Social Search: An Experimental Study*

Maria Bigoni† Michela Boldrini‡ Niccolò Lomys§

Emanuele Tarantino¶

May 14, 2022

Extended Abstract

Individuals seldom make search to inform a decision in isolation. In contrast, especially when facing different alternatives, they often rely on the information on others' choices and experiences (*social information*), to decide what and how much they should search about available alternatives before making their ultimate choice. In online environments, information about others' choices and experiences is abundant and readily available via online social networks, popularity rankings, and search engines. Offline, consumers rely on direct observation of, or communication with, their social connections. When individuals make search in presence of social information, many questions of interest arise. How does the information on others' choices affect individuals' search behavior? Does it alter individuals' decision on what to search for or their incentive to collect information? Do individuals over-exploit their social information or engage in excessive independent exploration? How do others' choices affect this trade-off?

In this paper, we experimentally investigate the interplay between social information and individual incentives to choose and acquire private information.[1]

*Niccolò Lomys acknowledges funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement Nº 714147) and from the ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program).

†University of Bologna; maria.bigoni@unibo.it.

‡Bocconi University; michela.boldrini@unibocconi.it.

§Toulouse School of Economics, University of Toulouse Capitole; niccolo.lomys@tse-fr.eu.

¶LUISS University, EIEF, and CEPR; etarantino@luiss.it.

[1]To our knowledge, this is the first paper to experimentally investigate information acquisition and choice in the presence of social learning.

We consider a simple variant of Weitzman (1979)'s sequential search model.[2] There are many agents. Each agent chooses between two actions. Actions' qualities are i.i.d. draws, about which agents are initially uninformed. Agents wish to choose the action with the highest quality. Agents acquire information about the qualities of the two actions via costly sequential search with recall. Searching for an action reveals its quality, but not the quality of the other action. After searching for the first action, an agent decides whether to search for the second action or not. Finally, the agent chooses an action among those he searched for.

Agents are matched in pairs. Within each pair, agents act in sequence, and we distinguish between first and second movers. First movers act in isolation, whereas second movers have social information. That is, before starting their search, second movers observe the final choice of the first mover they are matched with. First and second movers have the same preferences and, within each pair, they face the same draw of the two actions' qualities. Second movers observe neither first movers' search costs (which are i.i.d. across agents) nor search decisions.

Optimal search decisions of agents with social information (second movers) differ in three respects compared to those of isolated agents (first movers). First, second movers are not indifferent about which action to search first, whereas search movers are. In particular, second movers find it optimal to start searching from the action taken by their matched first movers. Second, the expected gain from the second search (and hence the incentive to explore) is lower for second movers compared to first movers. Third, whereas the expected gain from the second search for first movers increases as the quality of the first searched action decreases, this is not the case for second movers. In particular, for second movers, the expected gain from the second search is non-monotone (more precisely, inverted-U shaped) in the quality of the first action searched.

Our experimental design aims at identifying which are, if any, the systematic deviations of agents' actual search behavior from the theoretical benchmark. Do second movers start searching from the action taken by their matched first movers? Do agents engage in too much or too little information acquisition? Do under- or over-eager search always occurs or only in the presence/absence of social information? Under which conditions is agents' sequential search behavior efficient? What helps agents (to learn) to behave efficiently?

[2]Formally, the model is a simplified version of those in Mueller-Frank and Pai (2016), Lomys (2020), and Lomys and Tarantino (2021)

The experimental design consists of two parts. During Part 1 (lasting 10 rounds), all agents play the game in isolation, as first movers, to train. In Part 2 (lasting 20 rounds), we split agents into two groups: one group still plays the task as first movers, while the other group plays as second movers, having access to some social information. What type of social information is available to players in the second group varies across treatments. Second movers always observe the final choice of the first mover they are matched with. However, (1) matching can either be exogenous or endogenous, and (2) the information provided may, or may not, contain also some hints on the reputation of the matched first mover. First movers' reputation is measured in terms of agent's rating in Part 1: each agent's performance is rated based on the cumulative payoffs earned over the first ten rounds. We run four treatments (between-subjects design):

- *Benchmark Treatment: No Reputation and Exogenous Matching.* We do not communicate first movers' rating in Part 1 to second movers in Part 2; agents' matching across groups is exogenous (stranger matching protocol).
- *Treatment 1: No Reputation and Endogenous Matching.* We do not communicate first movers' rating in Part 1 to second movers in part 2; agents' matching across groups is endogenous (i.e., second movers can choose which first mover they want to be matched with).
- *Treatment 2: Reputation and Exogenous Matching.* We communicate first movers' rating in Part 1 to second movers in Part 2; agents' matching across groups is exogenous.
- *Treatment 3: Reputation and Endogenous Matching.* We communicate first movers' rating in Part 1 to second movers in Part 2; agents' matching across groups is endogenous.

Communicating matched first movers' rating allows second movers to make some inference regarding first movers' ability. The comparison between Treatment 2 and the Benchmark informs us regarding the impact of reputational information on first movers' ability on the second mover choice between searching and relying on social information. The comparison between Treatment 3 and the Benchmark, instead, allows us to study segregation, and test for behaviors similar to those operating in an "echo chamber" idea, keeping fixed the channel stemming from the observation of the first mover's ability. The comparison between Treatment 1 and the Benchmark allows us to study endogenous reputation formation, and check if the second movers tend to rely more on first movers who in the past proved to be reliable.

References

Lomys, Niccolò (2020), "Collective Search in Networks." *Working Paper*.

Lomys, Niccolò and Emanuele Tarantino (2021), "Identification and Estimation in Search Models with Social Information." *Working Paper*.

Mueller-Frank, Manuel and Mallesh M. Pai (2016), "Social Learning with Costly Search." *American Economic Journal: Microeconomics*, 8, 83–109.

Weitzman, Martin L. (1979), "Optimal Search for the Best Alternative." *Econometrica*, 43, 641–654.