Referee report for “The Free Installment Puzzle”

This paper obtains a fascinating data set on consumer credit usage and estimates how demand for credit responds to interest rates. Doing so is challenging because consumers appear to use non-rational behavior in making these decisions. In particular, the paper documents the *free installment puzzle,* that consumers offered interest rates of zero very often do not take the loan. The paper offers descriptive, reduced-form and structural analysis of these data, arguing that the structural model well replicates important features of the data set.

Overall however, I was left wondering how to think about the main contribution of this paper. Certainly, the data set is very interesting and the opportunity to study consumer credit choice at this level is unique. However, the basic point that there is a free installment puzzle could be made with a much simpler paper. The free installment puzzle is definitely interesting, but it is perhaps not that surprising. If the opportunity cost of passing on the free loan is to make whatever interest rate the consumer’s bank is paying on bank accounts for 3-12 months, the dollar numbers are likely to be pretty small. Computing the amount of money the consumer is leaving on the table might be an interesting complement to the bank rate-of-return calculations.

Thinking further about the contribution of the paper, there is an element of the paper that contrasts reduced-form techniques with structural techniques and argues that the structural model is necessary to properly estimate the effect of the interest rate on demand. However, the analysis mostly consisted of showing that reduced-form techniques got a different (wrong) result than the structural approach, without enough analysis on why that was the case. Both approaches utilize the same control variables. Both use static models. Why the difference? It should not be just that the structural model “got the right functional form.” Should we think of the bank and consumer decision as a form of simultaneous equations (like supply and demand) whereas the reduced-form technique just measured the equilibrium relationship between variables? Perhaps that is a good criticism of the matching model, but not the IV estimator. Also, I wasn’t sure if the IV estimator was implemented to take advantage of the same excluded variables as were used in structural estimator. The emphasis of the discussion of the IV estimator was on the bank’s cost of capital, whereas the emphasis of the excluded variable in the structural estimation was on merchant-type and time dummies. These dummies show up in the matching estimator, but I’m not sure if a matching estimator is sufficient to address a simultaneous equations problem.

At any rate, if the point of the paper is to contrast reduced-form and structural techniques, I think a more careful comparison is in order. My view is that should not be the focus of the paper. I felt the strongest element of the paper was its approach to estimating the causal coefficient of the interest rate on demand for credit when we do not observe bank offers, which requires integrating out the bank decision, as well as credible excluded variables between the consumer and bank decisions. Given the unusual length of this paper, I could imagine separating the paper into a “simple statistics” descriptive paper documenting the free-installment puzzle (and probably including the rate-of-return calculations) and a more econometrically ambitious paper on estimating the causal effect of the interest rate in this environment, where both the consumer and bank decision determine what we observe.

A few more issues to consider:

The authors estimate the relationship between consumer characteristics and interest rates. This is important for projecting the interest rate offers that the consumer received for options they did not take. This step is crucial to the main issues of the paper. This estimation is probably biased because it is selected – we observe interest rates only for consumers that take the loan, so perhaps they took a loan offer with an interest rate that was low for some unobserved reason. It would be interesting to discuss how this regression addresses the main concerns of the paper, namely that we have limited information on what banks offer consumers.

Interestingly, the paper begins the section entitled "Reduced-form approaches …" with a theoretical model. The point of the model is to explicate what we are getting when we focus on the decision to use credit conditional on the purchase choice. My impression is that we could start with a joint distribution of all four of the endogenous variables in Equation 1 (the number of shopping trips, the choice to use a credit card, the amount to purchase and whether to use credit) as a function of exogenous variables (r,x). Then, we could use Bayes Rule to condition one endogenous on the other in whatever order we wanted in order to focus on what we can estimate, which is perhaps a more direct way to get to Equation 2. This part makes clear that Equation 3 does not give us the full effect of the interest rate, since the interest rate affects the choice to use the card and the choice to use credit (r affects c and the subscript 1, as well as a).

The authors discuss identification extensively. While I focused on the exclusion restrictions, they discuss several more elements. For instance, they do a numerical analysis to show that the objective function is concave, which seems useful. They also offer some intuition that does not rely on exclusion restrictions, but I did not completely follow this argument (pg 51), although this could be my own failing.

One of the important variables in explaining the consumer decision is the share of purchases for which the consumer uses installments (*installshare*). However, this variable strikes me as very endogenous. The use of installments is the left-hand side variable in this regression, so it shows up both on the left and right-hand side. The variable is important, and plays a central role in the discussion of results and the counter-factual analysis. I am sure that *installshare* is a great predictor, and I am sympathetic to the problems described on page 48, but I don't know how to interpret this variable appropriately.

I found the distinction between "fixed costs" (\lambda) and "option value" (\rho) somewhat arbitrary. In practice, the first are the set of parameters that enter utility directly and the second are those that enter in interaction with transaction size. In Table 3, we see that a few variables enter one and not the other, but there was no discussion of how that choice was made, either in terms of option value or in terms of my more mechanical interpretation of the two sets of parameters.

Also, I was surprised to see several variables in Table 3 reported with values at the level of e-14 or even lower. One is e-25. That normally suggests to me that my maximizer did not operate properly. Are these variables constrained to be positive? Perhaps it would be better to just report them to be zero and say that the constraint was binding. I don't think that many optimization packages can well distinguish between 2.345e-16 and zero.

I wasn't sure what to make of the goodness of fit tests on page 62. Given all of the dummy variables in the estimation strategy, I would expect the model to fit the average incidence of these variables just about perfectly. A tougher test would be to predict individual choices. For instance, look at purchases at which a consumer used an installment plan. What is the odds ratio relative to no installment plan from the model for that observation? How does it compare to what I would get if I just used the simple frequency of installments in the whole data set?

Why is sigma identified? That is not normally the case in a discrete choice model. Is it because c(a,r,d) enters the utility function without a parameter?

On page 43, what is the difference between T_{I} and T_0? Both are described as the subset of purchase dates at which the consumer purchased under installment.