

Report on the PhD Thesis:
Essays on Endogenous Information Acquisition in
Economics
Pavel Ilinov (Center for Economic Research and Graduate
Education, Charles University Prague)

February 28, 2023

The thesis presents three chapters. The common theme to all chapters is that economic agents can become better informed at a cost, and that such decision is modeled within the framework developed by the literature on rational inattention. The first chapter is more methodological and establishes an equivalence between a rational inattention game and a conformity game. The second and third chapters are more applied and both look at the problem of a principal selecting an agent. In the second chapter the principal selects and agent to delegate a decision over which the agent can acquire better information. In the third chapter the agent takes no decision and the information acquisition is on the part of the principal that seeks to select the better agent.

The three chapters are original piece of research that bring new insights to an already well developed literature. The relation with the existing literature is well described and the thesis well written. The first chapter is already published in Economics letter and the other two should be publishable in good/very good theory/applied theory journals. Based on the above, I have no hesitation in recommending the thesis to be presented in front of the defense committee in its current form. In the remaining, I review the key contributions and provide some suggestions that may be of use in future revisions or for future research.

1 Chapter

The first chapter (joint with Ole Jann) is a technical note that connects the literature on rational inattention to the one on social norms and conformity. It establishes an equivalence between the two problem when the cost of deviation from conformity is given by the Kullback-Leibler divergence measure. The paper is already published in Economics letters. Indeed, it provides an interesting result .Not only it connects very relevant/popular literatures, but also brings in the message of being cautious when interpreting data. In fact, it might be hard to identify the underlying data generating process/model.

My only remark is that it would be important to understand whether people divergence from conformity follows the cost structure imposed by the Kullback-Leibler divergence. From the analysis its use appear crucial as it guarantees that externalities in the conformity game do not distort behavior from the welfare maximizing set of choices. Is there any existing evidence of it? If not, an avenue for future research could be to think of a lab experiment that could elicit the cost of deviating from conformity.

2 Chapter

The second chapter (joint with Matveenko, Senkov, and Starkov) looks at the optimal selection of an agent to whom to delegate a decision. The available agents are heterogeneous with respect to the prior they hold regarding the state of the world, whereas their preferences are aligned with the principal. In particular, they would like to match the chosen action to the true state of the world. Once selected the agent has the ability to acquire more information and, as for the other chapters, information acquisition is modeled within the framework of the rational inattention literature.

The main insight is that the optimal delegation consists in selecting an agent who is more uncertain about the state of the world than the principal is so as to ensure that he/she will acquire more information. The interesting point is that the principal does not go all the way to select the most uncertain agent. The reason is essentially that the principal prefers the agent to acquire more information about the states that are more supported by her prior. The paper starts with an example, it moves to describe a simplified setting with two actions/two states, then it presents the general model with more actions/states. Finally, it compares the optimal delegation in this setting where misalignments is on beliefs with misalignment in preferences and with contracting on actions/ on outcomes. It also discusses extensions such as varying the underlying preferences.

Main Comments:

- The paper provides an interesting contribution that can aim to be published in a good theory outlet. It is well written and seems at a fairly advanced stage. I only provide comments with the intention to help to submission (or resubmission) to a good journal.

- Starting a paper with a nice simple example is a very good idea, but the example should provide the main insights and/or clarify the workings of the more general setting. The example provided fails to do so. It is not useful to simplify the understanding of the technical parts of the analysis, and it does not convey the main insights either as the main trade off is absent. I would suggest if possible to come up with a better example. At least one where the principal does not pick the ex-ante more uncertain agent. That should be possible within an example where agents' information after paying the cost is not perfect (a simple modification could be that an agent upon paying the cost gets either a perfectly informative signal or a mildly informative one). If a better example is hard to find, I would present the insight conveyed by the present example in a short paragraph (the point made is quite direct) rather than putting it on the spotlight.

- The main economic insight of this paper is perhaps not as novel given the cited results by Che and Kartik (2009). The setting is obviously rather different so there is no issue this work covers a gap. Yet, to clarify and increase its value more space could be spent clarifying the differences with Che and Kartik (2009) and why a priori one might expect different insights to arise. Also, you could argue that your more flexible information acquisition structure could fit better some applications.

- I find that the reader could benefit from some guidance at the conceptual level. How should we think about the heterogeneous priors? It could be that the agents and the principal have different ways/models of interpreting the world. Is any of these correct? (clearly a belief $\mu_p = 1$ is incompatible with a belief $\mu = 0$ for an agent). Given that the priors are observable should we think of this as the remaining differences after having updated observing the opinions/priors of the others? A way that seems compatible with the model presented is that the principal holds the correct prior (or believes that for that matter) and benefits from the availability of agents who, from his point of view, have incorrect priors. This clarification is important not only conceptually but also for the results. In fact, the expected utility of the principal is evaluated assuming she has the correct prior. Take figure 2.2. Hiring an agent with a belief of 0 or 1 (or close to them) does not mean that the agent knows the true state, rather that the agent is so biased that he always selects the action he believes to be true. The

expected utility for the principal in those cases is respectively .3 and .7 because she knows that those agents do not acquire information and always choose their ex-ante optimal action whose relative likelihood is based on the prior of the principal (by the way it would be more clear to mark .3 and .7 on the y axes).

- Another conceptual point that could be covered in greater details is an understanding/clarification of the broader picture one should have in mind. If agents benefit from being selected, and they are so only based on their prior (which is publicly observable) we should probably think that they have some control over of it and they might select/influence it to increase the chances of being selected more often. What can explain their heterogeneity then? One perhaps could hint at a dynamic story where the principal is also appointed over time and that the prior of the selected principal might therefore vary over time. If agents are long lived having a bias in the prior might guarantee being selected when a similarly biased principal is.

- The optimal delegation is compared to aligned delegation. Aligned delegation gives the same choice the principal deciding on her own (if she had access to information acquisition). This is just about framing, but perhaps it would highlight the benefit from misalignment even more.

Figure 2.3 and 2.4 naturally call for some comparative static exercise on μ_p . How does the gap in prior between principal and selected agent varies with μ_p ? When the gap in the prior is the larger, is it the case that the gain in utility from delegating is the larger or is there a trade off? Note that understanding the utility gain for the principal would also help rationalize where the agents would ideally like to set their priors.

The specific state matching preferences used are important for some of the results. In particular, as explained in at the top of page 42 and 44, they limit the incentive for the principal to override the selected agent. I find that these comments are quite useful to understand the setting correctly and therefore could be placed earlier on in the paper.

The conclusions try to emphasize the main point as surprising and novel. Given that some papers have already highlighted a similar message, I would perhaps try to be more specific about the new points made in the current paper.

Other comments:

- I find the section on the general case a bit less well presented. For instance, I would try to phrase differently the following sentences:

“Despite the potential confusion this enables...”

“... which do not necessarily give the reader a good idea of its features and the intuition behind this solution”

- P36: “the results of Matevenko and Mikhailishev (2021)” It might help to point out a specific proposition.

- You may consider adding a corollary to Proposition 12 that summarizes the main results of that section.

3 Chapter

The third chapter (single authored) complements the search literature to allow for endogenous information acquisition.

The main analysis deals with the case of only two options that ex-ante bring the same value, and to quadratic costs of learning. The main finding is that the decision maker incurs a fixed cost of information, and can perfectly substitute the acquisition of information across periods. The decision maker ends up selecting the two options with an equal probability. This result is labeled as no discrimination and related to the serial position effect in the psychological literature.

In the latter part of the paper, the author explores the implications of relaxing these assumptions.

Main comments:

- The introduction is fine for the purpose of the current thesis but would benefit from a revision before submitting the work to a journal. My suggestion is to start it with a more applied angle by describing the leading application of a manager that seeks to select a candidate, and to use the application as a motivation to look into endogenous information acquisition.

- The underlying search model (or category of search models) the author wishes to extend should be more clearly described/narrowed down in the introduction. For instance, in his analysis the author assumes that the decision maker needs to choose one of the options she investigates, and that the options already investigated remain available. These assumptions are key to the analysis and define the type of search and range of applications one can speak to.

- Related to the above, the author points out that what the decision maker can learn about the alternatives at each stage is different from other papers such as Steiner et al (2017) and

Miao and Xing (2020). Which applications fit better the current's paper? You could argue that it fits better the interviewing candidates example.

- I find that the take home message from this chapter is not as well formulated as the in the others. The author relates his work to the discussion in psychology about the serial order effect. What can we conclude? I do not see a clear enough message. My reading of the results is that the substitutability in the acquisition of information across periods seems special, while the author gives the impression that he wants to stress when it holds. As it is there is no clear message on the serial order effect from the main analysis because it does not characterize when the primacy or the recency effect occur.

- Based on the above it is important to identify an extension where the primacy vs recency effect race could be clearly settled. The author suggests modifying the cost function in the conclusion. This may be indeed worth exploring.

- Additionally, I find that the results in the section "an additional alternative" could be given more importance. As it is, it seems one of the many extensions. However, it seems to me that it bring some of the more valuable insights and therefore should be part of the core analysis. As a matter of fact, in this setting there is a serial order effect and the substitutability of information does not hold, which seems the more interesting and perhaps more general case.

- It would help if more intuitions as to why the substitutability in the acquisition of information holds for in the baseline. Explain clearly why the two agents case is special.

- I thought that on top of having two alternatives perhaps one would also need the options to be ex-ante symmetric, while Lemma 3 shows that this is not the case at least for small asymmetries. Why is that the case? Also, what happens if the condition in Lemma 3 does not hold?

- The fact that the two options are picked with equal probability is labeled as no discrimination and the following intuition is offered:

"Intuitively the no-discrimination result can be explained by the special intertemporal role of information in the solution."

Could you elaborate on that? One could have the intuition that it could always hold by law of iterated expectation or a martingale property. We see from the result of the analysis of the three alternatives' case that this intuition is indeed not correct.

- The analysis of the three alternative case gives the impression that it could be generalized to N options. If true, the author when revising the work should aim to have the N options

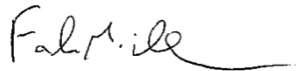
case. If not, he should point out the difficulties of extending the analysis.

Minor comments/suggestions:

- The index i is used both for the time/stage and the alternative, this might be confusing.
- Proposition 14 has 2 parts. They should be written in a more “symmetric” way. For instance, the first part mentions the manager choice, while the second does not. Also, the second part is written in a more technical way than the second.
- Pag 73: “obviously, in this scenario the manager chooses the better candidate more often”. The sentence is confusing because this ex-post is true also for the baseline case.
- The section on discounting is less interesting. It may be dropped and the result mentioned in a footnote.
- Theorem 3 is written in a bit convoluted manner, I suggest to write the statement more clearly.

I thank the thesis advisors and the Director of Graduate studies for having given me the chance to read this interesting thesis. Please do not hesitate to contact me in case I can provide any further useful information.

Yours Faithfully,



Fabio Michelucci
Professor of Economics
Department of Economics
Ca' Foscari University of Venice
Tel: +39 041 2349190
Email: fabio.michelucci@unive.it