Why Don't You Take a Leaf Out of Her Book? An Experiment on Social Search*

Maria Bigoni† Michela Boldrini‡ Niccolò Lomys§

Emanuele Tarantino¶

May 5, 2023

Abstract

We propose an experimental analysis to investigate the impact of social learning on individuals' acquisition of information before making a choice and how behavioral biases and the perceived reliability of the information source affect this process. We use a stylized version of the canonical sequential search model and solve it with and without social information. Agents without social information act in isolation, while those with social information observe a peer's choice but not their search process. We aim to explore the effects of social information on search behavior, the influence of the perceived quality of the information source, and the potential of alternative institutions to counteract behavioral biases and restore the benefits of social information. Our results can guide the development of policies that enhance the efficiency of decision-making processes by optimizing the use of social information.

Keywords: Sequential Search; Multi-Armed Bandits; Individual Decision Making; Social Information; Reputation.

JEL Classification: D8; C9; D1.

*Niccolò Lomys acknowledges funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreements N^{o} 714147 and N^{o} 714693). This work was done in part while Niccolò Lomys was visiting the Simons Institute for the Theory of Computing at UC Berkeley (program on Data-Driven Decision Processes).

†University of Bologna, CEPR, and IZA; maria.bigoni@unibo.it.

‡Bocconi University, IGIER, and RFF-CMCC EIEE; michela.boldrini@unibocconi.it.

§CSEF and Università degli Studi di Napoli Federico II; niccolomys@gmail.com.

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

1 Introduction

Social media have substantially increased the amount of information individuals have access to before they purchase a product. Information is also abundant offline by observing social connections' purchasing behavior. In principle, such a wealth of information should help people make better and more informed choices. Yet, behavioral biases may affect agents' (mis)use of the information available to them (Ambuehl and Li, 2018). In addition, the quality of the information source—which is not always transparent—appears to be a key determinant of its use (see, e.g., Reimers and Waldfogel 2021 and Aguiar, Waldfogel and Waldfogel 2021, for an analysis of online rating and ranking systems). Whereas experimental evidence documents how individuals acquire information through a dynamic search process (Gabaix, Laibson, Moloche and Weinberg, 2006; Brown, Flinn and Schotter, 2011; Caplin, Dean and Martin, 2011), to our knowledge, no experiment explores how information transmission via social learning shapes individuals' search behavior and how this process depends on the perceived reliability of the information source.

To fill this gap, we propose an experimental analysis of how agents use social learning when they acquire information about alternative options of unknown quality before a purchase. We consider a simple variant of Weitzman (1979)'s sequential search model and solve it with and without social information. Agents without social information act in isolation; agents with social information observe the choice, but not the search process, of a peer who does not derive any benefit from being imitated. Our goal is threefold. First, we investigate how social information affects search behavior and product discovery. Second, we analyze how the impact of social information depends on the perceived quality of the information source (i.e., on the reliability of the peer). Third, we study how alternative institutions helping agents to evaluate the information source's reliability can mitigate the effect of behavioral biases and restore the benefits of social information.

In our setup, without social information, an isolated agent must choose between two options whose qualities are drawn independently from the same distribution. Before searching, the agent knows the quality distribution but not options' realized qualities. Searching an option reveals its quality to the agent. After searching the first option for free, the agent decides whether to search the second option at a cost. With social information, before starting the search process, the agent observes the choice of one of her peers, but neither the peer's search behavior nor the peer's search cost. The peer acts in isolation and faces the same realization of options' qualities as the agent. Thus, social information helps the agent draw inferences about the options'

realized qualities from the peer's choice, but strategic considerations are absent.

Theory predicts that social information changes an agent's optimal search strategy in three ways. First, upon observing the peer's choice, the agent uses Bayes' rule to update her beliefs regarding how many searches the peer has conducted. Since the peer searched twice with positive probability, the quality distribution of the option chosen by the peer first-order stochastically dominates that of the other option. Thus, the agent begins searching from the option chosen by her peer. Second, the expected gain from a second search is lower than that of an isolated agent. The reason is that an agent with social information discounts the value of the second search by the probability that the peer searched only once, as only in this case the second search is valuable to the agent. Third, the expected gain from the second search is not decreasing in the quality of the first option searched, whereas it is so for an isolated agent. The reason is that if the quality of the option chosen by the peer is low, the peer must have searched twice: thus, the unchosen option's quality cannot be higher. As a result, the expected gain from the second search is nil, as when the first searched option has the highest possible quality.

We design our experiment to address the following questions. Do agents search optimally without and with social information? How do agents trade off independent exploration with the exploitation of social information? How does this depend on the reliability of their social information? Do reputational mechanisms help agents assess the reliability of their social information source and save on exploration costs?

In the experiment, participants act in sequence as either first or second movers. The role assignment (first mover vs. second mover) is exogenous and random. First movers are isolated. Second movers have social information: before searching, they observe the choice of a first mover. The experimental design consists of two parts. Part 1 lasts ten rounds and serves as a training stage for all participants: irrespective of role assignment, all participants act as first movers to familiarize themselves with the decision environment. Part 2 lasts twenty rounds: depending on their role, participants either continue playing as first movers or start playing as second movers, having access to some social information on the behavior of a first mover.

In the benchmark treatment, the second mover has no information on the quality of the search behavior of the first mover she observes, with whom she is randomly matched. In the other treatments, we manipulate the availability of this reputational information and the possibility for the second mover to choose which first mover to observe from a set of four.

As a result, the type of social information available to second movers in Part 2 varies across treatments. Second movers always observe the final choice and the

identity of the first mover they match with. However, (i) additional information containing hints about the reputation of their matched first mover may or may not be provided, and (ii) matching can be either exogenous or endogenous. The information about the first movers' reputation is about their overall performance in the task in Part 1: the rating depends on the level of the cumulative payoff accrued over the first ten rounds. Communicating first movers' ratings allows second movers to make some inferences about the first movers' reliability.

We implement a 2×2 (between-subjects) design and run four treatments:

- *NoRep-Exo (benchmark): No Reputation and Exogenous Matching.* We do not communicate first movers' ratings in Part 1 to second movers in Part 2; agents' matching across groups is exogenous (stranger matching protocol).
- *NoRep-Endo: No Reputation and Endogenous Matching.* We do not communicate first movers' ratings in Part 1 to second movers in Part 2; agents' matching across groups is endogenous (i.e., second movers can choose which first mover to match with).
- *Rep-Exo: Reputation and Exogenous Matching.* We communicate the first movers' ratings in Part 1 to the second movers in Part 2; agents' matching across groups is exogenous.
- *Rep-Endo: Reputation and Endogenous Matching.* We communicate the first movers' ratings in Part 1 to the second movers in Part 2; agents' matching across groups is endogenous.

The assignment of second movers to the four treatments is exogenous and random.

Across treatments, we manipulate the institutional settings that affect how agents assess the reliability of their social information source to examine how they influence the benefits of social information. We compare Rep-Exo with the benchmark to see how reputational information on the first movers' reliability affects the second movers' decision between searching (exploration) and using social information (exploitation). We compare NoRep-Endo with the benchmark to investigate endogenous reputation formation and test if the second movers tend to trust more first movers who were reliable before. Lastly, we compare Rep-Endo with the benchmark to observe how the two previous mechanisms interact. In particular, we check if they complement each other in a more realistic setting, where agents can choose which "expert" to follow based on experts' reputations.

Outline. In Section 2, we discuss the relation of our work to the existing literature. In Section 3, we present the theoretical model. In Section 4, we characterize optimal sequential search without and with social information and introduce the theoretical

predictions of interest. In Section 5, we present the experimental design. In Section 6, we present the hypothesis we aim to test, distinguishing between main and ancillary ones (thus providing a hierarchy of hypotheses). In Appendix A, we describe in detail the empirical analysis we plan to adopt for each of the hypotheses (detailing samples, variables, measurements, statistical tests, conditions, and power calculations). In Appendix B, we describe the implementation of the experiment and the relevant inclusion/exclusion criteria. In Appendix C, we provide information about the timeline of our next steps. Appendix D contains the experimental instructions.

2 Related Literature

Most of the work on how information about other agents' choices affects individual decision-making comes from the literature on social learning (see, e.g., Bikhchandani, Hirshleifer, Tamuz and Welch, 2022; Anderson and Holt, 2008, for surveys of the theoretical and experimental literature). Since the seminal work by Anderson and Holt (1997), social learning experiments mimic situations in which many agents act in sequence and try to learn about an underlying common state of the world. Each agent observes a private signal (i.i.d. across agents), which is informative about the state of the world, and all predecessors' choices. All information is exogenous and freely available. This literature focuses on how, in the long run, private and social information is aggregated and can affect agents' learning, hence inference accuracy. Evidence from social learning experiments shows that agents typically fail to process the available information optimally. In contrast to what Bayesian updating would predict, agents are prey to well-known behavioral biases such as, in particular, the overconfidence bias. That is, agents tend to overweight the informativeness of their private signals and/or their ability to interpret them and underweight those of others (see, e.g., Nöth and Weber, 2003; Çelen and Kariv, 2004; Goeree, Palfrey, Rogers and McKelvey, 2007; Kübler and Weizsäcker, 2004; Weizsäcker, 2010; Ziegelmeyer, March and Krügel, 2013; Eyster, Rabin and Weizsacker, 2018; Çelen, Geng and Li, 2020; De Filippis, Guarino, Jehiel and Kitagawa, 2022; Angrisani, Guarino, Jehiel and Kitagawa, 2021).[1]

[1]There is also empirical evidence in favor of alternative behavioral biases. Some work (see, e.g., Enke and Zimmermann, 2019) provides evidence of information redundancy neglect. According to this behavioral bias, agents fail to account that, like them, also their predecessors had access to multiple sources of information and tend to weigh early signals more than late ones. In slightly modified versions of standard social learning experiments, such as in Çelen, Kariv and Schotter (2010), agents have access to the choices and advice of other agents. Even when both pieces of information are equally informative, agents are more likely to rely on the advice of others, and the presence of advice improves the accuracy of agents' decisions and their payoffs. Duffy, Hopkins, Kornienko and Ma (2019) and Duffy, Hopkins and Kornienko (2021) show that, when called to choose between private or social information, most agents exhibit a suboptimal bias in favor of the latter,

Our experiment departs from this strand of literature in two ways. First, we do not focus on a long-run information aggregation problem. Second, our agents do not face a decision environment in which private and social information is free and exogenous, and the latter is abundant. In contrast, our agents face a decision environment in which private information is costly, whereas social information is free but limited (i.e., agents observe only one predecessor). In this context, our agents have to optimally trade off the exploitation of free social information with individual exploration via costly sequential search.

As such, the decision environment we model more closely resembles that of multi-armed bandit problems. In multi-armed bandit problems, agents repeatedly choose among different options—or bandit's arms—of initially unknown quality. At each point in time, agents face a trade-off between (i) optimally exploiting the information they have, choosing the arm they believe to be the best one given their current information set, and (ii) investing in costly exploration to unveil information about alternative arms and expand their information set, so to update their beliefs about the best arm and improve the quality of future decisions. Moving from the original description of the problem pioneered by Thompson (1933) and Robbins (1952), most of the literature on bandit problems considers settings in which decision-makers act in isolation, and different optimal exploration-exploitation policies have been derived, depending on the characteristics of the environment (see Bergemann and Valimaki, 2008; Hörner and Skrzypacz, 2017, for surveys of bandit problems in Economics). Accordingly, most experimental studies of these models focus only on the case in which agents solve the exploitation vs. exploration trade-off in isolation, without information about others' choices, to test the empirical relevance of the environmental game parameters, the search environment complexity, and the explanatory power of (alternative) optimal search rules derived by theoretical models (see, e.g, Meyer and Shi, 1995; Banks, Olson and Porter, 1997; Anderson, 2012; Banovetz, 2020; Hudja and Woods, 2022b,a; Banovetz and Oprea, 2023; Lykopoulos, Voucharas and Xefteris, 2022). Results show that players recognize the value of exploration and can play close to optimal. Agents' ambiguity aversion and risk preferences only partially explain deviations from optimal search behavior in the form of under- and over-exploration.

A few more recent experiments study multi-player (multi-armed) bandit problems, focusing on the role of strategic considerations. Building on the theoretical framework of Bolton and Harris (1999) and Keller, Rady and Cripps (2005), these papers study a strategic environment in which several players simultaneously solve a bandit problem: the best among many arms is the same for all players, but initially unknown; thus,

suggesting a tendency to underestimate the frequency of others' mistakes and the non-informativeness of social information.

players can learn from the publicly observable each other's experimentation decisions and/or their outcomes. These bandit problems become dynamic public-good dilemmas in which the public good is the dynamically evolving information about players' common state of the world. Boyce, Bruner and McKee (2016) and Hoelzemann and Klein (2021) rely on lab experiments to study this complex information transmission dilemma and investigate the empirical relevance of strategic exploration, free-riding, and coordination issues.

Similarly to Boyce et al. (2016) and Hoelzemann and Klein (2021), we focus on a multi-player bandit problem with a common state of the world. In contrast to these papers, however, in our setting, agents are short-lived, move sequentially, and there are informational but not payoff externalities. Lifting strategic considerations allows us to simplify the decision environment and focus exclusively on how agents solve the exploitation vs. exploration trade-off without and with social information. The main contribution of our paper is to analyze, in the cleanest possible framework, what institutions can nudge subjects into valuing and correctly exploiting social information, thus avoiding costly under- or over-exploration and ultimately making better choices. Second, we aim to study how, absent strategic considerations, social information affects individual search behavior by tilting the exploration-exploitation trade-off and test whether the exploitation of social information crowds out the incentive for independent exploration.