

Chakraborty, Subrata; Mengersen, Kerrie; Fidge, Colin; Ma, Lin; Lassen, David

Article

A Bayesian Network-based customer satisfaction model: A tool for management decisions in railway transport

Decision Analytics

Provided in Cooperation with:

Springer Nature

Suggested Citation: Chakraborty, Subrata; Mengersen, Kerrie; Fidge, Colin; Ma, Lin; Lassen, David (2016) : A Bayesian Network-based customer satisfaction model: A tool for management decisions in railway transport, Decision Analytics, ISSN 2193-8636, Springer, Heidelberg, Vol. 3, Iss. 1, pp. 1-24,
<https://doi.org/10.1186/s40165-016-0021-2>

This Version is available at:

<https://hdl.handle.net/10419/185051>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>

RESEARCH

Open Access



Models of sequential decision making in consumer lending

Kanshukan Rajaratnam^{1,2*†} , Peter A. Beling^{1,2†} and George A. Overstreet^{3†}

*Correspondence:

kanshukan.rajaratnam@uct.ac.za

[†]Kanshukan Rajaratnam, Peter A. Beling and George A. Overstreet contributed equally to this work

¹ Department of Finance and Tax, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa

Full list of author information is available at the end of the article

Abstract

In this paper, we introduce models of sequential decision making in consumer lending. From the definition of adverse selection in static lending models, we show that homogenous borrowers take-up offers at different instances of time when faced with a sequence of loan offers. We postulate that bounded rationality and diverse decision heuristics used by consumers drive the decisions they make about credit offers. Under that postulate, we show how observation of early decisions in a sequence can be informative about later decisions and can, when coupled with a type of adverse selection, also inform credit risk. We show through two examples how lenders may use such information in setting their offer rates.

Keywords: Sequential decision making, Consumer loans, Credit risk

Background

In consumer lending, portfolio managers typically have access to a scorecard used to forecast default probability for each applicant. Scorecards are built using historical data on loan accounts and their respective performance data. The inputs into the scorecard include financial, demographic and other personal information about each applicant. The output of the scorecard is a real-valued score for each applicant, which can then be mapped to a probability of default [see Hand and Henley (1997) for scorecard construction]. Similarly, portfolio managers may have access to a scorecard forecasting applicant responses to offers. In turn, forecasts of default probabilities and response probabilities serve as inputs to business metric functions such as expected profit.

In setting loan prices (or loan interest rates), loan portfolio managers face a trade-off between response and risk. Consumers prefer lower loan rates and hence lower loan rates results in higher take-up of the products, but lower profits for each account. Loan pricing is further complicated by the phenomenon of *adverse selection* in which the default rates of individuals who accept a loan offer may be higher than that of those who decline the offer, all other factors being equal (Phillips and Raffard 2009). Adverse selection is thought to be the result of information asymmetry. Credit bureau reports and public records, which lenders use as input for credit risk and response models, may not reflect the circumstances and immediate financial needs of the borrower. Additionally,

there may be subtle relationships between price elasticity and adverse selection (Oliver and Thaker 2013).

A portfolio manager may view a subset of the consumer population as homogenous as a result of the observable information available to him. However, private information held by individuals in the population subset differentiate their risk profiles. Portfolio managers do not have access to such private information. It would then seem obvious that reducing the information asymmetry between borrowers and lenders will result in more targeted marketing of credit products with appropriate rates. Typically, public information used by portfolio managers are those that are input into scorecards, i.e., financial, demographic and other personal information. The flat maximum effect indicates that, given roughly the same inputs, there is little difference in the performance of scorecards constructed using a variety of modeling approaches (Overstreet et al. 1992), i.e., new data sources or new variables need to be found in order to improve scorecard performance.

An important line of research in human decision making is bounded rationality. Bounded rationality describe how a decision is made rather than the outcome of that decision (Selten and Gigerenzer 2002, p. 4). It is the idea that decision makers are limited by the available information, time, cognitive ability, and the manageability of the problem. Humans use heuristics to make decisions, which are simple rules, but often lead to decision errors. Kahneman and Tversky were among the first to establish cognitive basis for errors arising from decision heuristics [see Tversky and Kahneman (1973) and Tversky and Kahneman (1974)]. Limited cognitive ability and incomplete information are some of the reasons for errors in decision making by human subjects. Many experiments have been conducted to reveal the decision heuristics used by human subjects in various classes of decision problems [see Winkler and Murphy (1973) for more]. Given the diversity in the decision heuristics used by human subjects, this forms a new source of data that may be used to improve scorecard performance.

One particular class of decision problems is the sequential decision problem in which agents are required to make a sequence of binary decisions. Sequential decision problems are of particular interest to consumer lending because consumers are often faced with a sequence of loan offers for which they make take/no take decisions. In this paper, we postulate that inference about the decision heuristics used by consumers when accepting or rejecting a loan offer may provide a new source of information for lenders. For example, decision heuristics could provide added information on borrower take/no-take behavior, thereby reducing the information asymmetry between lenders and borrowers. In particular, we show how observation of early decisions in a sequence can be informative about later decisions and can, when coupled with a type of adverse selection, also inform credit risk.

The paper is organized as follows. “[Adverse selection](#)” section extends the definition of adverse selection from Oliver and Thaker (2013) to a sequential offer setting that will serve as the basis for further study of the borrower’s decision process. “[Bounded rationality](#)” section discusses bounded rationality in human decision making and reviews literature on categorizing agents by their sequential decision making behavior. “[Problems involving sequential decisions in lending](#)” section introduces two sequential decision problems in the consumer lending space. The first decision problem relates to auction

mechanisms for peer-to-peer lending. Lenders cognizant of decision heuristics in the context of consumer lending may offer a lower bid rate and, hence, win the bidding process. We derive policy implications for a marketplace desirous of increasing borrowers' utility through lower interest rates. The second decision problem, set in the context of direct mail, is that of a lender required to choose when to market offers relative to the competition. We show how the lender may incorporate information learned about the decision heuristics of individual consumers. “**Conclusion**” section offers concluding remarks and suggestions for further research.

Adverse selection

In this section, we introduce basic notations, followed by the mathematical definition of adverse selection. We then extend notions of adverse selection to timing adverse selection (TAS) and provide motivation for the study of consumers' decision making process in the consumer lending space.

Suppose a portfolio manager has access to a homogenous population to which a credit product is marketed. We say a population is homogenous when members of the population have no observable differences between them. We use vector \bar{x} to denote past behavioral, financial and demographic information about each member of this population and only denotes information observable by the lender. We call \bar{x} , the characteristics vector. The portfolio manager makes an offer of credit with rate r . Once an offer is made, some subset of the population, the *Take* population, will accept the offer and open an account. Let T denote the event that an individual takes up an offer; so T^c is the event the individual declines the offer.

Suppose this is a simple loan account, where a unit of loan is lent to each account holder, and the account holder is required to repay the unit of loan plus the interest on the loan, $1 + r$, at the end of a specified time period. We assume there are only two mutually exclusive and exhaustive performance outcomes, G and B . The event G is associated with a *Good* customer, being one who does not default within the specified time period and repays the loan in full. The event B denotes the performance of a *Bad* customer, being one who is not Good. We denote the probability of default for a borrower with characteristic vector and offer rate r as $p(B|\bar{x}, r)$. The conditional probability of default for the Take population is then written as $p(B|T, \bar{x}, r)$.¹

Oliver and Thaker (2013) define adverse selection as

$$p(B|T, \bar{x}, r) > p(B|\bar{x}, r). \quad (1)$$

Equation 1 states the probability of a member in the Take population defaulting is higher than the probability of default in the general population, i.e., both the Take and Non-Take population. We use the total probability theorem to obtain the Bads among the Non-Takes, i.e.,

$$p(B|\bar{x}, r) = p(B|T, \bar{x}, r)p(T|\bar{x}, r) + p(B|T^c, \bar{x}, r)p(T^c|\bar{x}, r). \quad (2)$$

¹ At each instance of an offer, a homogenous subset of the total population (i.e., conditioned on the characteristic vector) offered the same rate (i.e., one rate at each instance of offer) will have some who take up the offer and some who do not. This take-up of an offer (or not taking up an offer) may be a result of unobservable information (i.e., unobservable to the portfolio manager). One may think of the random variable T as a function of some unobservable information provided by the take action that is independent of the characteristic vector and rate.

Following (Oliver and Thaker 2013), Bayes' rule can relate the conditional probability of Bad of a Take to the conditional probability of Take by a Bad, i.e.,

$$\frac{p(B|T, \bar{x}, r)}{p(B|\bar{x}, r)} = \frac{p(T|B, \bar{x}, r)}{p(T|\bar{x}, r)}. \quad (3)$$

Since, $p(T|\bar{x}, r) + p(T^c|\bar{x}, r) = 1$, combining Eqs. 1 and 2 results in the following inequality:

$$p(B|\bar{x}, r) > p(B|T^c, \bar{x}, r). \quad (4)$$

Equation 4 indicates the Non-Take population have a higher credit quality than the total population.

In defining Eq. 4, we assumed the portfolio manager makes a one-time offer of credit. Suppose instead of a one-time offer strategy, the portfolio manager markets repeatedly. At each of a finite number of epochs, the manager has the option to market to individuals who have not previously taken an offer. Below we show that, due to adverse selection, the credit quality of those not-taking up any prior offers improve after every marketing instance.

From Eqs. 1 and 3, it follows that,

$$\frac{p(B|T_1, \bar{x}, r_1)}{p(B|\bar{x}, r_1)} = \frac{p(T_1|B, \bar{x}, r_1)}{p(T_1|\bar{x}, r_1)} > 1, \quad (5)$$

where T_i is the random variable indicating take-up at the i th offer and r_i is the offer-rate in the i th marketing instance. It follows from Eqs. 4 and 5, that the credit quality of the non-take population after the first marketing instance is higher than the credit quality of the population prior to the first marketing instance, i.e., $p(B|T_1^c, \bar{x}, r_1) < p(B|\bar{x}, r_1)$, where T_i^c indicates the event a borrower declines the i th offer. Suppose the portfolio manager markets a second time to those who did not take up the offer in the first marketing instance. It follows from Eqs. 1 and 3 that,

$$\frac{p(B|T_2, T_1^c, \bar{x}, \bar{r}_2)}{p(B|T_1^c, \bar{x}, \bar{r}_2)} = \frac{p(T_2|B, T_1^c, \bar{x}, \bar{r}_2)}{p(T_2|T_1^c, \bar{x}, \bar{r}_2)} > 1, \quad (6)$$

where \bar{r}_i is a vector of all past and current offers, i.e., $\bar{r}_i = \{r_1, r_2, \dots, r_i\}$.

Equations 4 and 6 can both be generalized for the i th marketing instance, i.e.,

$$p(B|T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i) > p(B|T_1^c, \dots, T_{i-1}^c, T_i^c, \bar{x}, \bar{r}_i) \quad (7)$$

and

$$\frac{p(B|T_i, T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i)}{p(B|T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i)} = \frac{p(T_i|B, T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i)}{p(T_i|T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i)} > 1. \quad (8)$$

Note that Eqs. 7 and 8 are extensions of Eqs. 1 and 3. Equation 7 implies that due to adverse selection, the credit quality of successive non-take population improves after each marketing instance. This is due to the higher probability of Take among Bads than

the general population at each marketing instance. Furthermore, Eq. 8 indicates there is a time component to adverse selection. We call this time dependent characteristic of adverse selection, *timing adverse selection* (or TAS). Note, thus far the vector of offer rates in \bar{r}_i has not been specified.

Suppose some members of a marketed population decline all prior offers, it follows from Eq. 8,

$$p(T_i|B, T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i) > p(T_i|T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i). \quad (9)$$

It follows since $p(T_i|T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i) \geq 0$, the probability of take-up for Bads is strictly positive and is greater than the probability of take-up among the general population. In such a scenario where some members of a marketed population decline all prior offers, there is a positive probability at each marketing instance of a Bad declining all prior offers and taking up the latest offer, i.e.,

$$p(T_i|B, T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i) > 0 \quad \forall i. \quad (10)$$

Thus far, we have shown that given sequential offers, there is a positive probability at each marketing instance of Bads declining all prior offers and taking up the latest offer. This is derived from notions of adverse selection (see Eq. 1) as discussed in Oliver and Thaker (2013). However, Eq. 2 may be shown through an alternative classifier such as logistic regression built using the characteristic vector, past offers and performance data from a dataset built from an experiment. In order to create this data set, a portfolio manager requires access to a homogenous population. Given a homogenous population with characteristic vector, \bar{x} , a portfolio manager may make a sequence of offers, where each offer in the sequence is only made to the subset of the population who decline all previous offers. Using this data, the portfolio manager may estimate the probability of Take and Bad given past declines, i.e., $p(T_i \cap B|T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i)$ through a logistic regression scorecards.² If $p(T_i \cap B|T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i) > 0$, it follows that $p(T_i|B, T_1^c, \dots, T_{i-1}^c, \bar{x}, \bar{r}_i) > 0$ (i.e., Eq. 2).

Equation is in line with observations of real-life subjects faced with sequential decision making problems in the context of consumer lending. Observations have shown homogenous subjects taking up offers at different instances of a sequence. In the next section, we take a borrower's view of receiving a sequence of offers and discuss the decision heuristic observed in similar sequential decision making problems.

Bounded rationality

In “Adverse selection” section, we considered the case of a portfolio manager repeatedly marketing a credit product to a non-take population. The non-take population is updated after every offer is made. We showed under Oliver and Thaker's definition of adverse selection [see Oliver and Thaker (2013)], the credit quality of the non-take population improves monotonically with marketing instance and that there is a timing aspect to adverse selection.

² In order to estimate this probability, the portfolio manager requires all information regarding every offer made to each potential borrower. This requires each potential borrower to reveal every offer made by all institutions. The portfolio manager may make credit cards offers to a set of potential borrowers through direct mail channel, while incentivizing these potential borrowers to reveal offers from competitive institutions, whether these competitors offers were taken up or not. Credit performance of taken-up offers may be observed in credit bureau data.

Suppose now we take a borrower's view. A borrower receives a sequence of offers from a lender. We assume the offer expires before the next offer arrives. As each offer arrives, the borrower is required to make a decision on whether to take the offer. If an offer not taken, the borrower waits for the next offer. When rejecting an offer, the borrower risks the chance of receiving only lower quality offers in the future. A broad set of literature indicates that humans do not necessarily make rational decisions when faced with sequential decision problems, in part because of our bounded ability to take in information and limited cognitive abilities. Such limitations are known as *bounded rationality*. Bounded rationality may explain why homogenous borrowers accept credit offers at different point of time when faced with a sequence of offers. We provide an example, in the form of a well-studied problem known as the secretary problem, of how human subjects make decisions in a sequential decision problem setting.

Decision heuristics

The secretary problem, also known as the dowry problem, is a well-studied sequential decision problem involving optimal stopping theory. The secretary problem in its simplest form is as follows (Ferguson 1989). Suppose a manager wishes to fill a secretarial position. There is only one such position available, for which there are N applicants. The manager is aware of the number of applicants. We assume the applicants can be ranked from best to worst candidates without ties. The applicants are then interviewed sequentially and in a random fashion. Once an applicant is interviewed, the manager is required to make a decision to hire the applicant or not. If the applicant is hired, no further interviews takes place. However if the applicant is not hired, the decision maker interviews the next candidate. Rejected applicants cannot be recalled. The objective of the manager is to hire the best possible applicant. After each interview, the manager faces a trade-off, i.e., the manager could hire the current interviewee and risk the chance that a better applicant would have arrived later on in the interview process, or not hire the current interviewee but no higher quality applicant arrives later.

The optimal solution can be described using the idea of a *candidate*. An applicant is a candidate if he or she is the best applicant interviewed thus far. The optimal solution is then for the manager to reject the first $h - 1$ applicants, some integer $h \geq 1$, and then choose the next candidate (Ferguson 1989). Let N denote the number of applicants. For $N > 1$, the probability of selecting the best applicant is,

$$\begin{aligned}\phi_N(h) &= \sum_{j=h}^N p(\text{jth applicant is the best applicant and is selected}) \\ &= \sum_{j=h}^N \left(\frac{1}{N}\right) \left(\frac{h-1}{j-1}\right) \\ &= \left(\frac{h-1}{N}\right) \sum_{j=h}^N \frac{1}{j-1}.\end{aligned}$$

The optimal solution is $h^* = \operatorname{argmax}_t \phi_N(t)$. This is easily solved for small values of N . As $N \rightarrow \infty$, $h^* = N/e$ (Ferguson 1989). It follows that for large values of N , it is approximately optimal for the manager to interview 36.8% of the applicants and then select the

next applicant better than all previously interviewed applicants. The probability of successfully choosing the best candidate is approximately 36.8% [see Ferguson (1989), and Gilbert and Mosteller (2006) for more]. Stewart (1981) extended the secretary problem to one where the number of options is unknown. Under the assumption arrival times of each option is independent and identically distributed exponential random variable, the probability of choosing the best candidate with such a policy is $1/e$, which is the asymptotic optimal probability value for when the length is known (Stewart 1981).

Because of bounded rationality and behavioral biases, humans do not necessarily make decisions in a rational manner. Experiments in decision making with real-life subjects have shown diverse decision making heuristics. A field experiment by Seale and Rapoport (1997) is particularly important because it demonstrates that when people were presented with the secretary problem, they did not generally behave optimally but rather in fashions that could be explained as mixtures of three decision heuristics, each with a parameter. The decision making strategies reported by Seale and Rapoport (1997) are:

1. *Cutoff rule* Reject the first $h - 1$ applicants and then hire the next candidate.
2. *Successive non-candidate rule* Hire the first candidate who follows h successive non-candidate applicants since the last candidate.
3. *Candidate counting rule* Hire the h th candidate.

Note that of the three decision rules, only the cutoff rule is optimal, and then only if the correct parameter is chosen. Seale and Rapoport (1997) observed that human subjects seemed to follow a mixture of rules, with mixture weights and parameter values varying across individuals.

We speculate that multiple decision heuristics are in use by individuals responding to sequential credit offers. Such decision heuristics found among borrowers might explain why timing adverse selection occurs in practice. Furthermore, values for heuristics parameters might correlate with notions of patience on the part of the borrow, an idea explored below.

Credit hunger

In a field experiment with low to moderate income households, Meier and Sprenger (2010) tested whether time preferences can explain credit behavior. They measured time preferences of individuals through choice experiments. The choice experiment outcomes were then matched to credit report and tax return data. After controlling for disposable income and other characteristics, less patient individuals were found to have lower credit scores and higher default rates. While Meier and Sprenger's field experiment did not control for credit score, we posit that even when individuals do not have any observable differences, impatient consumer behaviors lead to higher default risk. We call this *credit hunger*.

If, as we speculate, credit hunger exists in consumer credit populations, there would be value in recognizing individuals with that characteristic. Methods for learning decision strategies from the observation of actions could provide such an ability. In the next section, we introduce recent work in machine learning that addresses related problems.

Behavior-based agent recognition

Suppose there exists a set of decision heuristics, similar to those observed by Seale and Rapoport (1997), governing consumers' decision making processes. Any lender that could gain the ability identify the decision heuristics being used by individual borrowers might then be able to achieve an advantage in lending strategy, relative to competitors without that ability. Recent work in machine learning has addressed a class of problems called behavior-based agent recognition (BAR), which center on the recognition of decision strategies (or the identity of agents) based on observation of decisions made by agents in sequential problems.

Qiao and Beling (2013) address the BAR problem by modeling the decision problem faced by agents as a Markov decision process (MDP). They use inverse reinforcement learning (IRL) to learn the reward vector of the MDP from the observed actions of the agents. The reward vector is, in turn, used as the feature space for supervised and unsupervised learning of decision agent identities. On several problems, feature spaces constructed from rewards learned from IRL outperform those constructed directly from observed actions (Qiao and Beling 2013). For the secretary problem, Qiao and Beling (2013) conduct a simulation experiment in which a distinct base parameter value was applied to each heuristics rule from Seale and Rapoport (1997). In addition, random noise was added to actions of the decision agents. The feature space learned from IRL resulted in clusters with high-accuracy relative to ground truth. The method did not require inputs on any description of the decision heuristics as a basis for recognition.

Suppose historical data of consumers' accept/reject decisions for a sequence of offers was available, including related historical account performance. In such a scenario, using Qiao and Beling's IRL model-based method, it might be possible to cluster consumers based on their decision heuristics. In addition, using the historical account performance, one could relate risk and response behavior to individual decision heuristics as well as the historical proportion of the borrower population using each decision heuristics. Furthermore, in identifying decision heuristics of historical population, distribution of parameter values for each decision heuristics could be estimated. A portfolio manager with access to such information might then incorporate his knowledge of the borrowers' decision heuristics in the consumer loan offer strategy.

In the next section, we discuss the impact of decision heuristics and parameter value information on a lender's decision.

Problems involving sequential decisions in lending

In this section, we introduce two sequential decision problems found in consumer loan settings. In the first problem, we introduce the lending process in a social lending platform, where lenders offer loans to borrowers through a bidding process. We model the offer policies of portfolio managers cognizant of notions of credit hunger and the resulting impact on the final-rate offered to the borrower. Whereas in the first problem, lenders were merely cognizant of credit hunger, in the second problem we assume lenders have access to greater information such as the distribution of decision heuristics and the distribution of heuristics parameter values found in a borrower population. The portfolio manager is required to decide whether to market a credit product to a homogenous

population. We assume each member of the marketable population receives a sequence of offers until an offer is taken-up by that member. The portfolio manager is required to decide whether to market in one of the instances of the offer sequence, and is required to decide on the offer rate if the product is marketed.

Social lending

Social lending offers an avenue for consumers to borrow money outside of the traditional banking system, where money is borrowed from lenders wanting to earn higher rates on their investment than through other accessible investment vehicles. Typically, these lenders are consumers. An internet marketplace provides the platform to bring together borrowers and lenders in order to benefit both parties. Such lending practices have been growing in many markets. In the United States, Peer-to-peer lending crossed the \$1 billion in outstanding loan amount in 2012.³ In China, regulatory tightening of bank credit has resulted in the growth of peer-to-peer social lending.⁴ Generally, potential borrowers register on a social lending site, and list both their details and loan requirement. Lenders, then compete to provide the loan at a competitive rate. The marketplace specifies the mechanism from which the rates are set with different rate setting mechanisms used by different marketplaces. Until recently, an auction mechanism for rate setting was used by the largest marketplace in the United States, Prosper.com. In this section, we model the impact of credit hunger on the offer rates in a social lending setting with an auction mechanism.

Prosper is the first peer-to-peer lending marketplace in the United States, currently with over two million members and \$692 million of funded loans.⁵ Prosper offers unsecured loans with fixed rates. The loans are fully amortized over the lending periods of 3 or 5 years. Prosper's current mechanism works as follows. A borrower creates a loan request, specifying the purpose of loan and the loan amount. A customer's specific interest rate is calculated using Prosper's internal models and listed for potential lenders to view. Lenders then compete to provide portions of the loan on first-come basis. This is known as a posted-price mechanism.⁶ Prior to December 20, 2010, Prosper followed the auction model in setting the lending rate. In this mechanism, the borrower lists an amount and a reserve rate. The reserve rate is the maximum rate, she is willing to take on for the loan. Lenders then bid on both the loan amount and an offer rate. At the close of the bidding process, the loan application is considered successful if the total loan amount bid by the lenders is no less than the requested amount. Only lenders bidding lower offer rates than the requested reserve rate are considered in determining the total loan amount. This mechanism is a uniform price mechanism where each winning lender receives the same rate.

In Prosper's auction mechanism, the lending process has a 2-week bidding period. However, the borrower has the option to either close the bidding process once the total loan amount bid by the lenders reaches the requested amount level, or wait until the

³ <http://techcrunch.com/2012/05/29/peer-to-peer-lending-crosses-1-billion-in-loans-issued/>.

⁴ <http://www.bloomberg.com/news/2012-07-23/china-shadow-bankers-go-online-as-peer-to-peer-sites-boom.html>.

⁵ <http://www.prosper.com/about/>. Data accessed on the 14th of February, 2014.

⁶ <http://www.lendacademy.com/prosper-com-ending-their-auction-process-dec-19th/>. Accessed on the 28th of February, 2014.

end of the official bidding period. We note that there is a de-facto signalling process that occurs when a borrower does not close their position immediately after the requested loan amount level has been reached. The observation that a borrower has not closed the position allows for lenders to lower their bid in recognition of possible lower default risk of the borrower.

The Prosper auction environment provides a natural environment to study timing adverse selection, particularly credit hunger. We show credit hunger has policy implications for a social lending marketplace wishing to provide greater benefits to the borrower. Chen et al. (2013) analyzed Prosper's mechanism as a game of complete information that fully characterizes the Nash equilibria found in such mechanisms. In contrast, we demonstrate that the provision of providing the lender with incremental information relating to credit hunger may result in lower rates for lower risk borrowers.

Notation, bidding process and Nash equilibrium

In this section, we define notation and the basic model required to understand the Nash equilibrium rates when lenders are cognizant of credit hunger.

We follow Ceyhan et al. (2011) in modeling the bidding process as a three time step process. Suppose a borrower wants to borrow an amount D and specifies a reserve interest rate of R . Both D and R are publicly listed on the borrower's listing at time $t = 0$. Each competing lender, L_i , specifies the amount she is willing to lend, a_i and her bid rate, b_i . Once the total loan amount bid by the lenders with $b_i \leq R$, exceeds the requested amount, the borrower has the option to stop the bidding process. We consider this time $t = 1$. All winners are announced at this point. If the borrower does not stop the bidding process, the lending process carries on until time $t = 2$, which is the maximum allowable time specified by the marketplace. In our model, we restrict each lender to bid at most once between time step 0 and 1. Note, time length between $t = 0$ and $t = 1$ varies for each lending process. Between time $t = 1$ and $t = 2$, any lenders not in a winning position may lower their rate bid in order gain a winning position. As each lender bids a lower rate, the latest leading lenders are announced. We assume each lender, L_i has a private rate, r_i which is the lowest rate she is willing to bid based on the characteristics of the borrower. Between time $t = 1$ and $t = 2$, lenders may only lower their bid if each bid increases their utility $u_i = x_i(p - r_i)$, where x_i is the loan amount bid by lender L_i and p is the winning rate. In this section, we drop all notions of a characteristics vector since we talk of one borrower. We assume each lender's private rate is a function of the credit risk of the borrower. Once a lender is declared a winner any time during the process, she may not pull out of the bidding process. The goal of each lender is to increase their utility. Our goal is to characterize the Nash equilibrium price when lenders are cognizant of credit hunger.

For completeness, we provide the definition of Nash equilibrium.

Definition 1 (*Nash equilibrium*) (Chen et al. 2013) A bid profile $b = (b_1, \dots, b_n)$ 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.

In the bidding process, the following allocation rules specifies the amount of loan allocated to each lender.

Definition 2 (*Allocation rules, last winner and first loser*) (Chen et al. 2013) Given a bid profile $b = (b_1, \dots, b_n)$, order lenders such that $b_i \leq b_{i+1}$ for all i . Let $k = \min\{j | \sum_{i=1}^j a_i \geq D, j = 1, \dots, n\}$. The allocation 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 $i > k$, where x_i is the amount borrowed from lender L_i . We refer to L_k as the last winner and L_{k+1} as the first loser. Let Δ be the list of winners when all lenders bid their true rate. We denote α and $\alpha + 1$ as the index of the last winner and first loser in Δ .

Note that $\sum_i x_i = D$. In Prosper's mechanism, the lending rate, $p = r_k$ when the last winner L_k does not fully utilize her budget (i.e., $x_k < a_k$), and $p = r_{k+1}$ when the last winner L_{k+1} fully utilizes her budget.

In addition, we require the following definition in order to define the price p in a Nash equilibrium.

Definition 3 (β) (Chen et al. 2013) Suppose we order all lenders indexed in a non-decreasing order of their true interest rates. For each $L_j \in \Delta$, let L_β be the last winner in Δ 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 $\sum_{i=1, i \neq j}^k a_i \geq D$. Define $\beta = \max_{L_j \in \Delta} \beta_j$.

Chen et al. (2013) provide bounds for the final price p .

Lemma 1 (Chen et al. 2013) *The price p in any Nash equilibrium b satisfies $r_{\alpha+1} \leq p \leq r_\beta$. Furthermore, $p = r_j$ for some L_j with $r_{\alpha+1} \leq r_j \leq r_\beta$.*

Note that in a Nash equilibrium, there is a finite set of prices which p can hold.

Suppose Nash equilibrium is reached at time $t = 1$. Assuming no lender may pull out of a winning bid, a borrower's winning rate cannot worsen between time $t = 1$ and $t = 2$. Since the winning rate may not worsen, the borrower does not face any trade-off in terms of the loan-rate. It would seem obvious that a borrower should remain in the bidding process in the hope of a lower loan-rate. However, credit hungry borrowers may close the bidding process at $t = 1$ and the winning rate at $t = 1$ is the final loan rate. On the other hand, a borrower who does not close the bidding process provides an opportunity for the lenders to bid a lower rate.

Nash equilibrium rates

Suppose some of the participating lenders are cognizant of credit hunger. We assume such lenders perceive the borrower's action as a signal of a lower borrower risk profile and hence, revise their bidding rates.

Now, we consider the case of one participating lender who is cognizant of notions of credit hunger. Suppose lender L_k is aware of credit hunger and has an initial rate of r_k^0 at time $t = 0$. Suppose at time $t = 1$, the borrower does not close the bidding process. This signals to the lender L_k that the borrower is less risky than previously thought and hence, the cognizant lender revises her private rate downwards. Let r_k^1 denote the new

rate at time $t = 1$ with $r_k^1 < r_k^0$. Let p_1 and p_2 be the winning rate at time $t = 1$ and $t = 2$ respectively. We consider the following cases:

Suppose $r_k^1 < p_1 < r_k^0$. Since $p_1 < r_k^0$, the cognizant lender is not in a winning position at time $t = 1$. However, she may revise her rate down to $p_1 - \epsilon$ for some $\epsilon > 0$ in order to move into a winning position and hence increase her utility. This re-bidding process dislodges the last winner at time $t = 1$. If the previously winning lender had bid her true rate between time $t = 0$ and $t = 1$, the process stops and $p_2 = p_1 - \epsilon$. If the winning lender at time $t = 1$ had not bid her true rate, then the auction process continues between those who have not bid their true rate and the cognizant lender, resulting in lower final rate for the borrower, i.e., $p_2 < p_1$.

Suppose $p_1 = r_k^0$, then lender L_k is the last winner. She will then lower her bid only if it results in greater utility through greater allocation of the loan to her.

Suppose $p_1 \geq r_k^0$. In this scenario, the cognizant lender will not revise her rate as any revision will have no impact on her utility as the winning rate is set by another lender.

Suppose instead, $p_1 \leq r_k^1$. Since $r_k^1 < r_k^0$, the cognizant lender is in a winning position at time $t = 1$, but is not the lender setting the final rate. Any lowering of her bidding rate will not result in an allocation change. It follows that the winning lenders at time $t = 1$ remain the winning lenders at time $t = 2$.

In the above example, only one lender was cognizant of credit hunger. In such a case, notions of credit hunger will only affect the final rate if the cognizant lender gains by lowering her rate and increase her utility in the process. This has policy implications for marketplaces wanting to benefit borrowers through lower final rates. For example, such marketplaces can set training policies to increase awareness of credit hunger. If more lenders are cognizant of credit hunger, those lenders may revise their bidding rates down for a patient borrower. This will result in greater situations where $r_i^1 < p_1 \leq r_i^0$ and thereby increasing the probability of lowering the final lending rates between time $t = 1$ and $t = 2$.

Direct mail lender's decision problem

In "[Social lending](#)" section, we showed if a lender is cognizant of credit hunger, this may result in a lower winning loan rate. Suppose instead, lenders had access to greater information such as access to historical data on credit offer decisions for consumers who were required to make decisions on a sequence of credit offers as well as their respective account performance data. The lender may then forecast the distribution of decision heuristics found in a potential borrower population and incorporate such information in the lending decisions. In this section, we discuss the impact of such forecasts on a lender's decision in a direct mail setting.

Problem definition

We setup the lender's decision problem as follows. Suppose a lender has access to a homogenous population and wishes to market credit offers. Since the population is a homogenous population, we do not include the characteristic vector \bar{x} in our analysis. The lender is required to decide whether to market or not to the population on the i th instance in a sequence of offers. If the lender decides to market on the i th instance then the lender is required to set the i th offer rate with a single objective of maximizing expected profit. The lender also has access to a historical database of similar

homogenous consumer population with information on sequence of past offers and the respective consumers' decisions on those credit offers. In addition, the lender has access to related historical account performance data.

A lender with access to borrowers' historical offer decision data may apply the BAR method on the data and categorize borrowers into groups, with each group associated with a single decision heuristics. Suppose all decision heuristics are ordered in some manner. Let $\psi_j = \psi_j(h_j)$ denote decision heuristics j with heuristics parameter h_j . In order to simplify our model, we define a candidate offer. A candidate offer is one that consumers prefer to rates on past offers. Since consumers prefer lower loan offer rates to higher rates, candidate offers are lower rates than all past offers' rates. We assume under each decision heuristics, only candidate offers are taken up by borrowers. Given historical data, the portfolio manager is able to estimate in a population, the probability of a borrower in a population using decision rule j , i.e., $p(\psi_j)$; and the conditional probability of a borrower using parameter value $h_j = k$, i.e., $p(h_j = k|\psi_j)$. Similarly, the portfolio manager may estimate default probability of a borrower conditioned on decision heuristic ψ_j , parameter h_j and Take, i.e., $p(B|\psi_j, h_j, T)$, where the random variable T indicates the event an account was opened.

The decision to take up an offer by a member of the population is dependent on the decision heuristics, the associated parameter, and a history of past and current offer rates. The decision heuristics and the heuristics parameter drives the cognitive process, while the sequence of offers determines the offer experience of the decision maker. In order to estimate the take rates, a sequence of offers must be specified. Let \bar{r}_i indicate a sequence of i offers with r_i the last offer. Hence, the probability of a Take conditioned on decision heuristics, the associated parameter, and a history of past and current offer rates may be estimated from the historical database, i.e., $p(T_i|\bar{r}_i, \psi_j, h_j)$. Suppose the last offer r_i is not a candidate offer, i.e., $r_i \geq \min_i[r_1, r_2, \dots, r_{i-1}]$. Since all offers accepted by a consumer are candidates, it follows that $p(T_i|\bar{r}_i, \psi_j, h_j) = 0 \forall \psi_j, h_j$ and $r_i \geq \min_i[r_1, r_2, \dots, r_{i-1}]$.

Lender's decision

The lender is required to make a decision on whether to make an i th offer in a sequence of offers. We assume the lender is only able to market on the i th offer. If the lender makes an offer, he is then required to set the offer rate. In order to optimize on the lender's objective of maximizing expected profit, we require an expected profit function. Suppose the lender markets an offer rate of r_i in a sequence of offers \bar{r}_i . For each loan unit, the lender earns $(1 + r_i)$ from each borrower who does not default. It recovers $D_e(1 - f_D)$ from each bad account, where D_e is the exposure of default and f_D is the fractional loss given default, with $f_D > 0$. Following (Rajaratnam et al. 2010; Beling et al. 2010), we make the simplifying assumption $D_e = 1$ in order to ensure the recovered amount is less than the original unit loan. We assume the lender funds the loan through debt with rate r_B . Furthermore, the cost of marketing to a single borrower is C . It follows that the expected profit conditional on take is,

$$E[P|\bar{r}_i, T_i, \psi_j, h_j] = (1 + r_i)[1 - p(B|\psi_j, h_j, T_i)] \\ + (1 - f_D)[p(B|\psi_j, h_j, T_i)] - (1 + r_B) - C.$$

Therefore,

$$E[P|\bar{r}_i, T_i, \psi_j, h_j] = (r_i - r_B - C) [1 - p(B|\psi_j, h_j, T_i)] - (r_B + f_D + C) [p(B|\psi_j, h_j, T_i)].$$

Since the cost of marketing to a consumer who does not take-up the offer is C , it follows that the expected profit conditioned on a decision heuristics, ψ_j and parameter h_j is,

$$E[P|\bar{r}_i, \psi_j, h_j] = E[P|\bar{r}_i, T_i, \psi_j, h_j] p(T_i|\bar{r}_i, \psi_j, h_j) - C [1 - p(T_i|\bar{r}_i, \psi_j, h_j)].$$

It follows that, the expected profit condition on a sequence of offer \bar{r}_i is,

$$E[P|\bar{r}_i] = \sum_j \sum_{h_j} E[P|\bar{r}_i, \psi_j(h_j), h_j] p(h_j|\psi_j) p(\psi_j),$$

where $p(h_j|\psi_j)$ and $p(\psi_j)$ are estimated from the historical database. Note, borrowers only take-up candidate offers. Since the lender's objective is to maximize expected profit and since all offers that are taken-up are candidate offers, the lender solves the following maximization problem,

$$\begin{aligned} \max_{r_i} & E[P|\bar{r}_i] \\ \text{s.t.} & r_i < \min_i [r_1, r_2, \dots, r_{i-1}]. \end{aligned} \tag{11}$$

Suppose $r_i = r^*$ is the profit maximizing rate, i.e., $r^* = \text{argmax}_{r_i} [E[P|\bar{r}_i]]$. The lender markets in the i th instance if $E[P|\bar{r}_{i-1}, r^*] \geq 0$.

Suppose r' and r'' are two possible candidate offers with $r' < r''$. Since $E[P|\bar{r}_i]$ increases monotonically with respect to all candidate offers r_i , it follows that $E[P|\bar{r}_{i-1}, r_i = r'] < E[P|\bar{r}_{i-1}, r_i = r'']$ for $r' < r''$, i.e., the lender prefers borrowers to accept a product with a higher offer rate given all else equal. Given borrowers will reject a non-candidate offer then the portfolio manager will offer a rate $r_i = \min_i [r_1, r_2, \dots, r_{i-1}] - \epsilon$ for some $\epsilon > 0$. We may think of ϵ as the marginal improvement in offers.

Note, since borrowers only take-up candidate offers, the rate set by the lender maximizing on expected profit is a candidate offer. We did not relate this to the quality of the candidate offer, i.e., the sequential decision model presented in this section seems to indicate a lower candidate offer rates does not attract more customers as would be expected in practice. The scenario presented here assumes the lender is cognizant of the order in which his offer lies within a sequence of offers. In our model, we assumed all consumers had received the same past offers prior to the current offer. In practice, when an offer is made, a lower offer rate attracts a higher take-up rate, i.e., the quality of the candidate offer affects the response rate. This may be attributed to each borrower receiving multiple offers from different lenders and hence, each borrower is at a different point on their sequences of offers.

Conclusion

This article extends the line of research in consumer lending to include the notion of sequential decision making. We use the definition of adverse selection presented by Oliver and Thaker (2013) to introduce the notion of timing adverse selection. We explain

this phenomenon through bounded rationality resulting in diverse decision heuristics used by consumers. Along with decision heuristics in consumer lending space, we introduced the notion of credit hunger. This was followed by introducing a method used to cluster agents based on their decision heuristics—behavior-based agent recognition. Finally, the paper illustrated the impact of credit hunger and decision heuristics on two decision problems in the consumer lending space.

As with adverse selection, timing adverse selection is not easy to measure in practice. However, sequential decision making experiments in other settings provide evidence that humans employ diverse decision heuristics, and this in turn suggests the existence of timing adverse selection. In order to categorize historical borrowers into clusters of decision heuristics, both take and non-take decision information is required. While offers taken-up by borrowers are found in credit bureau records, to our knowledge lenders do not share information on past declined offers. In addition, determining consumer lending decision heuristics requires an audit of offers, take behavior for all offers, and account performance for every offer accepted by a consumer.

While we have shown examples of the impact of timing adverse selection on a lender's decision, field experiments and further research is required in order to understand the phenomena described in this paper. In testing timing adverse selection and in determining consumers' decision heuristics, a sequence of offers need to be made and consumers' decisions recorded. Such a sequence of offers may be disrupted by other lenders marketing their own products. The set of potential borrowers must then be incentivized to reveal all competitors' offers whether they are taken up or not. Direct mail channels for credit cards lends itself to such an experiment.

Authors' contributions

All authors contributed equally. The idea was generated by Peter Beling. The work was written by all three co-authors. All authors read and approved the final manuscript.

Author details

¹ Department of Finance and Tax, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa. ² Department of Systems and Information Engineering, University of Virginia, 151 Engineer's Way, Charlottesville, VA 22904, USA. ³ McIntire School of Commerce, University of Virginia, Charlottesville, VA 22904, USA.

Acknowledgements

This work is based on the research supported in part by the National Research Foundation (NRF) of South Africa for the Grant No. 93649. Any opinion, finding and conclusion or recommendation expressed in this material is that of the author(s) and the NRF does not accept any liability in this regard.

Competing interests

The authors declare that they have no competing interests.

Received: 20 October 2015 Accepted: 6 November 2016

Published online: 24 November 2016

References

- Beling P, Overstreet G, Rajaratnam K. Estimation error in regulatory capital requirements: theoretical implications for consumer bank profitability. *J Oper Res Soci.* 2010;61(3):381–492.
- Ceyhan S, Shi X, Leskovec J. Dynamics of bidding in a P2P lending service: effects of herding and predicting loan success. In: Proceedings of the 20th international conference on world wide web. ACM; 2011. p. 547–56.
- Chen N, Ghosh A, Lambert NS. Auctions for social lending: a theoretical analysis. *Games Econ Behav.* 2013;86:367–91.
- Ferguson TS. Who solved the secretary problem? *Stat Sci.* 1989;4(3):282–9.
- Selten R, Gigerenzer G. Bounded rationality: the adaptive toolbox. Cambridge: MIT Press; 2002.
- Gilbert JP, Mosteller F. Recognizing the maximum of a sequence. In: Fienberg SE, Hoaglin DC, editors. Selected papers of Frederick Mosteller. Berlin: Springer; 2006. p. 355–98.
- Hand DJ, Henley WE. Statistical classification methods in consumer credit scoring. *J R Stat Soc Ser A.* 1997;160:523–41.
- Meier S, Sprenger C. Present-biased preferences and credit card borrowing. *Am Econ J Appl Econ.* 2010;2(1):193–210.

- Oliver R, Thaker A. Adverse selection and non-take inference with coherent risk and response scoring. *J Oper Res Soc.* 2013;64(1):70–85.
- Overstreet GA, Bradley EL, Kemp RS. The flat-maximum effect and generic linear scoring models: a test. *IMA J Manag Math.* 1992;4(1):97–109.
- Phillips R, Raffard R. Theory and empirical evidence for price-driven adverse selection in consumer lending. In: *Proceedings of the XI credit scoring conference 2009.*
- Qiao Q, Beling PA. Recognition of agents based on observation of their sequential behavior. In: Blockeel H, Kersting K, Nijssen S, Zelezny F, editors. *Machine learning and knowledge discovery in databases.* Berlin: Springer; 2013. p. 33–48.
- Rajaratnam K, Beling P, Overstreet G. Scoring decisions in the context of economic uncertainty. *J Oper Res Soc.* 2010;61(3):421–9.
- Seale DA, Rapoport A. Sequential decision making with relative ranks: an experimental investigation of the secretary problem. *Organ Behav Hum Decis Process.* 1997;69(3):221–36.
- Stewart T. The secretary problem with an unknown number of options. *Oper Res.* 1981;29(1):130–45.
- Tversky A, Kahneman D. Availability: a heuristic for judging frequency and probability. *Cogn Psychol.* 1973;5(2):207–32.
- Tversky A, Kahneman D. Judgment under uncertainty: heuristics and biases. *Science.* 1974;185:1124–31.
- Winkler RL, Murphy AH. Experiments in the laboratory and the real world. *Organ Behav Hum Perform.* 1973;10(2):252–70.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Immediate publication on acceptance
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ springeropen.com
