Comments on the Dissertation of Maxim Senkov

Summary

This is a great dissertation and obviously suffices for formal and content requirements for a PhD thesis in Economics; therefore, I recommend it for a defense. There are no binding recommendations that have to be changed.

The three chapters vary in scope and polish. However, this is to be expected and, in my opinion, all of them hold enough value to be good chapters of a thesis. The first chapter "Optimally Biased Expertise" is both the most ambitious and the most developed. It uses existing tools in information economics to show an interesting economic insight. It is an excellent work and I fully expect it to become well published.

The second chapter "Setting Interim Deadlines to Persuade" studies an interesting problem of optimal dynamic persuasion. While the chapter can still benefit from polishing its exposition, I believe that it can be published in a reasonable field journal.

The third chapter "Form of Preference Misalignment Linked to State-Pooling Structure in Bayesian Persuasion" is a short chapter, which is quite technical. This chapter has the narrowest scope of the three, but I think that its results can be of an interest to a specialized audience of a rapidly growing field of Bayesian persuasion.

Below, I offer several suggestions and comments specific for each chapter. 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.

Chapter 1: Optimally Biased Expertise

This chapter provides a rationale for why a principal might benefit from hiring an agent who is more uncertain than the principal about what the optimal course of action is.

This chapter has clearly benefited from an enormous amount of work, and as a result is very polished. My comments are therefore limited. I have a few suggestions that might broaden the appeal of the chapter:

1. The paper makes number of specific assumptions to derive its results. Examples include:
   (a) The number of actions is the same as the number of states.
   (b) Very specific (and symmetric) state-matching preferences.
   (c) There is no 'safe' action: action that is never optimal to be taken under full information, but which is optimal to be taken under uncertainty.
   (d) Finite number of actions and states.

It would be of interest to understand 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 a principal prefers to hire an agent with a more uncertain belief about the optimal action, because such an agent acquires 'more' information. It is therefore of interest to go further and to characterize the sets of conditions under which these results hold.

2. I am not sure how to think about the Shannon cost in this setting. The Shannon cost depends on the prior belief of the agent, but the prior belief of the agent is, in some sense, 'wrong'. Is it that the Shannon cost of an agent is a 'perceived' cost, given her prior belief, but the actual cost would then be different? This is rather a philosophical question and does not change the derivations, just I would like to understand how to think about this cost in this particular setting.

Comments on exposition:

1. At several instances we see statements about 'more information' (e.g. page 19, 21). How should we understand this? Is it in Blackwell sense?
2. When introducing sets $\mathcal{A}$ and Ω, you should state that you assume that these sets are finite.
3. I think it would be good to add a paragraph in the introduction that describes the particular setup of the model. As I was reading the introduction, it was quite abstract and a bit hard to grasp the description of the results.
4. I think that the paragraph (derivations) preceding Proposition 1.2 does not add much in terms of understanding Proposition 1.2 and could be left for appendix.

Chapter 2: Setting Interim Deadlines to Persuade

This chapter provides a rationale for setting up interim deadlines at which an agent reports back to an investor about a progress of a project in order to persuade the investor to invest in the project for as long as possible. Despite the enormous number of possible reporting schemes, the resulting optimal reporting scheme turns out to have a very simple and intuitive structure. I have one main comment and several minor comments, which are more directed on the exposition.

MAIN COMMENT

1. In the Full-information benchmark (Section 2.4.2) on page 79, it is written the following: "Consider now the case in which the principal knows that the first stage has not yet been completed. [...] I denote the continuation value of the principal at time t under full information conditional on the completion of first stage of the project by $V^{FI}_{t|1}$. [...] The principal's expected benefit from postponing the stopping for Δt is given by $V^{FI}_{t|1}\lambda\Delta t$."

- A minor point: From the notational perspective, I think that the amount of time should be denoted by Δ, and not Δt (the amount of time is not proportional to time t).
- A major point: I am a bit confused about these derivations. First of all, I would not call it 'postponing the stopping', because it would imply that after Δ amount of time, the investor is going to stop, which might not be the optimal continuation behaviour, even when no stage has been completed after that time. Second, my intuition about how it should work is as follows (and I am not sure whether it corresponds to the derivations made in the paper):

  For $t \in [0, T]$ and $n \in \{0, 1, 2\}$, let $V^{FI}_{t|n}$ denote the continuation value of the principal at time t under full information and under his optimal continuation behavior, conditional on $x_t = n$.

  Suppose that at time t, we have $x_t = 0$, i.e., we know that at t, no stage has been completed. Denote $t' = t + \Delta$, where $\Delta > 0$ is a (small) amount of time.

  At t, let us compare the expected benefit from investing for Δ amount of time (given the optimal continuation behaviour at t') versus the cost of investing for Δ amount of time.

  The cost of investing is $c\Delta$. The expected benefit of investing is a sum of different parts:

$$Pr[x_{t'} = 0|x_t = 0]V^{FI}_{t'|0} + Pr[x_{t'} = 1|x_t = 0]V^{FI}_{t'|1} + Pr[x_{t'} = 2|x_t = 0]V^{FI}_{t'|2} \qquad (1)$$

  One can probably argue that $Pr[x_{t'} = 2|x_t = 0] \to 0$ as $\Delta \to 0$ (probability of completing both stages within a very small amount of time is negligible), so the last term drops out. However, unless $V^{FI}_{t'|0} = 0$, the first term does not drop out. Then, the expression (1) is different from $V^{FI}_{t|1}\lambda\Delta t$, which is the expression used in the paper.
- If I am right, it is a computational mistake that should be corrected for before submitting the paper for publication in a journal. However, I do not think it changes the qualitative properties of the results. I might also be wrong. In any case, I think that the derivations should be made clearer to avoid any confusion.

MINOR COMMENTS

1. What is the importance of the exogenous time T at which the project ends for sure? Would the model have qualitatively different behaviour if no such external deadline existed and the project would stop only when the principal decides not to invest anymore?
2. I think that the setup of the model (Section 2.3.1) would benefit from a formal definition of the objective functions (or utility functions) of the agents. At the moment, the objectives are only stated verbally and I find it a bit confusing.
3. Figure 2.2: From the title of the Figure, it is not clear that it depicts the No-information benchmark.

4. On p. 81, it is written that "$P(x_\tau = 2)$ captures the belief about two stages of the project completed by τ, the random time of stopping in the future". I would avoid using the word 'belief', because it can bring confusion. Instead, I suggest to describe the term in the following spirit: "$P(x_\tau = 2)$ is the ex-ante probability that the two stages are going to be completed when the principal follows the investment schedule τ."

    - Similarly, it should be explained what $P(x_\tau|\mathcal{F}_t)$ is.

5. I cannot find anywhere the proper definition of what Ω is.

6. Since Assumption 2.1 is meant to express a lower bound on T, I think it would be better to rewrite the condition in a way that gives a clearer understanding to what that lower bound is (and provide an explanation on how we arrive to that assumption).

7. I find the statement of Lemma 2.7 unclear. What exactly does 'conditional on no completion (one completed stage/two completed stages) of the project' mean?

    - Lemma 2.7, point 2. "Conditional on one completed stage of the project, stopping the funding happens with probability zero". Does it mean that there is a zero ex-ante probability that stopping happens at $t = T$ when $x_T = 1$?

Chapter 3: Form of Preference Misalignment Linked to State-Pooling Structure in Bayesian Persuasion

This is a very theoretical paper which derives the characterization of the solution in the model of Bayesian Persuasion for a specific class of quadratic utility loss functions. Since the paper is quite technical, my comments are mainly aimed to improving the formalism.

1. I find the description in Figure 3.1 unclear. At first, I was not sure what the figures depict: whether they depict the optimal signals or something else.

2. Looking at expression (3.1) in the setup of the model, it is implicitly assumed that the signal π has a finite support (excluding continuous distributions). Is this assumption without loss of generality? If yes, it should be argued why it is so. If not, it should be mentioned why you need this assumption.

3. In the second paragraph on p. 132, it would be better to write $\rho'(\cdot) = b + \rho(\cdot)$ to make clear what elements are functions and what element is a constant. Similarly, it would be better to write $\rho(\cdot)$ instead of ρ throughout the entire text of the paper.

4. I think that it would be good to provide some intuition about the bounds on the slope of the function $\rho(\cdot)$, at which the sender strictly prefers full revelation. More specifically, where does the number '2' come from in Propositions 3.1, 3.2 and Lemma 3.1? Can that be explained on some more intuitive level?

5. I was missing some proposition that would state that an optimal signal has a state-pooling structure (the structure described in Chapter 3.5). Only by reading throughout the whole text and some proof, I realized that it is an implication of Proposition 3.3. I

think that it would be good to have a Corollary that would state that an optimal signal must follow the state-pooling structure.

6. To better understand Definition 3.2, I would appreciate for it to be accompanied by some figures showing an example of those defined terms.