer the funds to the borrower, from which Prosper deducts transaction fees. The borrower, in turn, sends repayments to the lenders, from which also Prosper deducts servicing fees.[7] Repayments are made on a monthly basis and are calculated so as to amortize the loan amount, over the term of the loan, and at the loan interest rate.

2.2. Modeling the Prosper auction

In most of this paper we consider a simple static model of the Prosper auction in a complete information setting, which offers a tractable alternative to the dynamic auction game. We note here that it is common to analyze open auctions under complete information (for example, the analysis of Ausubel and Schwartz, 1999; Grimm et al., 2003 and Riedel and Wolfstetter, 2006 assumes complete information). Complete information reflects the idea that agents learn through interactions, and in the semi-open auction format used by Prosper, it is arguably a better choice of information structure—as claimed, for example, by Varian (2007) and Edelman et al. (2007) who observe a similar situation regarding auctions for online advertising. The claim is supported, in part, by the fact that even with minimal knowledge about one's opponent, strategies of rational players in a repeated game converge to a Nash equilibrium of the one-shot game (Kalai and Lehrer, 1993). Here our auction game can be interpreted as a limit case of the dynamic Prosper mechanism, in which lenders will have learned entirely the private information of their opponents by the end of the 14-day period, and submit their last bid all at once strategically given full information. We extend this static model in Section 5 to an incomplete information setting, and in Section 6 to a dynamic setting.

To support some of our modeling assumptions, we retrieved a dataset provided by Prosper and extracted bidding data of all the auctions that took place since the company's debut until the abandonment of the auction model.[8] The data extracted consists of 28,135 auctions and 4.7 million bids. All Prosper statistics that we refer to in the remainder of this section are derived from this data.

We focus the analysis on a single auction for a loan of one year with a single repayment at the end of the term, with participating lenders who have private interest rates but fixed, publicly known amounts, or budgets, that they are willing to lend to this borrower. A possible interpretation of the fixed-budget assumption is that lenders do not behave strategically with respect to the amounts they bid in the course of the auction—every lender decides the amount he is willing to lend in this auction once and for all, and always bids this amount in every subsequent round. Because Prosper uses a semi-open format which makes public all the amounts bid, lenders will have learned each other's budgets after the first round of bids. Budgets then become common knowledge among the participants. The assumption does not preclude that the amount first bid be strategically chosen, for example as a function of the borrower's information, the lender's portfolio and the risk involved. What is important is that this amount becomes known and fixed through the auction process.

That lenders do not behave strategically with respect to budgets allows considerable simplifications and makes the analysis tractable. Yet the assumption is not innocuous. The major reason why lenders might engage in strategic behaviors is the phenomenon of demand reduction described by Ausubel and Cramton (2002). Ausubel and Cramton examine the case

[6] Because Prosper considers every bid independently, in theory lenders need not wait to be outbid to submit new bids, but in practice such cases are rare.

[7] Fees are in percentage of the loan amount. In 2010, a transaction fee of 3.0% of the loan amount was charged to the borrower, and a servicing fee equal to an annualized rate of 1.0% of the outstanding principal balance was deducted from each lender's share of loan repayments.

[8] The dataset is publicly accessible at http://www.prosper.com/about/academics.aspx. We extracted the auctions which (i) were successfully completed and ended up in a funded loan, and (ii) which had at least one losing bid. This second condition allows us to separate the auctions from Prosper's pre-set rate alternative which was first made available in 2009, as the dataset does not distinguish between the two.

of sealed-bid uniform-price auctions to sell multiple identical units of an item to bidders. Under fairly general conditions and incomplete information, they show that every Bayesian–Nash equilibrium yields an inefficient outcome with positive probability, because some bidders have incentive to bid a reduced demand. The idea is that by reducing demand, a bidder can also reduce the price she pays for the items won. For example, a bidder with a demand for two units which she values the same can end up bidding her true value for the first unit, but bid less for the second unit, in effect bidding a reduced demand function. What drives the result is that while shading her bid reduces the chances to be allocated the second unit, it also increases the probability of a low price for the first unit. Demand reduction is also present in the open-format counterpart of the sealed-bid uniform-price auction, the open ascending-price auctions (Ausubel and Schwartz, 1999; Grimm et al., 2003; Jun and Wolfstetter, 2004; Riedel and Wolfstetter, 2006).

Demand-reduction strategies can also be profitable in the Prosper (reverse) auction. Recall that in the actual Prosper auction, a bid is simply an amount and an interest rate. In principle, Prosper does not prevent a lender from submitting several bids for different amounts and interest rates, so that a lender can essentially bid a full demand (actually supply) schedule all at once via an appropriate combination of bids, as in the sealed-bid auctions considered in Ausubel and Cramton (2002). Even if limited to placing a single bid, reducing one's budget can potentially increase a lender's profit—by reducing the amount, a lender generally increases the final loan rate because more lenders will be needed to clear the loan so that the auctioneer will have to look further up for its last winner or first loser. If that increase is significant enough, it may compensate for the loss of revenue due to lesser demand.

Because we assume that lenders do not attempt to manipulate their budget, our framework does not capture a potential loss of utility due to demand reduction. Instead we explain the borrower's payment via competition on interest rates. It is worth noting that, according to bidding data, demand-reduction strategies seldom occur in Prosper auctions. Indeed, we observe that a lender wins with multiple bids about 2% of the time, so that lenders submit a single winning bid in the vast majority of the cases. This indicates that lenders do not submit multiple bids at once, rather, they submit a new bid only after being outbid. This also helps justify that in our model, each lender places a single bid.

A possible explanation is that, in the context of social lending, the cost of employing such demand-reduction strategies outweighs the benefits. From a lender's perspective, Prosper is an investment platform that offers a large number of loans to invest in. While Prosper limits every borrower to a maximum of two loans at a time, lenders can participate in as many auctions as they want. If lenders are at all risk-averse, they will prefer to diversify their risk through a portfolio of small loans. Indeed, the data shows that on average, a lender participates in about 70 auctions. If a strategy is costly in terms of time and effort, then that cost is multiplied by the number of auctions. Demand-reduction strategies can be costly, because they require a careful assessment of an indirect price impact. And given the low amounts bid on individual loans—on average, lenders bid \$80 per auction—the cost of lender participation is not negligible. In fact, as it can be cumbersome to scrutinize so many auctions, Prosper even provides an automated portfolio plan that facilitates the process and places bids for the lender.[9] On the other hand, demand-reduction strategies and the like may not be very profitable, since these strategies are most effective when (1) the lender bids a sufficiently large amount and (2) the competition is restrained. To take advantage of risk-sharing lenders should, and do bid small amounts for each loan—indeed, over 95% of the bids are for an amount of \$200 or less—and for that same reason many lenders should share the loan—and again, on average, a borrower requests \$6200 and the loan is divided among 85 lenders. On the contrary, strategies for bidding interest rates need not be complex, and as we will see in the next sections, there are indeed simple equilibrium strategies that can generate relatively large revenue for the lenders, when compared to the Vickrey auction.

If lenders choose not to strategize over budgets to increase their average profits, they might still want to do so to manage their risk. Individuals are risk neutral for small amounts but become risk-averse as the stakes increase. To the extent that rates are an indicator of loan quality, it may be advisable to lend different amounts when the interest rate is high, say 25%, compared to when the interest rate is low, say 5%. Our model remains compatible with this behavior, because most of the risk incurred can be assessed at the moment of deciding whether to participate in the auction, via the information available from the borrower. So even if lenders may want to vary the amount they bid *across* auctions, they may want to bid a fixed amount *within* auctions. The data supports this, demonstrating that large differences in rates are unlikely to appear within the same auction—the average difference between the highest rate bid over the 14-day period and the loan rate at the end of the period in Prosper is less than 3 percentage points.

Formal model. We now present a formal model for an auction in the Prosper marketplace. There is a single borrower who wants to borrow an amount of money D, hereafter referred to as the *demand*, and specifies a *reserve interest rate* R. Both R and D are publicly known before the auction starts. Multiple lenders, denoted by $L_1, \dots, L_n$, compete to finance the loan. Each lender L_i specifies her *budget* a_i, which is the amount she's willing to lend, and her *bid* b_i, which is the interest rate she seeks from her loan. The demand D and the budgets a_i are integers, *i.e.*, they are expressed in cents (or the smallest unit of currency). Each lender is limited to a single bid. We assume that the budgets a_i are exogenously fixed and common

[9] The system selects the auctions to participate in and the budget to bid as a function of criteria specified by the lender; these criteria concern borrower characteristics such as credit score and income level. The system tries to sustain the highest possible bids without losing to the competition.

knowledge. We also restrict attention to the case where no lender is indispensable for the borrower's loan to be funded in full, so that no lender can "clinch" part of the loan before any bidding takes place[10]: for any j, $\sum_{i \neq j} a_i \geqslant D$.

A *mechanism* for this setting computes an *allocation* and "*price*" for each lender, given the lenders' budgets $a_1, \ldots, a_n$, and bids $b_1, \ldots, b_n$. The allocation for lender L_i, $0 \leqslant x_i \leqslant a_i$, is the amount borrowed from lender L_i, and the price[11] p_i is the effective interest rate[12] at which the lender will be paid back by the borrower. As in the actual Prosper auction, we require that the total allocation exactly funds the loan, that is, $\sum_i x_i = D$. To ensure voluntary participation, we will require that $p_i \geqslant b_i$. We say that lender L_i is a *winner* if she receives a positive allocation $x_i > 0$.

We suppose that every lender L_i has a private true interest rate r_i, which is the interest rate equivalent to her best outside option. That rate r_i makes lender L_i the indifferent between getting her best outside option or investing her money in a Prosper loan with rate r_i.

Lenders are rational, which means they act to maximize their *utility* given their true interest rates. The utility of lender L_i is defined as[13]

$$u_i = x_i(p_i - r_i). \tag{1}$$

Almost all the mechanisms we consider in this paper will use the following allocation rule.

Definition 1 *(*Allocation $\mathcal{A}(\mathbf{b})$, Last Winner and First Loser*).* Given a bid profile $\mathbf{b} = (b_1, \ldots, b_n)$, order lenders so that $b_1 \leqslant b_2 \leqslant \cdots \leqslant b_n$. Let $k = \min\{j \mid \sum_{i=1}^{j} a_i \geqslant D, j = 1, \ldots, n\}$. Then the allocation $\mathcal{A}(\mathbf{b})$ is defined as $x_i = a_i$ for $i < k$, $x_k = D - \sum_{i=1}^{k-1} a_i$, and $x_i = 0$ for $i > k$. We refer to L_k as the last winner and L_{k+1} as the first loser.

Note that there is *at most one lender*—the last winner—who might not exhaust her budget.

Throughout the paper we will use k as index for the last winner and $k+1$ as index for the first loser. Note also that the ordering and index of lenders can change from one bid profile to another.

When multiple lenders bid the same interest rates, a fixed, preannounced tie-breaking rule is necessary. To maximize clarity of presentation, we use a tie-breaking rule which has oracle access to lenders' true interest rates. (Ties between lenders with the same true interest rate are broken arbitrarily.) While oracle access might appear to be a very strong assumption, all of the results in our paper hold, up to a modification by ϵ, for *any fixed preannounced tie-breaking rule*.[14]

For completeness, we recall the definition of a Nash equilibrium.

Definition 2 *(Nash equilibrium).* A bid profile $\mathbf{b} = (b_1, \ldots, b_n)$ is a Nash equilibrium if no lender can increase her utility by unilaterally changing her bid, that is, keeping the bids of other lenders fixed.

Given a set of bids from lenders, how should one select the winners and decide their respective interest rates? While most mechanisms investigated in this paper have the same allocation rule, they differ in their payments.

---

[10] See Ausubel (2004). If a lender can clinch part of the loan, he can set the interest rate to be the maximum possible rate allowed by the borrower. In practice, competition among a large number of lenders prevent such situations: Prosper counts over 140 bidding lenders per auction on average, only 85 of whom end up sharing the loan.

[11] Note the distinction between the payment and the price: The payment is the product of the allocation and the price.

[12] The effective interest rate is the ratio of the payment to the allocation.

[13] This linear form of utility is a consequence of the simple loan in our model. In the more general situation, the term of a Prosper loan is $n = 12$, 36 or 60 months, with monthly repayments. If the monthly lender rate is r and the servicing fee rate is f, the monthly repayment of the borrower is

$$\pi = \frac{(f+r)x}{1-(1+(f+r))^{-n}}$$

where x is the amount lent. Assuming the borrower neither repays early nor defaults, the utility of the loan for a lender is the discounted stream of cash flow

$$\begin{aligned} U &= \sum_{k=1}^{n} \frac{1}{(1+\delta)^k} \left[\pi - f \left((1+f+r)^{k-1} x - \sum_{i=1}^{k-1} (1+f+r)^{i-1} \pi \right) \right] \\ &= \frac{x}{1-(1+f+r)^{-n}} \left[\frac{r}{\delta} \left(1-(1+\delta)^{-n} \right) + \frac{f}{f+r-\delta} \left((1+\delta)^{-n} - (1+f+r)^{-n} \right) \right] \end{aligned}$$

where δ is the lender's discount rate. Of course the utility remains linear in x but loses its linearity in r. It can be seen that it is, however, well approximated by its first order Taylor expansion which gives the form of utility (1). Alternatively, one can interpret the bids and variables p_i and r_i in our model as an overall rate of return on the loan: Assuming that lenders are risk neutral for small amounts, whether lenders bid on the rate of return or the interest rate does not impact the analysis.

[14] Two alternative treatments are to discretize the bidding space to multiples of ϵ (as is actually the case in Prosper, where $\epsilon = 0.05$), or to consider ϵ-Nash equilibria. With the first, every Nash equilibrium with price p (using oracle access) translates to a Nash equilibrium with price p or $p - \epsilon$, depending on the particular tie-breaking rule used; with the second, a Nash equilibrium at price p translates to either a Nash or an ϵ-Nash equilibrium at the same price. See Immorlica et al. (2005), Karlin et al. (2005) for more discussion on tie-breaking rules.

VCG mechanism. The VCG mechanism is *incentive-compatible*, *i.e.*, it is a dominant strategy for every lender to report her true interest rate r_i.

Definition 3 *(Set Δ and bid profile* $\mathbf{r}$*).* Define a bid profile $\mathbf{r} = (r_1, \ldots, r_n)$ (*i.e.*, everyone bids truthfully). The VCG allocation is computed according to $\mathcal{A}(\mathbf{r})$ (as in Definition 1). We denote by Δ the set of VCG winners.

The VCG payments are computed as follows. Let $\Delta(j)$ be the set of winners in VCG after removing lender L_j from the group of lenders, and let $x_i(j)$ be the allocation of each $L_i \in \Delta(j)$. Observe that $\Delta \subset \Delta(j) \cup \{L_j\}$. The net payment to lender L_j in the VCG mechanism is

$$\sum_{L_i \in \Delta(j)} b_i x_i(j) - \sum_{L_i \in \Delta} b_i x_i + b_j x_j.$$

Note that the VCG mechanism is not a uniform-price mechanism. Indeed, the (effective) prices associated with the VCG payments given above are not necessarily the same for all winning lenders.

First price auction. Another natural mechanism is the "first price" auction: The allocation is according to $\mathcal{A}(\mathbf{b})$, and each winner is paid his offered interest rate b_i. In fact, the social lending Web site Zopa has recently introduced Listings, where lenders can bid on individual borrowers' loans as in Prosper, which uses a first price auction. The first price auction is clearly not incentive-compatible; it also need not have a Nash equilibrium. In fact, unlike other settings where the first price mechanism admits an ϵ-Nash equilibrium (such as single-item auctions or path auctions, see Immorlica et al., 2005), in our setting it need not even have an ϵ-Nash equilibrium, as the following example shows. Note that the choice of tie-breaking rule is not the reason the first price mechanism does not have an ϵ-Nash equilibrium; it is easy to see that this example does not have an ϵ-Nash equilibrium for any tie-breaking rule.

Example 1. Let $D = 15$ and suppose there are three lenders L_1, L_2, L_3, with budgets $a_1 = a_2 = a_3 = 10$, and interest rate $r_1 = r_2 = 0.1$, $r_3 = 0.5$. Assume ties are broken according to $L_1 > L_2 > L_3$. For any given $\epsilon > 0$, consider the bid b_1 made by L_1. If $b_1 = 0.1$, then L_2 will bid $b_2 = 0.5$ to obtain 5 units of allocation with a total utility of $5 \cdot (0.5 - 0.1) = 2$. If $b_1 = 0.5$, then L_2 will bid $b_2 = 0.5 - \epsilon$, to obtain 10 units of allocation with a total utility of $4 - 10\epsilon$. If $0.1 < b_1 < 0.5$, then L_2 will set either $b_2 = 0.5$ to obtain 5 units of allocation or $b_2 = b_1 - \epsilon$ to obtain 10 units of allocation, whichever utility is larger. Given the interest b_2 set by L_2, lender L_1 will set either $b_1 = 0.5$ or $b_1 = b_2$, whichever utility is larger. Thus, the strategies of L_1 and L_2 form a loop and there is no ϵ-Nash equilibrium.

We will consider the following one-shot model of the Prosper auction:

The Prosper mechanism. Given a bid profile $\mathbf{b}$, the mechanism used by Prosper, denoted by PROSPER, is the following.

- *Allocation*: PROSPER computes the allocation according to $\mathcal{A}(\mathbf{b})$ (recall Definition 1).
- *Pricing*: If $x_k = a_k$, *i.e.*, the last winner exhausts her budget, the price to each winner is the bid of the first loser, *i.e.* $p_i = b_{k+1}$ for $i = 1, \ldots, k$. We will refer to this interest rate as the *price* throughout. If $x_k < a_k$, *i.e.*, the last winner does not exhaust her budget, the price to all winners is the bid of the last winner, *i.e.* $p_i = b_k$ for $i = 1, \ldots, k$.

PROSPER is not truthful for the same reason that the first price auction is not: For example, suppose there are two lenders L_1 and L_2 with $a_1 = D + 1$, $r_1 = 1$ and $a_2 = D + 1$, $r_2 = 2$. Then L_1's utility is greater when bidding $b_1 = 2$ than when bidding (truthfully) $b_1 = 1$ (note that ties are broken by $L_1 > L_2$). However, as we will see shortly, PROSPER always has a Nash equilibrium, unlike the first price auction.

3. Equilibrium analysis

In this section, we will analyze the Nash equilibria of PROSPER. We are interested in the borrower's total payment, and PROSPER is a uniform-price mechanism. We therefore focus on characterizing the set of *prices* that can arise at an equilibrium. We will refer to the equilibrium with the smallest price as the *cheapest* Nash equilibrium, and to that with the largest price as the *worst* Nash equilibrium.

We start by showing that Prosper always admits a Nash equilibrium. We then provide a complete characterization of all equilibria in Section 3.1 where *losers* bid their true interest rate. Restricting ourselves to this subset of bid profiles renders possible an analysis that is otherwise intractable, and it does so without sacrificing too much of the realism of bidding behaviors, as we argue below. Finally we show in Section 3.2 that our characterization can be sharpened even further when, in addition, we assume that *winners* bid at least their true rate.

Our first result is that PROSPER always admits a Nash equilibrium, in contrast to the first price auction.

```
ALG-GREEDY
1. Define a bid profile b = r = (r_1, ..., r_n).
2. Repeat the following until convergence:
   • If there is L_i ∈ Δ such that L_i can increase utility by increasing her bid unilaterally
     in the profile b, increase b_i to the utility-maximizing point and update b.
3. Output the bid profile b.
```

Proposition 1. *ALG-GREEDY returns a Nash equilibrium of PROSPER.*

Proof. Denote by $\mathbf{b}$ the profile generated by ALG-GREEDY. Note that any lender $L_i \in \Delta$ who increases her bid in Step 2 becomes the last winner in the current bid profile. Let L_k denote the last winner in $\mathbf{b}$. Observe that all lenders in Δ are winners and $L_k \in \Delta$, since ALG-GREEDY starts with the profile of true rates, and only lenders in Δ can increase their bids to become the last winner.

First, no lender can obtain more utility by increasing her bid in $\mathbf{b}$—by definition, no winner in Δ wants to increase her bid; if winners not in Δ increase their bid higher than b_k, their allocation falls to zero. On the other hand, all losers bid their true interest (since we start with the profile $\mathbf{b} = \mathbf{r}$) and all winners other than L_k exhaust their budget, so they do not have an incentive to decrease their bid. For L_k, by the rule of ALG-GREEDY, L_k is the last lender in Δ who moves her bid up to the point where her utility is maximized. Hence, she cannot obtain more utility by decreasing her bid, and so $\mathbf{b}$ is a Nash equilibrium. $\square$

3.1. Characterizing equilibria

We now move on to characterizing the equilibria of PROSPER, which will help us analyze the borrower's payment in the cheapest and worst equilibria. In this subsection, we consider a subset of bid profiles: We assume that losers always bid their true interest rate. This assumption, crucial for tractability,[15] is not unrealistic.

One can first argue that losers should bid at least their true interest rate. A loser who bids less is exposed to the risk of obtaining a negative utility if the other lenders were to unexpectedly deviate from the equilibrium—for example if for some reason a winning lender were forced to exit the market (although very unlikely, Prosper occasionally withdraw bids). Bidding at least her true interest is, however, a safe strategy robust to any equilibrium deviation. That losers bid at least their true rate seems fairly reasonable.

One can argue further in favor of bidding *exactly* her true rate. Bidding above the true rate can potentially impact the lender's change of obtaining positive utility if some lenders were to deviate from the equilibrium, such as a winning lender forced to exit the market. A more compelling argument is perhaps that a loser who bid exactly her true rate minimizes her "envy" towards the winning bidders, in the sense that it potentially reduces the utility of the winning bidders, and does so at no cost since the losing bidder always gets zero utility. Such behavioral biases are not accounted for in the utility function, and we will not attempt to rationalize them in this paper. They are plausible nevertheless, and such ideas have been used in other auction settings to characterize bidding behavior (see, for example, Othman and Sandholm, 2010).

Under this assumption, we characterize the set of all prices and equilibrium bid profiles. Note that *winners* can, and indeed do, bid less than their true value in certain equilibria (see Examples 2 and 3 below). We begin by introducing some notation that will be used through the remainder of this paper.

Definition 4 *(Index α, $\alpha+1$ and β).* Given bid profile $\mathbf{r}$, we use α to denote the index of the last VCG winner and $\alpha+1$ to denote the index of the first VCG loser.

Assume lenders are indexed in a non-decreasing order of their true interest rates. For each $L_j \in \Delta$, let L_{β_j} be the last VCG winner when the set of lenders is restricted to $\{L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_n\}$, *i.e.*, it is the smallest index k such that $\sum_{i=1, i\neq j}^{k} a_i \geqslant D$. Define $\beta = \max_{L_j \in \Delta} \beta_j$.

Here β is the index of the last lender whose bid affects the VCG payment of some VCG winner. (Alternately, it is the largest index of lenders who enter the VCG solution if any VCG winner is removed from the market.) For example, suppose that $D = 11$, with 5 lenders $L_1, \dots, L_5$, with respective budgets $a_1 = 6$, $a_2 = 5$, $a_3 = 4$, $a_4 = 3$, $a_5 = 12$, and interest rates $r_1 < \dots < r_5$. Then $\Delta = \{L_1, L_2\}$, $\alpha = 2$, $\alpha + 1 = 3$ and $\beta = 4$. Now if instead $a_3 = 2$, then $\beta = 5$.

To characterize the equilibria, we need the following series of lemmas, whose proof is relegated to Appendix C.

Lemma 1. *In any Nash equilibrium $\mathbf{b}$ with price p, any lender L_i with $r_i < p$ is a winner.*

[15] We show in Appendix A that computing the cheapest and worst Nash equilibria is, without any restriction on bid profiles *a priori*, NP-hard. Perhaps surprisingly, this hardness result arises entirely because losers can bid strictly higher than their true values.

We point out that the reductions used to prove Theorems 9 and 10 illustrate that the above lemma need not hold for arbitrary equilibria of PROSPER: While no lender with $r_i > p$ can be a winner in a Nash equilibrium with price p, the converse is not true: Not every lender with $r_i < p$ need be a winner in an arbitrary equilibrium. That is, the assumption about losers bidding the true values is crucial this lemma.

Lemma 2. *The price p in any Nash equilibrium $\mathbf{b}$ of PROSPER satisfies $r_{\alpha+1} \leqslant p \leqslant r_\beta$. Furthermore, $p = r_j$ for some L_j with $r_{\alpha+1} \leqslant r_j \leqslant r_\beta$.*

The lemma above is crucial and characterizes the set of possible prices that can arise in an equilibrium. The next lemma follows easily from the previous ones.

Lemma 3. *For any Nash equilibrium with price p, there exists a Nash equilibrium with the same price and where all lenders in Δ are winners.*

It is tempting to think that we may assume without loss of generality that every winner other than the last winner bids her true rate, since the actual bid value of such a winner affects neither the set of winners nor the price. However, as the example below shows, this is not true: While increasing the bid to the true rate indeed does not change the allocation or the price, the resulting profile is no longer an equilibrium.

Example 2. Suppose $D = 11$. There are four lenders with $a_1 = 5$, $r_1 = 0.5$; $a_2 = 10$, $r_2 = 1$; $a_3 = 5$, $r_3 = 2$ and $a_4 = 10$, $r_4 = 7.1$. For example, bid profile $\mathbf{b} = (2, 0, 2, 7.1)$ is a Nash equilibrium with allocation $x_1 = 1$ and $x_2 = 10$. Actually, it can be seen that it is a cheapest Nash equilibrium as well. Note that if L_2 was to bid at least her true value 1, L_1 has no incentive to increase her bid any more, and thus L_2 will have to increase her bid to 7.1 to maximize her utility.

Note that in the above example, there is no equilibrium at price 2 with every winner bidding at least her true rate. While winners' bids in a Nash equilibrium can, in general, be quite complicated, the following simple equivalence holds.

Lemma 4. *Suppose that $\mathbf{b} = (b_1, \ldots, b_k, r_{k+1}, \ldots, r_n)$ is a Nash equilibrium where $L_1, \ldots, L_k$ are winners and L_k is the last winner. Then the profile $\mathbf{b}' = (0, \ldots, 0, b_k, r_{k+1}, \ldots, r_n)$ (i.e., every winner except the last winner bids 0) constitutes a Nash equilibrium with the same allocation and price.*

The lemmas above give us the following algorithm to compute Nash equilibria. The algorithm essentially checks all possible *pairs of last winners and prices*, for every lender in Δ and every price r_j, $r_{\alpha+1} \leqslant r_j \leqslant r_\beta$—note that the price in each $\mathbf{b}(k, r_j) \in S$ is exactly r_j.

ALG-PROSPER

1. Let $S = \emptyset$.
2. For each $L_k \in \Delta$ and $r_j \in \{r_\ell \mid L_\ell\colon r_{\alpha+1} \leqslant r_\ell \leqslant r_\beta\}$
   - define a bid profile $\mathbf{b}(k, r_j)$ where
     - $b_i = 0$ for each L_i with $r_i < r_j$, $i \neq k$;
     - $b_k = r_j$;
     - $b_i = r_i$ for each L_i with $r_i \geqslant r_j$, $i \neq k$;
   - If L_k is the last winner in $\mathbf{b}(k, r_j)$ and it is a Nash equilibrium, let $S \leftarrow S \cup \mathbf{b}(k, r_j)$.
3. Output S.

As there are n lenders in the market, $\beta - \alpha < n$, ALG-PROSPER runs in polynomial time.

Theorem 1. *The collection S of bidding profiles returned by ALG-PROSPER consists of Nash-equilibrium bidding profiles. Moreover, every Nash-equilibrium bidding profile is price- and allocation-equivalent to some profile returned by ALG-PROSPER.*

Proof. Consider any Nash equilibrium $\mathbf{b} = (b_1, \ldots, b_k, r_{k+1}, \ldots, r_n)$ where $L_1, \ldots, L_k$ are winners and L_k is the last winner. Let p be the price of $\mathbf{b}$ and assume without loss of generality that $b_k \leqslant r_{k+1} \leqslant \cdots \leqslant r_n$. Note that $r_{\alpha+1} \leqslant p \leqslant r_\beta$. By Lemma 4, it suffices to consider Nash equilibrium $\mathbf{b}' = (0, \ldots, 0, b_k, r_{k+1}, \ldots, r_n)$. Note that the lender L_k belongs to Δ: By Lemma 3, all lenders of Δ are winners. Since those lenders are enough to fulfill the demand, the last winner must be one of them. As $\mathbf{b}$ is a Nash equilibrium, $p = r_{k+1}$ anyway (no matter whether L_k exhausts her budget or not). Hence, $\mathbf{b}(k, r_{k+1})$ is a Nash equilibrium as well, *i.e.* $\mathbf{b}(k, r_{k+1}) \in S$. □

Why do we need to check all possible pairs of last winners and prices? The example below shows that this is necessary in general.

Example 3. Let $D = 12$. There are five lenders with $r_1 = 0$, $a_1 = 10$; $r_2 = 1$, $a_2 = 3$; $r_3 = 2$, $a_3 = 1$; $r_4 = 2.9$, $a_4 = 1$ and $r_5 = 4$, $a_5 = 12$, respectively. In this example, it can be seen that the cheapest Nash equilibrium is $\mathbf{b} = (0, 2.9, 0, 2.9, 5)$ with a total payment of $12 \cdot 2.9 = 34.8$, where L_1, L_2, L_3 are winners with allocation 10, 1, 1, respectively. An interesting fact of $\mathbf{b}$ is that L_3 is even not a VCG winner. To obtain positive utility, L_3 reduces her bid to 0, which drives L_2 to increase her bid to 2.9. However, if winners have to bid at least their true value, L_2 only wants to increase her bid to $r_3 = 2$ with an allocation of 2. In this case, the utility of L_1 is $10 \cdot 2 = 20$ and she will increase her bid to $r_5 = 4$ to obtain a utility of $(D - a_2 - a_3 - a_4) \cdot 4 = 7 \cdot 4 = 28$.

This example shows an interesting form of equilibrium where a lender with high interest rates (L_3) submits a very low bid which forces another lender (L_2) to *hold up* an interest rate that would not, otherwise, be utility-maximizing for her. Restricting winners to bid at least their true value removes such equilibria and allows us to provide a sharper characterization, as we will see in the following subsection.

The characterization in Theorem 1 easily gives the cheapest and worst Nash equilibria as the smallest and highest prices in S. In fact, the worst Nash equilibrium always has price r_β.

Theorem 2. *Let $L_j \in \Delta$ be the lender where L_β is a VCG winner by the lenders in $\{L_1, \ldots, L_{j-1}, L_{j+1}, \ldots, L_n\}$. Then the bid profile $\mathbf{b} = (b_1, \ldots, b_n)$, where $b_i = 0$ if $r_i < r_\beta$ and $i \neq j$, $b_j = r_\beta$ and $b_i = r_i$ if $r_i \geqslant r_\beta$ and $i \neq j$, is a worst Nash equilibrium with price r_β.*

Proof. We first prove that $\mathbf{b}$ is a Nash equilibrium with price r_β. Let L_k be the last winner and p be the price of $\mathbf{b}$, respectively. By the definition of β and selection of L_j, we know that $p \geqslant r_\beta$. If $p > r_\beta$, by the construction of $\mathbf{b}$, all lenders in $\Delta \cup \{L_\beta\}$ are winners. Since $L_\beta \notin \Delta$, we know that $L_k \in \Delta$ and L_k does not exhaust her budget, which implies that $b_k = p > r_\beta$, which is impossible. Hence $p = r_\beta$.

It is easy to see that no winner in $\mathbf{b}$ can obtain more utility by increasing her bid (this is because, for any winner L_i, if moving her bid to a value higher than $p = r_\beta$, she will not get any allocation). Additionally, all winners with either bid 0 or true interest $p = r_\beta$ and all losers do not have an incentive to decrease their bid. The only lender we need to consider is L_j. If $j \neq k$, L_j exhausts her budget in $\mathbf{b}$, and thus has no incentive to reduce her bid. If $j = k$, as $L_j = L_k$ does not exhaust her budget, then when reducing her bid, the price will be decreased as well. Hence, L_j is not willing to reduce her bid. Therefore, $\mathbf{b}$ is a Nash equilibrium.

From Lemma 2, we know that the price in any Nash equilibrium lies between $r_{\alpha+1}$ and r_β. Therefore this must be the worst Nash equilibrium. $\square$

3.2. Winners bid at least their private rates

We have, until now, placed no restrictions on the winners' bids. For example, it is possible that a winner wants to bid strictly less than her true interest rate to increase her allocation and utility in an equilibrium. For instance, in Example 2, L_2 can obtain a utility of 10 by bidding 0. However, such a bidding strategy also carries the risk of negative utility: Suppose a new lender L_5 with budget $a_5 = 5$ and bid $b_5 = 0.5$ enters the market. Then L_2 remains a winner but receives price 0.5, which is less than her true interest rate $r_2 = 1$, leading to a loss.

In this subsection we continue to assume that losers bid truthfully, and in addition we assume that winners bid no less than their true interest rates, for example to avoid the potential loss incurred by playing such strategies. This restricts the set of bid profiles to explore even further, and allows us to sharpen our equilibria characterization.

Starting with the profile $\mathbf{r}$ of true interest rates, let V_i be the subset of values to which each lender $L_i \in \Delta$ is willing to increase her bid to be the last winner, given the bids of other lenders (note that $V_i \in \{r_{\alpha+1}, \ldots, r_\beta\}$ when lenders are indexed in a non-decreasing order of interest rates). We will show that if any of these lenders deviates to bid one of these values, the set of values to which other lenders want to move up does not change much at all—it either shrinks or stays the same, but no new value is added to it. This will allow us to show that the set of prices in equilibria is a subset of $\bigcup_{L_i \in \Delta} V_i$. Our first lemma is analogous to Lemma 4.

Lemma 5. *Suppose that $\mathbf{b} = (b_1, \ldots, b_k, r_{k+1}, \ldots, r_n)$ is a Nash equilibrium where $L_1, \ldots, L_k$ are winners and L_k is the last winner.*

- *If L_k does not exhaust her budget, then the profile $\mathbf{b}' = (r_1, \ldots, r_{k-1}, b_k, r_{k+1}, \ldots, r_n)$ (i.e., every lender except L_k bids her true interest) is a Nash equilibrium with the same allocation and price.*
- *If L_k exhausts her budget, then $\mathbf{b}' = (r_1, \ldots, r_k, r_{k+1}, \ldots, r_n)$ (i.e., every lender bids her true interest) is a Nash equilibrium with the same allocation and price.*

Now consider a bid profile $\mathbf{b} = (b_1, \ldots, b_n)$. For each lender L_i, define a set of values $V_i(\mathbf{b})$ where $b'_i \in V_i(\mathbf{b})$ if (i) $b'_i > b_i$, (ii) when bidding b'_i, L_i is the last winner and does not exhaust her budget, and (iii) b'_i is the bid that maximizes her utility (*i.e.* $u_i(b'_i) \geqslant u_i(b)$ for any b, and in particular, $u_i(b'_i) \geqslant u_i(b_i)$). If $V_i(\mathbf{b}) \neq \emptyset$, we say that L_i is *weakly* willing to increase her bid. Intuitively, if L_i is weakly willing to increase her bid, then either she can obtain more utility by increasing her bid or the utility of other lenders can be increased without hurting her own utility. We can then establish the following result.

Lemma 6. *Given a bid profile $\mathbf{r} = \{r_1, \ldots, r_n\}$, let $T = \{L_i \mid V_i(\mathbf{r}) \neq \emptyset\}$. For each $L_j \in T$, let $b_j \in V_i(\mathbf{r})$ be an arbitrary value in $V_i(\mathbf{r})$. Define a profile*

$$\mathbf{b}(j, b_j) = (r_1, \ldots, r_{j-1}, b_j, r_{j+1}, \ldots, r_n)$$

and let

$$T(j, b_j) = \{L_i \mid V_i(\mathbf{b}(j, b_j)) \neq \emptyset\}$$

be the set of lenders who are weakly willing to increase their bids in $\mathbf{b}(j, b_j)$. Then for any $L_j \in T$,

(a) $T(j, b_j) \subseteq T$.
(b) *$L_i \notin T(j, b_j)$ for any $L_i \in T$, $i \neq j$, with $b_i \leqslant b_j$.*
(c) *$V_i(\mathbf{b}(j, b_j)) = V_i(\mathbf{r})$ for any $L_i \in T(j, b_j)$, $i \neq j$.*

Fact (a) tells us that the set of lenders who are weakly willing to increase their bid does not expand if some winner increases her bid. Fact (b) says if a lender $L_j \in T$ with larger b_j increases her bid, then all other lenders $L_i \in T$ with smaller b_i will not increase their bids any more. Fact (c) says that if L_i is still weakly willing to increase her bid after another lender $L_j \in T$ increases her bid, the set of values to which L_i is weakly willing to move will not change. The following corollary follows immediately.

Corollary 1. *Any sequence of moves starting with $\mathbf{r} = \{r_1, \ldots, r_n\}$, where a move consists of a lender $L_i \in T(\mathbf{b})$ increasing her bid to any $b \in V_i(\mathbf{b})$ where $\mathbf{b}$ is the current bid profile, converges to a Nash equilibrium.*

Let $V = \{v_1, v_2, \ldots, v_m\} = \bigcup_{L_i \in T} V_i(\mathbf{r})$ where values are ordered by $v_1 < v_2 < \cdots < v_m$. Note that $m \leqslant n$. For each $v_j \in V$, let $f(j)$ be the index of the lender where $v_j \in V_{f(j)}(\mathbf{r})$. (If there are multiple such lenders, pick one arbitrarily.) That is, v_j is a value to which lender $L_{f(j)}$ is weakly willing to move. Lemmas 5 and 6 can now be employed to derive the cheapest and the worst Nash equilibrium.

Theorem 3. *The following results hold*:

- *The cheapest Nash equilibrium is either $\mathbf{r}$ or given by the smallest index ℓ, $1 \leqslant \ell \leqslant m$, where $\mathbf{b}(f(\ell), v_\ell)$ is a Nash equilibrium.*
- *The worst Nash equilibrium is either $\mathbf{b}(f(m), v_m)$ or $\mathbf{r}$ (if $V = \emptyset$).*

Proof. For the simplicity of the proof, we denote $\mathbf{b}(f(j), v_j)$ by $\mathbf{b}(j)$.

Let $\mathbf{b}^* = (b_1^*, \ldots, b_n^*)$ be a cheapest Nash equilibrium with last winner L_k. If L_k exhausts her budget in $\mathbf{b}^*$, by Lemma 5, we know that $\mathbf{r}$ is a cheapest Nash equilibrium as well. If L_k does not exhaust her budget, again by Lemma 5, it is safe to assume that all other lenders bid their true interest in $\mathbf{b}^*$. As the utility of L_k is maximized by bidding b_k^*, we know that $L_k \in T$ and $b_k^* \in V_k(\mathbf{r}) \subseteq V$ (otherwise, $\mathbf{r}$ is a Nash equilibrium). Hence, any profile $\mathbf{b}(j)$, where $v_j < b_k^*$, is not a Nash equilibrium. Consider the profile $\mathbf{b}(j)$, where $v_j = b_k^*$. By Lemma 6, we know that $\mathbf{b}(j)$ is a Nash equilibrium, which is the cheapest Nash equilibrium as well. Similarly, we can prove that $\mathbf{b}(f(m), v_m)$ is a worst Nash equilibrium. □

4. Comparing Prosper and VCG

We now compare the total payment of VCG with that of PROSPER, in a setting where losers bid at least their true interest rate and with no restrictions on winners' bids. For any given instance, let CNE(PROSPER) and WNE(PROSPER) denote the total payment of the cheapest and worst Nash equilibrium of PROSPER, respectively. Let VCG denote the total payment of VCG.

Theorem 4. *The following inequalities hold*

$$\frac{1}{D} \cdot \mathrm{CNE(PROSPER)} \leqslant \mathrm{VCG} \leqslant O(\log D) \cdot \mathrm{CNE(PROSPER)}.$$

Proof. Consider any cheapest Nash equilibrium $\mathbf{b}^* = (b_1^*, \ldots, b_k^*)$ of PROSPER. Let x^* be the vector of allocations and p^* be the price to all winners in $\mathbf{b}^*$. Assume without loss of generality that lenders are ordered by $r_1 \leqslant \cdots \leqslant r_n$.

By Lemma 2, we know $\mathrm{CNE(PROSPER)} \leqslant D \cdot r_\beta$. On the other hand, by the definition of β, r_β will be counted in the total payment of some VCG winner in Δ (recall that Δ is the set of winners by VCG), which implies that $r_\beta \leqslant \mathrm{VCG}$. Hence, $\mathrm{CNE(PROSPER)} \leqslant D \cdot r_\beta \leqslant D \cdot \mathrm{VCG}$. It remains to prove the second inequality of the claim.

Let x_i be the allocation of each L_i by VCG. As D is the total demand, we know $x_i \leqslant D$. For each $L_i \in \Delta \setminus \{L_\alpha\}$, define an ordered multi-set S_i with $|S_i| = x_i$ by

$$S_i = \{\overbrace{r_\alpha, \ldots, r_\alpha}^{\text{total of } a_\alpha - x_\alpha}, \overbrace{r_{\alpha+1}, \ldots, r_{\alpha+1}}^{\text{total of } a_{\alpha+1}}, \overbrace{r_{\alpha+2}, \ldots, r_{\alpha+2}}^{\text{total of } a_{\alpha+2}}, \ldots\}.$$

That is, S_i contains $a_\alpha - x_\alpha$ many r_α's, $a_{\alpha+1}$ many $r_{\alpha+1}$'s, $a_{\alpha+2}$ many $r_{\alpha+2}$'s, and so on, until the size of S_i is x_i. Define an ordered multi-set S_α with $|S_\alpha| = x_\alpha$ by

$$S_\alpha = \{\overbrace{r_{\alpha+1}, \dots, r_{\alpha+1}}^{\text{total of } a_{\alpha+1}}, \overbrace{r_{\alpha+2}, \dots, r_{\alpha+2}}^{\text{total of } a_{\alpha+2}}, \dots\}.$$

That is, S_α contains $a_{\alpha+1}$ many $r_{\alpha+1}$'s, $a_{\alpha+2}$ many $r_{\alpha+2}$'s, and so on, until the size of S_α is x_α. By the rule of VCG, we know that the payment to each winner $L_i \in \Delta$ is the sum of elements in S_i.

For each S_i, denote its σ-th element by $f_i(\sigma)$, $\sigma = 1, \dots, x_i$, and let $\phi_i = \arg\max_{\sigma=1,\dots,x_i}(x_i - \sigma + 1) f_i(\sigma)$. That is, ϕ_i is the index that gives the maximum payment (not utility!) obtainable if L_i increases her interest unilaterally. Let $\lambda_i = (x_i - \phi_i + 1) f_i(\phi_i)$. Hence, the total payment of VCG satisfies

$$\begin{aligned}
\mathtt{VCG} &= \sum_{L_i \in \Delta} \sum_{\sigma=1}^{x_i} f_i(\sigma) \\
&= \sum_{L_i \in \Delta} \left(1 \cdot f_i(x_i) + \frac{1}{2} \cdot 2 f_i(x_i - 1) + \dots + \frac{1}{x_i} \cdot x_i f_i(1)\right) \\
&\leqslant \sum_{L_i \in \Delta} \left(\left(1 + \frac{1}{2} + \dots + \frac{1}{x_i}\right) \cdot \max_{\sigma=1,\dots,x_i} (x_i - \sigma + 1) f_i(\sigma)\right) \\
&= \sum_{L_i \in \Delta} O(\log x_i) \cdot \lambda_i \\
&\leqslant O(\log D) \sum_{L_i \in \Delta} \lambda_i.
\end{aligned}$$

On the other hand, consider the price p^* to all winners in $\mathbf{b}^*$. We divide the lenders in Δ into two groups:

$$\Delta_1 = \{L_i \in \Delta \mid p^* < f_i(\phi_i)\}$$

and

$$\Delta_2 = \{L_i \in \Delta \mid p^* \geqslant f_i(\phi_i)\}.$$

For each $L_i \in \Delta_1$, we claim that the total payment (not utility!) that L_i obtains in $\mathbf{b}^*$ is at least λ_i. Otherwise, assume that there is $L_j \in \Delta_1$ such that the total payment that L_j obtains in $\mathbf{b}^*$ is smaller than λ_j. Note that the utility of L_j in $\mathbf{b}^*$ is $u_j(\mathbf{b}^*) = x_j^* \cdot (p^* - r_j) = x_j^* p^* - x_j^* r_j$, where $x_j^* p^*$ is the total payment that L_j obtains. By the assumption, we know $x_j^* p^* < \lambda_j$. As $b_i^* \geqslant r_i$ for any loser L_i, when L_j increases her bid to $f_j(\phi_j)$ in $\mathbf{b}^*$, the total payment she obtains is at least λ_j. On the other hand, when the bid of L_j increases, the allocation of L_j, x_j^*, is not increasing. Since

$$\text{utility} = \text{total payment} - \text{allocation} \times \text{true interest}$$

we know that the utility of L_j increases, a contradiction to the fact that $\mathbf{b}^*$ is a Nash equilibrium. It follows that

$$\mathtt{CNE(PROSPER)} \geqslant \sum_{L_i \in \Delta_1} \text{total payment to } L_i \geqslant \sum_{L_i \in \Delta_1} \lambda_i.$$

For lenders in Δ_2, observe that

$$\begin{aligned}
\mathtt{CNE(PROSPER)} &= D \cdot p^* \\
&\geqslant p^* \cdot \sum_{L_i \in \Delta_2} x_i \\
&\geqslant p^* \cdot \sum_{L_i \in \Delta_2} (x_i - \phi_i + 1) \\
&\geqslant \sum_{L_i \in \Delta_2} (x_i - \phi_i + 1) f_i(\phi_i) \\
&= \sum_{L_i \in \Delta_2} \lambda_i.
\end{aligned}$$

Therefore,

$$\mathtt{CNE(PROSPER)} \geqslant \frac{1}{2} \sum_{L_i \in \Delta} \lambda_i \geqslant \frac{1}{O(\log D)} \mathtt{VCG}$$

which completes the proof of the theorem. $\square$

The inequalities in the above theorem are tight. Consider the following two examples:

- Let $D = 10m + 1$, where m is an arbitrary positive integer. There are six lenders with budget $a_0 = 5m + 1$, $a_1 = m$, $a_2 = a_3 = a_4 = 3m$, $a_5 = 11m$ and interest $r_0 = r_1 = r_2 = r_3 = r_4 = 0$ and $r_5 = 1$. It is easy to see that $\texttt{VCG} = 1$, where the only lender that obtains positive utility is L_0 who gets a payment of $r_5 = 1$. Note that $a_1 + a_2 + a_3 + a_4 = 10m$ and there is no way to partition $\{L_1, L_2, L_3, L_4\}$ into two groups such that the sum of budgets of each partition is $5m$. Using the argument from the proof of Theorem 9 presented in Appendix A, we obtain that the cheapest Nash equilibrium has a price 1 to each winner, which implies that the total payment is D. Hence, $\texttt{CNE}(\texttt{PROSPER}) = D \cdot \texttt{VCG}$.
- Let $D = n$. There are $n + 1$ lenders $L_0, L_1, \ldots, L_n$ with budget $a_0 = n$ and $a_i = 1$ for $i = 1, \ldots, n$. Let $r_0 = 0$ and $r_i = \frac{n}{n-i+1}$ for $i = 1, \ldots, n$. In VCG, L_0 wins and the total payment is $\sum_{i=1}^{n} r_i a_i = \sum_{i=1}^{n} \frac{n}{n-i+1} = O(n \log n)$. It can be seen that the profile $(r_0, r_1, \ldots, r_n)$ is a Nash equilibrium. Therefore, the total payment is $D \cdot r_1 = n$ and $\texttt{VCG} = O(\log D) \cdot \texttt{CNE}(\texttt{PROSPER})$.

For this general setting, *i.e.*, where losers can bid higher than their true value, the worst-case ratio between the worst Nash equilibrium of PROSPER and VCG can be arbitrarily large as the reduction in the proof of Theorem 10 shows. If losers bid their interest rate truthfully, we have the following result, similar to Theorem 4 (again, the bounds are tight).

Theorem 5.

$$\frac{1}{D} \cdot \texttt{WNE}(\texttt{PROSPER}) \leqslant \texttt{VCG} \leqslant O(\log D) \cdot \texttt{WNE}(\texttt{PROSPER}).$$

5. Bayesian analysis

The analysis of the preceding sections concentrates on the complete information setting. The actual Prosper auction is dynamic, so that the lenders learn about each other through the course of the bidding process. The complete information setting corresponds to the case in which all bidders come to fully learn each other's private interest rates. For completeness, in this section we look at the other extreme alternative: that bidders learn nothing about each other. This setting is captured by a sealed-bid auction format analyzed under incomplete information. Computing the *outcomes* of the Prosper auction in closed form is not tractable (see Hu, 2009, chapter 4); however the incomplete information setting does allow us to derive the auction that is optimal for the borrower, and draw comparisons between this optimal auction and the Prosper auction.

As before, the auction is initiated by one borrower for a particular loan, who demands D units of currency with a reserve rate R. There are n lenders $L_1, \ldots, L_n$. Each lender L_i has a true budget a_i, which is the maximum amount she is willing to put in this particular loan, and a true private interest rate r_i, which is the rate that makes this loan utility-equivalent to her best outside option. Let A be the largest amount lenders can bid. For every lender i, the pair (a_i, r_i) is identically and independently distributed according to F, with a density f and with full support on $[0, A] \times [0, R]$. (Note that this distribution applies to one particular listing, since f will depend in general on the borrower's characteristics.) We denote by $F(\cdot|a)$ and $f(\cdot|a)$ respectively the conditional distribution and density of rates for the lenders with budget a. We restrict ourselves to distributions such that the function

$$\Phi(a, r) = r + \frac{F(r \mid a)}{f(r \mid a)} \tag{2}$$

is non-decreasing in r for all budget amounts a, and is non-increasing in a for every rate r. The first condition is the reverse-auction analog to the monotone hazard rate assumption used in the derivation of optimal auctions (Myerson, 1981; Krishna, 2010); the second condition, discussed at the end of this section, ensures that lenders report their budgets truthfully.

As before, each lender i is allowed to make only one bid, $b_i = (\hat{a}_i, \hat{r}_i)$, which now includes both an amount $\hat{a}_i$ and interest rate $\hat{r}_i$. Here we no longer make the assumption that the true budgets a_i are common knowledge. Recall that the assumption was motivated by the fact that, if lenders do not change the amount bid through the course of the auction, they will have learned each other's "budget" after the first round of bids. Since there is now only one round, it is natural to assume that the amounts bid remain private information, and so the model we consider now allows lenders to choose these amounts strategically.

An auction mechanism is thus defined by an allocation function $\mathbf{x}(\cdot)$ and price function $\mathbf{p}(\cdot)$. Both take as input a bid profile: Given a bid profile $\mathbf{b} = (b_1, \ldots, b_n)$, $x_i(\mathbf{b}) \geqslant 0$ is the portion of the loan allocated to lender L_i, and $p_i(\mathbf{b})$ the corresponding loan rate. We restrict attention to the class of auctions $\mathcal{C}$ that satisfy the following conditions, which are both required by Prosper and natural in the social lending context:

1. The allocation to a lender must not exceed the amount bid.
2. The interest rate set for a lender cannot be less than the rate bid, and cannot be greater than the borrower's reserve rate R.
3. The total amount allocated must be either 0 (*i.e.*, the loan is not funded at all) or D (*i.e.*, the loan is funded in full).

The auction mechanism generates the following utility for the borrower:

$$v(\mathbf{b}) = \sum_i x_i(\mathbf{b})\big[R - p_i(\mathbf{b})\big],$$

and the following utility for each lender L_i:

$$u_i(\mathbf{b}) = x_i(\mathbf{b})\big[p_i(\mathbf{b}) - r_i\big].$$

A (pure) strategy for L_i is a function σ_i that maps her type to her bid: $\sigma_i(a_i, r_i) = (\hat{a}_i, \hat{r}_i)$, subject to the constraint that the amount bid does not exceed the budget: $\hat{a}_i \leqslant a_i$. A strategy profile $(\sigma_1, \dots, \sigma_n)$ is an equilibrium if each lender L_i maximizes her (ex-interim) expected utility when she bids $\sigma_i(a_i, r_i)$: for every alternative bid $(\hat{a}_i, \hat{r}_i)$,

$$\begin{aligned} &\mathop{\mathsf{E}}_{\substack{(a_j,r_j)\\ j\neq i}}\big[u_i\big(\sigma_1(a_1,r_1),\dots,\sigma_n(a_n,r_n)\big)\big] \\ &\quad\geqslant \mathop{\mathsf{E}}_{\substack{(a_j,r_j)\\ j\neq i}}\big[u_i\big(\sigma_1(a_1,r_1),\dots,\sigma_{i-1}(a_{i-1},r_{i-1}),(\hat{a}_i,\hat{r}_i),\sigma_{i+1}(a_{i+1},r_{i+1}),\dots,\sigma_n(a_n,r_n)\big)\big]. \end{aligned}$$

The auction is incentive-compatible when it is an equilibrium that each lender bid her full budget and true interest rate.

It can be seen that the Prosper auction is not incentive-compatible. This is because, even though the auction uses a uniform price, the bid of a lender can impact her price. Moreover, both components of the bids can influence the price and so lenders may strategize on both the rate they report, and the budget. Lenders can strategize on the interest rate because if the bid is partially winning, the bid sets the price. If participation is low, a lender with a large budget can find it profitable to bid a rate that is greater than her true rate, which decreases her likelihood of winning but increases her chances to be the last partial winner and set a high final loan rate. Lenders can also strategize over their declared budgets because the final rate awarded by the Prosper auction is generally decreasing with the size of the winner's allocations. A lender who gets allocated many units of loan deprives other lenders from winning any part of the loan. These additional losers drive the final loan rate down, because Prosper has to look further down to find the last partial winner or first loser. Therefore a lender might benefit by bidding a smaller budget, reducing her allocation but increasing the final rate for her net benefit. This suggests that the Prosper auction loses efficiency (and revenue for the borrower) by "demand reduction", in the same way as do the standard multi-unit uniform-price auctions (Ausubel and Cramton, 2002). But what causes demand reduction is somewhat different. In Ausubel and Cramton (2002), demand reduction is due to bidders submitting several bids, in the hope that the high bids lose and set the rate/price for the low (winning) bids, whereas here demand reduction in the reverse auction is due to bidders submitting a single high bid with low budget.

It is well known that revenue equivalence continues to hold for multi-unit auctions (Krishna, 2010). Revenue equivalence states that the utilities for the borrowers and lenders in an incentive-compatible auction depend solely on how the auction allocates the loan. By the Revelation Principle, every equilibrium of an auction is outcome-equivalent to the equilibrium of some incentive-compatible auction. Consequently all auctions which, at equilibrium, give the same allocations to the same lenders yield the same expected utility for the borrower and for each lender, irrespective of payments. In particular, to compare Prosper and another auction, we only need to compare their respective allocations. This will be useful later to understand why the Prosper auction does not minimize the borrower's payments in general. In the current framework revenue equivalence takes the following form:

Lemma 7. *At equilibrium, an incentive-compatible auction with allocation function* $\mathbf{x}(\cdot)$ *and price function* $\mathbf{p}(\cdot)$ *yields an* (*ex-ante*) *expected utility for the borrower equal to*

$$\mathop{\mathsf{E}}_{\substack{(a_i,r_i)\\ \forall i}} \sum_{i=1}^{n}\left(x_i\big((a_1,r_1),\dots,(a_n,r_n)\big)(R-r_i) - \int_{r_i}^{R} x_i\big((a_1,r_1),\dots,(a_i,\rho_i),\dots,(a_n,r_n)\big)\,\mathrm{d}\rho_i\right) \tag{3}$$

and an (*ex-interim*) *expected utility of each lender* L_i *equal to*

$$\mathop{\mathsf{E}}_{\substack{(a_j,r_j)\\ j\neq i}} \int_{r_i}^{R} x_i\big((a_1,r_1),\dots,(a_i,\rho_i),\dots,(a_n,r_n)\big)\,\mathrm{d}\rho_i. \tag{4}$$

Proof. Let $U_i(a_i, r_i)$ be the expected utility of lender L_i at equilibrium conditional on her true budget a_i and rate r_i. Let $X_i(\hat{a}_i, \hat{r}_i)$ (resp. $M_i(\hat{a}_i, \hat{r}_i)$) be the expected allocation (resp. total interest payments) to L_i when she bids amount $\hat{a}_i$ and rate $\hat{r}_i$. Incentive compatibility implies

$$U_i(a_i, r_i) = \max_{\hat{r}_i \in [0,R]} M_i(a_i, \hat{r}_i) - X_i(a_i, \hat{r}_i) r_i$$

and by the Envelope Theorem (*e.g.*, Milgrom and Segal, 2002), X_i is non-decreasing in its second argument and for every a_i we can write

$$\begin{aligned} U_i(a_i, r_i) &= U_i(a_i, R) + \int_{r_i}^{R} X_i(a_i, \rho_i)\,\mathrm{d}\rho_i \\ &= \int_{r_i}^{R} X_i(a_i, \rho_i)\,\mathrm{d}\rho_i \\ &= \mathop{\mathsf{E}}_{\substack{(a_j, r_j)\\ j\neq i}} \int_{r_i}^{R} x_i\big((a_1, r_1), \ldots, (a_i, \rho_i), \ldots, (a_n, r_n)\big)\,\mathrm{d}\rho_i \end{aligned}$$

where we first observe that a lender whose true rate is R can only receive rate R and makes a zero utility, and then use Fubini's theorem to interchange the integral signs. □

We now introduce an auction that is optimal from the borrower's viewpoint in that it maximizes her expected utility. Let us call the value $\Phi(a_i, r_i)$ the virtual interest rate of lender L_i, where Φ is defined by (2); these rates play a role analogous to the virtual valuations in Myerson (1981). Consider the following allocation rule $\mathbf{x}^{\mathrm{opt}}(\mathbf{b})$:

- If the total amount bid $\sum_i \hat{a}_i$ is less than the borrower's demand D, allocate zero to every lender.
- Otherwise, give temporary allocations to lenders in order of increasing value of $\Phi(\hat{a}_i, \hat{r}_i)$ up to exhaustion of the borrower's demand (with ties broken arbitrarily). Then verify that the allocation-weighted average of declared virtual interest rates is less the borrower's maximal rate R:

  $$\frac{1}{D}\sum_{i=1}^{n} \Phi(\hat{r}_i, \hat{a}_i) x_i < R,$$

  where x_i is the temporary allocation to lender L_i. If the inequality is satisfied, then keep these allocations, otherwise allocate zero to every lender.

Define the price function $\mathbf{p}^{\mathrm{opt}}(\mathbf{b})$ as

$$p_i^{\mathrm{opt}}(\mathbf{b}) = \begin{cases} \hat{r}_i + \frac{1}{x_i^{\mathrm{opt}}(\mathbf{b})} \int_{\hat{r}_i}^{R} x_i^{\mathrm{opt}}((\hat{a}_i, \rho_i), \mathbf{b}_{-i})\,\mathrm{d}\rho_i & \text{if } x_i^{\mathrm{opt}}(\mathbf{b}) > 0, \\ 0 & \text{if } x_i^{\mathrm{opt}}(\mathbf{b}) = 0. \end{cases}$$

It is easily verified that the auction belongs to the class $\mathcal{C}$ defined earlier in this section. It is also incentive-compatible: Suppose L_i bets an arbitrary amount $\hat{a}_i$, then bidding her true interest rate r_i instead of $\hat{r}_i$, lender L_i gains utility

$$\begin{aligned} (r_i - \hat{r}_i) x_i^{\mathrm{opt}}(r_i, \mathbf{b}_{-i}) + \int_{r_i}^{\hat{r}_i} x_i^{\mathrm{opt}}\big((a_i, \rho_i), \mathbf{b}_{-i}\big)\,\mathrm{d}\rho_i &= \int_{r_i}^{\hat{r}_i} \big[x_i^{\mathrm{opt}}\big((a_i, r_i), \mathbf{b}_{-i}\big) - x_i^{\mathrm{opt}}\big((a_i, \rho_i), \mathbf{b}_{-i}\big)\big]\,\mathrm{d}\rho_i \\ &\geqslant 0 \end{aligned}$$

since, by our distributional assumption, allocations are non-increasing in the bid interest rate. Also, if L_i bids interest rate r_i, then bidding her true budget a_i instead of $\hat{a}_i \leqslant a_i$, she gains utility

$$\int_{r_i}^{R} \big[x_i^{\mathrm{opt}}\big((a_i, \rho_i), \mathbf{b}_{-i}\big) - x_i^{\mathrm{opt}}\big((\hat{a}_i, \rho_i), \mathbf{b}_{-i}\big)\big]\,\mathrm{d}\rho_i \geqslant 0$$

since, by our distributional assumption, allocations are non-decreasing in amount bid.

In our simple setting the derivation of an optimal auction is fairly standard. The proposition below adapts the construction of Krishna (2010); the proof is relegated to Appendix C.[16]

Proposition 1. *The incentive-compatible auction defined by $\mathbf{x}^{\mathrm{opt}}(\cdot)$ and $\mathbf{p}^{\mathrm{opt}}(\cdot)$ yields a utility that is maximal for the borrower, among all equilibrium outcomes of all auctions of class $\mathcal{C}$.*

[16] For a design of optimal auctions and reverse auctions with multidimensional types under general assumptions, see Laffont et al. (1987), Maskin and Riley (2000), Vohra and Malakhov (2005), Goldberg et al. (2006), or Iyengar and Kumar (2008).

All auctions that allocate the same amount to the same lenders yield the same payments for the borrower. So why is the Prosper auction not optimal?

First, the Prosper auction funds the borrower's loan as long as the total amount bid covers the loan. In contrast, in order to minimize the borrower's payments, it is sometimes best to reject all bids and not fund the loan if too many lenders have high interest rates, *i.e.*, interest rates that are close to the borrower's maximum rate. The argument is similar to the case of single-item optimal auctions. Recall that by revenue equivalence (Lemma 7) the interest payments made to a lender L_i in an incentive-compatible auction equals the interest payments due to her true rate r_i, plus an information rent

$$\int_{r_i}^{R} x_i\big((a_1,r_1),\dots,(a_i,\rho_i),\dots,(a_n,r_n)\big)\,\mathrm{d}\rho_i$$

which is what it costs the borrower to incite L_i to reveal her private rate. By rejecting all bids when the lender rates are above a certain threshold, the borrower loses the chance of having her loan funded at a rate less than her reserve rate R, to gain the reduction in interest payments obtained by decreasing the information rent of the lenders.

The other distinctive feature of the optimal auction is the selection of lenders to whom it allocates the loan. While it is efficient to allocate the loan to the lenders with the lowest interest rate, the optimal allocation chooses lenders with the lowest *virtual* interest rate. It turns out that the true interest rate and the virtual interest rate can differ if there exist correlations between rates and budgets. For example, if lenders who have a high true interest rate also tend to have a large participation in the auctions, virtual rates can be strictly decreasing in the budget. In such situations, it is possible that the optimal auction selects lenders with large budgets but also high rates, instead of lenders with lower rates but low budget. While it is may appear surprising not to select the lenders with the lowest rates, it is merely a consequence of the idea behind Bayesian optimal auction design. Recall that the optimal auction sets a bidder's allocation as a function of how her true value (here, her true interest rate) compares with the distribution of values for this bidder. This is because, to ensure incentive compatibility, the price set for the bidder can only depend on the competing bids, it cannot depend on the bidder's self-reported value. Assuming every lender L_i reports truthfully her budget a_i, the distribution of true rates for each of these lenders is not F, but $F_i(\cdot) = F(\cdot|a_i)$. Therefore whether a lender should receive a positive allocation depends on how her interest rate compares to the rate of her population F_i: A lender L_i with a given true rate has a greater chance to receive a positive allocation if lenders of the same kind (with the same budget) tend to have larger rates, *i.e.*, if her distribution F_i is biased towards large rates. Of course to ensure that lenders report budgets truthfully and do not shade their bid, the distribution of rates associated with a large budget should not put them at disadvantage over the distribution of rates associated with small budget. This explains our requirement that the virtual interest rate be non-increasing in budgets.

Example 4. Suppose that the distribution of interest rates among the lenders with budget a follows a truncated normal distribution with mean $\mu_a = 3 + a/2$ and variance $\sigma^2 = 4$. Then, the virtual interest rate for a lender with true rate r and budget a is

$$\Phi(r,a) = r + \frac{\int_0^r e^{-(t-\mu_a)/(2\sigma^2)}\,\mathrm{d}t}{e^{-(r-\mu_a)/(2\sigma^2)}} = r + \int_0^r e^{(r-t)(r+t-2\mu_a)/(2\sigma^2)}\,\mathrm{d}t. \tag{5}$$

Let the borrower's demand be $D = 2$, with reserve interest rate $R = 10$. There are three lenders, L_1, L_2, and L_3, with $r_1 = 3$, $a_1 = 1$; $r_2 = 4$, $a_2 = 1$; $r_3 = 5$, $a_3 = 6$. The virtual interest rates for lenders L_1, L_2, and L_3 are respectively 4.34, 6.38, and 6.09. This means that an optimal auction will select lenders L_1 and L_3 and allocate to each of them one monetary unit. However, in the VCG auction, or in the Prosper auction in which lenders play identical bidding strategies increasing with respect to their true interest rates, the auction allocates one monetary unit to lenders L_1 and L_2 instead.

6. Dynamic Prosper mechanism

The actual mechanism used by Prosper is dynamic. In this section we extend our static auction model described in Section 2 to account for this dynamism, assuming that here also, the budgets are fixed once and for all and known to all bidders, and that a bidder can only submit new bids after being outbid.

We consider the following model of the dynamic auction. For simplicity we discretize the time, $t = 0, 1, \dots, T$. At $t = 0$, the borrower publicly announces the demand D and the reserve interest rate R. Each lender L_i submits an initial bid at $t = 1$, $b_i^{(1)}$; at subsequent times, lenders may decide to either lower their bid, or to maintain their most recent offer. We only allow bids that are multiples of the minimum increment $\epsilon > 0$ (equivalently, if a lender wants to reduce her bid, she must reduce it by at least ϵ). The budget a_i of each lender L_i is fixed and common knowledge. At every time t, the one-shot auction `PROSPER` is used to determine the allocation and price for the bid profile $(b_1^{(t)}, \dots, b_n^{(t)})$ (as before, we assume a fixed, preannounced tie-breaking rule). We assume without loss of generality that the bid profile must not be equal in two consecutive rounds, *i.e.* $(b_1^{(t)}, \dots, b_n^{(t)}) \neq (b_1^{(t+1)}, \dots, b_n^{(t+1)})$.[17] The winners at time t are announced publicly, as well as the price offered to winners and the bids of losers. The final outcome of the auction is the outcome of the last one-shot auction PROSPER at time T, and the price at that time is called *final price*. We refer to this dynamic process as PROSPER DYNAMICS. In this section, we index the lenders so that $r_1 \leqslant r_2 \leqslant \cdots \leqslant r_n$.

We first provide simple bounds on the final price under a very general and plausible assumption on bidding behavior: Consider a lender L_i whose utility is zero in the current round—if decreasing her bid will *strictly* increase her utility, assuming other lenders' bids remain unchanged, she will do so. This is a natural assumption, since if the other bids remain unchanged or decrease, L_i cannot get a positive utility anyway. Naturally lenders do not wish to end up with a negative utility, so similar to Section 3 we also assume that they never bid less than their true interest. Note that VCG winners are always among the winning lenders of PROSPER DYNAMICS, however there may be more winners. We bound the final price as stated below.

Theorem 6. *The final price of* PROSPER DYNAMICS *is between* r_α *and* r_β.

Proof. We start by showing that the final price p is no less than r_α. Assume by contradiction that $p < r_\alpha$. Then, as lenders never bid below their own interest rate, all winning lenders have an interest rate less than r_α, hence belong to the set $\{L_1, \dots, L_{\alpha-1}\}$. However, since $\sum_{i=1}^{\alpha-1} a_i < D$, the borrower's demand is not fulfilled by the winners of the auction, which is impossible as the total budget of lenders exceeds the demand. Hence $p \geqslant r_\alpha$.

We now show that $p \leqslant r_\beta$. Assume by contraction that $p > r_\beta$. Then, all lenders with a true interest no greater than r_β win the auction. Since the total budget of those lenders is greater than D, at least one lender has a budget that is not exhausted. The Prosper mechanism allows exactly one winner to have a non-exhausted budget, let L_j be this lender. Then $L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_\beta$ are all winners and exhaust their entire budget. Note that $L_j \in \Delta$ since the budget of those lenders is at least D. We remarked in Section 3 that β is the largest index of a VCG winner when the set of lenders is $\{L_1, \dots, L_{i-1}, L_{i+1}, \dots, L_n\}$ for any L_i. Therefore the total budget of the winners who completely exhaust their budget, $L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_\beta$, is at least D. This contradicts the fact that L_j is allocated a positive amount. Hence $p \leqslant r_\beta$. □

Note that the lower bound is slightly different from Lemma 2 for the equilibria of PROSPER, where the price is bounded below by $r_{\alpha+1}$.

We now consider lenders with two special types of bidding behaviors: *myopic greedy* behavior and *conservative* behavior. We will show that, for both cases, the final price is more constrained, to be either the lowest or highest possible value in the range of possible prices.

Myopic greedy lenders try to maximize their utility in the next round, under the assumption that the price of the current round remains unchanged. Such lenders choose to either keep the same bid for the next round when they cannot increase their utility, or decrease their bid just below the current price when this would allow them to get more allocation. Formally, a sequence of bids corresponds to a *myopic greedy* behavior when, for all L_i, $b_i^{(1)} = R$; and for all $t > 1$, if $x_i^{(t)} < a_i^{(t)}$ and $p^{(t)} - \epsilon \geqslant r_i$, then $b_i^{(t+1)} = p^{(t)} - \epsilon$; otherwise $b_i^{(t+1)} = b_i^{(t)}$. We show that, when all lenders follow myopic greedy strategies, the price converges to the lowest possible rate.

Theorem 7. *If all lenders are myopic greedy, then the winners are* Δ (*i.e. the set of VCG winners*) *and the final price is* $r_{\alpha+1}$ *when all winners exhaust their budget and* r_α *otherwise.*

Proof. Let p be the final price. By the above lemma, we know that $p \geqslant r_\alpha$.

All lenders in Δ are winners of PROSPER DYNAMICS. If there is $L_i \notin \Delta$ such that $x_i^{(T)} > 0$, then $p \geqslant r_{\alpha+1}$ and there is $L_j \in \Delta$ who does not exhaust her budget. This is, however, impossible as myopic greedy behavior would lead L_j to decrease her bid whenever receiving a partial allocation. Therefore the set of winners in PROSPER DYNAMICS with myopic greedy lenders is Δ.

For the same reason, no VCG winner whose budget is not exhausted should have a final bid higher than r_α. Hence, if the total budget of VCG winners exceeds D, the final price is r_α. If all VCG winners exhaust their budget, *i.e.* $\sum_{L_i \in \Delta} a_i = D$, then for any price above $r_{\alpha+1}$, lender $L_{\alpha+1}$ can bid just below the price and get a positive allocation, therefore the final bid of lender $L_{\alpha+1}$ is $r_{\alpha+1}$. Since no lender bid below their own interest rate, $r_{\alpha+1}$ is the lowest losing bid and is the final price. □

Interestingly, when all lenders are myopic greedy, the final outcome is exactly the same as that in PROSPER when all lenders bid their true interest rate (of course, this need not be an equilibrium bid profile in PROSPER).

Conservative lenders attempt to maximize their final utility at the last round *under worst-case assumptions* about other lenders' true interest rates. The worst-case scenario for a lender L_i occurs when every other winner has a true interest less

[17] Since bids are decreasing and bounded below by zero, they must converge in a finite number of rounds for any given $\epsilon > 0$, no matter what strategy lenders decide to follow; we assume that T is large enough so as to allow convergence in all cases.

than L_i's interest rate. When this is the case, a lender should never decrease her bid when she has a positive utility in the current round, otherwise she can decrease it by the minimum increment ϵ (as long as it is above her true interest rate). Formally, a sequence of bids corresponds to a *conservative* bidding behavior when, for all L_i, $b_i^{(1)} = R$, and for all $t > 1$, $b_i^{(t+1)} = b_i^{(t)}$ if $x_i^{(t)} > 0$ or $b_i^{(t)} = r_i$; and $b_i^{(t+1)} = b_i^{(t)} - \epsilon$ otherwise. When all lenders follow conservative bidding strategies, the final price is the maximum possible price.

Theorem 8. *If all lenders are conservative, then the final price is no less than $r_\beta - \epsilon$.*

Proof. Let $L_j \in \Delta$ be the lender in Δ whose VCG payment is affected by L_β. We remark that L_j must be a winner, since all VCG winners belong to the set of winners in the dynamic process. If the final price $p < r_\beta - \epsilon$, then there is a previous round where L_j bids $r_\beta - \epsilon$. We show that, when that is the case, L_j always get a positive allocation, so that L_j is not willing to lower her bid, which contradicts to $p < r_\beta - \epsilon$. Indeed, when L_j bids $r_\beta - \epsilon$, lenders who bid less than or equal to L_j belong to $S_j = \{L_1, \ldots, L_{\beta-1}\} \backslash \{L_j\}$. Since β is the index of the largest VCG winner when considering the set of lenders $\{L_1, \ldots, L_{j-1}, L_{j+1}, \ldots, L_n\}$, the total budget of the lenders of S_j is less than D. This implies that L_j must receive a positive allocation for that round. $\square$

A natural question to ask is how the borrower's payment in PROSPER DYNAMICS compares with that in VCG (the VCG mechanism also has a dynamic implementation for this setting, see Ausubel, 2004). However, it is easy to construct examples that show that neither mechanism dominates the other, in the sense that the total payment to the lenders in PROSPER DYNAMICS can be larger than that in VCG mechanism and vice versa, depending on the behavior of lenders.

- Let $D = 15$ and $R = 10$. There are three lenders L_1, L_2, L_3 with respective budgets $a_1 = 14$, $a_2 = 2$, $a_3 = 20$ and interest rates $r_1 = \delta$, $r_2 = 2\delta$, $r_3 = 1$. The VCG allocation is then $x_1 = 14$, $x_2 = 1$, $x_3 = 0$, and the VCG payment to L_1 is $2\delta \cdot 2 + 13 \cdot 1$, while the payment to L_2 is 1. In PROSPER DYNAMICS, if lenders play myopic greedy strategies, the price is δ and so the total payment is $15 \cdot 2\delta$, so that the VCG payments are arbitrarily many times larger than those of PROSPER DYNAMICS, when δ tends to 0.
- Assume that the demand is any $D > 0$, and consider $D + 1$ lenders, with respective budgets $a_1 = D$, $a_2 = \cdots = a_D = 1$, $a_{D+1} = 2D$, and interest rates $r_1 = \delta$, $r_2 = 2\delta = \cdots = r_D = 2\delta$, $r_{D+1} = 1$. The total VCG payment is then $2\delta \cdot (D - 1) + 1$, while in PROSPER DYNAMICS with conservative bidders, the total payment is $D \cdot 1$. As δ tends to 0, the ratio between PROSPER DYNAMICS and VCG payment tends to D. Note that the ratio cannot be greater than D as in all cases VCG allocates at least one unit of demand at rate r_β.

Developing a model for how lenders actually bid in the Prosper auction is essential to developing a more precise understanding of the equilibrium final price in the actual Prosper auction. We believe this is an interesting direction for further work.

7. Conclusion

In this paper, we took the first steps towards developing a theoretical understanding of social lending markets, investigating borrower payments in the mechanism used by Prosper, one of the largest social lending sites in the US, to auction off each borrower's loan. Our analysis allowed us to precisely characterize borrower's payments in equilibria of the Prosper auction, and we found that the Prosper auction does not provide the cheapest loans to the borrower in the presence of strategic lenders trying to maximize their profit from lending, in either a full information or incomplete information setting. Of course, while the basic premise of social lending is providing borrowers with cheap loans, Prosper's choice of mechanism is also likely driven by consideration of its own profits, which, given Prosper's business model where both lenders and borrowers are charged some fraction of their loan amounts, relies heavily on volume of trade.

The Prosper marketplace has since changed to a set rate per loan, rather than an auction process. This change is interesting in itself, and relates to questions regarding when auctions are the best mechanisms for selling goods. Note here that how auctions perform (specific auction formats, as well as auctions in general) also depends on how many and what kind of participants it is able to attract, since participation is not the same fixed quantity across all social lending platforms, but rather is *endogenous to the market design*. We note in this context that Prosper, which was the first social lending site in the US, began losing market share to Lending Club, which uses a simple pre-set rate mechanism, and overtook Prosper to become the market leader within 3 years of its entry into the social lending market. Arguably, this could be related to the fact that a system that is simple attracts more participants—a mechanism that people understand and is transparent might attract a much larger volume of participants and hence yield better outcomes (for some or all parties) than a complex mechanism such as the Prosper auction, or even an optimal auction that only manages to attract a small number of traders. This switch in market design effected by Prosper could lead to a wealth of interesting data for research questions regarding participation and behavior in the two different market designs, which could in turn form the basis for a theoretical study of the question of market design for social loans.

The online social lending market is clearly an exploding space, and one whose design has immediate financial consequences for the millions of borrowers and lenders trading in these markets. The switch effected by Prosper from auctions to a set rate, as well as the fact that other leading sites such as Zopa and Lending Club use rather different designs indicate both that the understanding of how to best design such markets is, at best, in a nascent phase, as well as offer the potential for data to observe the effect of these changes and differences on participation and trader behavior. We believe that the problem of market design for social lending, based on an understanding of the current social lending marketplace and the behavior of its participants, and addressing the benefits to all parties involved in the trade, is an important and exciting direction for future research.

Appendix A. NP-hardness of computing cheapest and worst equilibria of PROSPER

Theorem 9. *The computation of (the total payment of) a cheapest Nash equilibrium of PROSPER is NP-hard. Furthermore, for any polynomial time computable function $f(n)$, it does not admit any approximation algorithm within a ratio of $f(n)$, unless $P = NP$.*

Proof. We reduce from Partition: Given an instance of Partition with a set of integers $S = \{x_1, \dots, x_n\}$ where $\sum_{i=1}^n x_i = 2N$, we ask if S can be partitioned into two subsets such that the sum of the numbers in each subset is N. Assume without loss of generality that $1 \leqslant x_i \leqslant N$, for $i = 1, \dots, n$.

We construct an instance of our problem as follows: Let $M \triangleq f(n)$. For $i = 1, \dots, n$, there is a lender L_i with budget $a_i = x_i M$ and interest $r_i = 0$. Further, there are two extra lenders L_0 and L_{n+1} with budget $a_0 = MN + 1$, $a_{n+1} = 3MN$, and interest $r_0 = 0$, $r_{n+1} = 1$, respectively. Let $D = 2MN + 1$. We claim that it is NP-hard to distinguish whether the total payment of the cheapest Nash equilibrium is smaller than or equal to $\frac{2MN+1}{MN+1}$ or at least $2MN + 1$.

Assume that there is a partition of S into S_1 and S_2 such that the sum of the numbers in each subset is N. We construct a bid profile $\mathbf{b}$ as follows: Let $b_i = r_i$ for $i = 0, n + 1$, $b_i = 0$ if $x_i \in S_1$ and $b_i = \frac{1}{MN+1}$ if $x_i \in S_2$. Given $\mathbf{b}$, as $a_0 + \sum_{i:\ x_i \in S_1} a_i = 2MN + 1 = D$, the winners are L_0 and those corresponding to the set S_1 and all winners exhaust their budget. Thus, the price to each winner is $\frac{1}{MN+1}$ and the utility of L_0 is $a_0 \cdot \frac{1}{MN+1} = 1$. If L_0 increases her bid b_0 (to a point at most $b_{n+1} = 1$ to remain to be a winner), her payment is at most 1 and utility is at most $1 \cdot (D - \sum_{i:\ x_i \in S} a_i) = D - 2MN = 1$, which implies that L_0 has no incentive to change her bid. Further, it is easy to see that the lenders corresponding to the set S_1 do not have an incentive to change their bid as well. For each lender L_i, $x_i \in S_2$, although L_i can reduce her bid to 0 to be a winner, the price to winners becomes 0 as well, which leads to a 0 utility to L_i. Therefore, no lender can unilaterally increase her utility and $\mathbf{b}$ is a Nash equilibrium with a total payment of $D \cdot \frac{1}{MN+1} = \frac{2MN+1}{MN+1}$.

On the other hand, assume that there is no such a partition of S such that the sum of the numbers in each subset is N. Consider any Nash equilibrium $\mathbf{b} = \{b_0, b_1, \dots, b_n, b_{n+1}\}$. If L_{n+1} is a winner, then the price to each winner is at least 1 and we are done. Thus, it is safe to assume that L_{n+1} is not a winner. It follows that L_0 must be a winner. Let L_j be the last winner in $\mathbf{b}$. If L_j exhausts her budget, as $D - a_0 = MN$, the set of winners excluding L_0 defines a partition of S with sum N, a contradiction to our assumption. Hence, it suffices to consider the case where L_j does not exhaust her budget. By the rule of PROSPER, the price to each winner is b_j. It can be seen that $b_j > 0$ (otherwise, as argued above, L_0 can increase her bid to 1 to obtain a positive utility). In addition, if there is L_i, $x_i \in S$, such that L_i is not a winner, then L_i can reduce her bid to $b_j - \epsilon$ to be a winner with positive payment and utility, a contradiction. Thus, all lenders $L_1, \dots, L_n$ are winners. As $x_i \leqslant N$ for $i = 1, \dots, n$, the last winner L_j has to be L_0. In this case, by the property of Nash equilibrium, $b_j \geqslant 1$, which implies that the total payment to winners is at least $1 \cdot D = 2MN + 1$.

Hence, it is NP-hard to distinguish whether the total payment of the cheapest Nash equilibrium is smaller than or equal to $\frac{2MN+1}{MN+1}$ or at least $2MN + 1$. As $\frac{2MN+1}{(2MN+1)/(MN+1)} = MN + 1$, it is NP-hard to approximate the total payment of the cheapest Nash equilibrium within a ratio of $\Omega(M) = \Omega(f(n))$. $\square$

The computation of a worst Nash equilibrium is NP-hard as well, as the following result shows.

Theorem 10. *The computation of (the total payment of) a worst Nash equilibrium of PROSPER is NP-hard. Furthermore, it does not admit any approximation algorithm within any ratio, unless $P = NP$.*

Proof. We reduce from Partition: Given a set of integers $S = \{x_1, \dots, x_n\}$ where $\sum_{i=1}^n x_i = 2N$, can S be partitioned into two subsets such that the sum of the numbers in each subset is N? Assume without loss of generality that $1 \leqslant x_i \leqslant N$, for $i = 1, \dots, n$.

We construct an instance of our problem as follows: For $i = 1, \dots, n$, there is a lender L_i with budget $a_i = x_i$ and interest $r_i = 0$. Let $D = N$ and $R = 1$. We claim that it is NP-hard to distinguish whether the total payment of the worst Nash equilibrium is either 0 or N.

Assume that there is a partition of S into S_1 and S_2 such that the sum of the numbers in each subset is N. We construct a bid profile $\mathbf{b}$ as follows: Let $b_i = r_i$ if $x_i \in S_1$ and $b_i = R = 1$ if $x_i \in S_2$. Given $\mathbf{b}$, as $\sum_{i:\ x_i \in S_1} a_i = N = D$, the winners are those corresponding to the set S_1 and all winners exhaust their budget. Thus, the price to each winner is 1. It is easy to see that no winner is willing to increase her bid. Additionally, if any loser L_i reduces her bid to 0, even if L_i becomes a winner, since the price is reduced to 0 as well, L_i still obtains 0 utility. Hence, no lender can unilaterally increase her utility and $\mathbf{b}$ is a Nash equilibrium with a total payment of $D \cdot 1 = N$.

On the other hand, assume that there is no such a partition of S such that the sum of the numbers in each subset is N. Consider any Nash equilibrium $\mathbf{b} = \{b_1, \dots, b_n\}$. Let L_j be the last winner in $\mathbf{b}$. If L_j exhausts her budget, then the set of winners constitutes a partition of S with sum N, a contradiction to our assumption. Hence, L_j does not exhaust her budget. By the rule of PROSPER, the price to each winner is b_j. If $b_j > 0$, we claim that all lenders $L_1, \dots, L_n$ are winners. Otherwise, if L_i is not a winner (which implies that $b_i \geqslant b_j$), by reducing her bid to $b_j - \epsilon > 0$, L_i becomes a winner with price at least $b_j - \epsilon$, a contradiction to the fact that $\mathbf{b}$ is a Nash equilibrium. However, if all lenders are winners, we know that $\sum_{i:\ i \neq j} a_i = \sum_{i:\ i \neq j} x_i < D = N$. Because $\sum_{i=1}^{n} x_i = 2N$, we have $x_j > N$, a contradiction to our assumption. Hence, $b_j = 0$, which implies that the total payment to winners is 0.

Hence, it is NP-hard to distinguish whether the total payment of the worst Nash equilibrium is 0 or N, which implies that we do not have any approximation algorithm within any ratio, unless $P = \text{NP}$. □

Appendix B. Other uniform-price mechanisms

PROSPER is a uniform-price mechanism, meaning that all winning lenders receive the same price. How does it compare to other uniform-price mechanisms? Here there are two natural candidates—pay all winners the *bid of the last winner* (denoted by BLW), and pay all winners the *bid of the first loser* (denoted by BFL). Both mechanisms have the same allocation rule as PROSPER (as in Definition 1), but a slightly different pricing rule. The mechanism BLW offers a price equal to the bid of the last winner, while BFL offers a price equal to the bid of the first loser. Note that the price of PROSPER is either that of BLW or BFL, depending on whether or not the last winner exhausts her budget.

If all "items" were identical, meaning that every lender had a budget of one, BFL would be identical to VCG, which produces an *efficient allocation* amongst lenders. As the following example shows, however, BFL is in fact very different from VCG in terms of efficient allocation: Since the price is determined by the bid of the first loser, every winner has an incentive to bid as low as possible to increase her allocation when the total budget of winners is greater than the demand. Specifically, as long as the bid of the first loser is at least her true interest rate, a winner loses nothing by bidding as low as 0 to improve her allocation.

Example 5. Let $D = 11$. There are four lenders with budgets $a_1 = 3$, $a_2 = 4$, $a_3 = 5$, $a_4 = 10$ and interest rates $r_1 = 1$, $r_2 = 2$, $r_3 = 3$, $r_4 = 4$, respectively. The *only* equilibrium of BFL has L_1, L_2, L_3 bidding 0 and the losing lender L_4 bidding her true value, and the allocation determined according to the given tie-breaking rule. (As otherwise, the lender not exhausting her budget can always reduce her bid to 0 and increase her allocation, *with no change in price* which remains at $b_4 = 4$.)

Thus, which winner has leftover budget is entirely determined by the tie-breaking rule, which can be, in general, arbitrary: Even when the winners of BFL are exactly the VCG winners, lenders with lower interest rates do not necessarily receive a better allocation, leading to inefficiency. Note that this does not happen in either PROSPER or BLW, where the last winner never bids below her true interest rate.[18] Specifically, in the above example, both PROSPER and BLW have efficient equilibria, while BFL does not.

We now investigate BLW. As the following results show, BLW is very similar to PROSPER: The set of equilibria of BLW is a subset of that of PROSPER, and when losers are restricted to bid their true values (as in Section 3.1), the equilibria are identical.

Similar to Lemma 4 (and its proof), we have the following result.

Lemma 8. *Suppose* $\mathbf{b} = (b_1, \dots, b_k, b_{k+1}, \dots, b_n)$ *is a Nash equilibrium of* BLW*, where* $L_1, \dots, L_k$ *are winners and* L_k *is the last winner. Then the profile* $\mathbf{b}' = (0, \dots, 0, b_k, r_{k+1}, \dots, r_n)$ *(i.e., every winner bids* 0 *except* L_k*) constitutes a Nash equilibrium with the same allocation and price.*

Proof. In any Nash equilibrium $(b_1, \dots, b_k, b_{k+1}, \dots, b_n)$ of BLW, any lender L_j with $r_j < p = b_k$ must be a winner. Indeed, if not, L_j can bid $p - \epsilon \geqslant r_j$ and get positive utility. So all lenders whose private interest rate is less than p must be winners, and $r_i \geqslant p$ for all losers L_i. By assumption, losers bid greater or equal to their true value, so $b_i \geqslant r_i \geqslant p$ as well.

As in Lemma 4 winners L_i, $i < k$ can replace their bid by 0 and no one has an incentive to deviate. Suppose a loser L_j decreases her bid from b_j to $r_j \leqslant b_j$, no winner can profit from increasing her bid because the same increase would have profited her in the previous profile as well. □

Theorem 11. *Any bid profile that is an equilibrium of* BLW *is also an equilibrium in* PROSPER*. Moreover, if we restrict ourselves to bid profiles where losers bid their true value, every equilibrium in* PROSPER *is also one in* BLW*, so that both sets of equilibria are identical.*

[18] Strictly speaking, this holds for PROSPER only when the sum of the budgets of winners exceeds demand D. If it is equal to D, the last winner exhausts her budget anyway, and bidding lower than her true interest rate does not change anything.

Proof. Note that both mechanisms are identical when the last winner does not exhaust her budget. Let's consider an equilibrium $\mathbf{b} = (b_1, \ldots, b_n)$ of BLW where L_k is the last winner. By Lemma 8, we can assume without loss of generality that all lenders other than L_k bid her true interest rate.

Suppose otherwise that $\mathbf{b}$ is not an equilibrium in PROSPER. This means there is a profitable deviation for some lender L_i in PROSPER. There are the following cases.

Case 1. L_i is a winner, $i \neq k$. Her only possible profitable deviation is to increase her bid above b_k. Since $\mathbf{b}$ is an equilibrium in BLW, $b_k = r_{k+1}$, and therefore, by increasing her bid, at least one loser becomes a winner and L_i does not exhaust her budget anymore. In that case PROSPER and BLW compute the exact same allocation and price, which means that if L_i profits from that deviation in PROSPER, she would profit from the same deviation in BLW, which contradicts the fact that $\mathbf{b}$ is a Nash equilibrium in BLW.

Case 2. $L_i = L_k$. As she does not wish to deviate in BLW, $b_k = r_{k+1}$. In PROSPER, L_k cannot profit from decreasing her bid, since if she is exhausting her budget, her allocation and price remain the same. If she is not exhausting her budget, decreasing her bid can only decrease her price. In all cases, L_k cannot profit by decreasing. By an identical argument as the above case, she cannot profit by increasing her bid in PROSPER as well.

Case 3. L_i is a loser. We show that if there is a profitable deviation in PROSPER, there must be one in BLW as well. Increasing her bid does not increase her allocation (which is 0), and therefore does not increase her utility. Decreasing her bid leads to a price that is less than her true value, which decreases utility (since losers bid at least their true interest rate).

Therefore any Nash equilibrium in BLW is also a Nash equilibrium in PROSPER. When losers bid their true value, by Lemma 4, we can restrict ourselves to the Nash equilibria of PROSPER of the form $(0, \ldots, 0, b_k, r_{k+1}, \ldots, r_n)$, and the same reasoning as above can be used to show the converse. $\square$

When losers can bid any value greater equal their true interest rates (that is, they are not restricted to bidding truthfully), the cheapest Nash equilibrium of PROSPER can be a factor of D smaller than the cheapest equilibrium of BLW; conversely, the worst Nash equilibrium of PROSPER can be arbitrarily larger than the worst equilibrium of BLW. This follows directly from the examples of the hardness reductions in Theorems 9 and 10.

Appendix C. Omitted proofs

C.1. Proof of Lemma 1

The claim follows directly from the assumption that all losers bid their true interest and the definition of the mechanism: For any lender L_i with $r_i < p$, if L_i is not a winner, we know L_i bids her interest truthfully. This implies that the price p is at most r_i, a contradiction.

C.2. Proof of Lemma 2

(a) We start with the lower bound, $p \geqslant r_{\alpha+1}$. By contradiction, suppose that $p < r_{\alpha+1}$ for some Nash equilibrium. Note that since any VCG loser has an interest rate greater than p, any VCG loser is a loser. It is easy to see that $p \geqslant r_\alpha$ as otherwise, by the definition of α, the total demand cannot be fulfilled. If $p = r_\alpha$, the last VCG winner L_α obtains zero utility, and thus is willing to increase her bid to $r_{\alpha+1}$ to obtain positive utility. If $r_\alpha < p < r_{\alpha+1}$, by Lemma 1, all lenders of Δ are winners. Hence, the last winner does not exhaust her budget and would profit from increasing her bid to $r_{\alpha+1}$, a contradiction.

(b) We now deal with the upper bound, $p \leqslant r_\beta$. By contradiction, suppose that $p > r_\beta$ for some Nash equilibrium. Let L_k be the last winner. By Lemma 1, we know that all lenders L_i with an interest $r_i \leqslant r_\beta < p$ are winners. In particular, lenders in $\Delta \cup \{L_\beta\}$ are winners, implying by Definition 3 that $L_k \in \Delta$. This contradicts to Definition 4 of β.

(c) It remains to prove that $p = r_j$ for some L_j with $r_{\alpha+1} \leqslant r_j \leqslant r_\beta$. If $p \neq r_j$ for any L_j with $r_{\alpha+1} \leqslant r_j \leqslant r_\beta$, as all losers bid their true interest, the only possible case is that the last winner L_k does not exhaust her budget and hence $p = b_k$. In this case, however L_k will always increase her bid to the point which equals the bid of the first loser, a contradiction.

C.3. Proof of Lemma 3

From Lemma 2, $p \geqslant r_{\alpha+1}$. If $p > r_{\alpha+1}$, then by Lemma 1 all lenders with an interest rate less than or equal to $r_{\alpha+1}$ are winners, and so by definition of Δ all lenders of Δ are winners.

If $p = r_{\alpha+1}$, then by Lemma 1, all lenders $L_i \in \Delta$ with $r_i < r_{\alpha+1} = p$ are winners. If a VCG winner is a loser, she must bid her true interest rate by assumption, which must then be greater than or equal to p since she loses the auction, but also no greater than $r_{\alpha+1}$ since she belongs to Δ. Hence all lenders of Δ that are losers bid exactly $r_{\alpha+1}$. For any winner $L_i \notin \Delta$, $r_i \geqslant r_{\alpha+1}$, but since L_i gets a nonnegative utility, $r_i \leqslant r_{\alpha+1}$, and so $r_i = r_{\alpha+1}$. In other words, any winner that does not belong to Δ gets a zero utility, and so can increase her bid to her true interest rate $r_{\alpha+1}$ without changing her utility. This gives a Nash equilibrium with the same price, and by the tie-breaking rule, the winners are exactly the lenders of Δ.

C.4. Proof of Lemma 4

It is easy to see that both $\mathbf{b}$ and $\mathbf{b}'$ have the same allocation and price to winners, which is at least b_k. By Lemma 1, we know that for any L_i, $i = k+1, \ldots, n$, $r_i \geqslant b_k$. As all losers and the last winner L_k bid the same value in $\mathbf{b}$ and $\mathbf{b}'$, it is easy to see that no lender has an incentive to increase her bid in $\mathbf{b}'$ to obtain more utility. In addition, by the fact that $\mathbf{b}$ is a Nash equilibrium, it can be seen that lenders $L_1, \ldots, L_{k-1}, L_{k+1}, \ldots, L_n$ cannot obtain more utility by decreasing their bid. For L_k, even if L_k is able to reduce her bid down to 0, either the price drops to 0 as well (if L_k does not exhaust her budget in $\mathbf{b}$) or her allocation does not change (if L_k exhausts her budget in $\mathbf{b}$). Thus, L_k cannot obtain more utility by decreasing her bid. Therefore, $\mathbf{b}'$ is a Nash equilibrium with the same allocation and price as $\mathbf{b}$.

C.5. Proof of Lemma 5

It is easy to see that both $\mathbf{b}$ and $\mathbf{b}'$ have the same allocation and price to winners, which is at least b_k. By Lemma 1, we know that for any L_i, $i = k+1, \ldots, n$, $r_i \geqslant b_k$.

If L_k does not exhaust her budget, as all losers and the last winner L_k bid the same value in $\mathbf{b}$ and $\mathbf{b}'$, it is easy to see that no lender has an incentive to increase her bid in $\mathbf{b}'$ to obtain more utility. In addition, by the fact that $\mathbf{b}$ is a Nash equilibrium, it can be seen that lenders $L_1, \ldots, L_{k-1}, L_{k+1}, \ldots, L_n$ cannot obtain more utility by decreasing her bid. For L_k, if L_k can obtain more utility by decreasing her bid in $\mathbf{b}'$, she is willing to decrease her bid in $\mathbf{b}$ as well, a contradiction. Therefore, $\mathbf{b}'$ is a Nash equilibrium with the same allocation and price as $\mathbf{b}$.

The case where L_k exhausts her budget is similar. Hence, the claim follows.

C.6. Proof of Lemma 6

For simplicity, in the proof, we denote $\mathbf{b}(j, b_j)$ by $\mathbf{b}(b_j)$ and $T(j, b_j)$ by $T(b_j)$.

(a) Consider any $L_i \in T(b_j)$. Note that the only difference between $\mathbf{r}$ and $\mathbf{b}(b_j)$ is that L_j increases her bid from r_j to b_j. Hence, the utility of all other lenders does not decrease, which implies that $u_i(\mathbf{r}) \leqslant u_i(\mathbf{b}(b_j))$. Furthermore, by the definition of T and $T(j)$, L_i is weakly willing to increased her bid in $\mathbf{b}(b_j)$, which implies she is weakly willing to increase her bid in $\mathbf{r}$ as well. Hence, $L_i \in T$ and $T(b_j) \subseteq T$.

(b) Assume otherwise that there is $L_i \in T$ with $b_i \leqslant b_j$ such that $L_i \in T(b_j)$. Further, by the definition of $V_i(\mathbf{r})$ and $\mathbf{b}(b_j)$, L_j is the last winner in $\mathbf{b}(b_j)$ and her budget is not exhausted. Hence, L_i exhausts her budget with price b_j in $\mathbf{b}(b_j)$. In $\mathbf{b}(b_i)$, however, L_i does not exhaust her budget with price b_i. As $b_j \geqslant b_i > r_i$, we have $u_i(\mathbf{b}(b_j)) > u_i(\mathbf{b}(b_i))$. Therefore, if $L_i \in T(b_j)$, *i.e.* L_i is weakly willing to increase her bid in $\mathbf{b}(b_j)$, the utility of L_i can be strictly increased in $\mathbf{b}(b_i)$, which contradicts the definition of b_i—the point where the utility of L_i is maximized.

(c) As $L_i \in T(b_j)$, by Fact (a), we know that $L_i \in T$. If there is $b \in V_i(\mathbf{r})$ such that $b \leqslant b_j$, by picking $b_i = b$ for L_i, we know that $L_i \notin T(b_j)$ by Fact (b), a contradiction. Hence, in $\mathbf{r}$ and $\mathbf{b}(b_j)$, L_i is weakly willing to increase her bid to the same points. (Note that when L_i bids a value higher than b_j, it does not matter if L_j bids r_j or b_j.)

C.7. Proof of Proposition 1

First observe that by the Revelation Principle, if an auction yields an optimal utility for the borrower among all incentive-compatible auctions of class $\mathcal{C}$, it also yields an optimal utility among all outcomes of all equilibria of all auctions of class $\mathcal{C}$.

Let us consider an arbitrary incentive-compatible auction of class $\mathcal{C}$, with allocation function $\mathbf{x}(\cdot)$ and price function $\mathbf{p}(\cdot)$. We use same notation as in the proof of Lemma 7. The expected utility of the borrower is

$$\sum_{i=1}^{n} \mathop{\mathsf{E}}_{(a_i, r_i)} \big[R X_i(a_i, r_i) - M_i(a_i, r_i)\big],$$

and recall that by the envelope theorem

$$M_i(a_i, r_i) = r_i X_i(a_i, r_i) + \int_{r_i}^{R} X_i(a_i, \rho_i)\, \mathrm{d}\rho_i.$$

Interchanging the integrals yields

$$\int_0^A \int_0^R \int_{r_i}^R X_i(\rho_i) f(a_i, r_i)\, \mathrm{d}\rho_i\, \mathrm{d}r_i\, \mathrm{d}a_i = \int_0^A \int_0^R X_i(r_i) F(r_i \mid a_i) f(a_i)\, \mathrm{d}r_i\, \mathrm{d}a_i.$$

After substitution,

$$\mathop{\mathsf{E}}_{(a_i,r_i)}\bigl[M_i(a_i,r_i)\bigr] = \int_0^A \int_0^R r_i X_i(a_i,r_i) f(a_i,r_i)\,\mathrm{d}r_i\,\mathrm{d}a_i + \int_0^A \int_0^R X_i(a_i,r_i) F(r_i \mid a_i) f(a_i)\,\mathrm{d}r_i\,\mathrm{d}a_i$$

$$= \int_0^A \int_0^R \left[r_i + \frac{F(r_i \mid a_i)}{f(r_i \mid a_i)}\right] X_i(a_i,r_i) f(a_i,r_i)\,\mathrm{d}r_i\,\mathrm{d}a_i$$

$$= \int_{[0,A]^n} \int_{[0,R]^n} \Phi(r_i,a_i) x_i\bigl((a_1,r_1),\ldots,(a_n,r_n)\bigr) \prod_{i=1}^n f(a_i,r_i)\,\mathrm{d}\mathbf{r}\,\mathrm{d}\mathbf{a}.$$

Hence the expected utility of the borrower is

$$\sum_{i=1}^n \mathop{\mathsf{E}}_{(a_i,r_i)}\bigl[R X_i(a_i,r_i) - M_i(a_i,r_i)\bigr] = \int_{[0,A]^n} \int_{[0,R]^n} \left[\sum_{i=1}^n \bigl(R - \Phi(r_i,a_i)\bigr) x_i\bigl((a_1,r_1),\ldots,(a_n,r_n)\bigr)\right] \prod_{i=1}^n f(a_i,r_i)\,\mathrm{d}\mathbf{r}\,\mathrm{d}\mathbf{a}.$$

An auction of class $\mathcal{C}$ can only make positive allocations when these allocations total the borrower's demand D, there are thus two cases to consider: Given $(a_1,r_1),\ldots,(a_n,r_n)$, either $\sum_i x_i = 0$, in which case

$$\sum_{i=1}^n (R-\Phi)x_i = 0 \leqslant \sum_{i=1}^n (R-\Phi)x_i^{\mathrm{opt}}$$

by definition of $\mathbf{x}^{\mathrm{opt}}(\cdot)$, or $\sum_i x_i = D$, in which case

$$\sum_{i=1}^n (R-\Phi)x_i \leqslant \sum_{i=1}^n (R-\Phi)x_i^{\mathrm{opt}},$$

since $\mathbf{x}^{\mathrm{opt}}(\cdot)$ allocates the loan in priority to the lenders with the lowest virtual interest rate.

References

Archer, A., Tardos, E., 2002. Frugal path mechanisms. In: Proceedings of the 13th Annual Symposium on Discrete Algorithms, pp. 991–999.
Ausubel, L., 2003. Auction theory for the new economy. In: Jones, D.C. (Ed.), New Economy Handbook. Academic Press, pp. 124–162.
Ausubel, L., 2004. An efficient ascending-bid auction for multiple objects. Amer. Econ. Rev. 94, 1452–1475.
Ausubel, L., Cramton, P., 2002. Demand reduction and inefficiency in multi-unit auctions. Mimeo.
Ausubel, L., Schwartz, J., 1999. The ascending auction paradox. Mimeo.
Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R.V., 2011. An ascending Vickrey auction for selling bases of a matroid. Oper. Res. 59, 400–417.
Bresky, M., 2009. Revenue and efficiency in multi-unit uniform-price auctions. Mimeo.
Brusco, S., Lopomo, G., 2002. Collusion via signalling in simultaneous ascending bid auctions with heterogeneous objects, with and without complementarities. Rev. Econ. Stud. 69, 407–436.
Brusco, S., Lopomo, G., 2008. Budget constraints and demand reduction in simultaneous ascending-bid auctions. J. Ind. Econ. 56, 113–142.
Chakraborty, I., Engelbrecht-Wiggans, R., 2005. Asymptotic prices in uniform-price multi-unit auctions. Econ. Theory 26, 983–987.
Chen, N., Ghosh, A., 2011. A market clearing solution for social lending. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence, pp. 152–157.
Chen, N., Karlin, A., 2007. Cheap labor can be expensive. In: Proceedings of the 18th Annual Symposium on Discrete Algorithms, pp. 707–715.
Cramton, P., Day, R., 2008. The quadratic core-selecting payment rule for combinatorial auctions. Mimeo.
Day, R., Milgrom, P., 2008. Core-selecting package auctions. Int. J. Game Theory 36, 393–407.
Day, R., Raghavan, S., 2007. Fair payments for efficient allocations in public sector combinatorial auctions. Manage. Sci. 53, 1389–1406.
Draaisma, T., Noussair, C., 1997. Optimal bidding in a uniform price auction with multiunit demand. Econ. Letters 56, 157–162.
Edelman, B., Ostrovsky, M., Schwarz, M., 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. Am. Econ. Rev. 97, 242–259.
Elkind, E., Sahai, A., Steiglitz, K., 2004. Frugality in path auctions. In: Proceedings of the 15th Annual Symposium on Discrete Algorithms, pp. 701–709.
Engelbrecht-Wiggans, R., Kahn, C., 1998. Multi-unit auctions with uniform prices. Econ. Theory 12, 227–258.
Freedman, S., Jin, G., 2011. Learning by doing with asymmetric information: evidence from Prosper.com. Mimeo.
Goldberg, A., Hartline, J., Karlin, A., Saks, M., Wright, A., 2006. Competitive auctions. Games Econ. Behav. 55, 242–269.
Grimm, V., Riedel, F., Wolfstetter, E., 2003. Low price equilibrium in multi-unit auctions: the GSM spectrum auction in Germany. Int. J. Ind. Organ. 21, 1557–1569.
Hu, S., 2009. Unit-price procurement auctions: structural results and comparisons. Ph.D. thesis. University of Michigan.
Immorlica, N., Karger, D., Nikolova, E., Sami, R., 2005. First-price path auctions. In: Proceedings of the 6th ACM Conference on Electronic Commerce. ACM, New York, NY, USA, pp. 203–212.
Iyengar, G., Kumar, A., 2008. Optimal procurement mechanisms for divisible goods with capacitated suppliers. Rev. Econ. Design 12, 129–154.
Iyer, R., Khwaja, A., Luttmer, E., Shue, K., 2010. Screening in new credit markets. Mimeo.
Jun, B., Wolfstetter, E., 2004. Signaling equilibria in a multi-unit English clock auction. Mimeo.
Kagel, J., Levin, D., 2001. Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic Vickrey auctions. Econometrica 69, 413–454.
Kalai, E., Lehrer, E., 1993. Rational learning leads to Nash equilibrium. Econometrica 61, 1019–1045.

Karlin, A., Kempe, D., Tamir, T., 2005. Frugality of truthful mechanisms. In: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, pp. 615–626.
Katzman, B., 2009. Asymptotic properties of equilibrium in discriminatory and uniform price in multi-unit auctions. Econ. Bull. 29, 834–846.
Krishna, V., 2010. Auction Theory, 2nd ed. Academic Press.
Laffont, J., Maskin, E., Rochet, J., 1987. Optimal nonlinear pricing with two-dimensional characteristics. In: Hurwicz, L., Groves, T., Radner, R., Reiter, S. (Eds.), Information, Incentives and Economic Mechanisms. B. Blackwell, pp. 256–266.
List, J., Lucking-Reiley, D., 2000. Demand reduction in multiunit auctions: Evidence from a sportscard field experiment. Am. Econ. Rev. 90, 961–972.
Maskin, E., Riley, J., 2000. Optimal multi-unit auctions. Int. Libr. Crit. Writ. Econ. 113, 5–29.
Milgrom, P., 2006. Incentives in core-selecting auctions. Mimeo.
Milgrom, P., Segal, I., 2002. Envelope theorems for arbitrary choice sets. Econometrica 70, 583–601.
Myerson, R., 1981. Optimal auction design. Math. Oper. Res. 6, 58–73.
Nautz, D., 1995. Optimal bidding in multi-unit auctions with many bidders. Econ. Letters 48, 301–306.
Noussair, C., 1995. Equilibria in a multi-object uniform price sealed bid auction with multi-unit demands. Econ. Theory 5, 337–351.
Othman, A., Sandholm, T., 2010. Envy quotes and the iterated core-selecting combinatorial auction. In: Proceedings of AAAI, pp. 829–835.
Riedel, F., Wolfstetter, E., 2006. Immediate demand reduction in simultaneous ascending bid auctions. Econ. Theory 29, 721–726.
Sano, R., 2013. Vickrey-reserve auctions and an equilibrium equivalence. Math. Soc. Sci. 65, 112–117.
Swinkels, J., 2001. Efficiency of large private value auctions. Econometrica 69, 37–68.
Talwar, K., 2003. The price of truth: Frugality in truthful mechanisms. In: Proceedings of the 20th Annual Symposium on Theoretical Aspects of Computer Science, pp. 608–619.
Varian, H., 2007. Position auctions. Int. J. Ind. Organ. 25, 1163–1178.
Vohra, R., Malakhov, A., 2005. Optimal auction for capacitated bidders—a network approach. Technical Report. Working Paper. Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University.
Weber, R., 1997. Making more from less: Strategic demand reduction in the FCC spectrum auctions. J. Econ. Manag. Strategy 6, 529–548.