The Dissertation review report on the dissertation "Essays in Strategic Information Provision" by Maxim Senkov

Maxim Ivanov, McMaster University

September 2023

Assessment

After reviewing the dissertation, I am broadly pleased with the contribution of each chapter (out of three) of the thesis. In my opinion, it certainly satisfies all necessary requirements for a Ph.D. thesis in economics and is ready for a defense. Thus, I recommend that the Ph.D. candidate defend his dissertation.

Please find below a list of my comments and suggestions to improve each chapter of the dissertation. I would like to emphasize that none of these comments is substantial for the defense and thus can be addressed when the candidate prepares his papers for submission to a publication.

Comments on Chapter 1

- The results of the papers hinge on a few crucial assumptions regarding the players' preferences and the cardinality of the state and action spaces. First, the state and action spaces are finite. Second, there is no ex-post conflict of interests between the players. Together, these assumptions preclude the conflict of interest about the optimal action driven by the difference in the players' posterior beliefs (whenever this difference is small enough). Furthermore, the player's preferences take a 'hit-or-miss' form; that is, both players get a positive payoff if and only if the action matches the state. On one side, such a specification of the model allows to isolate the effect of the player's misaligned prior beliefs on the principal's ex-ante utility from that of conflicting preferences, and the paper should emphasize it more. On the other side, the paper does not justify this form of preferences well enough, either by providing real-life examples or citing literature that uses similar preferences. In my opinion, it is one of the major weak spots of the paper and has to be addressed seriously.

- The paper has to clearly explain the driving forces behind the main results. Specifically, the example of the binary case (which aims to serve that purpose) does not explain them well. The key explanations are stated on pp. 18–19 and are rather vague. First, the paper there (and in some other places) often uses the word 'different' to explain the difference between some variables without explaining the direction of the difference (increase or decrease) or its intensity. Here are a couple

of examples:
(p. 18): "*Further, he prefers to learn more about the more probable event: the higher is the probability that the agent's prior belief assigns to $\omega = r$, the more important is precision $\pi(R|r)$ for his payoff, compared to $\pi(L|l)$. Therefore, two agents with different beliefs would acquire different information, leading to different precisions $\pi(R|r)$ and $\pi(L|l)$.*"
– First, it is not clear what 'important' means in this context. Is it the marginal effect of $\pi(R|r)$ in the agent's ex-ante payoff or something else? Second, it is important to explain how exactly the agent's prior beliefs affect his incentives to acquire information about each state and why.

(p. 18–19): "*To summarize, the agent's belief μ affects his optimal decision precisions in two ways: a more uncertain agent acquires more information (and hence makes a better decision on average) than an agent who believes one state is more likely. However, the latter is more concerned with choosing the correct action in the ex ante more likely state, while neglecting the other state.*"
– What is the overall effect of this trade-off and why? Also, what is meant by "*the latter is more concerned with choosing the correct action*"?

To sum up, given the importance of this example, the author has to carefully address the language and clarity of the content in this and other parts of the paper that explain the main results.

- I don't think that the closest paper is Ball and Gao (2021), which is an extension of Szalay (2005). In these papers, the players share common beliefs, and the agent's incentives to acquire more information (which results in more dispersed optimal actions) are driven by excluding safe (i.e., ex-ante optimal) actions from the feasible action set. In the current paper, there are no safe actions. On top of that, there is an ex-post conflict of interest in Ball and Gao (2021). In my opinion, the closest paper is Che and Kartik (2009), and the paper has to emphasize and clearly explain the differences between both the model specifications and the results in the two papers. Again, the current explanation (on p. 10) is rather brief and does not provide the full picture of the relationship between the works.

- Regarding contracting on actions/misaligned preferences, there is a paper "Resisting Persuasion" by Tsakas et al. (2021), which seems to be a counterpart of the contracting problem considered in the current paper. In their paper, the principal can impose additional penalties on herself after taking certain actions. As they show, such a modification of preferences affects the sender's incentives to disclose information and may eventually benefit the principal.

- The description of the model does not say anything about the conditions on the action and the state spaces.

- I would suggest to remove the negative result on alternative preferences specifications and focus more on justifying the current ones.

Comments on Chapter 2

This part of the thesis considers a principal-agent problem, which is a combination of the dynamic information design and the optimal stopping time problems. Specifically, the paper considers a three-stage Poisson process with known and exogenous parameters: the rate $\lambda > 0$ and the finite horizon $T > 0$. The agent runs this process, observes its realization, and ex-ante commits to the disclosure policy about the realization, whereas the principal decides at each moment of time $t \in [0, T]$ whether to provide the instantaneous payment $c > 0$ to the agent for running the process. The principal receives a fixed positive benefit v if and only if the process runs and achieves the final stage at time $\tau \leq T$.

The established results—especially those about setting the endogenous interim deadline—are interesting and by no means trivial (though some of them are rather natural). Thus, my comments are mostly about the key assumptions of the model.

- In the current specification, the most of the model parameters are exogenous. We can start with the Poisson process. Its characteristics are known to both parties (that is, there is no ex-ante uncertainty about the properties of the process) and exogenous (that is, no party can modify the rate λ and the deadline T). At the same time, this process is interpreted as a project that is run, observed, and reported by the agent. This raises numerous questions about such an interpretation and the role of the agent. Specifically, the agent's role in controlling the project is rather passive, as he cannot affect the chance or timing of its success. Thus, it is not clear what is implied by the agent's ability to 'run the project', as he cannot affect the process in any way. In my opinion, the paper has to provide convincing real-life examples of such a setup. The current example of the relationship between a start-up (the agent) and an investor (the principal) does not seem suitable for the model. First, the success of the startup critically depends on the expertise, experience, and amount of effort put by its founder and employees. Second, typically there is a substantial amount of ex-ante uncertainty about the chance of success for the start-up.

- Also, the cost of running the project c is assumed to be static, fixed, and unrelated to the rate λ (which determines the probability of the project's success). First, the paper does not provide any justification for this assumption. Second, even if we abstain from the relationship between the cost and the probability of success rate and admit the passive role of the agent (whose only function is to report the evolution of the process), it seems natural to assume that the project requires some fixed budget $B > 0$ to be potentially successful, and the principal can allocate this

budget over time to affect the agent's incentives to disclose his information over time. In other words, what is the optimal flow payoff c_t, such that $\int_0^\tau c_t dt \geq B$?

- As far as I understand (p. 89), there are multiple optimal disclosure rules for the agent. Clearly, a complete characterization of them might be intractably difficult. However, is it possible to characterize certain properties that any optimal rule must satisfy?

- To sum up, I think that the success of this part of the thesis will crucially depend on the justification of the model assumptions.

Comments on Chapter 3

This part of the thesis characterizes optimal signal structures in the standard Bayesian persuasion payoff for a specific setup: a finite state space, a continuous action space, and the quadratic preferences of players, such that the bias between the optimal actions of the sender and the receiver takes an arbitrary form. The proposed algorithm of pooling or separating states is nice and elegant and based on the observation that inducing interior posteriors is never optimal for the sender. As a result, it is sufficient to restrict attention only to the boundaries of the simplex of the posterior induced by subsets of states. As a result, it tremendously simplifies the sender's problem.

- It seems that Proposition 3.2 and Lemma 3.1 have a close relationship with the lobbying example in Gentzkow and Kamenica (2011, section V.A). This is because the underlying logic in all these results is very similar and does not depend on the cardinality of the state space.

- I would suggest a couple of other related papers that provide sufficient conditions on the sender's payoff function under which optimal signal structures are monotone partitional, which are extreme points in the space of posteriors and, hence, are not interior. The first one is "Censorship as optimal persuasion" by Kolotilin et al. (2022). The second one is my paper "Optimal monotone signals in Bayesian persuasion mechanisms" (2021). In that paper, Theorem 2 is also proved by using an algorithm of iterative pooling and separating the states. It would be interesting to compare the method proposed in this thesis to those two works.

- A comment related to the previous one is as follow. Suppose we impose some conditions on the behavior of the function $\rho(\omega)$. Is it possible to sharpen the characterization of optimal signal structure in this case?

On a separate note, I wish Maxim Senkov the best of luck and hope he will defend his thesis successfully.