The free installment puzzle

The authors analyze transaction data from a credit card company that follows individuals over time. The company offers clients installment credits, which are sometimes free. One puzzling feature is that customers pay back the free installments in advance. The authors essentially run a horse race between various reduced-form models and a more structural one. The argument is that the reduced-form models predict an upward slopping demand curve, whereas the more structural one does not.

The paper is way too long and goes into endless digressions, which just distract the reader from the main point. I believe the whole paper could fit in about 30/40 pages. There are just too many repetitions. Paradoxically, the paper lacks a clear economic motivation and should be expanded accordingly. The reference list is also particularly thin. The nature of the data makes it a hard problem to analyze, but the reader never gets a sense of why it is important to read through all this material. I found the policy experiments in the end rather disappointing and the results rather unsurprising. Given the model, there is not much of an installment puzzle. The model rationalizes most choices with iid preference shocks and fits the data rather well.

More fundamentally, there is an omitted variable problem underlying all of the impressive effort of the authors. The data comes from only one company, but customers could own multiple cards, or be offered installments by merchants in special deals. This means that the interest rate recorded in the data is not the one the customer uses to optimize. The conditional demand for credit should be downward sloping in the overall interest rate, but not necessarily in the one proposed by that particular company. It could be that a low interest rate offered by that company is matched by others and we would then see few purchases through this company, while a high observed interest rate could in fact be lower than the

competition, so we would see more purchases. I could not see this point addressed in the paper, and especially not in any formal sense. Obviously, it would require more data and of a different nature than the existing one as well as a more complex model. Given this problem, I wonder what the estimation exactly identifies.

It is unclear whether the customer knows the occurrence of free installments beforehand. One would think that for the purchase of larger items, the customer shops around and would choose the merchant with the best deal and installments would be part of it. Some part of the text appears to imply that this is not the case, but the authors should provide more evidence. If the customers were to sort and select the shop with the best deal, it would invalidate the econometric procedure.

The model incorporates utility shocks to the choice of installment terms. These are iid, extreme value distributed shocks. This is a strange choice of modeling. One would expect some correlation for these shocks, especially for installment terms, which are close. I understand that this would complicate the analysis, as this would depart from the multinomial logistic form. The authors should explain better the reason for this particular choice.
I found the discussion at the end of section 4.3 rather speculative, especially for a paper submitted at Econometrica.

Even though the data set is impressive, it only allows identifying the conditional choices, which makes the scope of the paper rather narrow. A number of important economic decisions are taken as given, for instance the timing of expenditures and their amount or the particular merchant chosen for the transaction. Given the fixed transaction cost, would that imply that consumers bunch their purchases? I understand that the current model is somehow the most complex one that can be identified by the data, but this is not a good excuse for a paper at a top journal.

I liked the discussion about the dominance assumption in its various forms. However, I found the discussion about the identification frustrating. Of course, given

all the parametric assumptions made, one would expect the likelihood to have a clear maximum, but this sounds a bit contrived. The real issue is what exactly identifies the probability of a free installment. It looks like identification only comes out of functional form, so that even the estimation of the conditional demand model appears to be somewhat arbitrary. The authors seem to have access to people within the credit card company. It would be of use to ask them to confirm that free installments are indeed offered around 30% of the time. This would strengthen the results.