

# Bayesian Heuristics for Group Decisions

Ali Jadbabaie & M. Amin Rahimian

[jadbabai,rahimian@mit.edu](mailto:jadbabai,rahimian@mit.edu)

We propose a model of inference for heuristic decision-making in groups that is rooted in the Bayes rule but avoids the complexities of rational inference in partially observed environments with incomplete information, which are characteristic of group interactions. Our model is also consistent with a dual-process psychological theory of thinking: the group members behave rationally at the initiation of their interactions with each other (the slow and deliberative mode); however, in the ensuing decision epochs, they rely on a heuristic that replicates their experiences from the first stage (the fast automatic mode). We specialize this model to a group decision scenario where private observations are received at the beginning, and agents aim to take the best action given the aggregate observations of all group members. We study the implications of the information structure and the choice of the probability distributions for signal likelihoods and beliefs. These factors also determine the structure of the so-called “Bayesian heuristics” that the agents follow in our model. We further analyze the group decision outcomes in two classes of linear action updates and log-linear belief updates and show that many inefficiencies arise in group decisions as a result of repeated interactions between individuals, leading to overconfident beliefs as well as choice-shifts toward extreme actions. Nevertheless, balanced regular structures demonstrate a measure of efficiency in terms of aggregating the initial information of individuals. These results not only verify some well-known insights about group decision-making, but also complement these insights by revealing additional mechanistic interpretations for the group decision-process, as well as psychological and cognitive intuitions about the group interaction model.

*Key words:* opinion dynamics; social learning; Bayesian learning; non-Bayesian learning; rational learning; observational learning; statistical learning; distributed learning; distributed hypothesis testing; distributed detection; DeGroot model; linear regression; conjugate priors; exponential families

*MSC2000 subject classification:* Primary: 91B06; secondary: 91A35, 62C10

*OR/MS subject classification:* Primary: Games/group decisions: Voting/committees; secondary: Organizational studies: Decision making Effectiveness/performance Information; Networks/graphs: Heuristics

*JEL:* D83, D85.

*History:* This version is dated August 17, 2017. A preliminary version was presented in NBER-NSF Seminar on Bayesian Inference in Econometrics and Statistics on April 30, 2016.

**1. Introduction, background & motivation of the work.** Daniel Kahneman in his highly acclaimed work, “*Thinking, Fast and Slow*”, points out that the proper way to elicit information from a group is not through a public discussion but rather confidentially collecting each person’s judgment [47, Chapter 23]. Indeed, decision making among groups of individuals exhibit many singularities and important inefficiencies that lead to Kahneman’s noted advice. As a team converges on a decision, expressing doubts about the wisdom of the consensus choice is suppressed; subsequently teams of decision makers are afflicted with *groupthink* as they appear to reach a consensus.<sup>1</sup> The mechanisms of uncritical optimism, overconfidence, and the illusions of validity in group interactions also lead to *group polarization*, making the individuals more amenable toward extreme opinions [84]. An enhanced understanding of decision making and learning in groups sheds light on the role of individuals in shaping public opinion and how they influence efficiency of information transmissions. These in turn help to improve the predictions about group behavior and provide guidelines for designing effective social and organizational policies. We contribute to this literature by offering an analytic model of decision making in groups subject to heuristics and biases. We further investigate the evolution of beliefs and properties of the group decision outcome as the agents in our model repeatedly interact with each other and deliberate their options.

The rational approach advocates application of Bayes rule to the entire sequence of observations successively at every step. However, such repeated applications of Bayes rule in networks become

<sup>1</sup>Gar Klein proposes a famous method of *project premortem* to overcome the groupthink through an exercise: imagining that the the planned decision was failed in implementation and writing a brief report of the failure [50].

very complex, especially if the agents are unaware of the global network structure; and as they use their local data to make inferences about all possible contingencies that can lead to their observations. While rational learning continues to receive quite a significant amount of attention [1], it has also been criticized in the literature due to its unrealistic computational and cognitive demand on the agents [42]. On the one hand, the properties of rational learning models are difficult to analyze beyond some simple asymptotic facts such as convergence. On the other hand, these models make unrealistic assumptions about the cognitive ability and amount of computations that agents perform before committing to a decision. To avoid these shortcomings, an alternative non-Bayesian approach relies on simple and intuitive “*heuristics*” that are descriptive of how agents aggregate the reports of their neighbors before coming up with a decision.

Heuristics are used widely in the literature to model social interactions and decision making [26, 27, 28]. They provide tractable tools to analyze boundedly rational behavior and offer insights about decision making under uncertainty. Hegselmann and Krause [37] investigate various ways of averaging to model opinion dynamics and compare their performance for computations and analysis. Using such heuristics one can avoid the complexities of fully rational inference, and their suitability are also verified in experimental studies by Grimm and Mengel [33] and Chandrasekhar, Larreguy and Xandri [10]. The study of such heuristics started in 1974 with the seminal work of DeGroot [12] in linear opinion pooling, where agents update their opinions to a convex combination of their neighbors’ beliefs and the coefficients correspond to the level of confidence that each agent puts in each of her neighbors. More recently, Jadbabaie, Molavi and Tahbaz-Salehi [40, 41, 60] consider a variation of this model for streaming observations, where in addition to the neighboring beliefs the agents also receive private signals. Despite their widespread applications, theoretical and axiomatic foundations of social inferences using heuristics and non-Bayesian updates have received attention only recently [60, 62]. The main goal of our paper is to analyze the connection between rational and heuristic approaches from a behavioral perspective.

Both the rational and heuristic approach have a deep history in the decision theory of groups. Sobel [81] provides a theoretical framework to study the interplay between individual recommendations and rationality of group decisions. The seminal work of Janis [44] provides various examples involving the American foreign policy in the mid-twentieth where the desire for harmony or conformity in the group have resulted in bad group decisions, a phenomenon that he coins *groupthink*. Various other works have looked at the choice shift toward more extreme options [16, 83] and group polarization [39, 73]. In [55] the authors investigate the effects of mistakes and biases that arise from the group members’ emphasis on their common information and their negligence of their private data leading to a hidden profile problem in group decision making.

A dual process theory for the psychology of mind and its operation identifies two systems of thinking [17]: one that is fast, intuitive, non-deliberative, habitual and automatic (system one); and a second one that is slow, attentive, effortful, deliberative, and conscious (system two).<sup>1</sup> Major advances in behavioral economics are due to incorporation of this dual process theory and the subsequent models of bounded rationality [46]. Reliance on heuristics for decision making is a distinctive feature of system one that avoids the computational burdens of a rational evaluation; but also subjects people to systematic and universal errors: the so-called “*cognitive biases*”. Hence, it is important to understand the nature and properties of heuristic decision making and its consequences to individual and organizational choice behavior. This premise underlies many of the recent advances in behavioral economics [85], and it motivates our work as well.

<sup>1</sup>While many decision science applications focus on developing dual process theories of cognition and decision making (cf. [18, 59] and the references therein); other researchers identify multiple neural systems that derive decision making and action selection: ranging from reflexive and fast (Pavlovian) responses to deliberative and procedural (learned) ones; and these systems are in turn supported by several motoric, perceptual, situation-categorization and motivational routines which together comprise the decision making systems [71, Chapter 6].

**2. Summary of contributions and organization of the paper.** In this work we are concerned with the operations of system one (the fast/automatic system): we aim to study heuristics for information aggregation in group decision scenarios when the relevant information is dispersed among many individuals. In such situations, individuals in the group are subjected to informational (but not strategic) externalities. By the same token, the heuristics that are developed for decision making in such situations are also aimed at information aggregation. In our model as the agent interacts with her environment, her initial response would engage her system two (the slow and deliberative system): she rationally evaluates the reports of her neighbors and uses them to make a decision. However, after her initial experience and by engaging in repeated interactions with other group members her system one takes over the decision processes, implementing a heuristic that imitates her (rational/Bayesian) inferences from her initial experience; hence avoiding the burden of additional cognitive processing in the ensuing interactions with her neighbors. This follows the propositions of Tversky and Kahneman in [87], who argue that humans have limited time and brainpower, therefore they rely on simple rules of thumb, i.e. heuristics, to help them make judgments under uncertainty. However, the use of such heuristics causes people to make predictable errors and subjects them to various cognitive biases. The specific cognitive bias that we formulate and analyze in the case of group decision-making is the human error in attributing recommendations of other people to their private information. In reality these recommendations are shaped not only by their private information, but also by other recommendations that are observable to them across the social network.

Our main contribution is to offer a behavioral framework for heuristic decision making, by relying on the time-one Bayesian update and using it for all future decision epochs. This model offers a behavioral foundation for non-Bayesian updating that is compatible with the dual-process psychological theory of decision making and the principles of judgment under uncertainty subject to heuristics and biases. Since our 2014 paper [69], various other authors have also developed results based on variations of this no-recall idea. In [54] the authors propose this model to analyze dynamic consensus and use it to study the effects of correlation neglect on voting behavior [54], as well as the persuasion power of media [52]. In [63], the authors use this model to analyze the structural and environmental conditions that are necessary for learning with binary actions. Where we depart from this body of work is in developing a behavioral rationale for no-recall updates. In particular, we show when such updates lead to inefficient information aggregation and leverage these insights to propose better team and organizational decision making strategies.

On the one hand, our model of inference based on *Bayesian* or *no-recall heuristics* is motivated by the real-world behavior of people induced by their system one and reflected in their spur-of-the-moment decisions and impromptu behavior: Basing decisions only on the immediately observed actions and disregarding the history of the observed actions or the possibility of correlations among different observations; i.e. “*what you see is all there is*” [47]. On the other hand, the proposed Bayesian (no-recall) heuristics offer a boundedly rational approach to model decision making over social networks. The latter is in the sense of the word as coined by Herbert A. Simon, i.e. “to incorporate modifications that lead to substantial simplifications in the original choice problem” [77].<sup>1</sup> By ignoring the history of interactions, the heuristic (no-recall) agents are left with a substantially simplified model of their environment that they can respond to optimally. This is in contrast with the Bayesian approach which is not only unrealistic in the amount of cognitive burden that it imposes on the agents, but also is often computationally intractable and complex to analyze [42].

<sup>1</sup>Simon advocates “bounded rationality” as compatible with the information access and the computational capacities that are actually possessed by the agents in their environments. Most importantly he proposes the use of so-called “*satisficing*” heuristics; i.e. to search for alternatives that exceed some “*aspiration levels*” by satisfying a set of minimal acceptability criteria [78, 79].

In Section 3, we describe the mathematical details of our model; in particular, we explain the mathematical steps for deriving the so-called Bayesian heuristics in a given decision scenario. Specific cases of Bayesian heuristics that we explore in the following sections are the log-linear (multiplicative) updating of beliefs over the probability simplex, and the linear (weighted arithmetic average) updating of actions over the Euclidean space. In Section 4, we specialize our group decision model to a setting involving exponential family of distributions for both signal likelihoods and agents' beliefs. The agents aim to estimate the expected values of the sufficient statistics for their exponential family signal structures. We show that the Bayesian heuristics in this case are affine rules in the self and neighboring actions, and we give explicit expressions for their coefficients. Subsequently, we provide conditions under which these action updates constitute a convex combination as in the DeGroot model, with actions converging to a consensus in the latter case. We also investigate the efficiency of the consensus action in aggregating the initial observations of all agents across the network.

Next in Section 5, we discuss a situation where agents exchange beliefs about a truth state that can take one of the finitely many possibilities. The Bayesian heuristics in this case take the form of log-linear rules that set the updated beliefs proportionally to the product of self and neighboring beliefs in every decision epoch. We investigate the evolution of beliefs under the prescribed “no-recall” update rules and compare the asymptotic beliefs with that of a Bayesian agent with direct access to all the private information; thus characterizing the inefficiencies of the asymptotic beliefs, in particular, their redundancy. We next propose a coordination scheme for teaming agents across several time-steps without exposing them to redundant beliefs. We also briefly discuss possible extensions of our framework to belief exchange settings with increased recall.

We summarize our findings from the analysis of linear action and log-linear belief updates in Section 6, where we reiterate the assumptions that are implicit in the adoption of popular aggregation heuristics such as the DeGroot model; moreover, we discuss the inefficiencies that arise as a result of their application. Such heuristics allow us to aggregate the information in our environment, and provide for desirable asymptotic properties such as consensus; however, this consensus often fails as an efficient group aggregate for the individuals' private data. We provide the mathematical proofs and the relevant details for many of the results in the appendices at end of the paper.

**3. The model.** Consider a group of  $n$  agents that are labeled by  $[n]$  and interact according to a digraph  $\mathcal{G} = ([n], \mathcal{E})$ .<sup>1</sup> The neighborhood of agent  $i$  is the set of all agents whom she observes including herself, and it is denoted by  $\mathcal{N}_i = \{j \in [n]; (j, i) \in \mathcal{E}\} \cup \{i\}$ ; every node having a self-loop:  $(i, i) \in \mathcal{E}$  for all  $i$ . We refer to the cardinality of  $\mathcal{N}_i$  as the degree of node  $i$  and denote it by  $\text{deg}(i)$ . There is a state  $\theta \in \Theta$  that is unknown to the agents and it is chosen arbitrarily by nature from an underlying state space  $\Theta$ , which is measurable by a  $\sigma$ -finite measure  $\mathcal{G}_\theta(\cdot)$ . For example if a space  $(\Theta$  or  $\mathcal{S})$  is a countable set, then we can take its  $\sigma$ -finite measure ( $\mathcal{G}_\theta$  or  $\mathcal{G}_s$ ) to be the counting measure, denoted by  $\mathcal{K}(\cdot)$ ; and if the space is a subset of  $\mathbb{R}^k$  with positive Lebesgue measure, then we can take its  $\sigma$ -finite measure to be the Lebesgue measure on  $\mathbb{R}^k$ , denoted by  $\Lambda_k(\cdot)$ . Associated with each agent  $i$ ,  $\mathcal{S}_i$  is a measurable space called the signal space of  $i$ , and given  $\theta$ ,  $\mathcal{L}_i(\cdot | \theta)$  is a probability measure on  $\mathcal{S}_i$ , which is referred to as the *signal structure* of agent  $i$ . Furthermore,  $(\Omega, \mathcal{F}, \mathcal{P}_\theta)$  is a probability triplet, where  $\Omega = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$  is a product space, and

<sup>1</sup> Some notations: Throughout the paper,  $\mathbb{R}$  is the set of real numbers,  $\mathbb{N}$  denotes the set of all natural numbers, and  $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$ . For  $n \in \mathbb{N}$  a fixed integer the set of integers  $\{1, 2, \dots, n\}$  is denoted by  $[n]$ , while any other set is represented by a capital Greek or calligraphic letter. For a measurable set  $\mathcal{X}$  we use  $\Delta\mathcal{X}$  to denote the set of all probability distributions over the set  $\mathcal{X}$ . Furthermore, any random variable is denoted in boldface letter, vectors are represented in lowercase letters and with a bar over them, measures are denoted by upper case Greek or calligraphic Latin letters, and matrices are denoted in upper case Latin letters. For a matrix  $A$ , its spectral radius  $\rho(A)$  is the largest magnitude of all its eigenvalues.

$\mathcal{F}$  is a properly defined sigma field over  $\Omega$ . The probability measure on  $\Omega$  is  $\mathcal{P}_\theta(\cdot)$  which assigns probabilities consistently with the signal structures  $\mathcal{L}_i(\cdot|\theta), i \in [n]$ ; and in such a way that with  $\theta$  fixed, the random variables  $\mathbf{s}_i, i \in [n]$  taking values in  $\mathcal{S}_i$ , are independent. These random variables represent the private signals that agents  $i \in [n]$  observe at time 0. Note that the private signals are independent across the agents. The expectation operator  $\mathbb{E}_\theta\{\cdot\}$  represents integration with respect to  $\mathcal{P}_\theta(d\omega), \omega \in \Omega$ .

**3.1. Beliefs, actions, and rewards.** An agents' belief about the unknown allows her to make decisions even as the outcome is dependent on the unknown value  $\theta$ . These beliefs about the unknown state are probability distributions over  $\Theta$ . Even before any observations are made, every agent  $i \in [n]$  holds a prior belief  $\mathcal{V}_i(\cdot) \in \Delta\Theta$ ; this represents her subjective biases about the possible values of  $\theta$ . For each time instant  $t$ , let  $\mathcal{M}_{i,t}(\cdot)$  be the (random) probability distribution over  $\Theta$ , representing the *opinion* or *belief* at time  $t$  of agent  $i$  about the realized value of  $\theta$ . Moreover, let the associated expectation operator be  $\mathbb{E}_{i,t}\{\cdot\}$ , representing integration with respect to  $\mathcal{M}_{i,t}(d\theta)$ . We assume that all agents share the common knowledge of signal structures  $\mathcal{L}_i(\cdot|\theta'), \forall \theta' \in \Theta$ , their priors  $\mathcal{V}_i(\cdot)$ , and their corresponding sample spaces  $\mathcal{S}_i$  and  $\Theta$  for all  $i \in [n]$ .<sup>1</sup>

Let  $t \in \mathbb{N}_0$  denote the time index; at  $t=0$  the values  $\theta \in \Theta$  followed by  $s_i \in \mathcal{S}_i$  of  $\mathbf{s}_i$  are realized and the latter is observed privately by each agent  $i$  for all  $i \in [n]$ . Associated with every agent  $i$  is an action space  $\mathcal{A}_i$  that represents all the choices available to her at every point of time  $t \in \mathbb{N}_0$ , and a utility  $u_i(\cdot, \cdot) : \mathcal{A}_i \times \Theta \rightarrow \mathbb{R}$  which in expectation represents her von Neumann-Morgenstern preferences regarding lotteries with independent draws from  $\mathcal{A}_i$  and/or  $\Theta$ . These utilities are additive over time corresponding to successive independent draws. The utility functions and action spaces are common knowledge among the agents. Subsequently, at every time  $t \in \mathbb{N}_0$  each agent  $i \in [n]$  chooses an action  $\mathbf{a}_{i,t} \in \mathcal{A}_i$  and is rewarded  $u_i(\mathbf{a}_{i,t}, \theta)$ .

**3.2. The group decision problem and aggregation heuristics.** Given  $\mathbf{s}_i$ , agent  $i$  forms an initial Bayesian opinion  $\mathcal{M}_{i,0}(\cdot)$  about the value of  $\theta$  and chooses her action  $\mathbf{a}_{i,0} \leftarrow \arg \max_{a_i \in \mathcal{A}_i} \int_{\Theta} u_i(a_i, \theta') \mathcal{M}_{i,0}(d\theta')$ , maximizing her expected reward. Here for a set  $\mathcal{A}$ , we use the notation  $\mathbf{a} \leftarrow \mathcal{A}$  to denote an arbitrary choice from the elements of  $\mathcal{A}$  that is assigned to  $\mathbf{a}$ . Not being notified of the actual realized value for  $u_i(\mathbf{a}_{i,0}, \theta)$ , she then observes the actions that her neighbors have taken. Given her extended set of observations  $\{\mathbf{a}_{j,0}, j \in \mathcal{N}_i\}$  at time  $t=1$ , she refines her opinion into  $\mathcal{M}_{i,1}(\cdot)$  and makes a second, and possibly different, move  $\mathbf{a}_{i,1}$  according to:

$$\mathbf{a}_{i,1} \leftarrow \arg \max_{a_i \in \mathcal{A}_i} \int_{\Theta} u_i(a_i, \theta) \mathcal{M}_{i,1}(d\theta), \quad (1)$$

maximizing her expected pay off conditional on everything that she has observed thus far; i.e. maximizing  $\mathbb{E}_{i,1}\{u_i(a_i, \theta)\} = \mathbb{E}_\theta\{u_i(\mathbf{a}_{i,1}, \theta) | \mathbf{s}_i, \mathbf{a}_{j,0} : j \in \mathcal{N}_i\} = \int_{\Theta} u_i(a_i, \theta) \mathcal{M}_{i,1}(d\theta)$ . Subsequently, she is granted her net reward of  $u_i(\mathbf{a}_{i,0}, \theta) + u_i(\mathbf{a}_{i,1}, \theta)$  from her past two plays. Following realization of rewards for their first two plays, in any subsequent time instance  $t > 1$  each agent  $i \in [n]$  observes the preceding actions of her neighbors  $\mathbf{a}_{j,t-1} : j \in \mathcal{N}_i$  and takes an option  $\mathbf{a}_{i,t}$  out of the set  $\mathcal{A}_i$ . Of particular significance in our description of the behavior of agents in the succeeding time periods  $t > 1$ , is the relation:

$$f_i(\mathbf{a}_{j,0} : j \in \mathcal{N}_i) := \mathbf{a}_{i,1} \leftarrow \arg \max_{a_i \in \mathcal{A}_i} \mathbb{E}_{i,1}\{u_i(a_i, \theta)\}. \quad (2)$$

<sup>1</sup>The signal structures  $\mathcal{L}_i(\cdot|\theta'), \forall \theta' \in \Theta$  and the priors  $\mathcal{V}_i(\cdot)$ , as well as the corresponding sample spaces  $\mathcal{S}_i$  and  $\Theta$  are common knowledge amongst the agents for all  $i \in [n]$ . The assumption of common knowledge in the case of fully rational (Bayesian) agents implies that given the same observations of one another's beliefs or private signals distinct agents would make identical inferences; in the sense that starting from the same belief about the unknown  $\theta$ , their updated beliefs given the same observations would be the same; in Aumann's words, rational agents cannot agree to disagree [5].

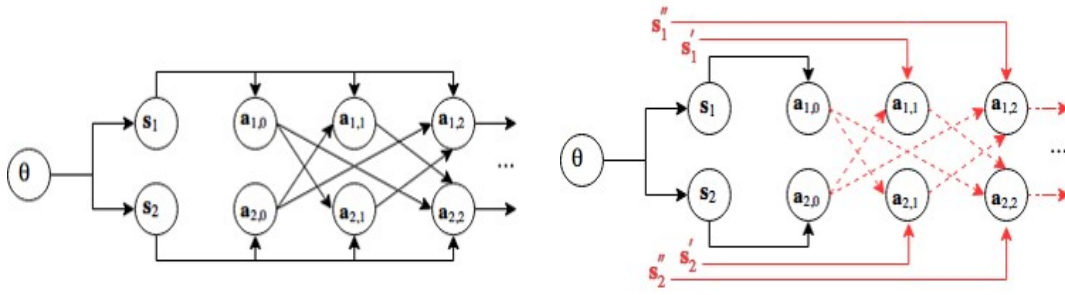


FIGURE 1. On the left, the decision flow diagram for the inferences of two rational agents who communicate their actions repeatedly. On the right, heuristic agents ignore the history of interactions (dotted arrows) in their inferences and always attribute actions to (imaginary) private signals.

derived in (1), which given the observations of agent  $i$  at time  $t = 0$ , specifies her (Bayesian) pay-off maximizing action for time  $t = 1$ . Once the format of the mapping  $f_i(\cdot)$  is obtained, it is then used as a heuristic for decision making in every future epoch. The agents update their action by choosing:  $\mathbf{a}_{i,t} = f_i(\mathbf{a}_{j,t-1} : j \in \mathcal{N}_i), \forall t > 1$ . We refer to the mapping  $f_i : \prod_{j \in \mathcal{N}_i} \mathcal{A}_j \rightarrow \mathcal{A}_i$  thus obtained, as a *Bayesian* or *no-recall heuristic*.

REMARK 1 (MYOPIC PREFERENCES & PURELY INFORMATIONAL EXTERNALITIES). In modeling the group decision process, we assume that the preferences of agents across time are myopic. At every time  $t \in \mathbb{N}$ , agents  $i$  takes action  $\mathbf{a}_{i,t}$  to maximize her expected utility,  $\mathbb{E}_{i,t}\{u_i(\mathbf{a}_{i,t}, \theta)\}$ . This myopia is rooted in the underlying group decision scenario that we are modeling: the agents goal for interacting with other group members is to come up with a decision that is more informed than if they were to act solely based on their own private data. Hence, by observing the recommendations of their neighboring agents  $\mathbf{a}_{j,t}$  they hope to augment their information with what their neighbors as well as other agents in the network know that they do not. In particular, the agent do not have the freedom to learn form consequences of their recommendations, not before committing to a choice. Expressly in the group decision scenario the agents do not learn from the realized values of the utilities of their previous recommendations (unless they commit to their choice), rather the purpose of the group discussion is to augment their information by learning from recommendations of other as much as possible before committing to a choice. Moreover, the network externalities that arise in above settings are purely informational. People are therefore interacting with each other, only to learn from one another, and to improve the quality of their decisions; for example, in jury deliberations after jurors are each individually and independently exposed to the court proceedings, the jury enters deliberations to decide on a verdict. In another case, several doctors may examine a patient and then engage in group discussions to determine the source of an illness; or a panel of judges may repeatedly evaluate the performance of contenders in weightlifting, figure skating, or diving competitions. Lack of strategic externalities is an important characteristic of the kind of human interactions that we investigate in this paper.

REMARK 2 (“WHAT YOU SEE IS ALL THERE IS!”). The heuristics thus obtained suffer from same fallacies of snap judgments that are associated with the recommendations of system one in “Thinking, Fast and Slow”; flawed judgments that rely on simplistic interpretations: “*what you see is all there is*”, in Kahneman’s elegant words [47]. Indeed, the use of the initial Bayesian update for future decision epochs entails a certain level of naivety on the part of the decision maker: she has to either assume that the structure of her neighbors’ reports have not departed from their initial format, or that they are not being influenced back by her own or other group members and can thus be regarded as independent sources of information; see Fig. 1. Such naivety in disregarding the history of interactions has been highlighted in our earlier works on Bayesian learning without recall

[67], where we interpret the use of time-one Bayesian update for future decision epochs, as a rational but memoryless behavior: by regarding their observations as being direct consequences of private signals, the agents reject any possibility of a past history beyond their immediate observations.

REMARK 3 (NAIVE INFERENCES). Similar and related forms of naivety have been suggested in the literature. Eyster and Rabin [19, 20] propose the autarkic model of naive inference, where players at each generation observe their predecessors but naively think that any predecessor's action relies solely on that player's private information, thus ignoring the possibility that successive generations are learning from each other. Bala and Goyal [6] study another form of naivety and bounded-rational behavior by considering a variation of observational learning in which agents observe the action and pay-offs of their neighbors and make rational inferences about the action/pay-off correspondences, based on their observations of the neighboring actions; however, they ignore the fact that their neighbors are themselves learning and trying to maximize their own pay-offs. Levy and Razin look at a particularly relevant cognitive bias called correlation neglect, which makes individuals regard the sources of their information as independent [53]; they analyze its implications to diffusion of information, and focus in particular, on the voting behavior.

**4. Affine action updates, linear updating, and DeGroot learning.** In this section we explore the essential modeling features that lead to a linear structure in the Bayesian heuristics (linear update rules). We present a general scenario that involves the exponential family of distributions and leads to linear action updates.

To describe the signal structures, we consider a measurable sample space  $\mathcal{S}$  with a  $\sigma$ -finite measure  $\mathcal{G}_s(\cdot)$ , and a parametrized class of sampling functions  $\{\mathcal{L}(\cdot|\theta; \sigma_i, \delta_i) \in \Delta\mathcal{S} : \sigma_i > 0 \text{ and } \delta_i > 0\}$  belonging to the  $k$ -dimensional exponential family as follows:

$$\ell(s|\theta; \sigma_i, \delta_i) := \frac{d\mathcal{L}(\cdot|\theta; \sigma_i, \delta_i)}{d\mathcal{G}_s} = \sigma_i \left| \frac{\Lambda_k(\xi(ds))}{\mathcal{G}_s(ds)} \right| \tau(\sigma_i \xi(s), \delta_i) e^{\sigma_i \eta(\theta)^T \xi(s) - \delta_i \gamma(\eta(\theta))}, \quad (3)$$

where  $\xi(s) : \mathcal{S} \rightarrow \mathbb{R}^k$  is a measurable function acting as a sufficient statistic for the random samples,  $\eta : \Theta \rightarrow \mathbb{R}^k$  is a mapping from the parameter space  $\Theta$  to  $\mathbb{R}^k$ ,  $\tau : \mathbb{R}^k \times (0, +\infty) \rightarrow (0, +\infty)$  is a positive weighting function, and

$$\gamma(\eta(\theta)) := \frac{1}{\delta_i} \ln \int_{s \in \mathcal{S}} \sigma_i \left| \frac{\Lambda_k(\xi(ds))}{\mathcal{G}_s(ds)} \right| \tau(\sigma_i \xi(s), \delta_i) e^{\sigma_i \eta(\theta)^T \xi(s)} \mathcal{G}_s(ds),$$

is a normalization factor that is constant when  $\theta$  is fixed, even though  $\delta_i > 0$  and  $\sigma_i > 0$  vary. This normalization constant for each  $\theta$  is uniquely determined by the functions  $\eta(\cdot)$ ,  $\xi(\cdot)$  and  $\tau(\cdot)$ . The parameter space  $\Theta$  and the mapping  $\eta(\cdot)$  are such that the range space  $\Omega_\theta := \{\eta(\theta) : \theta \in \Theta\}$  is an open subset of the natural parameter space  $\Omega_\eta := \{\eta \in \mathbb{R}^k : \int_{s \in \mathcal{S}} |\Lambda_k(\xi(ds))/\mathcal{G}_s(ds)| \tau(\xi(s), 1) e^{\eta^T \xi(s)} \mathcal{G}_s(ds) < \infty\}$ . In (3),  $\sigma_i > 0$  and  $\delta_i > 0$  for each  $i$  are scaling factors that determine the quality or informativeness of the random sample  $\mathbf{s}_i$  with regard to the unknown  $\theta$ : fixing either one of the two factors  $\sigma_i$  or  $\delta_i$ , the value of the other one increases with the increasing informativeness of the observed value  $\xi(\mathbf{s}_i)$ . The following conjugate family of priors<sup>1</sup> are associated with the likelihood structure (3). This family is determined uniquely by the transformation and normalization functions:  $\eta(\cdot)$  and  $\gamma(\cdot)$ , and it is parametrized through a pair of parameters  $(\alpha, \beta)$ ,  $\alpha \in \mathbb{R}^k$  and  $\beta > 0$ :

<sup>1</sup>Consider a parameter space  $\Theta$ , a sample space  $\mathcal{S}$ , and a sampling distribution  $\mathcal{L}(\cdot|\theta) \in \Delta\mathcal{S}$ ,  $\theta \in \Theta$ . Suppose that  $\mathbf{s}$  is a random variable which is distributed according to  $\mathcal{L}(\cdot|\theta)$  for any  $\theta$ . A family  $\mathcal{F} \subset \Delta\Theta$  is a conjugate family for  $\mathcal{L}(\cdot|\theta)$ , if starting from any prior distribution  $\mathcal{V}(\cdot) \in \mathcal{F}$  and for any signal  $s \in \mathcal{S}$ , the posterior distribution given the observation  $\mathbf{s} = s$  belongs to  $\mathcal{F}$ .

$$\mathcal{F}_{\gamma,\eta} := \left\{ \mathcal{V}(\theta; \alpha, \beta) \in \Delta\Theta, \alpha \in \mathbb{R}^k, \beta_i > 0 : \nu(\theta; \alpha, \beta) := \frac{d\mathcal{V}(\cdot; \alpha, \beta)}{d\mathcal{G}_\theta} = \left| \frac{\Lambda_k(\eta(d\theta))}{\mathcal{G}_\theta(d\theta)} \right| \frac{e^{\eta(\theta)^T \alpha - \beta \gamma(\eta(\theta))}}{\kappa(\alpha, \beta)}, \right. \\ \left. \kappa(\alpha, \beta) := \int_{\theta \in \Theta} \left| \frac{\Lambda_k(\eta(d\theta))}{\mathcal{G}_\theta(d\theta)} \right| e^{\eta(\theta)^T \alpha - \beta \gamma(\eta(\theta))} \mathcal{G}_\theta(d\theta) < \infty \right\}.$$

Furthermore, we assume that agents take actions in  $\mathbb{R}^k$ , and that they aim for a minimum variance estimation of the regression function or conditional expectation (given  $\theta$ ) of the sufficient statistic  $\xi(\mathbf{s}_i)$ . Hence, we endow every agent  $i \in [n]$  with the quadratic utility  $u_i(a, \theta) = -(a - m_{i,\theta})^T (a - m_{i,\theta})$ ,  $\forall a \in \mathcal{A}_i = \mathbb{R}^k$ , where  $m_{i,\theta} := \mathbb{E}_{i,\theta}\{\xi(\mathbf{s}_i)\} := \int_{s \in \mathcal{S}} \xi(s) \mathcal{L}(ds|\theta; \sigma_i, \delta_i) \in \mathbb{R}^k$ .

Our main result in this section prescribes a scenario in which each agent starts from a prior belief  $\mathcal{V}(\cdot; \alpha_i, \beta_i)$  belonging to  $\mathcal{F}_{\gamma,\eta}$  and she observes a fixed number  $n_i$  of i.i.d. samples from the distribution  $\mathcal{L}(\cdot | \theta; \sigma_i, \delta_i)$ . The agents then repeatedly communicate their actions aimed at minimum variance estimation of  $m_{i,\theta}$ . These settings are formalized under the following assumption that we term the *Exponential Family Signal-Utility Structure*.

**ASSUMPTION 1 (Exponential family signal-utility structure).**

- (i) Every agent  $i \in [n]$  observes  $n_i$  i.i.d. private samples  $\mathbf{s}_{i,p}, p \in [n_i]$  from the common sample space  $\mathcal{S}$  and that the random samples are distributed according to the law  $\mathcal{L}(\cdot | \theta; \sigma_i, \delta_i)$  given by (3) as a member of the  $k$ -dimensional exponential family.
- (ii) Every agent starts from a conjugate prior  $\mathcal{V}_i(\cdot) = \mathcal{V}(\cdot; \alpha_i, \beta_i) \in \mathcal{F}_{\gamma,\eta}$ , for all  $i \in [n]$ .
- (iii) Every agent chooses actions  $a \in \mathcal{A}_i = \mathbb{R}^k$  and bears the quadratic utility  $u_i(a, \theta) = -(a - m_{i,\theta})^T (a - m_{i,\theta})$ , where  $m_{i,\theta} := \mathbb{E}_{i,\theta}\{\xi(\mathbf{s}_i)\} := \int_{s \in \mathcal{S}} \xi(s) \mathcal{L}(ds|\theta; \sigma_i, \delta_i) \in \mathbb{R}^k$ .

The Bayesian heuristics  $f_i(\cdot), i \in [n]$  under the settings prescribed by the exponential family signal-utility structure (Assumption 1) are linear functions of the neighboring actions with specified coefficients that depend only on the likelihood structure parameters:  $n_i, \sigma_i$  and  $\delta_i$  as well as the prior parameters:  $\alpha_i$  and  $\beta_i$ , for all  $i \in [n]$ .<sup>1</sup>

**THEOREM 1 (Affine action updates).** Under the exponential family signal-utility structure specified in Assumption 1, the Bayesian heuristics describing the action update of every agent  $i \in [n]$  are given by:  $\mathbf{a}_{i,t} = f_i(\mathbf{a}_{j,t-1} : j \in \mathcal{N}_i) = \sum_{j \in \mathcal{N}_i} T_{ij} \mathbf{a}_{j,t-1} + \epsilon_i$ , where for all  $i, j \in [n]$  the constants  $T_{ij}$  and  $\delta_i$  are as follows:

$$T_{ij} = \frac{\delta_i \sigma_j (n_j + \delta_j^{-1} \beta_j)}{\sigma_i (\beta_i + \sum_{p \in \mathcal{N}_i} n_p \delta_p)}, \quad \epsilon_i = - \frac{\delta_i}{\sigma_i (\beta_i + \sum_{p \in \mathcal{N}_i} n_p \delta_p)} \sum_{j \in \mathcal{N}_i \setminus \{i\}} \alpha_j.$$

The action profile at time  $t$  is the concatenation of all actions in a column vector:  $\bar{\mathbf{a}}_t = (\mathbf{a}_{1,t}^T, \dots, \mathbf{a}_{n,t}^T)^T$ . The matrix  $T$  with entries  $T_{ij}, i, j \in [n]$  given in Theorem 1 is called the social influence matrix. The constant terms  $\epsilon_i$  in this theorem appear as the rational agents attempt to compensate for the prior biases of their neighbors when making inferences about the observations in their neighborhood; we denote  $\bar{\epsilon} = (\epsilon_1^T, \dots, \epsilon_n^T)^T$  and refer to it as the vector of neighborhood biases. The evolution of action profiles under conditions of Theorem 1 can be specified as follows:  $\bar{\mathbf{a}}_{t+1} = (T \otimes I_k) \bar{\mathbf{a}}_t + \bar{\epsilon}$ , where  $I_k$  is the  $k \times k$  identity matrix and  $(T \otimes I_k)$  is a Kronecker product.

<sup>1</sup>Some of the non-Bayesian update rules have the property that they resemble the replication of a first step of a Bayesian update from a common prior. For instance, DeMarzo, Vayanos and Zwiebel [13] interpret the weights in the DeGroot model as those assigned initially by rational agents to the noisy opinions of their neighbors based on their perceived precision. However, by repeatedly applying the same weights over and over again, the agents ignore the need to update these weights and to account for repetitions in their information sources (the so-called persuasion bias); as one of our main objectives, we formalize this setup as a Bayesian heuristic.



Subsequently, the evolution of action profiles over time follows a non-homogeneous positive linear discrete-time dynamics, cf. [21]. If the spectral radius of  $T$  is strictly less than unity:  $\rho(T) < 1$ , then  $I - T$  is non-singular; there is a unique equilibrium action profile (the steady-state action profile that is regarded as the group decision outcome) given by  $\bar{a}_e = ((I - T)^{-1} \otimes I_k) \bar{c}$  and  $\lim_{t \rightarrow \infty} \bar{\mathbf{a}}_t = \bar{a}_e$ . If unity is an eigenvalue of  $T$ , then there may be no equilibrium action profiles or an infinity of them. If  $\rho(T) > 1$ , then the linear discrete-time dynamics is unstable and the action profiles may grow unbounded in their magnitude, cf. [43].

**EXAMPLE 1 (GAUSSIAN SIGNALS WITH GAUSSIAN BELIEFS).** Mossel and Tamuz [61] consider the case where the initial private signals as well as the unknown states are normally distributed and the agents all have full knowledge of the network structure. They show that by iteratively observing their neighbors' mean estimates and updating their beliefs using Bayes rule all agents converge to the same belief. The limiting belief is the same as what a Bayesian agent with direct access to everybody's private signals would have held; and furthermore, the belief updates at each step can be computed efficiently and convergence occurs in a number of steps that is bounded in the network size and its diameter. These results however assume complete knowledge of the network structure by all the agents.

Here, we consider the linear action updates in the Gaussian setting. Let  $\Theta = \mathbb{R}$  be the parameter space associated with the unknown parameter  $\theta \in \Theta$ . Suppose that each agent  $i \in [n]$  holds onto a Gaussian prior belief with mean  $\alpha_i \beta_i^{-1}$  and variance  $\beta_i^{-1}$ ; here,  $\gamma(\theta) = \theta^2/2$  and  $\eta(\theta) = \theta$ . Further suppose that each agent observes an independent private Gaussian signal  $\mathbf{s}_i$  with mean  $\theta$  and variance  $\sigma_i^{-1} = \delta_i^{-1}$ , for all  $i \in [n]$ ; hence,  $\xi(\mathbf{s}_i) = \mathbf{s}_i$  and  $\tau(\sigma_i \xi(\mathbf{s}_i), \delta_i) = \tau(\sigma_i \mathbf{s}_i, \sigma_i) = (2\pi/\sigma_i)^{-1/2} \exp(\sigma_i \mathbf{s}_i^2/2)$ . After observing the private signals, everybody engages in repeated communications with her neighbors. Finally, we assume that each agent is trying to estimate the mean  $m_{i,\theta} = \mathbb{E}_{i,\theta}\{\mathbf{s}_i\}$  of her private signal with as little variance as possible. Under the prescribed setting, Theorem 1 applies and the Bayesian heuristic update rules are affine with the coefficients as specified in the theorem with  $n_i = 1$  and  $\sigma_i = \delta_i$  for all  $i$ . In particular, if  $\alpha_i = (0, \dots, 0) \in \mathbb{R}^k$  and  $\beta_i \rightarrow 0$  for all  $i$ , then  $\epsilon_i = 0$  for all  $i$  and the coefficients  $T_{ij} = \sigma_j / \sum_{p \in \mathcal{N}_i} \sigma_p > 0$  specify a convex combination:  $\sum_{j \in \mathcal{N}_i} T_{ij} = 1$  for all  $i$ .

**EXAMPLE 2 (POISSON SIGNALS WITH GAMMA BELIEFS).** As the second example, suppose that each agent observe  $n_i$  i.i.d. Poisson signals  $\mathbf{s}_{i,p} : p \in [n_i]$  with mean  $\delta_i \theta$ , so that  $\Theta = \mathcal{A}_i = (0, +\infty)$  for all  $i \in [n]$ . Moreover, we take each agent's prior to be a Gamma distribution with parameters  $\alpha_i > 0$  and  $\beta_i > 0$ , denoted  $\text{Gamma}(\alpha_i, \beta_i)$ :

$$\nu_i(\theta) := \frac{d\mathcal{V}_i}{d\Lambda_1} = \frac{\beta_i^{\alpha_i}}{\Gamma(\alpha_i)} \theta^{\alpha_i-1} e^{-\beta_i \theta},$$

for all  $\theta \in (0, \infty)$  and each  $i \in [n]$ . Note that here  $\eta(\theta) = \log \theta$ ,  $\gamma(\eta(\theta)) = \exp(\eta(\theta)) = \theta$ ,  $\kappa(\alpha_i, \beta_i) = \Gamma(\alpha_i) \beta_i^{-\alpha_i}$ ,  $m_{i,\theta} = \delta_i \theta$ ,  $\xi(\mathbf{s}_{i,p}) = \mathbf{s}_{i,p}$ ,  $\sigma_i = 1$  and  $\tau(\sigma_i \xi(\mathbf{s}_{i,p}), \delta_i) = \delta_i^{\mathbf{s}_{i,p}} / (\mathbf{s}_{i,p}!)$ , for all  $i, p$ . This setting corresponds also to a case of Poisson observers with common rate  $\theta$  and individual exposures  $\delta_i, i \in [n]$ , cf. [25, p. 54]. The posterior distribution over  $\Theta$  after observing the sum of  $n_i$  Poisson mean  $\delta_i \theta$  samples is again a Gamma distribution with updated (random) parameters  $\sum_{p=1}^{n_i} \mathbf{s}_{i,p} + \alpha_i$  and  $n_i \delta_i + \beta_i$ , [25, pp. 52–53]. Using a quadratic utility  $-(a - \delta_i \theta)^2$ , the expected pay-off at time zero is maximized by the  $\delta_i$ -scaled mean of the posterior Gamma belief distribution [25, p. 587]:  $\mathbf{a}_{i,0} = \delta_i (\sum_{p=1}^{n_i} \mathbf{s}_{i,p} + \alpha_i) / (n_i \delta_i + \beta_i)$ . Given the information in the self-inclusive neighborhood:  $\sum_{p=1}^{n_j} \mathbf{s}_{j,p} = (n_j + \beta_j \delta_j^{-1}) \mathbf{a}_{j,0} - \alpha_j$ ,  $\forall j \in \mathcal{N}_i$ , agent  $i$  can refine her belief into a Gamma distribution with parameters  $\alpha_i + \sum_{j \in \mathcal{N}_i} [(n_j + \beta_j \delta_j^{-1}) \mathbf{a}_{j,0} - \alpha_j]$  and  $\beta_i + \sum_{j \in \mathcal{N}_i} n_j \delta_j$ . The subsequent optimal action at time 1 and the resultant Bayesian heuristics are as claimed in Theorem 1 with  $\sigma_i = 1$  for all  $i \in [n]$ . Here if we let  $\alpha_i, \beta_i \rightarrow 0$  and  $\delta_i = \delta > 0$  for all  $i$ , then  $\epsilon_i = 0$  for all  $i$  and the coefficients  $T_{ij} = n_j / \sum_{p \in \mathcal{N}_i} n_p > 0$  again specify a convex combination:  $\sum_{j \in \mathcal{N}_i} T_{ij} = 1$  for all  $i$  as in the DeGroot model. In the following two subsections, we shall further explore this correspondence with the DeGroot updates and the implied asymptotic consensus among the agents.

**4.1. Linear updating and convergence.** In general, the constant terms  $\epsilon_i$  in Theorem 1 depend on the neighboring prior parameters  $\alpha_j$ ,  $j \in \mathcal{N}_i \setminus \{i\}$  and can be non-zero. Accumulation of constant terms over time when  $\rho(T) \geq 1$  prevents the action profiles from converging to any finite values or may cause them to oscillate indefinitely (depending upon the model parameters). However, if the prior parameters are vanishingly small, then the affine action updates in Theorem 1 reduce to linear update and  $\epsilon_i = 0$ . This requirement on the prior parameters is captured by our next assumption.

**ASSUMPTION 2 (Non-informative priors).** For a member  $\mathcal{V}(\cdot; \alpha, \beta)$  of the conjugate family  $\mathcal{F}_{\gamma, \eta}$  we denote the limit  $\lim_{\alpha_i, \beta_i \rightarrow 0} \mathcal{V}(\cdot; \alpha, \beta)$  by  $\mathcal{V}_\emptyset(\cdot)$  and refer to it as the non-informative (and improper, if  $\mathcal{V}_\emptyset(\cdot) \notin \mathcal{F}_{\gamma, \eta}$ ) prior.<sup>1</sup> All agents start from a common non-informative prior:  $\mathcal{V}_i(\cdot) = \mathcal{V}_\emptyset(\cdot)$ ,  $\forall i$ .

As the name suggest non-informative priors do not inform the agent's action at time 0 and the optimal action is completely determined by the observed signal  $\mathbf{s}_{i,p} : p \in [n_i]$  and its likelihood structure, parameterized by  $\sigma_i$  and  $\delta_i$ . If we let  $\alpha_i, \beta_i \rightarrow 0$  in the expressions of  $T_{ij}$  and  $\epsilon_i$  from Theorem 1, then the affine action updates reduce to linear combinations and the succeeding corollary is immediate.

**COROLLARY 1 (Linear updating).** Under the exponential family signal-utility structure (Assumption 1) with non-informative priors (Assumption 2); the Bayesian heuristics describe each updated action  $\mathbf{a}_{i,t}$  as a linear combination of the neighboring actions  $\mathbf{a}_{j,t-1}$ ,  $j \in \mathcal{N}_i$ :  $\mathbf{a}_{i,t} = \sum_{j \in \mathcal{N}_i} T_{ij} \mathbf{a}_{j,t-1}$ , where  $T_{ij} = \delta_i \sigma_j n_j / (\sigma_i \sum_{p \in \mathcal{N}_i} n_p \delta_p)$ .

The action profiles under Corollary 1 evolve as a homogeneous positive linear discrete-time system:  $\bar{\mathbf{a}}_{t+1} = (T \otimes I_k) \bar{\mathbf{a}}_t$  and if the spectral radius of  $T$  is strictly less than unity, then  $\lim_{t \rightarrow \infty} \bar{\mathbf{a}}_t = \bar{\mathbf{0}}$ . For a strongly connected social network with  $T_{ii} > 0$  for all  $i$  the Perron-Frobenius theory [76, Theorems 1.5 and 1.7] implies that  $T$  has a simple positive real eigenvalue equal to  $\rho(T)$ . Moreover, the left and right eigenspaces associated with  $\rho(T)$  are both one-dimensional with the corresponding eigenvectors  $\bar{\mathbf{l}} = (l_1, \dots, l_n)^T$  and  $\bar{\mathbf{r}} = (r_1, \dots, r_n)^T$ , uniquely satisfying  $\|\bar{\mathbf{l}}\|_2 = \|\bar{\mathbf{r}}\|_2 = 1$ ,  $l_i > 0$ ,  $r_i > 0$ ,  $\forall i$  and  $\sum_{i=1}^n l_i r_i = 1$ . The magnitude of any other eigenvalue of  $T$  is strictly less than  $\rho(T)$ . If  $\rho(T) = 1$ , then  $\lim_{t \rightarrow \infty} \bar{\mathbf{a}}_t = \lim_{t \rightarrow \infty} (T^t \otimes I_k) \bar{\mathbf{a}}_0 = (\bar{\mathbf{r}} \bar{\mathbf{l}}^T \otimes I_k) \bar{\mathbf{a}}_0$ ; in particular, the asymptotic action profile may not represent a consensus although every action converges to some point within the convex hull of the initial actions  $\{\mathbf{a}_{i,0}, i \in [n]\}$ . The asymptotic actions also deviate from the efficient actions defined as follows:

**DEFINITION 1 (EFFICIENCY OF THE GROUP DECISION OUTCOME).** The efficient action of an agent in the group decision process is her expected utility maximizing action, given her Bayesian posterior belief subject to all the private signals of all agents across the network.

If  $\rho(T) > 1$ , then the linear discrete-time dynamics is unstable and the action profiles may increase or decrease without bound, pushing the decision outcome to extremes. Thus, we can associate  $\rho(T) > 1$  to cases of polarizing group interactions.

<sup>1</sup>Conjugate priors offer a technique for deriving the prior distributions based on the sample distribution (likelihood structures). However, in lack of any prior information it is impossible to justify their application on any subjective basis or to determine their associated parameters for any agent. Subsequently, the use of non-informative priors is suggested by Bayesian analysts and various techniques for selecting non-informative priors is explored in the literature [49]. Amongst the many proposed techniques for selecting non-informative priors, Jeffery's method sets its choice proportional to the square root of Fisher's information measure of the likelihood structure [72, Section 3.5.3], while Laplace's classical principle of insufficient reason favors equiprobability leading to priors which are uniform over the parameter space.

**4.2. DeGroot updates, consensus, and efficiency.** In order for the linear action updates in Corollary 1 to constitute a convex combination as in the DeGroot model,<sup>1</sup> we need to introduce some additional restrictions on the likelihood structure of the private signals.

**ASSUMPTION 3 (Locally balanced likelihoods).** *The likelihood structures given in (3) are called locally balanced if for all  $i \in [n]$ ,  $(\delta_i/\sigma_i) = (\sum_{j \in \mathcal{N}_i} \delta_j n_j) / \sum_{j \in \mathcal{N}_i} \sigma_j n_j$ .*

Assumption 3 signifies a local balance property for the two exponential family parameters  $\sigma_i$  and  $\delta_i$ , across every neighborhood in the network. In particular, we need for the likelihood structures of every agent  $i$  and her neighborhood to satisfy:  $\delta_i \sum_{j \in \mathcal{N}_i} \sigma_j n_j = \sigma_i \sum_{j \in \mathcal{N}_i} \delta_j n_j$ . Since parameters  $\sigma_i$  and  $\delta_i$  are both measures of accuracy or precision for private signals of agent  $i$ , the balance condition in Assumption 3 imply that the signal precisions are spread evenly over the agents; i.e. the quality of observations obey a rule of social balance such that no agent is in a position of superiority to everyone else. Indeed, fixing  $\delta_i = \delta$  for all  $i$ , the latter condition reduces to a harmonic property for the parameters  $\sigma_i$ , when viewed as a function of their respective nodes (cf. [57, Section 2.1] for the definition and properties of harmonic functions):

$$\sigma_i = \sum_{j \in \mathcal{N}_i} \frac{n_j \sigma_j}{\sum_{k \in \mathcal{N}_i} n_k}, \delta_i = \delta, \forall i. \quad (4)$$

However, in a strongly connected social network (4) cannot hold true unless  $\sigma_i$  is a constant:  $\sigma_i = \sigma$  for all  $i$ . Similarly, when  $\sigma_i = \sigma$  is a constant, then under Assumption 3  $\delta_i$  is spread as a harmonic function over the network nodes, and therefore can only take a constant value in a strongly connected network:  $\delta_i = \delta$  for all  $i$ , cf. [57, Section 2.1, Maximum Principle]. In particular, fixing either of the parameters  $\sigma_i$  or  $\delta_i$  for all agents and under the local balance condition in Assumption 3, it follows that the other parameter should be also fixed across the network; hence, the ratio  $\sigma_i/\delta_i$  will be a constant for all  $i$ . Later when we consider the efficiency of consensus action we introduce a strengthening of Assumption 3, called globally balanced likelihood (cf. Assumption 4), where the ratio  $\delta_i/\sigma_i$  should be a constant for all agents across the network. Examples 1 and 2 above provide two scenarios in which the preceding balancedness conditions may be satisfied: (i) having  $\sigma_i = \delta_i$  for all  $i$ , as was the case with the Gaussian signals in Example 1, ensures that the likelihoods are globally balanced; (ii) all agents receiving i.i.d. signals from a common distribution in Examples 2 (Poisson signals with the common rate  $\theta$  and common exposure  $\delta$ ) makes a case for likelihoods being locally balanced.

**THEOREM 2 (DeGroot updating and consensus).** *Under the exponential family signal-utility structure (Assumption 1), with non-informative priors (Assumption 2) and locally balanced likelihoods (Assumption 3); the updated action  $\mathbf{a}_{i,t}$  is a convex combination of the neighboring actions  $\mathbf{a}_{j,t-1}, j \in \mathcal{N}_i$ :  $\mathbf{a}_{i,t} = \sum_{j \in \mathcal{N}_i} T_{ij} \mathbf{a}_{j,t-1}$ ,  $\sum_{j \in \mathcal{N}_i} T_{ij} = 1$  for all  $i$ . Hence, in a strongly connected social network the action profiles converge to a consensus, and the consensus value is a convex combination of the initial actions  $\mathbf{a}_{i,0} : i \in [n]$ .*

In light of Theorem 2, it is of interest to know if the consensus action agrees with the minimum variance unbiased estimator of  $m_{i,\theta}$  given all the observations of every agent across the network, i.e. whether the Bayesian heuristics *efficiently* aggregate all the information amongst the networked agents. Our next result addresses this question. For that to hold we need to introduce a strengthening of Assumption 3:

<sup>1</sup>The use of linear averaging rules for modeling opinion dynamics has a long history in mathematical sociology and social psychology [24]; their origins can be traced to French's seminal work on "A Formal Theory of Social Power" [23]. This was followed up by Harary's investigation of the mathematical properties of the averaging model, including the consensus criteria, and its relations to Markov chain theory [36]. This model was later generalized to belief exchange dynamics and popularized by DeGroot's seminal work [12] on linear opinion pools.

**ASSUMPTION 4 (Globally balanced likelihoods).** *The likelihood structures given in (3) are called globally balanced if for all  $i \in [n]$  and some common constant  $C > 0$ ,  $\delta_i/\sigma_i = C$ .*

In particular, under Assumption 4,  $\sigma_i\delta_j = \sigma_j\delta_i$  for all  $i, j$ , and it follows that the local balance of likelihoods is automatically satisfied. According to Definition 1, the consensus action is *efficient* if it coincides with the minimum variance unbiased estimator of  $m_{i,\theta}$  for all  $i$  and given all the observations of every agent across the network. Our next result indicates that global balance is a necessary condition for the agents to reach consensus on a globally optimal (efficient) action. To proceed, let the network graph structure be encoded by its adjacency matrix  $A$  defined as  $[A]_{ij} = 1 \iff (j, i) \in \mathcal{E}$  for  $i \neq j$ , and  $[A]_{ij} = 0$  otherwise. Following the common convention, we consider the adjacency matrix  $A$  with zero diagonals. To express the conditions for efficiency of consensus, we need to consider the set of all agents who listen to the beliefs of a given agent  $j$ ; we denote this set of agents by  $\mathcal{N}_j^{\text{out}} := \{i \in [n] : [I + A]_{ij} = 1\}$  and refer to them as the out-neighborhood of agent  $j$ . This is in contrast to her neighborhood  $\mathcal{N}_j$ , which is the set of all agents whom she listens to. Both sets  $\mathcal{N}_j$  and  $\mathcal{N}_j^{\text{out}}$  include agent  $j$  as a member.

**THEOREM 3 ((In-)Efficiency of consensus).** *Under the exponential family signal-utility structure (Assumption 1) and with non-informative priors (Assumption 2); in a strongly connected social network, the agents achieve consensus at an efficient action if, and only if, the likelihoods are globally balanced and  $\sum_{p \in \mathcal{N}_j^{\text{out}}} n_p \delta_p = \sum_{p \in \mathcal{N}_i} n_p \delta_p$ , for all  $i$  and  $j$ . The efficient consensus action is then given by  $\mathbf{a}^* = \sum_{j=1}^n \left( \delta_j n_j \mathbf{a}_{j,0} / \sum_{p=1}^n n_p \delta_p \right)$ .*

Following our discussion of Assumption 3 and equation (4), we pointed out that if either of the two parameters  $\sigma_i$  and  $\delta_i$  that characterize the exponential family distribution of (3) are held fixed amongst the agents, then the harmonicity condition required for the local balancedness of the likelihoods implies that the other parameter is also fixed for all the agents. Therefore the local balancedness in many familiar cases (see Example 2) restricts the agents to observing i.i.d. signals: allowing heterogeneity only in the sample sizes, but not in the distribution of each sample. This special case is treated in our next corollary, where we also provide the simpler forms of the Bayesian heuristics and their linearity coefficients in the i.i.d. case:

**COROLLARY 2 (DeGroot learning with i.i.d. samples).** *Suppose that each agent  $i \in [n]$  observes  $n_i$  i.i.d. samples belonging to the same exponential family signal-utility structure (Assumption 1 with  $\sigma_i = \sigma$  and  $\delta_i = \delta$  for all  $i$ ). If the agents have non-informative priors (Assumption 2) and the social network is strongly connected, then the (no-recall) heuristic agents update their action according to the linear combination:  $\mathbf{a}_{i,t} = \sum_{j \in \mathcal{N}_i} T_{ij} \mathbf{a}_{j,t-1}$ , where  $T_{ij} = n_j / \sum_{p \in \mathcal{N}_i} n_p$ , and reach a consensus. The consensus action is efficient if, and only if,  $\sum_{p \in \mathcal{N}_j^{\text{out}}} n_p = \sum_{p \in \mathcal{N}_i} n_p$  for all  $i$  and  $j$ , and the efficient consensus action is given by  $\mathbf{a}^* = \sum_{j=1}^n \left( \mathbf{a}_{j,0} n_j / \sum_{p=1}^n n_p \right)$ .*

It is notable that the consensus value pinpointed by Theorem 2 does not necessarily agree with the MVUE of  $m_{i,\theta}$  given all the private signals of all agents across the network; in other words, by following Bayesian heuristics agents may not aggregate all the initial data efficiently. As a simple example, consider the exponential family signal-utility structure with non-informative priors (Assumptions 1 and 2) and suppose that every agent observes an i.i.d. sample from a common distribution  $\mathcal{L}(\cdot|\theta; 1, 1)$ . In this case, the action updates proceed by simple iterative averaging:  $\mathbf{a}_{i,t} = (1/|\mathcal{N}_i|) \sum_{j \in \mathcal{N}_i} \mathbf{a}_{j,t-1}$  for all  $i \in [n]$  and any  $t \in \mathbb{N}$ . For an undirected graph  $\mathcal{G}$  it is well-known that the asymptotic consensus action following simple iterative averaging is the degree-weighted average  $\sum_{i=1}^n (\text{deg}(i)/|\mathcal{E}|) \mathbf{a}_{i,0}$ , cf. [31, Section II.C]; and the consensus action is different from the global MVUE  $\mathbf{a}^* = (1/n) \sum_{i=1}^n \mathbf{a}_{i,0}$  unless the social network is a regular graph in which case,  $\text{deg}(i) = d$  is fixed for all  $i$ , and  $|\mathcal{E}| = n \cdot d$ .

REMARK 4 (EFFICIENCY OF BALANCED REGULAR STRUCTURES). In general, if we assume that all agents receive the same number of i.i.d. samples from the same distribution, then the condition for efficiency of consensus,  $\sum_{p \in \mathcal{N}_j^{\text{out}}} n_p = \sum_{p \in \mathcal{N}_i} n_p$ , is satisfied for balanced regular structures. In such highly symmetric structures, the number of outgoing and incoming links are the same for every node and equal to a fixed number  $d$ .

Our results shed light on the deviations from the globally optimal (efficient) actions, when consensus is being achieved through the Bayesian heuristics. This inefficiency of Bayesian heuristics in globally aggregating the observations can be attributed to the agents' naivety in inferring the sources of their information, and their inability to interpret the actions of their neighbors rationally, [31]; in particular, the more central agents tend to influence the asymptotic outcomes unfairly. This sensitivity to social structure is also due to the failure of agents to correct for the repetitions in the sources of their information: agent  $i$  may receive multiple copies that are all influenced by the same observations from a far away agent; however, she fails to correct for these repetitions in the sources of her observations, leading to the co-called persuasion bias, [13].

**5. Log-Linear belief updates.** When the state space  $\Theta$  is finite, the action space is the probability simplex and the agents have a quadratic utility that measures the distance between their action and the point mass on the true state, the communication structure between the agents is rich enough for them to reveal their beliefs at every time period. The Bayesian heuristic in this case leads to a log-linear updating of beliefs similar to what we analyze in [65, 67] under the Bayesian without recall model; the chief question of interest to the latter works is whether the agents, after being exposed to sequence of private observations and while communicating with each other, can learn the truth using the Bayesian without recall update rules. The learning framework of [65, 66, 67, 70] in which agents have access to a stream of new observations is in contrast with the group decision model of this paper; the difference being in the fact that here the agents have a single initial observation and engage in group decision making to come up with the best decision that aggregates their individual private data with those of the other group members.

Consider an environment where the state space is a finite set of cardinality  $m$  and agents take actions over the  $(m - 1)$ -simplex of probability measures while trying to minimize their distance to a point mass on the true state. Specifically, let  $\Theta = \{\theta_1, \dots, \theta_m\}$ ,  $\mathcal{A}_i = \Delta\Theta$  and for any probability measure  $\mathcal{M}(\cdot) \in \mathcal{A}_i$  with probability mass function<sup>1</sup>  $\mu(\cdot) := d\mathcal{M}/d\mathcal{K}$ , suppose that:

$$u_i(\mathcal{M}, \theta_j) = -(1 - \mu(\theta_j))^2 - \sum_{\substack{k=1, \\ k \neq j}}^m \mu(\theta_k)^2.$$

Subsequently,  $\mathcal{M}_{i,t} = \arg \max_{\mathcal{M} \in \Delta\Theta} \mathbb{E}_{i,t}\{u_i(\mathcal{M}, \theta)\}$ , and the agents proceed by truthfully announcing their beliefs to each other at every time step. In particular, if we denote the belief probability mass functions  $\nu_i(\cdot) := d\mathcal{V}_i/d\mathcal{K}$  and  $\mu_{i,t}(\cdot) := d\mathcal{M}_{i,t}/d\mathcal{K}$  for all  $t$ , then we can follow the steps of [69]

<sup>1</sup>Recall that  $\mathcal{K}(\cdot)$  denotes the counting measure.

to derive the Bayesian heuristic  $f_i$  in (2) by replicating the time-one Bayesian belief update for all future time-steps:<sup>1</sup>

$$\boldsymbol{\mu}_{i,t}(\theta') = \frac{\boldsymbol{\mu}_{i,t-1}(\theta') \left( \prod_{j \in \mathcal{N}_i \setminus \{i\}} \frac{\boldsymbol{\mu}_{j,t-1}(\theta')}{\nu_j(\theta')} \right)}{\sum_{\tilde{\theta} \in \Theta} \boldsymbol{\mu}_{i,t-1}(\tilde{\theta}) \left( \prod_{j \in \mathcal{N}_i \setminus \{i\}} \frac{\boldsymbol{\mu}_{j,t-1}(\tilde{\theta})}{\nu_j(\tilde{\theta})} \right)}, \text{ for all } \theta' \in \Theta \text{ and at any } t > 1. \quad (5)$$

REMARK 5 (HISTORY-NEGLECT AND NO-RECALL). In writing (5), every time agent  $i$  regards each of her neighbors  $j \in \mathcal{N}_i$  as having started from some prior belief  $\nu_j(\cdot)$  and arrived at their currently reported belief  $\boldsymbol{\mu}_{j,t-1}(\cdot)$  upon observing their private signals, hence rejecting any possibility of a past history, or learning and correlation between their neighbors. Such a rule is of course not the optimum Bayesian update of agent  $i$  at any step  $t > 1$ , because the agent is not taking into account the complete observed history of beliefs and is instead, basing her inference entirely on the initial signals and the immediately observed beliefs.

REMARK 6 (GENERALIZATION TO ONE-STEP RECALL BELIEF UPDATES). In updating her belief at time  $t$  according to (5), agent  $i$  is implicitly assuming that each of her neighbors,  $j \in \mathcal{N}_i \setminus \{i\}$ , have inferred their reported beliefs at time  $t-1$ ,  $\boldsymbol{\mu}_{j,t-1}(\cdot)$ , from a fixed prior,  $\nu_j(\cdot)$ . To improve on this assumption, one can instead use each neighbor's own beliefs at time  $t-2$  for the denominators of the multiplicative belief ratio terms in (5). This modifications requires only a single unit of memory and leads to the following update rule:

$$\boldsymbol{\mu}_{i,t}(\theta') = \frac{\boldsymbol{\mu}_{i,t-1}(\theta') \left( \prod_{j \in \mathcal{N}_i \setminus \{i\}} \frac{\boldsymbol{\mu}_{j,t-1}(\theta')}{\boldsymbol{\mu}_{j,t-2}(\theta')} \right)}{\sum_{\tilde{\theta} \in \Theta} \boldsymbol{\mu}_{i,t-1}(\tilde{\theta}) \left( \prod_{j \in \mathcal{N}_i \setminus \{i\}} \frac{\boldsymbol{\mu}_{j,t-1}(\tilde{\theta})}{\boldsymbol{\mu}_{j,t-2}(\tilde{\theta})} \right)}, \text{ for all } \theta' \in \Theta \text{ and at any } t > 1. \quad (6)$$

This requires agent  $i$  to remember the penultimate beliefs of her neighbors when interpreting their most recent reports; hence, relaxing the no-recall constraint. In Appendix D, we analyze the evolution of beliefs when the agents interact and update their beliefs according to (6). In particular, we show that (6) implements the Bayesian (rational) belief update for agent  $i$ , where the network structure is a directed rooted tree, rooted at node  $i$ . Hence, the update rule in (6) can also describe the behavior of an agent who updates her belief as if the group structure is a rooted directed tree. The bounded rationality, in this case, arise from the agent's "naivety" in regarding the actual complex network structure as a rooted directed tree. Other authors have also pointed out the grounds for such naivety by processing the streaming information from neighboring agents as independent sources [48, Remark V.4].

REMARK 7 (ALGEBRA OF BELIEFS). Both the linear action updates studied in the previous chapter as well as the weighted majority update rules that arise in the binary case and are studied in [67, 68] have a familiar algebraic structure over the respective action spaces (the Euclidean space  $\mathbb{R}^k$  and the Galois field  $\mathcal{GF}(2)$ ). In Appendix E, we develop similar structural properties for belief updates in (5) and over the space  $\Delta\Theta$ , i.e. the points of the standard  $(m-1)$ -simplex. In particular,  $\Delta\Theta$  can be endowed with an addition and a subtraction operation as well as an identity element (the uniform distribution over the state states  $\Theta$ ). In the resultant abelian group, the updated

<sup>1</sup>It is notable that the Bayesian heuristic in (5) has a log-linear structure. Geometric averaging and logarithmic opinion pools have a long history in Bayesian analysis and behavioral decision models [29, 74] and they can be also justified under specific behavioral assumptions [60]. In the past, we have investigated the implications of such log-linear behavior and properties of convergence and learning when agents are exposed to a stream of private observation [65, 66, 69, 70]. We show that the naivety of agents in these cases impedes their ability to learn; except in simple social structures such as cycles or rooted trees (cf. [66]). In [70], we show that learning in social networks with complex neighborhood structures can be achieved if agents choose a neighbor randomly at every round and restrict their belief update to the selected neighbor each time.

belief in (5) can be expressed as an addition of the self and neighboring beliefs subtracted by their priors (cf. (23) in Appendix E).<sup>1</sup>

**5.1. Becoming certain about the group aggregate.** We begin our investigation of the evolution of beliefs under (5) by considering the optimal response (belief) of an agents who has been given access to the set of all private observations across the network; indeed, such a response can be achieved in practice if one follows Kahneman’s advice and collect each individual’s information privately before combining them or allowing the individuals to engage in public discussions [47, Chapter 23]. Starting from the uniform prior and after observing everybody’s private data our aggregate belief about the truth state is given by the following implementation of the Bayes rule:

$$\boldsymbol{\mu}^*(\theta') = \frac{\prod_{j=1}^n \ell_j(\mathbf{s}_j|\theta')}{\sum_{\tilde{\theta} \in \Theta} \prod_{j=1}^n \ell_j(\mathbf{s}_j|\tilde{\theta})}. \quad (7)$$

Our next theorem describes the asymptotic outcome of the group decision process when the agents report their beliefs and follow the Bayesian heuristic (5) to aggregate them. The outcome indicated in Theorem 4 departs from the global optimum  $\boldsymbol{\mu}^*$  in two major respects. Firstly, the agents reach consensus on a belief that is supported over  $\Theta^\diamond := \arg \max_{\tilde{\theta} \in \Theta} \sum_{i=1}^n \alpha_i \log(\ell_i(\mathbf{s}_i|\tilde{\theta}))$ , as opposed to the global (network-wide) likelihood maximizer  $\Theta^* := \arg \max_{\tilde{\theta} \in \Theta} \boldsymbol{\mu}^*(\tilde{\theta}) = \arg \max_{\tilde{\theta} \in \Theta} \sum_{i=1}^n \log(\ell_i(\mathbf{s}_i|\tilde{\theta}))$ ; note that the signal log-likelihoods in the case of  $\Theta^\diamond$  are weighted by the centralities ( $\alpha_i$ ) of their respective nodes. Secondly, the consensus belief is concentrated uniformly over  $\Theta^\diamond$ , its support does not include the entire state space  $\Theta$  and those states which score lower on the centrality-weighted likelihood scale are asymptotically rejected as a candidate for the truth state; in particular, if  $\{\theta^\diamond\} = \Theta^\diamond$  is a singleton, then the agents effectively become certain about the truth state of  $\theta^\diamond$ , in spite of their essentially bounded aggregate information and in contrast with the rational (optimal) belief  $\boldsymbol{\mu}^*$  that is given by the Bayes rule in (7) and do not discredit or reject any of the less probable states. This unwarranted certainty in the face of limited aggregate data is a manifestation of the group polarization effect that derive the agents to more extreme beliefs, rejecting the possibility of any alternatives outside of the most probable states  $\Theta^\diamond$ .

**THEOREM 4 (Certainty about the group aggregate).** *Under the no-recall belief update (5), with  $\Theta^\diamond := \arg \max_{\tilde{\theta} \in \Theta} \sum_{i=1}^n \alpha_i \log(\ell_i(\mathbf{s}_i|\tilde{\theta}))$ , we have that  $\lim_{t \rightarrow \infty} \boldsymbol{\mu}_{i,t}(\theta') = 1/|\Theta^\diamond|$  for  $\theta' \in \Theta^\diamond$  and  $\lim_{t \rightarrow \infty} \boldsymbol{\mu}_{i,t}(\theta') = 0$  for  $\theta' \notin \Theta^\diamond$ . In particular, if the sum of signal log-likelihoods weighted by node centralities is uniquely maximized by  $\theta^\diamond$ , i.e.  $\{\theta^\diamond\} = \Theta^\diamond$ , then  $\lim_{t \rightarrow \infty} \boldsymbol{\mu}_{i,t}(\theta^\diamond) = 1$ .*

**REMARK 8 (RELATIVE EFFICIENCY OF BALANCED REGULAR NETWORKS).** The fact that log-likelihoods in  $\Theta^\diamond$  are weighted by the node centralities is a source of inefficiency for the asymptotic outcome of the group decision process. This inefficiency is warded off in especially symmetric typologies, where in and out degrees of all nodes in the network are the same. In these so-called balanced regular digraphs, there is a fixed integer  $d$  such that all agents receive reports from exactly

<sup>1</sup>It is instructive to also point out the propinquity to “cognitive algebras” that arise in information integration theory. Indeed, cognitive and psychological roots of the Bayesian heuristics as aggregation rules can be traced to Anderson’s seminal theory of information integration, developed throughout 1970s and 1980s [3]. Accordingly, a so-called “value function” assigns psychological values to each of the stimuli and these psychological values are then combined into a single psychological (and later an observable) response through what is called the “integration function”. A fundamental assumption is that valuation can be represented at a higher (molar) level as a value on the response dimension for each stimulus, as well as a weight representing the salience of this stimulus in the overall response. These valuations and weights are themselves the result of integration processes in the lower (molecular) level. At the heart of information integration theory is the “cognitive algebra” which describes the rules by which the values and weights of stimuli are integrated into an overall response [4].

$d$  agents, and also send their reports to some other  $d$  agents;  $d$ -regular graphs are a special case, since all links are bidirectional and each agent sends her reports to and receive reports from the same  $d$  agents. In such structures  $\bar{\alpha} = (1/n)\mathbb{1}$  so that  $\Theta^* = \Theta^\diamond$  and the support of the consensus belief identifies the global maximum likelihood estimator (MLE); i.e. the maximum likelihood estimator of the unknown  $\theta$ , given the entire set of observations from all agents in the network.

**5.2. Teaming for efficient belief exchange.** Studies in committee mechanism design [30, 51], and Group Decision Support Systems (GDSS) [14], strive to construct mechanisms for information aggregation so that the group members arrive at optimal results. The overconfidence that Theorem 4 predicts as a result of redundancy in aggregated beliefs has many adverse consequences to the health, wealth and welfare of the general public: overconfident investors take on excessively risky projects, overconfident doctors reject otherwise beneficial alternative treatments, overconfident voters are susceptible to polarization and be enticed by political extremes, overconfident jurors impose excessive penalties on presumed culprits in terms of fines and damages, etc.

We begin by the observation that the no-recall belief updates in (5) are efficient in aggregating the observations of neighboring agents after hearing their reports at time one. Hence, one may propose an efficient deliberation mechanism by having all agents hear each other’s initial beliefs only once; and subsequently expect them all to hold efficient beliefs at time one as predicted by (5). Indeed, such a scheme achieves efficient belief aggregation, and the time-one beliefs of all agents would coincide with their Bayesian posteriors given all the private signals in the network (cf. (22)). However, for large groups the requirement to listen to reports of all other  $n - 1$  agents, all at once, may impose excessive cognitive and communicative burdens on the agent and would thus be difficult to implement in practice.

To mitigate the burdens of meeting in large groups, we propose the deliberations to take place through a sequence of coordinated meetings, such that at any given period agents meet in groups of size at most  $D$ .<sup>1</sup> Let  $T_{n,D} := \lceil \log_D(n) \rceil$ , where  $\lceil \cdot \rceil$  denotes the smallest integer greater than or equal to its argument. Then a total of  $T_{n,D}$  time-steps would be enough to coordinate all the necessary meetings for all the private information to be aggregated in everybody’s beliefs. To achieve efficient belief aggregation, the meetings should be coordinated such that people are regrouped at the end of each round, after reporting their beliefs to and hearing the beliefs of the other group members at that round. The regrouping should take into account the history of interactions among agents to avoid redundancies in the aggregate belief following no-recall heuristics. The procedure is not sensitive to the initial assignment of agents to different subgroups, but a random assignment at  $t = 0$  is a reasonable design choice. At the ensuing periods  $t > 0$ , a set of  $D$  people will be grouped

<sup>1</sup>The question of the optimal group size  $D$  for efficient group performance has a long history in social psychology, going back to mid-twentieth century [80, 82]. A classic study of group size in 1958 by Slater [80] concludes that groups of size five are most efficient for dealing with intellectual tasks that involve the collection or exchange of information and decision-making based on the aggregate information. The conventional wisdom is that there is an inverted-U relationship between team size and group performance, but the exact nature of the relationship depends on the task types and specific coordination requirements. The inverted U-relationship has been pointed out by Steiner and others [82, 56], who articulate the trade-offs between additions of individual skills with the increasing group size, and faculty losses that are due to motivational or coordinational shortcomings of large groups. The combination of such conflicting effects leads to an inverted-U relationship between the team size and group performance: on the one hand, larger groups have more potential for productivity but these potentials are compromised by the difficulties of coordination, communication and sustained motivation in large groups [8]. Determining an optimal group size in each case depends, not only, on the task type and task requirements, but also on the social relations between group members and other intra-group factors [35, 22]. A good rule is to have sufficiently many members, but not greater than that [45, 34]. Overall, team size is considered to be important in determining group performance by both theoretical and experimental studies [86, 32]. It continues to attract attention in social and organizational psychology for both business and management applications [11, 38, 75, 58].



together (and hear each other's beliefs) only if no two of them have met with each other or with a same person or with other people who have met between themselves or with a same person, or other people who have met with other people who have met between themselves or with a same person, and so on and so forth. This no-redundancy requirement can be expressed inductively as follows:

**CONDITION 1 (No-redundancy).**

- (i) Two people cannot meet if they have met with each other or with a same person.
- (ii) Two people cannot meet if they have met with people who themselves cannot meet with each other.

The following procedure ensures that the coordinated meetings satisfy the requirements of Condition 1 (no-redundancy). To describe the procedure formally, let  $\bar{n} := D^{T_{n,D}}$ ,  $X_{n,D,t} = D^{T_{n,D}-t-1} = \bar{n}/D^{t+1}$ ,  $Z_{n,D} = \bar{n}/D = D^{T_{n,D}-1} = D^t X_{n,D,t}$ , where  $t$  indexes the rounds of communications and  $T_{n,D} = \lceil \log_D(n) \rceil$  is as defined above.

**(A1: NO-REDUNDANCY-COORDINATION).** For  $n$  people to exchange beliefs in  $T_{n,D}$  rounds, with  $Z_{n,D}$  parallel meetings during each round and at most  $D$  people in each meeting:

1. Impose an arbitrary (possibly random) ordering  $\prec$  on the agents and label them accordingly:  $1 \prec 2 \prec \dots \prec n$ .
2. Add  $\bar{n} - n$  dummy agents labeled by  $\{n+1, n+2, \dots, \bar{n}\}$  to the group (if  $\log_D(n)$  is an integer, then  $\bar{n} = n$  is a power of  $D$  and no dummy agents are added).
3. For  $t = 0, \dots, T_{n,D} - 1$ , organize  $Z_{n,D} = D^t X_{n,D,t}$  meetings in parallel, where the participants in each meeting are indexed as follows:

$$\{\tau + \tau' D^t + \tau'' D^{t+1} : \tau' = 0, 1, 2, \dots, D-1\}, \quad (8)$$

and the range of the constants  $\tau$  and  $\tau''$  are given by:  $\tau = 1, 2, \dots, D^t$ , and  $\tau'' = 0, 1, 2, \dots, X_{n,D,t} - 1$ , generating the desired  $Z_{n,D} = D^t X_{n,D,t}$  meetings at round  $t$ .

According to (A1), at  $t=0$  people meet and learn each other's private signals in groups of size  $D$ . At  $t=1$ ,  $D$  groups of size  $D$  are "combined" according to the ordering that has been imposed on the agents from step (1) of (A1): highest ranked members from each of the  $D$  groups meet each other in a new group of size  $D$ , second highest ranked members also meet each other, and the third highest ranked members, and so on and so forth until the least ranked members. At the end of  $t=1$ , there are  $X_{n,D,1} = \bar{n}/D^2$  groups of size  $D^2$  each. The beliefs of the members of each subgroup aggregates all the information pertaining to the  $D^2$  private signals that are available in the respective subgroups after the second round of meetings ( $t=1$ ). In general, at round  $t$ ,  $D$  subgroups of size  $D^t$  are combined and new meetings are coordinated among their members (according to the ordering of the agents) such that at the end of round  $t$  there are  $X_{n,D,t}$  subgroups of size  $D^{t+1}$  and the information in each subgroup is fully aggregated among its members. In (8),  $\tau$  indexes the members of each subgroup of size  $D^t$  according to their rankings and  $\tau''$  indexes the  $X_{n,D,t}$  different subgroups whose information will be aggregated at the end of round  $t$ . This aggregation continues to propagate by combining each  $D$  subgroups at every round, until the final round ( $t = T_{n,D} - 1$ ) where  $D$  subgroups of size  $D^{T_{n,D}-1}$  will be combined. At the end of the final round, there would be  $X_{n,D,T_{n,D}-1} = 1$  group of size  $\bar{n} = D^{T_{n,D}}$  with fully aggregated information among all members.

**EXAMPLE 3 (SIX PEOPLE COORDINATED TO MEET IN PAIRS).** In Fig. 2, we provide an example implementation of (A1) with  $D=2$  and  $n=6$ . The top diagram in Fig. 2 depicts the propagation




FIGURE 2. Implementation of (A1) for a group with  $D = 2$  and  $n = 6$ . Agents 7 and 8 are dummies and whoever is paired with them at a particular round, will be idle at that round.

of information with advancing time steps, as groups are merged and agents are regrouped in new pairs. The bottom diagrams of Fig. 2 depict the flow of information in the network and across time as agents are regrouped according to (A1). The no-redundancy requirement of Condition 1 implies that there is a unique path connecting each of the agents at  $t = 0$  to another agent at  $t = 2$ . Subsequently, the network is a directed rooted tree from the viewpoint of each of the agents at the end of the meetings. The coordination imposed by (A1) ensures that after  $T_{n,D}$  rounds there are  $\bar{n}$  rooted spanning trees, one for each agent, giving them access to the aggregate information with no redundancy (see Fig. 2, the diagram on the bottom left).

It is worth highlighting that if we do not follow a coordinated schedule, then situations may arise where reaching an optimal belief becomes impossible, given the history of past meetings, even if that history has not expose the agents to any redundancy hitherto. In the example above with  $n = 6$  and  $D = 2$ , suppose we pair the agents such that  $\{1, 2\}$ ,  $\{3, 4\}$ ,  $\{5, 6\}$  meet at time one and  $\{2, 3\}$ ,  $\{1, 5\}$ ,  $\{4, 6\}$  meet at time two. Then there is no way to pair the agents for time three without exposing them to redundancy, and if we stop at time two, then the beliefs are sub-optimal (there is some information missing from each person's belief). For example, on the one hand, agent 1 has not learned about the private information of agents 3 or 4 and on the other hand, pairing her with any of the other agents 2, 3, 4, 5 or 6 at time three would expose her (and the agent with whom she is paired) to some redundancy.

If the agents start from uniform common priors, then the no-recall belief update in (5) simplifies as follows:

$$\mu_{i,t}(\theta') = \frac{\mu_{i,t-1}(\theta') \prod_{j \in G_{i,t}} \mu_{j,t-1}(\theta')}{\sum_{\tilde{\theta} \in \Theta} \mu_{i,t-1}(\tilde{\theta}) \prod_{j \in G_{i,t}} \mu_{j,t-1}(\tilde{\theta})}. \quad (9)$$

where  $G_{i,t} \subset [n]$  is the group of people with whom agent  $i$  is scheduled to meet at time  $t$ . Under (9) and after attending  $T_{n,D}$  meetings all agents will hold the (common) Bayesian aggregate belief in (7). However, if different agents start with different priors then after no-recall updating they will each hold the Bayesian posterior belief given all the information available in the network subject to their own priors (cf. (22) in Appendix D).

**6. Concluding remarks and discussions.** We propose the Bayesian heuristics framework to address the problem of information aggregation and decision making in groups. Our model is consistent with the dual process theory of mind with one system developing the heuristics through deliberation and slow processing, and another system adopting the heuristics for fast and automatic decision making: once the time-one Bayesian update is developed, it is used as a heuristic for all future decision epochs. On the one hand, this model offers a behavioral foundation for non-Bayesian updating; in particular, linear action updates and log-linear belief updates. On the other hand, its deviation from the rational choice theory captures common fallacies of snap-judgments and history neglect that are observed in real life. Our behavioral method also complements the axiomatic approaches which investigate the structure of belief aggregation rules and require them to satisfy specific axioms such as label neutrality and imperfect recall, as well as independence or separability for log-linear and linear rules, respectively [60].

We showed that under a natural quadratic utility and for a wide class of distributions from the exponential family the Bayesian heuristics correspond to a minimum variance Bayes estimation with a known linear structure. If the agents have non-informative priors, and their signal structures satisfy certain homogeneity conditions, then these action updates constitute a convex combination as in the DeGroot model, where agents reach consensus on a point in the convex hull of their initial actions. In case of belief updates (when agents communicate their beliefs), we showed that the agents update their beliefs proportionally to the product of the self and neighboring beliefs. Subsequently, their beliefs converge to a consensus supported over a maximum likelihood set, where the signal likelihoods are weighted by the centralities of their respective agents.

Our results indicate certain deviations from the globally efficient outcomes, when consensus is being achieved through the Bayesian heuristics. This inefficiency of Bayesian heuristics in globally aggregating the observations is attributed to the agents' naivety in inferring the sources of their information, which makes them vulnerable to structural network influences, in particular, redundancy and multipath effects: the share of centrally located agents in shaping the asymptotic outcome is more than what is warranted by the quality of their data. Another source of inefficiency is in the group polarization that arise as a result of repeated group interactions; in case of belief updates, this is manifested in the structure of the (asymptotic) consensus beliefs. The latter assigns zero probability to any alternative that scores lower than the maximum in the weighted likelihoods scale: the agents reject the possibility of less probable alternatives with certainty, in spite of their limited initial data. This overconfidence in the group aggregate and shift toward more extreme beliefs is a key indicator of group polarization and is demonstrated very well by the asymptotic outcome of the group decision process.

We pinpoint some key differences between the action and belief updates (linear and log-linear, respectively): the former are weighted updates, whereas the latter are unweighted symmetric updates. Accordingly, an agent weighs each neighbor's action differently and in accordance with the quality of their private signals (which she expects in them and infers from their actions). On the other hand, when communicating their beliefs the quality of each neighbor's signal is already internalized in their reported beliefs; hence, when incorporating her neighboring beliefs, an agent regards the reported beliefs of all her neighbors equally and symmetrically. Moreover, in the case of linear action updates the initial biases are amplified and accumulated in every iteration. Hence, the interactions of biased agents are very much dominated by their prior beliefs rather than their observations. This issue can push their choices to extremes, depending on the aggregate value of their initial biases. Therefore, if the Bayesian heuristics are to aggregate information from the observed actions satisfactorily, then it is necessary for the agents to be unbiased, i.e. they should hold non-informative priors about the state of the world and base their actions entirely on their observations. In contrast, when agents exchange beliefs with each other the multiplicative belief update can aggregate the observations, irrespective of the prior beliefs. The latter are asymptotically canceled; hence, multiplicative belief updates are robust to the influence of priors.

The Bayesian heuristics approach is strongly motivated by the behavioral processes that underlie human decision making. These processes often deviate from the predictions of the rational choice theory, and our investigation of the Bayesian heuristics highlights both the mechanisms for such deviations and their ramifications. In our ongoing research, we expand this behavioral approach by incorporating additional cognitive biases such as inattentiveness, and investigate how the decision processes are affected. On the one hand, the obtained insights highlight the value of educating the public about benefits of rational decision making and unbiased judgment, and how to avoid common cognitive errors when making decisions. On the other hand, by investigating the effect of cognitive biases, we can improve the practice of social and organizational policies, such that new designs can accommodate commonly observed biases, and work well in spite of them.

## Appendices

**Appendix A: Proof of Theorem 1.** If agent  $i$  starts from a prior belief  $\mathcal{V}_i(\cdot) = \mathcal{V}(\cdot; \alpha_i, \beta_i) \in \mathcal{F}_{\gamma, \eta}$ , then we can use the Bayes rule to verify that, cf. [72, Proposition 3.3.13], the Radon-Nikodym derivative of the Bayesian posterior of agent  $i$  after observing  $n_i$  samples  $\mathbf{s}_{i,p} \in \mathcal{S}$ ,  $p \in [n_i]$ , with likelihood (3) is  $\nu(\cdot; \alpha_i + \sigma_i \sum_{p=1}^{n_i} \xi(\mathbf{s}_{i,p}), \beta_i + n_i \delta_i)$ , and in particular the Bayesian posterior at time zero belongs to the conjugate family  $\mathcal{F}_{\gamma, \eta}$ :  $\mathcal{M}_{i,0}(\cdot) = \mathcal{V}(\cdot; \alpha_i + \sigma_i \sum_{p=1}^{n_i} \xi(\mathbf{s}_{i,p}), \beta_i + n_i \delta_i)$ .

Subject to the quadratic utility  $u_i(a, \theta) = -(a - m_{i,\theta})^T (a - m_{i,\theta})$ , the expected pay-off at any time  $t$  is maximized is by choosing [7, Lemma 1.4.1]:

$$\mathbf{a}_{i,t} = \mathbb{E}_{i,t}\{m_{i,\theta}\} := \int_{\theta \in \Theta} m_{i,\theta} \mathcal{M}_{i,t}(d\theta),$$

which coincides with her minimum variance unbiased estimator (Bayes estimate) for  $m_{i,\theta}$ . The members of the conjugate family  $\mathcal{F}_{\gamma, \eta}$  satisfy the following linearity property of the Bayes estimates that is key to our derivations.<sup>1</sup>

**LEMMA 1 (Proposition 3.3.14 of [72]).** *Let  $\zeta \in \mathbb{R}^k$  be a parameter and suppose that the parameter space  $\Omega_\zeta$  is an open set in  $\mathbb{R}^k$ . Suppose further that  $\zeta \in \Omega_\zeta$  has the prior distribution  $\mathcal{W}(\cdot; \alpha, \beta)$  with density  $\kappa'(\alpha, \beta) e^{\zeta^T \alpha - \beta \gamma(\zeta)}$  w.r.t.  $\Lambda_k$  where  $\kappa'(\alpha, \beta)$  is the normalization constant. If  $\mathbf{s}' \in \mathcal{S}' \subset \mathbb{R}^k$  is a random signal with distribution  $\mathcal{D}(\cdot; \zeta)$  and density  $\tau'(\mathbf{s}) e^{\zeta^T \mathbf{s} - \gamma'(\zeta)}$  w.r.t.  $\Lambda_k$ , then*

$$\int_{\zeta \in \Omega_\zeta} \int_{\mathbf{s}' \in \mathcal{S}'} \mathbf{s}' \mathcal{D}(d\mathbf{s}'; \zeta) \mathcal{W}(d\zeta; \alpha, \beta) = \frac{\alpha}{\beta}.$$

Hence for any  $\mathcal{V}(\cdot; \alpha, \beta) \in \mathcal{F}_{\gamma, \eta}$  we can write

$$\begin{aligned} \int_{\theta \in \Theta} m_{i,\theta} \mathcal{V}(d\theta; \alpha, \beta) &= \int_{\theta \in \Theta} \int_{s \in \mathcal{S}} \xi(s) \mathcal{L}(ds|\theta; \sigma, \delta) \mathcal{V}(d\theta; \alpha, \beta) \\ &= \int_{\theta \in \Theta} \left| \frac{\Lambda_k(\eta(d\theta))}{\mathcal{G}_\theta(d\theta)} \right| \frac{e^{\eta(\theta)^T \alpha - \beta \gamma(\eta(\theta))}}{\kappa(\alpha, \beta)} \mathcal{G}_\theta(d\theta) \int_{s \in \mathcal{S}} \xi(s) \sigma \left| \frac{\Lambda_k(\xi(ds))}{\mathcal{G}_s(ds)} \right| \tau(\sigma \xi(s), \delta) e^{\sigma \eta(\theta)^T \xi(s) - \delta \gamma(\eta(\theta))} \mathcal{G}_s(ds) \\ &= \int_{\zeta \in \Omega_\zeta} \left| \frac{\Lambda_k(\eta(d\theta))}{\mathcal{G}_\theta(d\theta)} \right| \frac{e^{\eta(\theta)^T \alpha - \beta \gamma(\eta(\theta))}}{\kappa(\alpha, \beta)} \mathcal{G}_\theta(d\theta) \int_{s \in \mathcal{S}} \xi(s) \sigma \left| \frac{\Lambda_k(\xi(ds))}{\mathcal{G}_s(ds)} \right| \tau(\sigma \xi(s), \delta) e^{\sigma \eta(\theta)^T \xi(s) - \delta \gamma(\eta(\theta))} \mathcal{G}_s(ds) \\ &= \int_{\zeta \in \Omega_\zeta} \frac{e^{\zeta^T \alpha - \frac{\beta}{\sigma} \gamma'(\zeta)}}{\kappa(\alpha, \beta)} \Lambda_k(d\zeta) \int_{s' \in \mathcal{S}'} \frac{s' \tau'(s')}{\sigma} e^{\zeta^T s' - \gamma'(\zeta)} \Lambda_k(ds') = \frac{\alpha \delta}{\sigma \beta}, \end{aligned} \quad (10)$$

<sup>1</sup>In fact, such an affine mapping from the observations to the Bayes estimate characterizes the conjugate family  $\mathcal{F}_{\gamma, \eta}$  and every member of this family can be uniquely identified from the constants of the affine transform [15].

where in the penultimate equality we have employed the following change of variables:  $\zeta = \eta(\theta)$ ,  $s' = \sigma\xi(s)$ ,  $\gamma'(\zeta) = \delta\gamma(\zeta)$ ,  $\tau'(s') = \tau(s', \delta)$ ; and the last equality is a direct application of Lemma 1. In particular, given  $\mathcal{M}_{i,0}(\cdot) = \mathcal{V}(\cdot; \alpha_i + \sigma_i \sum_{p=1}^{n_i} \xi(\mathbf{s}_{i,p}), \beta_i + n_i \delta_i)$ , the expectation maximizing action at time zero coincides with:

$$\mathbf{a}_{i,0} = \frac{\sum_{p=1}^{n_i} \xi(\mathbf{s}_{i,p}) + \sigma_i^{-1} \alpha_i}{n_i + \delta_i^{-1} \beta_i}. \quad (11)$$

Subsequently, following her observations of  $\mathbf{a}_{j,0}$ ,  $j \in \mathcal{N}_i$  and from her knowledge of her neighbor's priors and signal likelihood structure, agent  $i$  infers the observed values of  $\sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p})$  for all her neighbors. Hence, we get

$$\sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}) = (n_j + \delta_j^{-1} \beta_j) \mathbf{a}_{j,0} - \sigma_j^{-1} \alpha_j, \forall j \in \mathcal{N}_i. \quad (12)$$

The observations of agent  $i$  are therefore augmented by the set of independent samples from her neighbors:  $\{\sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}) : j \in \mathcal{N}_i\}$ , and her refined belief at time 1 is again a member of the conjugate family  $\mathcal{F}_{\gamma,\eta}$  and is give by:

$$\mathcal{M}_{i,1}(\cdot) = \mathcal{V}(\cdot; \alpha_i + \sum_{j \in \mathcal{N}_i} \sigma_j \sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}), \beta_i + \sum_{j \in \mathcal{N}_i} n_j \delta_j).$$

We can again invoke the linearity of the Bayes estimate for the conjugate family  $\mathcal{F}_{\gamma,\eta}$  and the subsequent result in (10), to get that the expected pay-off maximizing action at time 1 is given by:

$$\mathbf{a}_{i,1} = \frac{\delta_i \left( \alpha_i + \sum_{j \in \mathcal{N}_i} \sigma_j \sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}) \right)}{\sigma_i \left( \beta_i + \sum_{j \in \mathcal{N}_i} n_j \delta_j \right)}. \quad (13)$$

Finally, we can use (12) to replace for the neighboring signals and derive the expression of the action update of agent  $i$  at time 1 in terms of her own and the neighboring actions  $\mathbf{a}_{j,0}$ ,  $j \in \mathcal{N}_i$ ; leading to the expression of linear Bayesian heuristics as claimed in Theorem 1.  $\square$

**Appendix B: Proof of Theorem 2.** The balancedness of likelihoods (Assumption 3) ensures that the coefficients of the linear combination from Corollary 1 sum to one:  $\sum_{j \in \mathcal{N}_i} T_{ij} = 1$ , for all  $i$ ; thus forming a convex combination as in the DeGroot model. Subsequently, the agents begin by setting  $\mathbf{a}_{i,0} = \sum_{p=1}^{n_i} \xi(\mathbf{s}_{i,p})/n_i$  according to (11), and at every  $t > 1$  they update their actions according to  $\bar{\mathbf{a}}_t = (T \otimes I_k) \bar{\mathbf{a}}_{t-1} = (T^t \otimes I_k) \bar{\mathbf{a}}_0$ , where  $\bar{\mathbf{a}}_t = (\mathbf{a}_{1,t}^T, \dots, \mathbf{a}_{n,t}^T)^T$  and  $T$  is the  $n \times n$  matrix whose  $i, j$ -th entry is  $T_{ij}$ . Next note from the analysis of convergence for DeGroot model, cf. [31, Proposition 1], that for a strongly connected network  $\mathcal{G}$  if it is aperiodic (meaning that one is the greatest common divisor of the lengths of all its circles; and it is the case for us, since the diagonal entries of  $T$  are all non-zero), then  $\lim_{\tau \rightarrow \infty} T^\tau = \mathbf{1} \bar{\mathbf{s}}^T$ , where  $\bar{\mathbf{s}} := (s_1, \dots, s_n)^T$  is the unique left eigenvector associated with the unit eigenvalue of  $T$  and satisfying  $\sum_{i=1}^n s_i = 1$ ,  $s_i > 0$ ,  $\forall i$ . Hence, starting from non-informative priors agents follow the DeGroot update and if  $\mathcal{G}$  is also strongly connected, then they reach a consensus at  $\sum_{i=1}^n s_i \mathbf{a}_{i,0} = \sum_{i=1}^n s_i (\sum_{p=1}^{n_i} \xi(\mathbf{s}_{i,p})/n_i)$ .  $\square$

**Appendix C: Proof of Theorem 3.** We begin by a lemma that determines the so-called global MVUE for each  $i$ , i.e. the MVUE of  $m_{i,\theta}$  given all the observations of all agents across the network.

LEMMA 2 (**Global MVUE**). *Under the exponential family signal-utility structure (Assumption 1), the (global) MVUE of  $m_{i,\theta}$  given the entire set of observations of all the agents across the network is given by:*

$$\mathbf{a}_i^* = \frac{\delta_i \left( \alpha_i + \sum_{j=1}^n \sigma_j \sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}) \right)}{\sigma_i \left( \beta_i + \sum_{j=1}^n n_j \delta_j \right)}. \quad (14)$$

If we further impose non-informative priors (Assumption 2), then the global MVUE for each  $i$  can be rewritten as

$$\mathbf{a}_i^* = \frac{\delta_i \left( \sum_{j=1}^n \sigma_j \sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}) \right)}{\sigma_i \left( \sum_{j=1}^n n_j \delta_j \right)} = \frac{\delta_i}{\sigma_i} \sum_{j=1}^n \frac{\sigma_j n_j}{\sum_{p=1}^n n_p \delta_p} \mathbf{a}_{j,0}. \quad (15)$$

This lemma can be proved easily. Following the same steps that lead to (13), yields (14). Next by making the necessary substitutions under Assumption 2, (15) is obtained. From (15), it is immediately clear that if some consensus action is to be the efficient estimator (global MVUE) for all agents  $i \in [n]$ , then we need  $\delta_i \sigma_j = \sigma_i \delta_j$  for all  $i, j$ ; hence, the global balance is indeed a necessary condition. Under this condition, the local balance of likelihoods (Assumption 3) is automatically satisfied and given non-informative priors Theorem 2 guarantees convergence to consensus in a strongly connected social network. Moreover, we can rewrite (15) as  $\mathbf{a}_i^* = \mathbf{a}^* = (\sum_{j=1}^n \delta_j n_j \mathbf{a}_{j,0}) / \sum_{p=1}^n n_p \delta_p$ , for all  $i$ . Hence, if the consensus action  $(\sum_{i=1}^n s_i \mathbf{a}_{i,0})$  in the proof of Theorem 2, Appendix B) is to be efficient then we need  $s_i = \delta_i n_i / \sum_{j=1}^n n_j \delta_j$  for all  $i$ ;  $\bar{s} = (s_1, \dots, s_n)$  being the unique normalized left eigenvector associated with the unit eigenvalue of  $T$ :  $\bar{s}^T T = \bar{s}^T$ , as defined in Appendix B. Using  $\delta_i \sigma_j = \sigma_i \delta_j$ , we can also rewrite the coefficients  $T_{ij}$  of the DeGroot update in Theorem 2 as  $T_{ij} = \delta_j n_j / (\sum_{p \in \mathcal{N}_i} n_p \delta_p)$ .

Therefore, by expanding the eigenvector condition  $\bar{s}^T T = \bar{s}^T$  we obtain that in order for the consensus action  $\bar{s}^T \bar{\mathbf{a}}_0$  to agree with the efficient consensus  $\mathbf{a}^*$ , it is necessary and sufficient to have that for all  $j$

$$\sum_{i=1}^n s_i T_{ij} = \sum_{i=1}^n \left( \frac{\delta_i n_i}{\sum_{j=1}^n \delta_j n_j} \right) \frac{\delta_j n_j [I + A]_{ij}}{\sum_{p \in \mathcal{N}_i} n_p \delta_p} = s_j = \frac{\delta_j n_j}{\sum_{j=1}^n \delta_j n_j}, \quad (16)$$

or equivalently,

$$\sum_{i: j \in \mathcal{N}_i} \frac{\delta_i n_i}{\sum_{p \in \mathcal{N}_i} n_p \delta_p} = \sum_{i \in \mathcal{N}_i^{out}} \frac{\delta_i n_i}{\sum_{p \in \mathcal{N}_i} n_p \delta_p} = 1, \quad (17)$$

for all  $j$ . Under the global balance condition (Assumption 4),  $\delta_i \sigma_j = \delta_j \sigma_i$ , the weights  $T_{ij} = \delta_j n_j / (\sum_{p \in \mathcal{N}_i} n_p \delta_p)$  as given above, correspond to transition probabilities of a node-weighted random walk on the social network graph, cf. [9, Section 5]; where each node  $i \in [n]$  is weighted by  $w_i = n_i \delta_i$ . Such a random walk is a special case of the more common type of random walks on weighted graphs where the edge weights determine the jump probabilities; indeed, if for any edge  $(i, j) \in \mathcal{E}$  we set its weight equal to  $w_{i,j} = w_i w_j$  then the random walk on the edge-weighted graph reduces to a random walk on the node-weighted graph with node weights  $w_i, i \in [n]$ . If the social network graph is undirected and connected (so that  $w_{i,j} = w_{j,i}$  for all  $i, j$ ), then the edge-weighted (whence also the node-weighted) random walks are time-reversible and their stationary distributions  $(s_1, \dots, s_n)^T$  can be calculated in closed form as follows [2, Section 3.2]:

$$s_i = \frac{\sum_{j \in \mathcal{N}_i} w_{i,j}}{\sum_{i=1}^n \sum_{j \in \mathcal{N}_i} w_{i,j}}. \quad (18)$$

In a node-weighted random walk we can replace  $w_{i,j} = w_i w_j$  for all  $j \in \mathcal{N}_i$  and (18) simplifies into

$$s_i = \frac{w_i \sum_{j \in \mathcal{N}_i} w_j}{\sum_{i=1}^n w_i \sum_{j \in \mathcal{N}_i} w_j}.$$

Similarly to (16), the consensus action will be efficient if and only if

$$s_i = \frac{w_i \sum_{j \in \mathcal{N}_i} w_j}{\sum_{i=1}^n w_i \sum_{j \in \mathcal{N}_i} w_j} = \frac{w_i}{\sum_{k=1}^n w_k}, \forall i, \text{ or equivalently: } \left( \sum_{k=1}^n w_k \right) \sum_{j \in \mathcal{N}_i} w_j = \sum_{i=1}^n (w_i \sum_{j \in \mathcal{N}_i} w_j), \forall i,$$

which holds true only if  $\sum_{j \in \mathcal{N}_i} w_j$  is a common constant that is the same for all agents, i.e.  $\sum_{j \in \mathcal{N}_i} w_j = \sum_{j \in \mathcal{N}_i} \delta_j n_j = C' > 0$  for all  $i \in [n]$ . Next replacing in (17) yields that, in fact,  $C' = \sum_{i \in \mathcal{N}_j^{\text{out}}} \delta_i n_i$  for all  $j$ , completing the proof for the conditions of efficiency.  $\square$

**Appendix D: Generalization to belief updates with one-step recall.** Defining the log-ratio of beliefs, priors, and likelihoods at different states facilitates the analysis of the one-step recall update rule in (6). For any pair of states  $\hat{\theta}$  and  $\check{\theta}$  and all  $i, t$ , let

$$\phi_{i,t}(\hat{\theta}, \check{\theta}) := \log \left( \boldsymbol{\mu}_{i,t}(\hat{\theta}) / \boldsymbol{\mu}_{i,t}(\check{\theta}) \right), \boldsymbol{\lambda}_i(\hat{\theta}, \check{\theta}) := \log \left( \ell_i(\mathbf{s}_i | \hat{\theta}) / \ell_i(\mathbf{s}_i | \check{\theta}) \right), \gamma_i(\hat{\theta}, \check{\theta}) := \log \left( \nu_i(\hat{\theta}) / \nu_i(\check{\theta}) \right).$$

and further define  $\hat{\phi}_{i,t}(\hat{\theta}, \check{\theta}) := \phi_{i,t}(\hat{\theta}, \check{\theta}) - \phi_{i,t-1}(\hat{\theta}, \check{\theta})$ . By concatenating the log-ratio statistics of the  $n$  networked agents, we obtain the following four vectorizations for the log-ratio statistics:  $\bar{\boldsymbol{\phi}}_t(\hat{\theta}, \check{\theta}) := (\phi_{1,t}(\hat{\theta}, \check{\theta}), \dots, \phi_{n,t}(\hat{\theta}, \check{\theta}))^T$ ,  $\bar{\boldsymbol{\phi}}_{\Delta,t} := (\hat{\phi}_{1,t}(\hat{\theta}, \check{\theta}), \dots, \hat{\phi}_{n,t}(\hat{\theta}, \check{\theta}))^T$ ,  $\bar{\boldsymbol{\lambda}}(\hat{\theta}, \check{\theta}) := (\boldsymbol{\lambda}_1(\hat{\theta}, \check{\theta}), \dots, \boldsymbol{\lambda}_n(\hat{\theta}, \check{\theta}))^T$ ,  $\bar{\boldsymbol{\gamma}}(\hat{\theta}, \check{\theta}) := (\gamma_1(\hat{\theta}, \check{\theta}), \dots, \gamma_n(\hat{\theta}, \check{\theta}))^T$ . Under (6) the log-belief ratios evolve as follows:

$$\boldsymbol{\phi}_{i,t}(\hat{\theta}, \check{\theta}) = \boldsymbol{\phi}_{i,t-1}(\hat{\theta}, \check{\theta}) + \hat{\boldsymbol{\phi}}_{i,t}(\hat{\theta}, \check{\theta}), \text{ where } \hat{\boldsymbol{\phi}}_{i,t}(\hat{\theta}, \check{\theta}) := \sum_{j \in \mathcal{N}_i \setminus \{i\}} \hat{\boldsymbol{\phi}}_{j,t-1}(\hat{\theta}, \check{\theta}), \quad (19)$$

initialized by:  $\boldsymbol{\phi}_{i,0}(\hat{\theta}, \check{\theta}) = \gamma_i(\hat{\theta}, \check{\theta}) + \boldsymbol{\lambda}_i(\hat{\theta}, \check{\theta})$  and  $\hat{\boldsymbol{\phi}}_{i,0}(\hat{\theta}, \check{\theta}) = \boldsymbol{\lambda}_i(\hat{\theta}, \check{\theta})$ , for all  $i$ . In [42, Proposition 4.3 and its following paragraph] we point out that (19) implements the rational (Bayesian) belief update if the network structure is a rooted directed tree, where there is a unique path connecting each ‘‘upstream’’ agent  $j$  to the agent  $i$ . We refer to the  $\hat{\boldsymbol{\phi}}_{i,t}(\hat{\theta}, \check{\theta})$  terms as innovations. Following (19), the total innovation in the belief of agent  $i$  at time  $t$  is set equal to the sum of innovations in the beliefs of her neighbors in the preceding time step. The uniqueness of paths in rooted directed trees ensures that innovations are not multiply counted and the updated beliefs are not subject to redundancy, as expected, for an optimal (Bayesian) belief. However if there are multiple paths between pairs of agents in the network, then the beliefs deviate from Bayesian rationality.

Using the adjacency matrix  $A$  and the vectorized notations for the log-ratio statistics, the belief dynamics under (6) (or equivalently (19)) can be analyzed as follows:

$$\begin{aligned} \bar{\boldsymbol{\phi}}_{\Delta,t}(\hat{\theta}, \check{\theta}) &= A \bar{\boldsymbol{\phi}}_{\Delta,t-1}(\hat{\theta}, \check{\theta}) = A^t \bar{\boldsymbol{\phi}}_{\Delta,0}(\hat{\theta}, \check{\theta}), \text{ and} \\ \bar{\boldsymbol{\phi}}_t(\hat{\theta}, \check{\theta}) &= \bar{\boldsymbol{\phi}}_{t-1}(\hat{\theta}, \check{\theta}) + \bar{\boldsymbol{\phi}}_{\Delta,t}(\hat{\theta}, \check{\theta}) = \bar{\boldsymbol{\phi}}_0(\hat{\theta}, \check{\theta}) + \sum_{\tau=1}^t A^\tau \bar{\boldsymbol{\phi}}_{\Delta,0}(\hat{\theta}, \check{\theta}) = \bar{\boldsymbol{\gamma}}(\hat{\theta}, \check{\theta}) + \sum_{\tau=0}^t A^\tau \bar{\boldsymbol{\lambda}}(\hat{\theta}, \check{\theta}) \end{aligned}$$

Note that  $C_t(i, j) := \sum_{\tau=0}^t [A^\tau]_{i,j}$  denotes the number of paths of length less than or equal to  $\tau$  (with possibly repeated edges and vertices) that start from node  $j$  and end at node  $i$  (using the convention  $C_0(i, j) = 1$  if  $i = j$  and  $C_0(i, j) = 0$ , otherwise) [64, Section 6.10]. Using the  $C_t(i, j)$  notation, the log-ratio of belief of agent  $i$  at time  $t$  can be expressed as:

$$\phi_{i,t}(\hat{\theta}, \check{\theta}) = \gamma_i(\hat{\theta}, \check{\theta}) + \sum_{j=1}^n C_t(i, j) \lambda_j(\hat{\theta}, \check{\theta}). \quad (20)$$

When the network structure is a directed acyclic graph (DAG), there are no paths of length greater than  $n$  connecting any two nodes. Hence,  $C_t(i, j) = C_n(i, j)$  for all  $t \geq n$ . As a consequence of (20), in such cases the asymptotic outcome will be reached in finite time (bounded by  $n$ ); i.e. agents become constant in their beliefs after  $n$  steps. This final belief is then given by:

$$\boldsymbol{\mu}_{i,t}(\theta') = \frac{\nu_i(\theta') \prod_{j=1}^n \ell_j(\mathbf{s}_j | \theta')^{C_n(i,j)}}{\sum_{\theta'' \in \Theta} \nu_i(\theta'') \prod_{j=1}^n \ell_j(\mathbf{s}_j | \theta'')^{C_n(i,j)}}, \forall t \geq n. \quad (21)$$

In the special case of directed rooted trees, agent  $i$  is connected to each of the agents whose decisions influence her beliefs (directly or indirectly), through a unique path. Without any loss of generality, we can assume that the directed tree is rooted at agent  $i$  (by possibly eliminating the agents whose decisions will not influence the beliefs of agent  $i$ ). Hence, in a directed rooted tree we have that  $C_n(i, j) = 1$  for all  $j$  and the final belief of agent  $i$  in a directed rooted tree is given by:

$$\boldsymbol{\mu}_{i,t}(\theta') = \frac{\nu_i(\theta') \prod_{j=1}^n \ell_j(\mathbf{s}_j | \theta')}{\sum_{\theta'' \in \Theta} \nu_i(\theta'') \prod_{j=1}^n \ell_j(\mathbf{s}_j | \theta'')}, \forall t \geq n. \quad (22)$$

The final belief (22) in a directed rooted tree coincides with the Bayesian posterior for agent  $i$  given all the private signals of all agents who influence her decisions (directly or indirectly). However, the final belief in general acyclic structure, deviates from the Bayesian optimum. In particular, the latter beliefs are subject to redundancy effects, as reflected in the exponential weights  $C_n(i, j)$  that appear in (21). The  $C_n(i, j)$  coefficients weigh the signal likelihoods of different agents in the final belief of agent  $i$ , according to the total number of paths that connects each of them to agent  $i$ . In presence of cycles between  $i$  and  $j$ , the total number of paths  $C_t(i, j)$  grows unbounded with increasing  $t$  since traversing the same cycles repeatedly yields an increasing number of longer paths. Hence, in loopy structures redundancy leads to overconfidence and the asymptotic beliefs concentrate on a subset of alternatives, rejecting others. We elaborate on such overconfident asymptotic beliefs in Subsection 5.1. In Subsection 5.2, we exploit the optimality of directed rooted trees for information aggregation (even with Bayesian heuristics) by proposing a coordination scheme that ensures all agents see a directed rooted information structure, following a prescribed schedule of meetings that take place in a few rounds.

**Appendix E: An algebra of beliefs.** Both the linear action updates studied in the previous chapter as well as the weighted majority update rules that arise in the binary case and are studied in [68, 67] have a familiar algebraic structure. It is instructive to develop similar structural properties for belief updates in (5) and over the space  $\Delta\Theta$ , i.e. the points of the standard  $(m-1)$ -simplex. Given two beliefs  $\mu_1(\cdot)$  and  $\mu_2(\cdot)$  over  $\Theta$  we denote their “addition” as

$$\mu_1 \oplus \mu_2(\theta') = \frac{\mu_1(\theta') \mu_2(\theta')}{\sum_{\theta'' \in \Theta} \mu_1(\theta'') \mu_2(\theta'')}.$$

Indeed, let  $\Delta\Theta^\circ$  denote the  $(m-1)$ -simplex of probability measure over  $\Theta$  after all the edges are excluded;  $\Delta\Theta^\circ$  endowed with the  $\oplus$  operation, constitutes a group (in the algebraic sense of the word). It is easy to verify that the uniform distribution  $\bar{\mu}(\theta') = 1/|\Theta|$  acts as the identity element



for the group; in the sense that  $\bar{\mu} \oplus \mu = \mu$  for all  $\mu \in \Delta\Theta^\circ$ , and given any such  $\mu$  we can uniquely identify its inverse as follows:

$$\mu^{inv}(\theta') = \frac{1/\mu(\theta')}{\sum_{\theta'' \in \Theta} 1/\mu(\theta'')}.$$

Moreover, the group operation  $\oplus$  is commutative and we can thus endow the abelian group  $(\Delta\Theta^\circ, \oplus)$  with a subtraction operation:

$$\mu_1 \ominus \mu_2(\theta') = \mu_1 \oplus \mu_2^{inv}(\theta') = \frac{\mu_1(\theta')/\mu_2(\theta')}{\sum_{\tilde{\theta} \in \Theta} \mu_1(\tilde{\theta})/\mu_2(\tilde{\theta})}.$$

We are now in a position to rewrite the Bayesian heuristic for belief updates in terms of the group operations  $\oplus$  and  $\ominus$  over the simplex interior:

$$\boldsymbol{\mu}_{i,t} = \bigoplus_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t-1} \ominus_{j \in \mathcal{N}_i \setminus \{i\}} \boldsymbol{\nu}_j. \quad (23)$$

The above belief update has a structure similar to the linear action updates studied in (1): the agents incorporate the beliefs of their neighbors while compensating for the neighboring priors to isolate the observational parts of the neighbors' reports. A key difference between the action and belief updates is in the fact that action updates studied in Section 4 are weighted in accordance with the observational ability of each neighbor, whereas the belief updates are not. Indeed, the quality of signals are already internalized in the reported beliefs of each neighbor; therefore there is no need to re-weight the reported beliefs when aggregating them.

Given the abelian group structure we can further consider the ‘‘powers’’ of each element  $\mu^2 = \mu \oplus \mu$  and so on; in general for each integer  $n$  and any belief  $\mu \in \Delta\Theta^\circ$ , let the  $n$ -th power of  $\mu$  be denoted by  $n \odot \mu := \mu^n$ , defined as follows:<sup>1</sup>

$$\mu^n(\theta') = \frac{\mu^n(\theta')}{\sum_{\tilde{\theta} \in \Theta} \mu^n(\tilde{\theta})}.$$

Using the  $\oplus$  and  $\odot$  notations, as well as the adjacency matrix  $A$  we get:

$$\boldsymbol{\mu}_{i,t+1} = \bigoplus_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t} \ominus_{j \in \mathcal{N}_i \setminus \{i\}} \boldsymbol{\nu}_j = \bigoplus_{j \in [n]} ([I + A]_{ij} \odot \boldsymbol{\mu}_{j,t}) \ominus_{j \in [n]} ([A]_{ij} \odot \boldsymbol{\nu}_j). \quad (27)$$

With some abuse of notation, we can concatenate the network beliefs at every time  $t$  into a column vector  $\boldsymbol{\mu}_t = (\boldsymbol{\mu}_{1,t}, \dots, \boldsymbol{\mu}_{n,t})^T$  and similarly for the priors  $\boldsymbol{\nu} = (\boldsymbol{\nu}_1, \dots, \boldsymbol{\nu}_n)^T$ ; thus (27) can be written in the vectorized format by using the matrix notation as follows:

$$\boldsymbol{\mu}_t = \{(I + A) \odot \boldsymbol{\mu}_{t-1}\} \ominus \{A \odot \boldsymbol{\nu}\} \quad (28)$$

Iterating over  $t$  and in the common matrix notation we obtain:

$$\boldsymbol{\mu}_t = \left\{ (I + A)^t \odot \boldsymbol{\mu}_0 \right\} \ominus \left\{ \left( \sum_{\tau=0}^{t-1} (I + A)^\tau A \right) \odot \boldsymbol{\nu} \right\}. \quad (29)$$

<sup>1</sup>This notation extends to all real numbers  $n \in \mathbb{R}$ , and it is easy to verify that the following distributive properties are satisfied:

$$\begin{aligned} n \odot (\mu_1 \oplus \mu_2) &= (n \odot \mu_1) \oplus (n \odot \mu_2), \\ (m + n) \odot \mu_1 &= (m \odot \mu_1) \oplus (n \odot \mu_1), \\ (m.n) \odot \mu_1 &= m \odot (n \odot \mu_1), \end{aligned}$$

for all  $m, n \in \mathbb{R}$  and  $\mu_1, \mu_2 \in \Delta\Theta^\circ$ .

The above is key to understanding the evolution of beliefs under the Bayesian heuristics in (5), as we will explore next. In particular, when all agents have uniform priors  $\nu_j = \bar{\mu}$  for all  $j$ , then (28) and (29) simplify as follows:  $\boldsymbol{\mu}_t = (I + A) \odot \boldsymbol{\mu}_{t-1} = (I + A)^t \odot \boldsymbol{\mu}_0$ . This assumption of a common uniform prior is the counterpart of Assumption 1 (non-informative priors) in Subsection 4.1, which paved the way for transition from affine action updates into linear ones. In the case of beliefs over a finite state space  $\Theta$ , the uniform prior  $\bar{\mu}$  is non-informative. If all agents start from common uniform priors, the belief update in (5) simplifies to:

$$\boldsymbol{\mu}_{i,t}(\theta^t) = \frac{\prod_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}(\theta^t)}{\sum_{\tilde{\theta} \in \Theta} \prod_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}(\tilde{\theta})}. \quad (30)$$

Our main focus in the Subsection 5.1 is to understand how the individual beliefs evolve under (5), or (30) which is a spacial case of (5). The gist of our analysis is encapsulated in the group theoretic iterations:  $\boldsymbol{\mu}_t = (I + A)^t \odot \boldsymbol{\mu}_0$ , derived above for the common uniform priors case. In particular, our understanding of the increasing matrix powers  $(I + A)^t$  plays a key role. When the network graph  $\mathcal{G}$  is strongly connected, the matrix  $I + A$  is primitive. The Perron-Frobenius theory [76, Theorems 1.5 and 1.7] implies that  $I + A$  has a simple positive real eigenvalue equal to its spectral radius  $\rho(I + A) = 1 + \rho$ , where we adopt the shorthand notation  $\rho := \rho(A)$ . Moreover, the left and right eigenspaces associated with this eigenvalue are both one-dimensional and the corresponding eigenvectors can be taken such that they both have strictly positive entries. The magnitude of any other eigenvalue of  $I + A$  is strictly less than  $1 + \rho$ . Hence, the eigenvalues of  $I + A$  denoted by  $\lambda_i(I + A)$ ,  $i \in [n]$ , can be ordered in their magnitudes as follows:  $|\lambda_n(I + A)| \leq |\lambda_{n-1}(I + A)| \leq \dots < \lambda_1(I + A) = 1 + \rho$ . Subsequently, we can employ the eigendecomposition of  $(I + A)$  to analyze the behavior of  $(I + A)^{t+1}$ . Specifically, we can take a set of bi-orthonormal vectors  $\bar{l}_i, \bar{r}_i$  as the left and right eigenvectors corresponding to the  $i$ th eigenvalue of  $I + A$ , satisfying:  $\|\bar{l}_i\|_2 = \|\bar{r}_i\|_2 = 1$ ,  $\bar{l}_i^T \bar{r}_i = 1$  for all  $i$  and  $\bar{l}_i^T \bar{r}_j = 0$ ,  $i \neq j$ ; in particular, the left eigenspace associated with  $\rho$  is one-dimensional with the corresponding eigenvector  $\bar{l}_1 = \bar{\alpha} = (\alpha_1, \dots, \alpha_n)^T$ , uniquely satisfying  $\sum_{i=1}^n \alpha_i = 1$ ,  $\alpha_i > 0$ ,  $\forall i \in [n]$ , and  $\bar{\alpha}^T A = (\rho + 1)\bar{\alpha}^T$ . The entry  $\alpha_i$  is called the centrality of agent  $i$  and as the name suggests, it measures how central is the location of agent in the network. We can now use the spectral representation of  $A$  to write [43, Section 6]:

$$(I + A)^t = (1 + \rho)^t \left( \bar{r}_1 \bar{\alpha}^T + \sum_{i=2}^n (\lambda_i(I + A)/(1 + \rho))^t \bar{r}_i \bar{l}_i^T \right). \quad (31)$$

Asymptotically, we get that all eigenvalues other than the Perron-Frobenius eigenvalue  $1 + \rho$  are subdominant; hence,  $(I + A)^t \rightarrow (1 + \rho)^t \bar{r}_1 \bar{\alpha}^T$  and  $\boldsymbol{\mu}_t = (1 + \rho)^t \bar{r}_1 \bar{\alpha}^T \odot \boldsymbol{\mu}_0$  as  $t \rightarrow \infty$ ; the latter holds true for the common uniform priors case and also in general, as we shall see in the proof of Theorem F.

**Appendix F: Proof of Theorem 4.** Using the log-ratio variables and their concatenations defined at the beginning of Appendix D, we can rewrite the log-linear belief updates of (5) in a linearized vector format as shown below:

$$\begin{aligned} \bar{\boldsymbol{\phi}}_{t+1}(\hat{\theta}, \check{\theta}) &= (I + A) \bar{\boldsymbol{\phi}}_t(\hat{\theta}, \check{\theta}) - A \bar{\boldsymbol{\gamma}}(\hat{\theta}, \check{\theta}) = (I + A)^{t+1} \bar{\boldsymbol{\phi}}_0(\hat{\theta}, \check{\theta}) - \sum_{\tau=0}^t (I + A)^\tau A \bar{\boldsymbol{\gamma}}(\hat{\theta}, \check{\theta}) \\ &= (I + A)^{t+1} \left( \bar{\boldsymbol{\lambda}}(\hat{\theta}, \check{\theta}) + \bar{\boldsymbol{\gamma}}(\hat{\theta}, \check{\theta}) \right) - \sum_{\tau=0}^t (I + A)^\tau A \bar{\boldsymbol{\gamma}}(\hat{\theta}, \check{\theta}) \\ &= (I + A)^{t+1} \bar{\boldsymbol{\lambda}}(\hat{\theta}, \check{\theta}) + \left( (I + A)^{t+1} - \sum_{\tau=0}^t (I + A)^\tau A \right) \bar{\boldsymbol{\gamma}}(\hat{\theta}, \check{\theta}). \end{aligned}$$

Next we use the spectral decomposition in (31) to obtain: <sup>1</sup>

$$\begin{aligned} \bar{\phi}_{t+1}(\hat{\theta}, \check{\theta}) &= (1 + \rho)^{t+1} \bar{r}_1 \mathbf{\Lambda}(\hat{\theta}, \check{\theta}) + ((1 + \rho)^{t+1} - \sum_{\tau=0}^t (1 + \rho)^\tau \rho) \bar{r}_1 \beta(\hat{\theta}, \check{\theta}) + o((1 + \rho)^{t+1}) \\ &= (1 + \rho)^{t+1} \left( \bar{r}_1 \mathbf{\Lambda}(\hat{\theta}, \check{\theta}) + (1 - \sum_{\tau=0}^t (1 + \rho)^{\tau-t-1} \rho) \bar{r}_1 \beta(\hat{\theta}, \check{\theta}) + o(1) \right) \rightarrow (1 + \rho)^{t+1} \bar{r}_1 \mathbf{\Lambda}(\hat{\theta}, \check{\theta}), \end{aligned} \quad (32)$$

where we adopt the following notations for the global log likelihood and prior ratio statistics:  $\beta(\hat{\theta}, \check{\theta}) := \bar{\alpha}^T \bar{\gamma}(\hat{\theta}, \check{\theta})$  and  $\mathbf{\Lambda}(\hat{\theta}, \check{\theta}) := \bar{\alpha}^T \bar{\mathbf{\Lambda}}(\hat{\theta}, \check{\theta})$ ; furthermore, in calculation of the limit in the last step of (32) we use the geometric summation identity  $\sum_{\tau=0}^{\infty} \rho(1 + \rho)^{\tau-1} = 1$ .

To proceed denote  $\mathbf{\Lambda}(\hat{\theta}) := \sum_{i=1}^n \alpha_i \ell_i(\mathbf{s}_i | \hat{\theta})$  so that  $\mathbf{\Lambda}(\hat{\theta}, \check{\theta}) = \mathbf{\Lambda}(\hat{\theta}) - \mathbf{\Lambda}(\check{\theta})$ . Since  $\Theta^\diamond$  consists of the set of all maximizers of  $\mathbf{\Lambda}(\hat{\theta})$ , we have that  $\mathbf{\Lambda}(\hat{\theta}, \check{\theta}) < 0$  whenever  $\check{\theta} \in \Theta^\diamond$  and  $\hat{\theta} \notin \Theta^\diamond$ . Next recall from (32) that  $\bar{\phi}_{t+1}(\hat{\theta}, \check{\theta}) \rightarrow (1 + \rho)^{t+1} \bar{r}_1 \mathbf{\Lambda}(\hat{\theta}, \check{\theta})$  and  $\bar{r}_1$  is the right Perron-Frobenius eigenvector with all positive entries; hence, for all  $\check{\theta} \in \Theta^\diamond$  and any  $\hat{\theta}$ ,  $\phi_{i,t}(\hat{\theta}, \check{\theta}) \rightarrow -\infty$  if  $\hat{\theta} \notin \Theta^\diamond$  and  $\phi_{i,t}(\hat{\theta}, \check{\theta}) = 0$  whenever  $\hat{\theta} \in \Theta^\diamond$ ; or equivalently,  $\mu_{i,t}(\hat{\theta})/\mu_{i,t}(\check{\theta}) \rightarrow 0$  for all  $\hat{\theta} \notin \Theta^\diamond$ , while  $\lim_{t \rightarrow \infty} \mu_{i,t}(\hat{\theta}) = \lim_{t \rightarrow \infty} \mu_{i,t}(\check{\theta})$  for any  $\hat{\theta} \in \Theta^\diamond$ . The latter together with the fact that  $\sum_{\theta \in \Theta} \mu_{i,t}(\theta) = 1$  for all  $t$  implies that with probability one:  $\lim_{t \rightarrow \infty} \mu_{i,t}(\check{\theta}) = 1/|\Theta^\diamond|, \forall \check{\theta} \in \Theta^\diamond$  and  $\lim_{t \rightarrow \infty} \mu_{i,t}(\tilde{\theta}) = 0, \forall \tilde{\theta} \notin \Theta^\diamond$  as claimed in the Theorem. In the special case that  $\Theta^\diamond$  is a singleton,  $\{\theta^\diamond\} = \Theta^\diamond$ , we get that  $\lim_{t \rightarrow \infty} \mu_{i,t}(\theta^\diamond) = 1$  almost surely for all  $i \in [n]$ .

**Acknowledgments.** The authors would like to thank Peter Kraftt, Weiwen Leung and Armand Makowski for insightful discussions and the participants at the 2016 NBER-NSF Seminar on Bayesian Inference in Econometrics and Statistics, and the University of Pennsylvania, University of Maryland, Princeton and MIT local seminars for their comments and feedback.

This work is supported by ARO MURI W911NF-12-1-0509.

## References

- [1] Acemoglu D, Dahleh MA, Lobel I, Ozdaglar A (2011) Bayesian learning in social networks. *The Review of Economic Studies* 78(4):1201–1236.
- [2] Aldous D, Fill J (2002) Reversible markov chains and random walks on graphs Book in progress.
- [3] Anderson NH (1971) Integration theory and attitude change. *Psychological review* 78(3):171.
- [4] Anderson NH (1981) *Foundations of information integration theory* (Academic Press).
- [5] Aumann RJ (1976) Agreeing to disagree. *The annals of statistics* 1236–1239.
- [6] Bala V, Goyal S (1998) Learning from neighbours. *The Review of Economic Studies* 65(3):595–621.
- [7] Bickel PJ, Doksum KA (2015) *Mathematical Statistics: Basic Ideas and Selected Topics, volume I* (CRC Press).
- [8] Bray NLK Robert M, Atkin RS (1978) Effects of group size, problem difficulty, and sex on group performance and member reactions. *Journal of Personality and Social Psychology* 36(11):1224.
- [9] Brightwell GR, Winkler P (1999) Graph homomorphisms and phase transitions. *Journal of Combinatorial Theory, Series B* 77(2):221–262.
- [10] Chandrasekhar AG, Larreguy H, Xandri JP (2015) Testing models of social learning on networks: Evidence from a lab experiment in the field. Technical report, National Bureau of Economic Research.

<sup>1</sup>Given two functions  $f(\cdot)$  and  $g(\cdot)$  we use the asymptotic notation  $f(t) = o(g(t))$  to signify the relations  $\lim_{t \rightarrow \infty} |f(t)/g(t)| = 0$ .

- 
- [11] Cural LA, Forrester RH, Dawson JF, West MA (2001) It's what you do and the way that you do it: Team task, team size, and innovation-related group processes. *European Journal of Work and Organizational Psychology* 10(2):187–204.
- [12] DeGroot MH (1974) Reaching a consensus. *Journal of American Statistical Association* 69:118 – 121.
- [13] DeMarzo PM, Vayanos D, Zwiebel J (2003) Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics* 118:909–968.
- [14] Desanctis G, Gallupe RB (1987) A foundation for the study of group decision support systems. *Management Science* 33(5):589–609.
- [15] Diaconis P, Ylvisaker D (1979) Conjugate priors for exponential families. *The Annals of statistics* 7(2):269–281.
- [16] Eliaz K, Ray D, Razin R (2006) Choice shifts in groups: A decision-theoretic basis. *The American economic review* 1321–1332.
- [17] Evans JSB (2003) In two minds: dual-process accounts of reasoning. *Trends in cognitive sciences* 7(10):454–459.
- [18] Evans JSB, Stanovich KE (2013) Dual-process theories of higher cognition advancing the debate. *Perspectives on psychological science* 8(3):223–241.
- [19] Eyster E, Rabin M (2010) Naive herding in rich-information settings. *American economic journal: microeconomics* 2(4):221–243.
- [20] Eyster E, Rabin M (2014) Extensive imitation is irrational and harmful. *The Quarterly Journal of Economics* qju021.
- [21] Farina L, Rinaldi S (2011) *Positive linear systems: theory and applications*, volume 50 (John Wiley & Sons).
- [22] Frank F, Anderson LR (1971) Effects of task and group size upon group productivity and member satisfaction. *Sociometry* 34(1):135–149.
- [23] French Jr JR (1956) A formal theory of social power. *Psychological review* 63(3):181–194.
- [24] Friedkin NE (1986) A formal theory of social power. *Journal of Mathematical Sociology* 12(2):103–126.
- [25] Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB (2014) *Bayesian data analysis*. Texts in Statistical Science Series (CRC Press, Boca Raton, FL), third edition.
- [26] Gigerenzer G, Gaissmaier W (2011) Heuristic decision making. *Annual review of psychology* 62:451–482.
- [27] Gigerenzer G, Goldstein DG (1996) Reasoning the fast and frugal way: models of bounded rationality. *Psychological review* 103(4):650.
- [28] Gigerenzer G, Todd PM (1999) *Simple heuristics that make us smart* (Oxford University Press, USA).
- [29] Gilardoni GL, Clayton MK (1993) On reaching a consensus using DeGroot's iterative pooling. *The Annals of Statistics* 391–401.
- [30] Glazer J, Rubinstein A (1998) Motives and implementation: On the design of mechanisms to elicit opinions. *Journal of Economic Theory* 79(2):157–173.
- [31] Golub B, Jackson MO (2010) Naïve Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics* 2(1):112–149.
- [32] Gooding RZ, Wagner JA (1985) A meta-analytic review of the relationship between size and performance: The productivity and efficiency of organizations and their subunits. *Administrative Science Quarterly* 30(4):462–481.
- [33] Grimm V, Mengel F (2014) An experiment on belief formation in networks. *Available at SSRN 2361007*.
- [34] Guzzo RA, Shea GP (1992) Group performance and intergroup relations in organizations. *Handbook of industrial and organizational*. Palo Alto, CA, US: Consulting Psychologists Press 269–313.
- [35] Hackman JR, Vidmar N (1970) Effects of size and task type on group performance and member reactions. *Sociometry* 37–54.

- 
- [36] Harary F (1959) A criterion for unanimity in french's theory of social power. *Studies in Social Power*, 168–182 (University of Michigan), d. Cartwright, Ed.
- [37] Hegselmann R, Krause U (2005) Opinion dynamics driven by various ways of averaging. *Computational Economics* 25(4):381–405.
- [38] Hoegl M (2005) Smaller teams-better teamwork: How to keep project teams small. *Business Horizons* 48(3):209–214.
- [39] Isenberg DJ (1986) Group polarization: a critical review and meta-analysis. *Journal of personality and social psychology* 50(6):1141.
- [40] Jadbabaie A, Molavi P, Sandroni A, Tahbaz-Salehi A (2012) Non-bayesian social learning. *Games and Economic Behavior* 76(1):210 – 225.
- [41] Jadbabaie A, Molavi P, Tahbaz-Salehi A (2013) Information heterogeneity and the speed of learning in social networks. *Revise and Resubmit, Review of Economic Studies* .
- [42] Jadbabaie A, Mossel E, Rahimian MA (2017) Bayesian group decisions: algorithms and complexity. *arXiv preprint* .
- [43] James G, Rumchev V (2005) Stability of positive linear discrete-time systems. *Bulletin of the Polish Academy of Sciences. Technical Sciences* 53(1):1–8.
- [44] Janis IL (1982) *Groupthink: Psychological studies of policy decisions and fiascoes* (Houghton Mifflin Boston).
- [45] JR H (1990) *Groups that work (and those that don't)* (San Francisco: Jossey-Bass).
- [46] Kahneman D (2003) Maps of bounded rationality: Psychology for behavioral economics. *The American economic review* 93(5):1449–1475.
- [47] Kahneman D (2011) *Thinking, fast and slow* (Farrar Straus Giroux).
- [48] Kanoria Y, Tamuz O (2013) Tractable bayesian social learning on trees. *Selected Areas in Communications, IEEE Journal on* 31(4):756–765.
- [49] Kass RE, Wasserman L (1996) The selection of prior distributions by formal rules. *Journal of the American Statistical Association* 91(435):1343–1370.
- [50] Klein G (2007) Performing a project premortem. *Harvard Business Review* 85(9):18–19.
- [51] Laslier JF, Weibull J (2008) Committee decisions: optimality and equilibrium Cahier de recherche 2008-24. v3 du halshs-00121741.
- [52] Levy G, de Barra IM, Razin R (2017) Persuasion with correlation neglect: Media bias and media power through correlation of news content. *preprint* .
- [53] Levy G, Razin R (2015) Correlation neglect, voting behavior, and information aggregation. *The American Economic Review* 105(4):1634–1645.
- [54] Levy G, Razin R (2016) Information diffusion in networks with the bayesian peer influence heuristic. *preprint* .
- [55] Lightle JP, Kagel JH, Arkes HR (2009) Information exchange in group decision making: The hidden profile problem reconsidered. *Management Science* 55(4):568–581.
- [56] Littlepage GE (1991) Effects of group size and task characteristics on group performance: A test of steiner's model. *Personality and Social Psychology Bulletin* 17(4):449–456.
- [57] Lyons R, Peres Y (2016) *Probability on Trees and Networks* (Cambridge Series in Statistical and Probabilistic Mathematics).
- [58] Mao A, Mason W, Suri S, Watts DJ (2016) An experimental study of team size and performance on a complex task. *PloS one* 11(4):e0153048.
- [59] Masicampo EJ, Baumeister RF (2008) Toward a physiology of dual-process reasoning and judgment: Lemonade, willpower, and expensive rule-based analysis. *Psychological Science* 19(3):255–260.
- [60] Molavi P, Tahbaz-Salehi A, Jadbabaie A (2016) Foundations of non-bayesian social learning. *Columbia Business School Research Paper* .

- [61] Mossel E, Tamuz O (2010) Efficient bayesian learning in social networks with gaussian estimators. *arXiv preprint arXiv:1002.0747* .
- [62] Mueller-Frank M, Neri C (2015) A general model of boundedly rational observational learning: Theory and experiment. *Available at SSRN 2566210* .
- [63] Mueller-Frank M, Neri C (2017) Quasi-bayesian updating in social networks. *Available at SSRN 2933411* .
- [64] Newman M (2010) *Networks: An Introduction* (Oxford University Press).
- [65] Rahimian MA, Jadbabaie A (2015) Learning without recall: A case for log-linear learning. *5th IFAC Workshop on Distributed Estimation and Control in Networked Systems* .
- [66] Rahimian MA, Jadbabaie A (2015) Learning without recall in directed circles and rooted trees. *American Control Conference* 4222–4227.
- [67] Rahimian MA, Jadbabaie A (2016) Bayesian learning without recall. *IEEE Transactions on Signal and Information Processing over Networks* .
- [68] Rahimian MA, Jadbabaie A (2016) Naive social learning in ising networks. *American Control Conference* .
- [69] Rahimian MA, Molavi P, Jadbabaie A (2014) (Non-) bayesian learning without recall. *IEEE Conference on Decision and Control (CDC)* 5730–5735.
- [70] Rahimian MA, Shahrampour S, Jadbabaie A (2015) Learning without recall by random walks on directed graphs. *IEEE Conference on Decision and Control (CDC)* .
- [71] Redish AD (2013) *The mind within the brain: How we make decisions and how those decisions go wrong* (Oxford University Press).
- [72] Robert C (2001) *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Springer Texts in Statistics (Springer).
- [73] Roux N, Sobel J (2015) Group polarization in a model of information aggregation. *American Economic Journal: Microeconomics* 7(4):202–32.
- [74] Rufo MJ, Martin J, Pérez CJ, et al. (2012) Log-linear pool to combine prior distributions: A suggestion for a calibration-based approach. *Bayesian Analysis* 7(2):411–438.
- [75] School W (2006) Is your team too big? too small? whats the right number? *Knowledge@Wharton* .
- [76] Seneta E (2006) *Non-negative matrices and Markov chains* (Springer).
- [77] Simon HA (1955) A behavioral model of rational choice. *The Quarterly Journal of Economics* 99–118.
- [78] Simon HA (1956) Rational choice and the structure of the environment. *Psychological review* 63(2):129.
- [79] Simon HA (1979) Rational decision making in business organizations. *The American economic review* 69(4):493–513.
- [80] Slater PE (1958) Contrasting correlates of group size. *Sociometry* 21(2):129–139.
- [81] Sobel J (2014) On the relationship between individual and group decisions. *Theoretical Economics* 9(1):163–185.
- [82] Steiner ID (1972) *Group Processes and Productivity, Social psychology: a series of monographs, treatises, and texts* (Academic Press).
- [83] Stoner JA (1968) Risky and cautious shifts in group decisions: The influence of widely held values. *Journal of Experimental Social Psychology* 4(4):442–459.
- [84] Sunstein CR (2000) Deliberative trouble? why groups go to extremes. *The Yale Law Journal* 110(1):71–119.
- [85] Thaler RH (2015) *Misbehaving: The making of behavioral economics* (WW Norton & Company).
- [86] Thomas CF Edwin J; Fink (1963) Effects of group size. *Psychological Bulletin* 60(4):371–384.
- [87] Tversky A, Kahneman D (1975) Judgment under uncertainty: Heuristics and biases. *Utility, probability, and human decision making*, 141–162 (Springer).