Comments on the Dissertation of Ludmila Matyskova

Summary

Overall I found this to be an excellent dissertation, which would be acceptable for a PhD with minimal adjustments at my home institution. The three parts clearly vary in scope and polish, but this is to be expected and, in my opinion all of them have enough interest to be good chapters of a thesis.

The second chapter "Bayesian Persuasion with Costly Information Acquisition" is both the most ambitious and the most complete. It is an excellent piece of work, offering both a technical advance and economic insight. I fully expected both to be well published and to make a lasting contribution to this important research area.

The first chapter "Manipulation of Cursed Beliefs in Online Reviews" has the narrowest scope of the three, but still manages to make an interesting point. If it were fleshed out a little I also think it could be published in a reasonable field journal.

The final chapter "Habit Formation: An Experimental Study" is the least developed of the three, but shows great promise. The results from the data collected so far are intriguing, and with more work could also be developed into a well published and important paper.

Below I make a number of suggestions and comments specific to each paper. Primarily these are intended to be suggestions on how these papers could be developed post-dissertation as the candidate thinks about how to convert these chapters into publishable papers. Indeed many of them are speculative, and certainly should not be taken as requirements that need to be fulfilled for the dissertation to be satisfactory. I have very little to offer in terms of changes that need to be made prior to accepting the dissertation. Essentially I think the document is acceptable for the purposes of the PhD with very little modification.

Chapter 1: "Manipulation of Cursed Beliefs in Online Reviews"

The first chapter of the dissertation deals with the economically important topic of online reviews. Specifically, it asks whether a company has the incentive to use prices to manipulate the type of people who buy their product, and so the resulting online reviews. The potential mechanism that the candidate considers is one in which companies begin by charging high prices, meaning only those that really like the good buy it, who then leave online reviews. Subsequent purchasers fail to take this selection effect into account when making use of online reviews (the 'cursedness' of the title), meaning that they overestimate the quality of the good, allowing the firm to subsequently charge higher prices.

This chapter is a neat application of the well-established notion of cursed equilibrium to an interesting domain. My main comment is that, at the moment it contains a rather limited set of results (essentially one proposition) which, while fine for a dissertation chapter, might make it hard to subsequently publish.

I have three sets of suggestions here. First, the paper makes a number of specific assumptions when proving this result. Examples include:

1. Fully cursed consumers
2. One period lived consumers

3. Two periods in total
4. Reviews are based only on product quality, not the surplus generated by the purchase
5. A monopolist supplier of the good
6. No other forms of learning.

It would be of interest to know which of these components is key to the result. I believe that doing so would be more than just a box ticking exercise: the paper establishes that there are conditions under which monopolists do use price screening in the first period and ones in which there are not: it is therefore of interest to go further and try to characterize the sets of conditions under which such manipulation is a potential concern.

Second (and relatedly), the candidate shows (I think) that, in their setting, cursedness is necessary for price manipulation. Can this be strengthened to a more general result: for example is it the case that with fully rational consumers the presence of online reviews does not affect the price path of the monopolist? Or affects it in a way which is not down the manipulation of reviews?

Finally, I wonder if the author could provide some welfare comparisons. What does price manipulation do to profits and consumer surplus? Perhaps more interestingly, are online reviews Pareto improving in the presence of cursed beliefs?

In addition I have two minor points:

1. The link between cursedness and the result the candidate cites from Li and Hitt [2008] was not clear to me. In principle it sounds like updating beliefs as if the reviewer had the same preferences as the potential buyer could be rather different thing to cursedness. Could this be fleshed out?
2. The difference between this work and the cited work by Papanastasiou, Bakshi and Savva [2014] could, I think, be further clarified.

Chapter 2: Bayesian Persuasion with Costly Information Acquisition

The second chapter works with the extremely influential framework of Bayesian Persuasion: A sender provides a verifiable signal to a receiver who then decides which action to take, the results of which affect both parties. The question is: what information structure should the sender use in order to maximize their payoff? This paper extends this framework by allowing the receiver to collect additional information at a cost following the receipt of the senders signal. The chapter shows how to solve this obvious and important extension to the basic model, as well as proving a number of results – perhaps the most interesting of which is that the ability to collect more information may in fact leave the receiver worse off than if they could not.

This chapter has clearly benefitted from an enormous amount of work, and as a result is very polished. My comments are therefore limited. First I had two proposals that might broaden the appeal of the chapter when it eventually becomes a paper:

1. It would be nice to have an application which illustrates how allowing for the receiver to collect information leads to qualitatively different predictions from not doing do. Of course this is easier to say than to do, but it might be worth having a look back over the applications of the ‘standard’ Bayesian persuasion model in the literature to see if there is anything fits the bill

2. The candidate presents a method which makes this difficult seeming problem very tractable. One of the benefits that is mentioned is that this approach might also make the standard Bayesian persuasion model more tractable. This sounds promising and it might be worth putting in an example of a case which is problematic to solve using standard techniques but becomes much simpler using this new approach (as far as I could tell the chapter does not presently contain such an example.)

Some minor, more presentational comments

1. The candidate tends to explain the intuition of the no learning result through the idea that receivers with this type of cost function would only ever want to indulge in one round of costly learning. I think this needs to be better explained: on the one hand how does the set up need to be adjusted to allow for multiple rounds of learning (for example lemma 2 doesn't seem to make much formal sense in the current set up). On the other what does this have to do with no learning in equilibrium?
2. Is lemma 4 tight – i.e. and if and only if condition?
3. I found the materiel round lemma 5 to be hard to understand
4. It might be worth putting in an example of how propositions 4 and 5 help in solving a more complex example of the model
5. Footnote 19 describes 'more informative', but this concept has already been used
6. I wasn't sure what the materiel directly above section 2.7 was referring to
7. In the discussion of alternative cost functions, and which results hold, propositions 4 and 5 are not discussed

Chapter 3: Habit Formation: an Experiment

The final chapter of the thesis reports the result of an experiment looking at the formation of habits – the idea that taking an action once makes it more likely that the same action will be taken again. The authors abstract from the 'standard' explanation of habits – state dependence of utility – an concentrate instead on two models: one in which information is noisy, and the other in which there is a literal switching cost to changing from one action to another. They run a laboratory experiment in which subject make choices in two round – the basic task being one in which an action must be correctly matched to a randomly determined state, the nature of which can be uncovered through costly effort. Four treatments independently vary whether the states are correlated across the two rounds, and whether information about the true state is provided after round one. The information based story suggests that 'habits' should only be formed in the case of correlated states and no feedback, which is indeed what the authors find.

I think this is an extremely interesting experiment. The current results offer food for thought, and I think that further analysis could yield further interesting findings. My comments are as follows

1. As a matter of personal taste I found the ordering of the chapter to be confusing – I spent a lot of sections 3.2 and 3.3 getting cross because I didn't know what the authors meant by the switching costs and information processing models. I would have like to have seen these formally defined much earlier, and used to justify the hypotheses in table 3.1
2. I also wonder if the 'switching costs' model is a bit of an implausible straw man. Here is (I think) a variant of the model which I would find more plausible: There is a mental cost of 'engaging'

with the problem in each time period. So the subject can either (i) not think about the problem at all (ii) make a plan of action for both periods at time period 1 based on information available in that time period or (iii) think about the problem in both periods. Would a model of this type act like the switching cost model? Perhaps one treatment which might differentiate between this and the switching cost model the authors discuss would be one in which the states are negatively correlated.

3. A minor point but Kwah, Stevens and Woodford {2017] do not find that a 'standard' application of the rational inattention model well explains their data.
4. Subjects display surprisingly high accuracy in the experiment which may make it hard to spot correlations in behavior. This is important because negative results are as interesting as positive ones in the correlation analysis. Possibly the reason for the high accuracy is the relatively small number of tasks they complete relative to other similar experiments. It may be worth making the problem harder, perhaps by making the number of balls on the screen larger.
5. I have a feeling that if costs are based on Shannon mutual information then the regularity condition on page 55 will automatically be satisfied (but I could be wrong).
6. The current set of hypothesis does not present an exhaustive list of the things that one could look at in each model. For example, I would guess that the information model would predict that more info should be acquired in period 1 in the correlated case than in the uncorrelated case. It would be good to test these other predictions
7. Is it a worry that the realized state correlations seems to be significantly above 0.5 in the FI case?
8. I don't understand the use of the score variable. Shouldn't robust standard errors take care of this?