

LEARNING AND DECISION MAKING IN GROUPS

M. Amin Rahimian

A DISSERTATION

in

Electrical and Systems Engineering

Presented to the Faculties of the University of Pennsylvania
in Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy

2017

Supervisor of Dissertation

Ali Jadbabaie, JR East Professor of Engineering, MIT

Graduate Group Chairperson

Alejandro Ribeiro, Associate Professor of Electrical and Systems Engineering

Dissertation Committee:

Rakesh Vohra, Professor of Economics & Electrical and Systems Engineering

Victor M. Preciado, Assistant Professor of Electrical and Systems Engineering

Alvaro Sandroni, Professor of Managerial Economics & Decision Sciences

LEARNING AND DECISION MAKING IN GROUPS

© COPYRIGHT

2017

M. Amin Rahimian

This work is licensed under the
Creative Commons Attribution
NonCommercial-ShareAlike 3.0
License

To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nc-sa/3.0/>

*To my mother,
for her love and sacrifice*

Acknowledgments

First and foremost, I would like to thank Ali Jadbabaie, for his support and guidance throughout my PhD studies. This past five years have been a very formative and constructive part of my life, both professionally and personally. This, to a large extent, is due to him. He was always ready to answer my questions, and extremely patient in explaining his reasons and helping me understand them.

I would also like to thank Victor Preciado. I interacted a lot with Victor during my PhD studies. Our discussions were often quite technical and I learned so much from him; not only his techniques, but also his organization skills and way of thinking.

I had the good pleasure of assisting Santosh Venkatesh in several probability courses that he was offering at Penn. We had many meetings during these years and we often ended up discussing many topics besides, but often somehow related to, probability, pedagogy, and problem-solving. Santosh was very generous in sharing his personal experiences and thoughts. I really miss our conversations.

In my fourth and fifth years, I was extremely fortunate to take two courses with Elchanan Mossel and to later collaborate with him on the results in Chapter two of my thesis (Bayesian learning). Elchanan, together with Dean Eckles, co-advises my postdoctoral research at MIT Institute for Data, Systems and Society, and I am really looking forward to our continued interactions.

I want to especially thank Rakesh Vohra and Alvaro Sandroni for being on my thesis committee. They took their time to discuss my research results with me, to provide their feedback, and eventually to evaluate my work.

Penn has a very collegial and friendly environment that fosters research and collaboration. It has been made possible through the efforts of its leadership and administration. I would like to particularly thank George Pappas, Daniel Lee, Lilian Wu, Charity Payne and other members of the leadership and administration at ESE, GRASP, and Wharton Statistics for their effective efforts in creating such a wonderful environment. It is perhaps also a good place to thank my colleagues, classmates, and friends. I am very grateful for the pleasure of their company throughout these years.

Some other individual scientific contributions are acknowledged in the footnotes at the beginning of each chapter.

The financial support through ARO grant MURI W911NF-12-1-0509 is gratefully acknowledged.

ABSTRACT
LEARNING AND DECISION MAKING IN GROUPS

M. Amin Rahimian
Ali Jadbabaie

Many important real-world decision-making problems involve group interactions among individuals with purely informational interactions. Such situations arise for example in jury deliberations, expert committees, medical diagnoses, etc. We model the purely informational interactions of group members, where they receive private information and act based on that information while also observing other people's beliefs or actions.

In the first part of the thesis, we address the computations that a rational (Bayesian) decision-maker should undertake to realize her optimal actions, maximizing her expected utility given all available information at every decision epoch. We use an approach called iterated eliminations of infeasible signals (IEIS) to model the thinking process as well as the calculations of a Bayesian agent in a group decision scenario. Accordingly, as the Bayesian agent attempts to infer the true state of the world from her sequence of observations, she recursively refines her belief about the signals that other players could have observed and beliefs that they would have hold given the assumption that other players are also rational. We show that IEIS algorithm runs in exponential time; however, when the group structure is a partially ordered set the Bayesian calculations simplify and polynomial-time computation of the Bayesian recommendations is possible. We also analyze the computational complexity of the Bayesian belief formation in groups and show that it is NP-hard. We investigate the factors underlying this computational complexity and show how belief calculations simplify in special network structures or cases with strong inherent symmetries. We finally give insights about the statistical efficiency (optimality) of the beliefs and its relations to computational efficiency.

In the second part, we propose the “no-recall” model of inference for heuristic decision-making that is rooted in the Bayes rule but avoids the complexities of rational inference in group interactions. Accordingly to this model, the group members behave rationally at the initiation of their interactions with each other; however, in the ensuing decision epochs, they rely on heuristics that replicate their experiences from the first stage and can be justified as optimal responses to simplified versions of their complex environments. We study the implications of the information structure, together with the properties of the probability distributions, which determine the structure of the so-called “Bayesian heuristics” that the agents follow in this model. We also 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. Finally, we extend this model to a case where agents are exposed to a stream of private data in addition to observing each other's actions and analyze properties of learning and convergence under the no-recall framework.

Contents

1	Introduction & Overview	1
1.1	Background & Motivation	2
1.1.1	Rational Decision-Making	4
1.1.2	Heuristic Decision-Making	5
1.2	Contributions & Organization	6
1.2.1	The Group Decision Setup	7
1.2.2	Algorithms and Complexity of Rational Group Decisions	8
1.2.3	Heuristics and Biases in Group Decision Making	12
1.2.4	Learning from Streaming Data	13
2	Rational Decision Making in Groups	16
2.1	The Bayesian Model	17
2.1.1	Iterated Elimination of Infeasible Signals (IEIS)	22
2.1.2	IEIS over POSETs	26
2.2	The Case of Revealed Beliefs	29
2.2.1	Belief Calculations in Transparent Structures	35

2.2.2	Belief Calculations in Nontransparent Structures	38
2.2.3	Hardness of GROUP-DECISION	40
3	Heuristic Decision Making in Groups	45
3.1	The No-Recall Model of Group Decision Making	46
3.2	Affine Action Updates, Linear Updating, and DeGroot Learning	49
3.2.1	Linear Updating and Convergence	52
3.2.2	DeGroot Updates, Consensus, and Efficiency	54
3.3	Log-Linear Belief Updates	57
3.3.1	An Algebra of Beliefs	59
3.3.2	Becoming Certain about the Group Aggregate	62
3.3.3	Teaming for Efficient Belief Exchange subject	63
4	Learning from Stream of Observations	68
4.1	The No-Recall Model Applied to Streaming Data	69
4.2	Weighted Majority and Threshold Rules	77
4.2.1	Analysis of Convergence and Learning in Ising Networks	79
4.2.2	Equilibrium, Consensus, and (Mis-)Learning	80
4.3	Log-Linear Learning	80
4.3.1	Analysis of Convergence and Log-Linear Learning	81
4.3.2	Learning in Circles and General Connected Topologies	84

4.3.3	Learning by Random Walks on Directed Graphs	87
4.3.4	Log-Linear Learning with Time-Varying Priors	90
5	Conclusions & Future Directions	97
5.1	Summary & Discussion of Part one	97
5.2	Summary & Discussion of Part Two	99
5.3	Directions for Future Work	102
A	Proofs & Mathematical Derivations	104
A.1	Complexity of Bayesian Decisions Using (A1: BAYES-GROUP)	104
A.2	Computational Complexity of (A2: BAYES-POSET)	105
A.3	Proof of Theorem 2.7: A SUBSET-SUM Reduction	106
A.4	Proof of Theorem 2.7: An EXACT-COVER Reduction	108
A.5	Belief Calculations in Bounded Neighborhoods with i.i.d. Signals	111
A.6	Proof of Theorem 3.3	113
A.7	Proof of Theorem 3.6	115
A.8	Proof of Theorem 3.7	115
A.9	Proof of Theorem 3.11	117
A.10	Proof of Lemma 4.1: Time-One Bayesian Actions	119
A.11	A Markov Chain on the Boolean Cube	121
A.12	Proof of Proposition 4.2: Equilibrium Action Profiles	122

A.13 Proof of Lemma 4.2: Time-One Bayesian Beliefs	122
A.14 Proof of Theorem 4.4: No Learning when $\rho > 1$	123
A.15 Proof of Theorem 4.7: Learning by Random Walks on Directed Graphs . . .	125
A.16 Analysis of Convergence when Agents Recall their Self Beliefs	129
B Rational Equilibrium and Consensus in Symmetric Groups	132
C Bayesian Calculations for Forward Reasoning	134
D Efficient Rational Choice in Symmetric Binary Environments	139
E Minimum Variance Unbiased Estimation and Online Learning	143
F Generalization to Belief Updates with One-Step Recall	149
G Distributed Estimation and Learning from Intermittent Data	152
Bibliography	157

List of Tables

2.1	List of the variables that play a role in the Bayesian calculations for group decision-making (BAYES-GROUP).	24
2.2	List of the variables for Bayesian calculations in POSET groups (BAYES-POSET).	27

List of Figures

- 1.1 On the left, “The Jury”, by John Morgan(1861) - public domain; on the right, “Wikimedia advisory board meeting, Taipei, 2007” - photo credit: Chih Hao, Taiwan, CC BY-SA 2.0 Creative Commons license. Many professional societies in engineering, science, and medicine rely on technical committees to deliberate on different issues and make important policy decisions. 4

- 1.2 POSET structures: on the left, in a POSET structure any neighbors of a neighbor of an agent are also among her immediate (directly observable) neighbors; in the middle, a round-table discussions; on the right, a vested agent (in blue) investigates the sources of her information in one-on-one meetings; imposing a POSET structure on her neighborhood. 10

- 1.3 On the left, bipartisanship is an example of a binary state space. On the right, people reveal their beliefs through status updates and what they share and post on various social media platforms. 14

- 2.1 The Decision Flow Diagram for Two Bayesian Agents 20

2.2	A structure with bounded growth: each agent is influenced by no more than three other agents even as the network (organization) size grows to infinity.	26
2.3	On the left, a directed path of length four. On the right, a directed graph is acyclic if and only if it has a topological ordering; a topological ordering of a DAG orders its vertices such that every edge goes from a lesser node (to the left) to a higher one (to the right).	29
2.4	The last three structures are transparent but the first one is not.	35
2.5	The graph structure on the left is used for the SUBSET-SUM reduction. The graph structure on the right is used for the EXACT-COVER reduction. In the particular instance of EXACT-COVER that is depicted on the right, we have that $j_1 \in l_1$ and $j_1 \in l_m$, as the links in the top two layers indicate the inclusion relations among subsets l_1, \dots, l_m and elements j_1, \dots, j_n .	42
3.1	Heuristic agents ignore the history of interactions (dotted arrows) in their inferences and always attribute actions to (imaginary) private signals; compare with the decision flow diagram for the inferences of two rational agents who communicate their actions repeatedly (Fig. 2.1).	48
3.2	Implementation of (A5) 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.	67

4.1	Evolution of the first agent's beliefs over time	75
4.2	The difference in agents' beliefs over time	76
4.3	A hybrid structure	86
4.4	Evolution of the third agent's beliefs over time	87
4.5	The difference between the first and eighth agents' beliefs over time	88
4.6	Network Structure for Example 1	90
4.7	Evolution of the second agents beliefs over time	91
4.8	The difference between the third and eighth agents beliefs over time	92
C.1	Learning about the private signals of neighbors, neighbors of neighbors, and so on.	135
D.1	On the left, the Bayesian computations of agent i in a rooted ordered tree reduces to those in the unique path connecting her to the root (the leftmost node); On the right, an oriented (singly connected) tree with depth three. . .	142

Chapter 1

Introduction & Overview

In this chapter, we give an overview of our contributions to the group decision making and learning literature, while highlighting the background of the study and our motivations for it. Individuals often exchange opinions with their peers in order to learn from their knowledge and experiences, and in making various decisions such as investing in stock markets, voting in elections, choosing their political affiliations, selecting a brand of a product or a medical treatment. These interactions occur in groups and through a variety of media which we collectively refer to as social networks. In this thesis, we examine the Bayesian and non-Bayesian models of decision making in groups or social networks by analyzing their computational properties and exploring their common behavioral foundations.¹

¹I would like to thank Mohsen Jafari Songhori for pointing me to the relevant literature in organization science, and Weiwen Leung for discussions about heuristic decision making and its relations to dual process theory.

In Section 1.1, we provide the background of our work and present our motivations; in particular, we explain the two major theoretical approaches: Bayesian (rational) and non-Bayesian (heuristics) for modeling group decision making. In Section 1.2, we present a summary of our contributions and explain how they are organized in the three chapters that follow.

1.1 Background & Motivation

Social learning or learning from actions of others is an important area of research in microeconomics. Author James Surowiecki in his popular science book on wisdom of crowds [1], provides well-known cases for information aggregation in social networks, and argues how under the right circumstances (diversity of opinion, independence, decentralization and aggregation) groups outperform even their smartest or best informed members. One historical example provided in [1] is the essentially perfect performance of the middlemost estimate at the weight-judging competition of the 1906, West of England Fat Stock and Poultry Exhibition studied by Francis Galton in his 1907, Nature article [2], entitled “Vox Populi” (The Wisdom of Crowds). Another historical case mentioned by James Surowiecki is the market’s reaction to the 1986 challenger disaster in [3], where it is pointed out that the main responsible company’s (Morton Thiokol) stock was hit hardest of all, even months before the cause of the accident could be officially established.

On the other hand, several studies point out that the evolution of people’s opinions and decisions in groups are subject to various kind of biases and inefficiencies [4–8]. 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 [9, 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.¹ 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 [11]. 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.

In decision theory, the seminal work of Aumann [12] studies the interactions of two

¹Gar 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 [10].

rational agents with common prior beliefs and concludes that if the values of their posterior beliefs are common knowledge between the two agents, then the two values should be the same: rational agents cannot agree to disagree. The later work of Geanakoplos and Polemarchakis [13] investigates how rational agents reach an agreement by communicating back and forth and refining their information partitions. Following [12] and [13], a large body of literature studies the strategic interaction of agents in a social network, where they receive private information and act based on that information while also observing each other's actions [14, 15]. These observations are in turn informative about other agents' private signals; information that can be then used in making future decisions. In this line of work, it is important to understand the effectiveness of information sharing/exchange through observed actions and the effectiveness of decision-making using the available information; indeed, the quality of decision-making depends on the quality of information exchange and vice versa.

In organization science, the economic theory of teams has a rich history devoted to choosing optimal information instruments subject to limited and dispersed resources in organizations [16]. Some of the main issues that arise in the study of decision-making organizations are information aggregation [17] and architecture [18].¹ The author in [21] compares the performance of hierarchical and polyarchical organization structures in a project selection task, where each agent possesses a private signal about the quality of the projects and acts rationally (maximizing the expected pay-off from subject to her information). Limiting attention to two decision-makers, the author shows how each agent's decision reflects the organizational structure while accounting for the rationality of the other actor. Algorithmic and complexity aspects of organizational decision-making are relatively unexplored. The author in [22] uses the formalism of constraint satisfaction problems to model the product development process in organizations. The author is thus able to identify some algorithmic and structural features that help reduce backtracking and rework costs of the design process in the organization. In this thesis, we provide new algorithmic and computational results about decision-making organizations. Addition of new results in this domain can further facilitate scalable and efficient cooperation among colleagues in large organizations (cf. Remark 2.2 and Subsection 3.3.3).

Throughout this thesis, we model the purely informational interactions of rational agents in a group, where they make private observations and act based upon that information while also observing other people's recommendations repeatedly; such lack of strategic externalities in group interactions arise since people are interested in each other's action, only to learn what others know which they do not know, for example, in jury deliberations, expert committees, medical diagnosis, etc. (cf. Fig. 1.1).

¹The organizational economics literature devotes considerable attention to incentive issues and agency problems that arise in organizational decision-making [19]; however, issues relating to distributed information processing and communication are less explored [20].



Figure 1.1: On the left, “The Jury”, by John Morgan(1861) - public domain; on the right, “Wikimedia advisory board meeting, Taipei, 2007” - photo credit: Chih Hao, Taiwan, CC BY-SA 2.0 Creative Commons license. Many professional societies in engineering, science, and medicine rely on technical committees to deliberate on different issues and make important policy decisions.

1.1.1 Rational Decision-Making

Some of the early results addressing the problem of social learning are due to Banerjee in [23], and Bikhchandani, Hirshleifer, and Welch in [24] where the authors consider a complete graph structure where the agent’s observations are public information and also ordered in time, such that each agent has access to the observations of all the past agents. These assumptions help analyze and explain the interplay between public and private information leading to fashion, fads, herding, etc. Later results by [14] relax some of these assumptions by considering the agents that make simultaneous observations of only their neighbors rather than the whole network, but the computational complexities limit the analysis to networks with only two or three agents. In more recent results, [25] provides a framework of rational learning that is analytically amenable and applies to general choice and information structures.

It is worth noting here that in this line of work and in ours as well, it is considered infeasible for the agents to share their private signals. Formally, sharing the signals may come at a prohibitive cost to the agents. Indeed, private signals might belong to different spaces and might not be interpretable by different agents. For example experts may observe signals that pertain to their domains of expertise and these signals may not be readable to people outside of their domains. In such a case, obtaining the required expertise for communicating signals come at a very high cost to non-experts. Moreover, signal spaces that represent personal experiences or individual perceptions might be far richer than action or belief spaces through which the agents communicate.

In general, when a rational agent observes her neighbors in a network, she should compensate for redundancies in information: the same neighbors’ actions are repeatedly observed and neighboring actions may be affected by the past actions of the agent herself. Hence major challenges of Bayesian inference for social learning are due to the private

signals and third party interactions that are hidden from the agent. Moreover, the existence of loops in the network causes dependencies and correlations in the information received from different neighbors, which further complicates the inference task. Failure to account for such structural dependencies subjects the agents to mistakes and inefficiencies such as redundancy neglect [26] (by neglecting the fact that several of the neighboring agents may have been influenced by the same source of information), and data incest [27] (by neglecting the fact that neighboring actions may have been affected by the past actions of the agent herself).

The agents form beliefs about unknown parameters of interest and their actions and decisions reflect their beliefs. The rational approach advocates formation of Bayesian posterior beliefs by application of Bayes rule to the entire sequence of observations successively at every step. However, such repeated applications of Bayes rule in networks become very complex, especially if the agents are unaware of the global network structure, and have to use their local data to make inferences about global contingencies that can lead to their observations. While rational learning continues to receive quite a significant amount of attention [28], it has also been criticized in the literature due to its unrealistic computational and cognitive demand on the agents [29]. Indeed, one of our goals in the first part of the thesis is to formalize this complexity notion and investigate the structural conditions around it. 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.

1.1.2 Heuristic Decision-Making

Heuristics are used widely in the literature to model social interactions and decision making [30–32]. They provide tractable tools to analyze boundedly rational behavior and offer insights about decision making under uncertainty. A dual process theory for the psychology of mind and its operation identifies two systems of thinking [33]: 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).¹ Major advances in behavioral economics are due to incorporation of this dual process theory and the subsequent models

¹While many decision science applications focus on developing dual process theories of cognition and decision making (cf. [34, 35] 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 [36, Chapter 6].

of bounded rationality [37]. 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 [38], and it also motivates our work in the second part of the thesis.

Hegselmann and Krause [39] 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 [40] and Chandrasekhar, Larreguy and Xandri [41]. The study of such heuristics took off with the seminal work of DeGroot [42] in 1974 on 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 [43–45] 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 [45, 46].

One of the main goals of this work is to analyze the connection between rational and heuristic approaches from a behavioral perspective. 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, and we aim to formalize such a setup. For instance [47] 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 with the increasing information.

1.2 Contributions & Organization

Both the rational and heuristic approach have a deep history in the decision theory of groups. Sobel [48] provides a theoretical framework to study the interplay between individual recommendations and rationality of group decisions. The seminal work of Janis [4] 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 [5, 6] and group polarization [7, 8]. In [49] 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.

1.2.1 The Group Decision Setup

Throughout this work, we model the group decision process (GDP) by postulating an underlying state of the world (denoted by θ) that is common but unknown to all agents. The state θ belonging to the finite set Θ models the topic of the discussion/group decision process. For example, in the course of a political debate, Θ can be the set of all political parties and it takes a binary value in a bipartisan system (cf. Fig. 1.3, on the left). The value/identity of θ is not known to the agents. To each agent, we assign an action space and a utility function that rewards their actions depending on the true state of the world and irrespective of the actions of others. Agents are interested in other people's actions only to the extent that other actions are informative about other people's private signals and allows them to make inferences about what other people know that they do not know. Such modeling of rewards to actions at successive time periods is common place in the study of learning in games and a central question of interest is that of regret which measures possible gains by the players if they were to play other actions from what they have chosen in a realized path of play; in particular, existence of update/decision rules that would guarantee a vanishing time-average of regret as $t \rightarrow \infty$, and possible equilibria that characterize the limiting behavior of agents under such rules have attracted much attention, cf. [50–52]. Under a similar framework Rosenberg et al. [15] consider consensus in a general setting with players who observe a private signal, choose an action and receive a pay-off at every stage, and pay-offs that depend only on an unknown parameter and players' actions. They show that in this setting with no pay-off externalities and interactions which are purely informational players asymptotically play their best-replies given their beliefs and will agree in their pay-offs; in particular, all motives for experimentation will eventually disappear.

In modeling the group decision process, we assume that the preferences of agents across time are myopic. At every time t , 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 does not have the freedom to learn from consequences of their recommendations, not before committing to a choice. Specifically 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

and end the group discussion); rather the purpose of the group discussion is to augment their information by learning from recommendations of others as much as possible before committing to a choice. 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 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 thesis.¹

Throughout this work, we analyze the repeated interactions of agents in a GDP following both the Bayesian and non-Bayesian frameworks. We first show how the computations of a Bayesian agent scale up with the increasing network size and give some structural conditions that can help curb this computational complexity. We then propose the no-recall model of belief formation and decision making in groups 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.

1.2.2 Algorithms and Complexity of Rational Group Decisions

In Chapter 2, we focus on the development of algorithms for Bayesian decision-making in groups and characterizing their complexity. We are interested in the computations that the Bayesian agent should undertake to achieve the goal of producing best recommendations at every decision epoch during a group discussion. We are further interested in determining how the complexity of these computations scale up with the increasing network size.²

¹Rosenberg et al. in their study of the emergence of consensus under purely informational externalities [15], show that even with forward-looking agents the incentives to experiment disappear, thus leading them to a consensus on their myopic best-responses subject to common utility and action structures. The authors in [53] also look at forward-looking agents with binary state and action space and propose an egalitarian condition on the topology of the network to guarantee learning in infinite networks. An egalitarian graph is one in which all degrees are bounded and every agent who is being observed by some agent i , observes her back, (possibly indirectly) through a path of bounded length.

²The authors in [54] analyze the problem of estimating a binary state of the world from a single initial private signal that is independent and identically distributed among the agents conditioned on the true state. The authors show that by repeatedly observing each other's best estimates of the unknown, as the size of the network increases, Bayesian agents asymptotically learn the true state with high probability. Hence, the agents are able to combine their initial private observations and learn the truth. This setting is very close to our formulation of group decision processes; however, rather than the asymptotic analysis of the probability of mistakes with the increasing network size, we are interested in the computations that each agent should undertake to realize her rational choices during the group decision process. In particular, we investigate how

We begin by explaining the Bayesian model of decision-making in groups, the so-called group decision process (GDP), and the kind of calculations that it entails. In recent works, recursive techniques have been applied to analyze Bayesian decision problems with partial success [55–57]. The importance of a recursive implementation becomes apparent in light of the fact that a forward reasoning approach is bound to scale terribly with the increasing group size (cf. Appendix C). In a forward reasoning approach, to interpret observations of the actions of others, the agent considers the causes of those actions and is able to form a Bayesian posterior by weighing all contingencies that could have lead to those actions according to their probabilities. This requires the rational agent to simulate the inferences of her neighbors at all possible actions that they could have observed, and which she cannot observe directly but can only learn about partially (and indirectly) after knowing what her neighbors do. Although this forward reasoning about causes of the actions is natural to human nature [58], it is extremely difficult to adapt to the complexities of a partially observed setting where hidden causes lead to a multiplicity of contingencies.

In Chapter 2, we will use the framework of iterated eliminations to model the thinking process of a Bayesian agent in a group decision-making scenario: As the Bayesian agent attempts to infer the true state of the world from her private signal and sequence of observations of actions of others, her decision problems at every epoch can be cast recursively, as a dynamic program [57]. By the same token, the private signals of all agents constitute the state space of the problem and with every new observation, the agent refines her knowledge about the private signals that other agents have observed, by eliminating all cases that are inconsistent with her observations under the assumption that other agents are acting rationally. The iterated elimination of infeasible signals (IEIS) approach curbs some of the complexities of the group decision process, but only to a limited extent.

In a group decision scenario, the initial private signals of the agents constitute a search space that is exponential in the size of the network. The ultimate goal of the agents is to get informed about the private signals of each other and use that information to produce the best actions. A Bayesian agent is initially informed of only her own signal; however, as the history of interactions with other group members becomes enriched, her knowledge of the possible private signals that others may have observed also gets refined; thus enabling her to make better decisions. While the search over the feasible signal profiles in the IEIS algorithm runs in exponential time, these calculations may simplify in special highly connected structures: in Subsection 2.1.2, we give an efficient algorithm that enables a Bayesian agent to compute her posterior beliefs at every decision epoch, where the graph structure is a partially ordered set (POSET) as in [59]. In such structures, any agent whose actions indirectly influences the observations of agent i is also directly observed by her. Hence, any neighbor of a neighbor of agent i is a neighbor of agent i as well; the same is

the complexity of these computations scale up with the increasing network size.




Figure 1.2: POSET structures: on the left, in a POSET structure any neighbors of a neighbor of an agent are also among her immediate (directly observable) neighbors; in the middle, a round-table discussions; on the right, a vested agent (in blue) investigates the sources of her information in one-on-one meetings; imposing a POSET structure on her neighborhood.

true for all neighbors of the neighbors of the neighbors of agent i , who would themselves be a neighbor of agent i , and so on and so forth. The complete graph, where everybody sees the actions of everybody else is an example of a POSET structure.¹ Such rich communication structures are in fact characteristic of round-table discussions, and they arise quite often in settings with purely informational externalities: for example in applications such as jury deliberations where jurors enter a group discussion after they each is independently exposed to the court proceedings. Other examples include group decision-makings among professionals such as medical doctors who have each made their own examination of a critically ill person and have come together to decide on the best course of treatment. A POSET structure may also arise as a result of a vested agent investigating the sources of her information; thus deliberately imposing a POSET structure on her neighborhood. Such scenarios arise when stakes are high enough as in gathering legal evidence, or documenting factual data for business decisions (cf. Fig. 1.2). Subsequently, in Subsection 2.1.2 we provide a partial answer to one of the open problems raised by [60], where the authors provide an efficient algorithm for computing the Bayesian binary actions in a complete graph: we show that efficient computation is possible for non-complete graphs (POSETs) with general finite action spaces.

IEIS approach curbs some of the complexities of the group decision process, but only to a limited extent. The Bayesian iterations during a group decision process can be cast into the framework of a partially observed Markov decision process (POMDP). Thereby, the private signals of all agents constitute the state space of the problem and the decision maker only has access to a deterministic function of the state, the so-called partial observations. In GDP the actions or beliefs of the neighbors constitute the partial observations. The partially observed problem and its relations to the decentralized and team decision problems

¹See Definition 2.1 for a formal statement of the POSET property and the respective constraints that are imposed on the network topology.

have been the subject of major contributions [61, 62]; in particular, the partially observed problem is known to be PSPACE-hard in the worst case [63, Theorem 6]. However, unlike the general POMDP, the state (private signals) in a GDP do not undergo Markovian jumps as they are fixed at the initiation of GDP. Hence, determining the complexity of GDP requires a different analysis. To address this requirement, in Section 2.2 we shift focus to a case where agents repeatedly exchange their beliefs (as opposed to announcing their best recommendations); subsequently in Section 2.2.3, we are able to show that computing the Bayesian posterior beliefs in a GDP is \mathcal{NP} -hard with respect to the increasing network size.¹ This result complements and informs the existing literature on Bayesian learning over networks; in particular, those which offer efficient algorithms for special settings such as Gaussian signals in a continuous state space [55], or with binary actions in a complete graph [60].

The conflict and interplay between rationality and computational tractability in economic models of human behavior has been a focus of attention by both the earlier and the contemporary scholars of the field: for example in the early works of Herbert Simon on bounded rationality, artificial intelligence and cognitive psychology [64], and in the contemporary research of Vela Velupillai on the computable foundations for economics [65]. Our work in Chapter 2 can be regarded as an effort in this direction; a particularly relevant recent study along these lines is due to [66] on complexity of agreement, who investigates the question of convergence of beliefs to a consensus and the number of messages (bits) that needs to be exchanged before one can guarantee that everybody’s beliefs are close to each other.

There is another relevant body of literature that is dedicated to computations of the Nash equilibria in games and characterizing their complexity [67]. Such equilibria of games are predictors for the behavior of rational players. The computational complexity results concerning Nash equilibria address the following problem (referred to as NASH): given a description of the game in terms of the strategy profiles and payoffs of its players (a normal-form game), compute a Nash equilibrium or an ϵ -approximation of it (a profile of mixed strategies where players cannot unilaterally improve their payoffs by more than ϵ for some positive rational ϵ). NASH is closely related to the problem of computing fixed points for certain maps, a class of problems which Etessami and Yannakakis [68] have coined as FIXP (for fixed points). The complexity of computing Nash equilibria is also tied with the PPAD complexity class (Polynomial Parity Arguments on Directed graphs), and NASH is known to be PPAD-complete [67, 69–71].² The authors in [73] study the related question

¹We provide two reductions to known \mathcal{NP} -complete problems. One reduction relies on the increasing number of different types of signals that are observed by different agents in the network. The other reduction relies on the increasing size of the agent’s neighborhood (with i.i.d signals).

²The PPAD formalism accounts for the fact that the existence of a solution is guaranteed for all instances of NASH by Brouwer’s fixed point theorem [72]. This is in sharp contrast with the majority of \mathcal{NP} -complete

of communication complexity for reaching a Nash equilibrium. They consider a setup, where each player is initially informed of only her own utility function. They show that the number of bits that need to be transmitted before reaching a pure or mixed Nash equilibrium increases exponentially with the number of players; however, a correlated equilibrium can be reached after exchanging a polynomial number of bit.

Our results also enrich the evolving body of literature on various inference problems over graphs and networks. The interplay between statistical and computational complexity in such problems, as well as their complexity and algorithmic landscapes, are interesting issues that we highlight along with other concluding remarks and discussions of future directions in Chapter 5.

1.2.3 Heuristics and Biases in Group Decision Making

In Chapter 3, we propose the no-recall model of belief formation and decision making in groups that avoids the computational drawbacks of the Bayesian model, by relying on the time-one Bayesian update and using it for all future decision epochs.

The “no-recall” 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 [74], various other authors have also developed results based on variations of this no-recall idea. In [75] the authors propose this model to analyze dynamic consensus and use it to study the effects of correlation neglect on voting behavior [75], as well as the persuasion power of media [76]. In [77], 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 *no-recall* heuristics is motivated by the real-world behavior of people induced by their system one (the fast/automatic system) 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*” [9]. On the other hand, the proposed Bayesian (no-recall) heuristics offer a boundedly rational approach to model decision making over social networks.¹ By ignoring the history of interactions, the heuristic (no-recall) agents are left with a substantially

problems, for which determining the existence of a solution (the feasibility problem) is just as hard.

¹This is bounded-rationality 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” [78]. 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

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 [29].

In Chapter 3, we analyze the format of the no-recall update rules under various structural conditions and investigate the evolution of beliefs and properties of the group decision outcome as the agents repeatedly interact with each other and deliberate their options under the no-recall model. Specific cases of Bayesian heuristics that we explore in Chapter 3 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.

1.2.4 Learning from Streaming Data

In Chapter 4, we focus on a case of repeated interactions for social learning when agents receive new private signals in addition to observing their neighboring decisions (actions) at every point in time. Such a model is a good descriptor for online reputation and polling systems such as Yelp[®] and TripAdvisor[®], where individuals' recommendations are based on their private observations and recommendations of their friends [81, Chapter 5]. The analysis of such systems is important not only because they play a significant role in generating revenues for the businesses that are being ranked [82], but also for the purposes of designing fair rankings and accurate recommendation systems.

Consider a Bayesian agent trying to estimate an unknown state of the world. She bases her estimation on a sequence of independent and identically distributed (i.i.d.) private signals that she observes, whose common distribution is determined by the unknown state. Suppose further that her belief about the unknown state is represented by a discrete probability distribution over the set of finitely many possibilities (denoted by Θ), and that she sequentially applies Bayes rule to her observations at each step, and updates her beliefs accordingly. It is a well-known consequence of the classical results in merging and learning theory [83, 84] that the beliefs formed in the above manner constitute a bounded martingale and converge to a limiting distribution as the number of observations increases. However, the limiting distribution may differ from a point mass centered at the truth, in which case the agent fails to learn the true state asymptotically. This may be the case, for instance if the agent faces an identification problem, that is when there are states other than the true state which are observationally equivalent to the true state and induce the same distribution on her sequence of privately observed signals. However, by communicating in a social network the agents can resolve their identification problems by relying on each other's observational abilities.

“satisficing” heuristics; i.e. to search for alternatives that exceed some *“aspiration levels”* by satisfying a set of minimal acceptability criteria [79, 80].




Figure 1.3: On the left, bipartisanship is an example of a binary state space. On the right, people reveal their beliefs through status updates and what they share and post on various social media platforms.

The fact that different people make independent observations about the underlying truth state, θ , gives them incentive to communicate in social networks, in order to benefit from each others' observations and to augment their private information. Moreover, different people differ in their observational abilities. For instance, let the signals and their likelihoods for some agent i be denoted by s_i and $\ell_i(s_i|\theta)$, respectively. Further suppose that the signal structure of agent i allows her to distinguish the truth θ and the false state $\check{\theta}$, while the two states $\hat{\theta}$ and θ are indistinguishable to her: i.e. $\ell_i(s_i|\check{\theta}) \neq \ell_i(s_i|\theta)$ for some $s_i \in \mathcal{S}_i$, whereas $\ell_i(s_i|\check{\theta}) = \ell_i(s_i|\theta)$ for all $s_i \in \mathcal{S}_i$, where \mathcal{S}_i is the signal space (the set of all signals that i may observe). In such circumstances, agent i can never resolve her ambiguity between θ and $\check{\theta}$ on her own; hence, she has no choice but to rely on other people's observations to be able to learn the truth state with certainty.

The chief question of interest in Chapter 4 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 Chapter 4 in which agents have access to a stream of new observations is in contrast with the group decision models of Chapter 2 and 3; the difference being in the fact that in Chapters 2 and 3 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.

We begin by specializing the no-recall model to a case where agents try to decide between one of the two possible states and are rewarded for every correct choice that they make; for example, when voting in a bipartisan political system (cf. Fig. 1.3, on the left).

When there are only finitely many states of the world and agents choose actions over the probability simplex, then the action spaces are rich enough to reveal the beliefs of every communicating agent. We show that under a quadratic utility and by taking actions over the probability simplex, agents announce their beliefs truthfully at every epoch; in practice, jurors may express their belief about the probability of guilt in a criminal case or more generally people may make statements that are expressive of their beliefs (cf. Fig. 1.3, on the right). It is explained in [26] that rich-enough action spaces can reveal

the underlying beliefs that lead to actions; subsequently, an individual's action is a fine reflection of her beliefs. The author in [85] characterizes the distinction between coarse and rich action spaces using the concept of “*responsiveness*”: the utility function is responsive, if a player with that utility chooses different actions at different beliefs (as is the case for the quadratic utility described above); the role of responsiveness in determining the observational learning outcome is also discussed in [85].

We show that the no-recall updates in the belief exchange case are *log-linear* in the reported beliefs of the neighbors and the likelihood of private signals. We investigate the properties of convergence and learning for such agents in a strongly connected social network, provided that the truth is identifiable through the aggregate observations of the agents across entire network. This is of particular interest, when the agents cannot distinguish the truth based solely on their private observations, and yet together they learn. Analysis of convergence and learning in this case reveals that almost-sure learning happens only if the agents are arranged in a directed circle. We characterize the rate of log-learning¹ in such cases as being asymptotically exponentially fast with an exponent that is linear in time and whose coefficient can be expressed as a weighted average of the relative entropies of the signal likelihoods of all agents.

¹In the literature on theory of learning in games [50], log-linear learning refers to a class of randomized strategies, where the probability of each action is proportional to an exponential of the difference between the utility of taking that action and the utility of the optimal choice. Such randomized strategies combine in a log-linear manner [86], and they have desirable convergence properties: under proper conditions, it can be shown that the limiting (stationary) distribution of action profiles is supported over the Nash equilibria of the game [87].

Chapter 2

Rational Decision Making in Groups

In this chapter, we analyze the decision problem of a Bayesian agent in the course of group interactions, as she attempts to infer the true state of the world from her sequence of observations of actions of others as well as her own private signal. Such an agent recursively refines her belief on the signals that other players could have observed and actions that they could have taken given the assumption that other players are also rational. The existing literature addresses asymptotic and equilibrium properties of Bayesian group decisions and important questions such as convergence to consensus and learning. In this work, we address the computations that the Bayesian agents should undertake to realize the optimal actions at every decision epoch. We use a scheme called iterated eliminations of infeasible signals (IEIS) to model the thinking process as well as the calculations of a Bayesian agent in a group decision scenario. Following IEIS, with every new piece of information the agent refines her knowledge about the private signals that other agents have observed, by eliminating all cases that are inconsistent with her observations given that other agents are acting rationally. We show that IEIS algorithm runs in exponential time; however, when the group structure is a partially ordered set the Bayesian calculations simplify and polynomial-time computation of the Bayesian recommendations is possible. We next shift attention to the case where agents reveal their beliefs (instead of actions) at every decision epoch. We analyze the computational complexity of the Bayesian belief formation in groups and show that it is \mathcal{NP} -hard. We also investigate the factors underlying this computational complexity and show how belief calculations simplify in special network structures or cases with strong inherent symmetries. We finally give insights about the statistical efficiency (optimality) of the beliefs and its relations to computational efficiency.^{1,2}

¹The following papers cover the results of this chapter: [29, 88]. Preliminary versions of the results were also presented in the following non-archival venues: 2017 Statistics & Data Science Center Conference, 2017 New York Computer Science and Economics Day, 2017 ACM Conference on Economics and Computation (poster session), and 2017 Symposium on the Control of Network Systems.

²I would like to thank Mina Karzand for discussions about the decision flow diagram of two agents (Fig. 2.1), Ankur Moitra for discussions about the symmetric binary environment (Appendix D), Pooya Molavi for private communication about Bayesian learning in another information structure, and Rasul Tutunov and Jonathan Weed for discussions about the \mathcal{NP} -hardness reductions.

We begin Section 2.1 by presenting the model of decision making in groups that will be applied throughout this work. In Subsection 2.1.1, we formalize the calculations of a Bayesian agent as an iterated elimination of infeasible signals (IEIS) algorithm. The IEIS approach curbs some of the complexities of the group decision process, but only to a limited extent. While the search over the feasible signal profiles in the IEIS algorithm runs in exponential time, these calculations may simplify in special highly connected structures. In Subsection 2.1.2, we give an efficient algorithm that enables a Bayesian agent to compute her posterior beliefs at every decision epoch, where the graph structure is a partially ordered set (POSET), cf. Definition 2.1 for the POSET property and the respective constraints that are imposed on the network topology. In Section 2.2 we shift focus to a case where agents repeatedly exchange their beliefs (as opposed to announcing their best recommendations); subsequently in Subsection 2.2.3, we are able to show that computing the Bayesian posterior beliefs in a GDP is \mathcal{NP} -hard with respect to the increasing network size. We provide two reductions to known \mathcal{NP} -complete problems. One reduction relies on the increasing number of different types of signals that are observed by different agents in the network. The other reduction relies on the increasing size of the agent’s neighborhood (with i.i.d signals).

2.1 The Bayesian Model

We consider¹ a group of n agents, labeled by $[n] = \{1, \dots, n\}$, and interact according to a fixed directed graph \mathcal{G} . For each agent $i \in [n]$, \mathcal{N}_i denotes a neighborhood $\mathcal{N}_i \subset [n]$, whose actions are observed by agent i . We use $\nabla(j, i)$ to denote the length (number of edges) of the shortest path in \mathcal{G} that connects j to 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 Θ , which is measurable by a σ -finite measure $\mathcal{G}_\theta(\cdot)$. For example if a space (Θ or \mathcal{S}) is a countable set, then we can take its σ -finite measure (\mathcal{G}_θ 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 σ -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 θ , $\mathcal{L}_i(\cdot | \theta)$ 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

¹ Some notations: \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 $\nabla \mathcal{X}$ 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.

triplet, where $\Omega = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ is a product space, and \mathcal{F} is a properly defined sigma field over Ω . The probability measure on Ω 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 θ 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 $\mathcal{P}_\theta(d\omega)$, $\omega \in \Omega$.

An agents' belief about the unknown allows her to make decisions even as the outcome is dependent on the unknown value θ . These beliefs about the unknown state are probability distributions over Θ . Even before any observations are made, every agent $i \in [n]$ holds a prior belief $\mathcal{V}_i(\cdot) \in \nabla\Theta$; this represents her subjective biases about the possible values of θ . For each time instant t , let $\mathcal{M}_{i,t}(\cdot)$ be the (random) probability distribution over Θ , representing the *opinion* or *belief* at time t of agent i about the realized value of θ . Moreover, let the associated expectation operator be $\mathbb{E}_{i,t}\{\cdot\}$, representing integration with respect to $\mathcal{M}_{i,t}(d\theta)$.

Let $t \in \mathbb{N}_0$ denote the time index; at $t = 0$ the values $\theta \in \Theta$ followed by $\mathbf{s}_i \in \mathcal{S}_i$ 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 Θ .¹ 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, $\mathbb{E}_{i,t}\{u_i(\mathbf{a}_{i,t}, \theta)\}$, where the expectation is with respect to $\mathcal{M}_{i,t}$.

We now proceed to present the elements of the rational model for decision-making in a group. We assume that the signal, state, and action spaces are finite sets. Because some of our algorithms rely critically on the ability of the Bayesian agent to enumerate all possible private signals that the other network agents may have observed. We may relax this assumption in special cases where calculations are possible without resorting to exhaustive enumerations. The ultimate goal of a Bayesian agent can be described as learning enough about the private signals of all other agents in the network to be able to compute the Bayesian posterior belief about the true state, given her local observations; this, however, can be extremely complex, if not impossible.

In a rational group decision process (GDP), each agent $i \in [n]$ receives a private signal \mathbf{s}_i at the beginning and then engages in repeated interactions with other group members

¹The signal, action, and utility structures, as well as the priors, are all common knowledge among the players; this assumption of common knowledge, in particular, implies that given the same access to each other's behavior or private information distinct agents would make identical inferences; in the sense that starting from the same belief about the unknown θ , their updated beliefs given the same observations would be the same: rational agents cannot agree to disagree, in Aumann's words [12].

in the ensuing decision epochs: choosing actions and observing neighbors' choices every time. The agents start with a prior belief about the value of θ , which is a probability distribution over the set Θ with probability mass function $\nu(\cdot) = d\mathcal{V}_i/d\mathcal{K} : \Theta \rightarrow [0, 1]$ for all i . Throughout Chapter 2, we assume that this prior is common to all agents.¹ At each time t , we denote the Bayesian posterior belief of agent i given her history of observations by its probability mass function $\mu_{i,t}(\cdot) = d\mathcal{M}_{i,t}/d\mathcal{K} : \Theta \rightarrow [0, 1]$. Initially, every agent receives a private signal about the unknown θ . Each signal \mathbf{s}_i belongs to a finite set \mathcal{S}_i and its distribution conditioned on θ is given by $\ell_i(\cdot | \theta) := d\mathcal{L}_i(\cdot | \theta)/d\mathcal{K}$ which is referred to as the *signal structure* of agent i . We use $\mathcal{L}(\cdot | \theta)$ to denote the joint distribution of the private signals of all agents, signals being independent across the agents.²

Accordingly, at every time t , agent i observes the most recent actions of her neighbors, $\{\mathbf{a}_{j,t-1} \text{ for all } j \in \mathcal{N}_i\}$, and chooses an action $\mathbf{a}_{i,t} \in \mathcal{A}_i$, maximizing her expected utility given all her observations up to time t , $\{\mathbf{a}_{j,\tau} \text{ for all } j \in \mathcal{N}_i, \text{ and } \tau \leq t - 1\}$. For example in the case of two communicating agents the action of agent one at time two $\mathbf{a}_{1,2}$ is influenced by own private signal \mathbf{s}_1 as well as the neighboring action at times zero and one; part of the difficulty of the analysis is due to the fact that the action of agent two at time one is shaped not only by the private information of agent two but also by the action of agent one at time zero, cf. Fig. 2.1. A heuristic behavior may be justified as a mistake (cognitive bias) by interpreting actions of others as consequences of their private information, thus ignoring the history of observations when making inferences about the actions of others; in Fig. 2.1 this corresponds to ignoring all the arrows except those which are exiting the signal and state nodes: \mathbf{s}_1 , \mathbf{s}_2 , and θ (cf. Fig. 3.1). Such biased inferences are the focus of our results in the next two chapters, where we study the heuristics for decision making and learning in groups.

In more general structures, there are also unobserved third party interactions that influence the decisions of agent two but are not available to agent one (and therefore should be inferred indirectly).

¹The assumption of common priors in our framework is without loss of generality, as long as each agents' prior is known to all people who observe her directly or indirectly in the network. This is due to the fact that agents can account for their respective priors when making inferences about each other's actions. In particular, all of the algorithms presented throughout the paper can be adapted to work with personalized priors. This is true because when simulating other agents' actions in (2.1.4) or computing their beliefs in (2.2.5), the agents can use each individuals' own prior beliefs in their calculations. The same principle applies to special structures where computations simplify and exhaustive simulations can be avoided; for example in step (1) of Algorithm (A2) or step (2) of Algorithm (A3). Finally, our \mathcal{NP} -hardness results naturally carry through since they are concerned with the worst-case structures.

²This independence of private signals allows us to exploit a decomposition property of feasible signal profiles in case of POSETs in Subsection 2.1.2 to achieve polynomial-time Bayesian computations.




Figure 2.1: The Decision Flow Diagram for Two Bayesian Agents

For each agent i , her history of observations $\mathbf{h}_{i,t}$ is an element of the set:

$$\mathcal{H}_{i,t} = \mathcal{S}_i \times \left(\prod_{j \in \mathcal{N}_i} \mathcal{A}_j \right)^{t-1}.$$

At every time t , the expected reward to agent i given her choice of action \mathbf{a}_i and observed history $\mathbf{h}_{i,t}$ is given by the expected reward function $r_{i,t} : \mathcal{A}_i \times \mathcal{H}_{i,t} \rightarrow \mathbb{R}$, as follows:

$$r_{i,t}(\mathbf{a}_i, \mathbf{h}_{i,t}) = \mathbb{E}_{i,t} \{u_i(\mathbf{a}_i, \theta) \mid \mathbf{h}_{i,t}\} = \sum_{\hat{\theta} \in \Theta} u_i(\mathbf{a}_i, \hat{\theta}) \mu_{i,t}(\hat{\theta}),$$

for all $\mathbf{h}_{i,t} \in \mathcal{H}_{i,t}$, where $\mu_{i,t}(\hat{\theta})$ is the Bayesian posterior of agent i about the truth θ given the observed history $\mathbf{h}_{i,t}$. The (myopic) optimal action of agent i at time t is then given by $\mathbf{a}_{i,t} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} r_{i,t}(\mathbf{a}_i, \mathbf{h}_{i,t})$. Here for a set \mathcal{A} , we use the notation $\mathbf{a} \leftarrow \mathcal{A}$ to denote a deterministic choice from the elements of \mathcal{A} that is assigned to \mathbf{a} . The deterministic choice is dictated by a tie-breaking rule that is known to all agents; specifically, we assume that all of the action spaces are ordered (arbitrarily) and the ordering of each action space \mathcal{A}_i is known to all agents who observe agent i , directly or indirectly. According to the tie-breaking rule \leftarrow , whenever an agent is indifferent between a multitude of options she will choose the one that ranks lowest in their ordering. The restriction to deterministic tie-breaking rules is not without loss of generality. Because in the case of randomized tie-breaking rules, rational agents would have to make inferences about how past occurrences of ties have been resolved by other agents, whom they observe directly or indirectly. This is in addition to their inferences about private signals and other unknown random quantities whose values they are trying to learn. Thus the agent's problem is to calculate her Bayesian posterior belief $\mu_{i,t}$, given her history of past observations: $\mathbf{h}_{i,t} := \{\mathbf{s}_i, \mathbf{a}_{j,\tau}, j \in \mathcal{N}_i, \tau \in [t-1]\}$. Asymptotic properties of Bayesian group decisions, including convergence of the actions to a consensus and learning (convergence to an "optimal" aggregate action), can be studied using the Markov Bayesian equilibrium as a solution concept (cf. Appendix B); however, our main focus in this paper is on the computational and algorithmic aspects of

the GDP rather than its asymptotic properties.

Refinement of information partitions with the increasing observations is a key feature of rational learning problems and it is fundamental to major classical results that establish agreement [13] or learning [83, 84] among rational agents. Several follow-up works of [13] have extended different aspects of information exchange among rational agents. In this line of work, it is of particular interest to derive conditions that ensure the refinement of information partitions would lead to the consensus on and/or the common knowledge of an aggregate decision.¹ In particular, the author in [25] points out that rational social learning requires all agents in every period to consider the set of possible information partitions of other agents and to further determine how each choice would impact the information partitions of others in the subsequent periods.

In the GDP setting, the list of feasible signals can be regarded as the information set representing the current understanding of the agent about her environment and the way additional observations are informative is by trimming the current information set and reducing the ambiguity in the set of initial signals that have caused the agent’s history of past observations. In Section 2.1.1, we describe a recursive implementation² for the refinement of the information sets (partitions) that relies on iterated elimination of infeasible signals (IEIS) for all the agents. The IEIS calculations scale exponentially with the network size; this is true with the exception of some very well-connected agents who have, indeed, direct access to all the observations of their neighbors and can thus analyze the decisions of each of their neighbors based on their respective observations. We expand on this special case (called POSETs) in Subsection 2.1.2 and explain how the Bayesian calculations simplify as a result.

¹Some notable examples include the works of [89–91], which consider information exchange by repeatedly reporting the values of a general set function $f(\cdot)$ over the state space (rather than the conditional probabilities, which are the Bayesian beliefs). The authors in [89, 90] propose a condition of union consistency on $f(\cdot)$ and the authors in [91] replace this union consistency condition with a convexity property for $f(\cdot)$, all ensuring that the value of $f(\cdot)$ become common knowledge among the agents after repeated exchanges.

²The recursive implementation is an important factor in reducing the complexity of Bayesian calculations. In Appendix C, we describe the Bayesian calculations for causal (forward) reasoning over a path. The Bayesian calculations in Appendix C involve simulation of the network behavior for all the feasible signal profiles of the neighboring agents contingent on all of their possible “observations” (i.e. the set of all possible unobserved histories). Such a need arise as a result of forward reasoning on the part of the agents to try and combine the incoming information with what they have already learned: agents isolate all the new information about the private signals of a far away agent that has reached them for the first time and use that to refine what they already know about previously discovered agents. Such a forward implementation is very cumbersome for complex network structures; in other words, forward simulation of other people’s inferences at all their possible observations is extremely inefficient and scales very poorly when applied to general group structures.

2.1.1 Iterated Elimination of Infeasible Signals (IEIS)

Building on the prior works [25, 57], we implement the refinement of information partitions for rational agents in a group decision process as an iterated elimination of infeasible signals. Accordingly, at every decision time, the signal profiles that are inconsistent with the most recent observations are removed, leading to a refined information set for next period. In this section, we analyze the Bayesian calculations that take place among the group members as they calculate their refined information partitions and the corresponding beliefs. To calculate their Bayesian posteriors, each of the agents keeps track of a list of possible combinations of private signals of all the other agents. At each iteration, they refine their list of feasible signal profiles in accordance with the most recent actions of their neighbors.

To proceed, let $\bar{s} = (s_1, \dots, s_n) \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ be any profile of initial signals observed by each agent across the network, and denote the set of all private signal profiles that agent i regards as feasible at time t , i.e. her information set at time t , by $\mathcal{I}_{i,t} \subset \mathcal{S}_1 \times \dots \times \mathcal{S}_n$; this set is a random set, as it is determined by the random observations of agent i up to time t . Starting from $\mathcal{I}_{i,0} = \{s_i\} \times \prod_{j \neq i} \mathcal{S}_j$, at every decision epoch agent i removes those signal profiles in $\mathcal{I}_{i,t-1}$ that are not consistent with her history of observations $\mathbf{h}_{i,t}$ and comes up with a trimmed set of signal profiles $\mathcal{I}_{i,t} \subset \mathcal{I}_{i,t-1}$ to form her Bayesian posterior belief and make her decision at time t . The set of feasible signals $\mathcal{I}_{i,t}$ is mapped to a Bayesian posterior for agent i at time t as follows:

$$\boldsymbol{\mu}_{i,t}(\theta) = \frac{\sum_{\bar{s} \in \mathcal{I}_{i,t}} \mathcal{L}(\bar{s}|\theta) \nu(\theta)}{\sum_{\hat{\theta} \in \Theta} \sum_{\bar{s} \in \mathcal{I}_{i,t}} \mathcal{L}(\bar{s}|\hat{\theta}) \nu(\hat{\theta})}, \quad (2.1.1)$$

which in turn enables the agent to choose an optimal (myopic) action given her observations:¹

$$\mathbf{a}_{i,t} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \sum_{\hat{\theta} \in \Theta} u_i(\mathbf{a}_i, \hat{\theta}) \boldsymbol{\mu}_{i,t}(\hat{\theta}). \quad (2.1.2)$$

For example at time zero agent i learns her private signal s_i , this enables her to initialize her list of feasible signals: $\mathcal{I}_{i,0} = \{s_i\} \times \prod_{k \in [n] \setminus \{i\}} \mathcal{S}_k$. Subsequently, her Bayesian posterior at time zero is given by:

$$\boldsymbol{\mu}_{i,0}(\theta) = \frac{\sum_{\bar{s} \in \mathcal{I}_{i,0}} \mathcal{L}(\bar{s}|\theta) \nu(\theta)}{\sum_{\hat{\theta} \in \Theta} \sum_{\bar{s} \in \mathcal{I}_{i,0}} \mathcal{L}(\bar{s}|\hat{\theta}) \nu(\hat{\theta})} = \frac{l_i(\mathbf{s}_i|\theta) \nu(\theta)}{\sum_{\hat{\theta} \in \Theta} l_i(\mathbf{s}_i|\hat{\theta}) \nu(\hat{\theta})}$$

¹In this sense, the Bayesian posterior is a sufficient statistic for the history of observations and unlike the observation history, it does not grow in dimension with time.

and her optimal action (recommendation) at time one is as follows:

$$\mathbf{a}_{i,0} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \frac{\sum_{\hat{\theta} \in \Theta} \mathbf{u}_i(\mathbf{a}_i, \hat{\theta}) \ell_i(\mathbf{s}_i | \hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \ell_i(\mathbf{s}_i | \tilde{\theta}) \nu(\tilde{\theta})} = \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \sum_{\hat{\theta} \in \Theta} \mathbf{u}_i(\mathbf{a}_i, \hat{\theta}) \mu_{i,1}(\hat{\theta}). \quad (2.1.3)$$

In IEIS, the agent not only needs to keep track of the list of private signals that are consistent with her observations, denoted by $\mathcal{I}_{i,t}$, but also she needs to consider what other agents regard as consistent with their own observations under the particular set of initial signals. The latter consideration enables the decision maker to calculate actions of other agents under any circumstances that arise at a fixed profile of initial signals, as she tries to evaluate the feasibility of that particular signal profile given her observations. In other words, the neighbors are acting rationally in accordance with what they regard as being a feasible set of initial signal profiles. Hence, with every new observation of the neighboring actions, agent i not only refines her knowledge of other people's private signals but also her knowledge of what signal profiles other agents would regard as feasible.

For any agent $j \neq i$ and at every signal profile \bar{s} , let $\mathcal{I}_{j,t}^{(i)}(\bar{s})$ be the set of all signal profiles that agent i believes have not yet been rejected by agent j , given all her observations and conditioned on the initial private signals being \bar{s} . Consider the feasible action calculated by agent i for agent j under the assumption that the initial private signals are prescribed by $\bar{s} = (s_1, \dots, s_n)$, i.e.

$$\mathbf{a}_{j,\tau}^{(i)}(\bar{s}) \leftarrow \arg \max_{\mathbf{a}_j \in \mathcal{A}_j} \sum_{\hat{\theta} \in \Theta} \mathbf{u}_j(\mathbf{a}_j, \hat{\theta}) \frac{\sum_{\bar{s}' \in \mathcal{I}_{j,\tau}^{(i)}(\bar{s})} \mathcal{L}(\bar{s}' | \hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \sum_{\bar{s}' \in \mathcal{I}_{j,\tau}^{(i)}(\bar{s})} \mathcal{L}(\bar{s}' | \tilde{\theta}) \nu(\tilde{\theta})}, \forall \tau \in [t], \quad (2.1.4)$$

where $\mathcal{I}_{j,\tau}^{(i)}(\bar{s})$ is defined in Table 2.1. Given $\mathbf{a}_{j,t}^{(i)}(\bar{s})$ for all $\bar{s} \in \mathcal{I}_{i,t-1}$ and every $j \in \mathcal{N}_i$, the agent can reject any \bar{s} for which the observed neighboring action $\mathbf{a}_{j,t}$ for some $j \in \mathcal{N}_i$ does not agree with the simulated feasible action conditioned on \bar{s} : $\mathbf{a}_{j,t} \neq \mathbf{a}_{j,t}^{(i)}(\bar{s})$. To proceed, we introduce the notation \mathcal{N}_i^τ as the τ -th order neighborhood of agent i comprising entirely of those agents who are connected to agent i through a walk of length τ : $\mathcal{N}_i^\tau = \{j \in [n] : j \in \mathcal{N}_{i_1}, i_1 \in \mathcal{N}_{i_2}, \dots, i_{\tau-1} \in \mathcal{N}_{i_\tau}, i_\tau = i, \text{ for some } i_1, \dots, i_{\tau-1} \in [n]\}$; in particular, $\mathcal{N}_i^1 = \mathcal{N}_i$ and we use the convention $\mathcal{N}_i^0 = \{i\}$. We further denote $\bar{\mathcal{N}}_i^t := \cup_{\tau=0}^t \mathcal{N}_i^\tau$ as the set of all agents who are within distance t of or closer to agent i ; we sometimes refer to $\bar{\mathcal{N}}_i^t$ as her t -radius ego-net. For example, we use $\bar{\mathcal{N}}_i^1 = \bar{\mathcal{N}}_i := \{i\} \cup \mathcal{N}_i$ to denote the self-inclusive neighborhood of agent i . We refer to the cardinality of $\bar{\mathcal{N}}_i^t$ as the degree of node i and denote it by $\deg(i)$.

We now describe the calculations that agent i undertakes at every time t to update her list of feasible signal profiles from $\mathcal{I}_{i,t}$ to $\mathcal{I}_{i,t+1}$: agent i initializes her list of feasible signals $\mathcal{I}_{i,0} = \{\mathbf{s}_i\} \times \prod_{j \neq i} \mathcal{S}_j$; at time t she would have access to $\mathcal{I}_{i,t}$, the list of feasible signal profiles that are consistent with her observations, as well as all signal profiles that

Table 2.1: List of the variables that play a role in the Bayesian calculations for group decision-making (BAYES-GROUP).

\bar{s}	$\bar{s} = (s_1, s_2, \dots, s_n) \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ is a profile of initial private signals.
$\mathcal{I}_{i,t}$	$\mathcal{I}_{i,t} \subset \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ is the list of all signal profiles that are deemed feasible by agent i , given her observations up until time t .
$\mathcal{I}_{j,t}^{(i)}(\bar{s})$	$\mathcal{I}_{j,t}^{(i)}(\bar{s}) \subset \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ is the list of all signal profiles that agent i believes are deemed feasible by agent j , given what agent i believes agent j could have observed up until time t conditioned on the event that the initial signals of all agents are prescribed according to \bar{s} .
$\alpha_{j,t}^{(i)}(\bar{s})$	$\alpha_{j,t}^{(i)}(\bar{s}) \in \mathcal{A}_j$ is the action that agent i deems optimal for agent j , given what agent i believes agent j could have observed up until time t conditioned on the event that the initial signals of all agents are prescribed according to \bar{s} .

she thinks each of the other agents would regard as feasible conditioned on any profile of initial signals: $\mathcal{I}_{j,t-\tau}^{(i)}(\bar{s})$ for all $\bar{s} \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$, all $j \in \mathcal{N}_i^\tau$, and all $\tau \in [t]$. Calculations of agent i at time t enables her to update her information at time t to incorporate the newly obtained data which constitute her observations of neighbors' most recent actions $\mathbf{a}_{j,t}$ for all $j \in \mathcal{N}_i$; whence she refines $\mathcal{I}_{i,t}$ to $\mathcal{I}_{i,t+1}$ and updates her belief and actions accordingly, cf. (2.1.1) and (2.1.2). This is achieved as follows (recall that we use $\nabla(j, i)$ to denote the length of the shortest path connecting j to i):

(II: BAYES-GROUP). The information available to agent i at time t :

- $\mathcal{I}_{j,t-\tau}^{(i)}(\bar{s})$ for all $\bar{s} \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$, all $j \in \mathcal{N}_i^\tau$, and all $\tau \in [t]$.
- $\mathcal{I}_{i,t}$, i.e. all signal profiles that she regards as feasible given her observations.

(A1: BAYES-GROUP). Calculations of agent i at time t for deciding $\mathbf{a}_{i,t+1}$:

1. For all $\bar{s} := (s_1, \dots, s_n) \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ and all $j \in \mathcal{N}_i^{t+1}$ do:
 - If $\nabla(j, i) = t + 1$, initialize $\mathcal{I}_{j,0}^{(i)}(\bar{s}) = \{s_j\} \times \prod_{k \neq j} \mathcal{S}_k$.
 - Else initialize $\mathcal{I}_{j,t+1-\nabla(j,i)}^{(i)}(\bar{s}) = \mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$ and for all $\tilde{s} \in \mathcal{I}_{j,t+1-\nabla(j,i)}^{(i)}(\bar{s})$ do:
 - For all $k \in \mathcal{N}_j$ if $\mathbf{a}_{k,t-\tau}^{(i)}(\tilde{s}) \neq \mathbf{a}_{k,t-\tau}^{(i)}(\bar{s})$, then $\mathcal{I}_{j,t+1-\tau}^{(i)}(\bar{s}) = \mathcal{I}_{j,t+1-\tau}^{(i)}(\bar{s}) \setminus \{\tilde{s}\}$, where $\mathbf{a}_{k,t-\tau}^{(i)}(\tilde{s})$ and $\mathbf{a}_{k,t-\tau}^{(i)}(\bar{s})$ are calculated using (2.1.4), based on $\mathcal{I}_{k,t-\tau}^{(i)}(\tilde{s})$ and $\mathcal{I}_{k,t-\tau}^{(i)}(\bar{s})$.
2. Initialize $\mathcal{I}_{i,t+1} = \mathcal{I}_{i,t}$ and for all $\bar{s} \in \mathcal{I}_{i,t+1}$ do:
 - For all $j \in \mathcal{N}_i$ if $\mathbf{a}_{j,t} \neq \mathbf{a}_{j,t}^{(i)}(\bar{s})$, then $\mathcal{I}_{i,t+1} = \mathcal{I}_{i,t+1} \setminus \{\bar{s}\}$.

In Appendix A.1 we describe the complexity of the computations that the agent should undertake using (A1) at any time t in order to calculate her posterior probability $\mu_{i,t+1}$ and Bayesian decision $\mathbf{a}_{i,t+1}$ given all her observations up to time t . Subsequently, we prove that:

Theorem 2.1 (Complexity of IEIS). *There exists an IEIS algorithm with an $O(n^2 M^{2n-1} m_A)$ running time, which given the private signal of agent i and the previous actions of her neighbors $\{\mathbf{a}_{j,\tau} : j \in \mathcal{N}_i, \tau < t\}$ in any network structure, calculates $\mathbf{a}_{i,t}$, the updated action of agent i at time t .*

Remark 2.2 (Structure and Complexity in Decision Making Organizations). Suppose the cardinality of the set of agents who influence the decisions of agent i (her cone of influence) remains bounded with the network size: $\bar{\mathcal{N}}_i^n \leq D$ for some fixed $D \in \mathbb{N}$. In such structures, where the growth is bounded, the Bayesian computations using (A1) become polynomial, upon replacing n with fixed D in (A.1.2). Such bounded structures can, for example, arise as a result of horizontal growth in organizations as shown in Fig. 2.2. The question of structure and its relation to performance receive considerable attention in organization studies. Through a series of seminal papers [92–94], Sah and Stiglitz popularized a model of project selection in organizations to study the effect of their structures, and in particular to compare the performance of hierarchies and polyarchies. The authors in [95] consider the




Figure 2.2: A structure with bounded growth: each agent is influenced by no more than three other agents even as the network (organization) size grows to infinity.

optimal decision making structures for reducing the probability of two error types in project evaluation tasks (rejecting profitable projects, type I error, or accepting unprofitable ones, type II error). They point out that either of the hierarchical or polyarchical organization structures are suitable for reducing one error type and they can be combined optimally to produce good overall performance. They further study the incremental improvement from the addition of new decision-makers and point out that polyarchical structures allow for the information to propagate throughout the organization, while in hierarchical organizations most information is filter out on the way to the top. Therefore, from a complexity standpoint, extending hierarchies to accommodate new members can lead to better tractability with the increasing organization size.

2.1.2 IEIS over POSETs

We now shift focus to the special case of POSET networks. A partially ordered set (POSET) consists of a set together with an order which is a reflexive, antisymmetric, and transitive binary relation (indicating that, for certain pairs of elements, one of them precedes the other in the ordering).

Definition 2.1 (POSET Networks). *We call a network structure a POSET if the directed neighborhood relationship between its nodes satisfies the reflexive and transitive properties (note that we relax the anti-symmetric property). In particular, the transitive property implies that anyone whose actions indirectly influences the observations of agent i is also directly observed by her, i.e. any neighbor of a neighbor of agent i is a neighbor of agent i as well.*

In a POSET network it is always true that $\mathcal{N}_i^t \subset \mathcal{N}_i^\tau \subset \mathcal{N}_i$ for all $t \leq \tau$; and in particular, $\mathcal{N}_i^t = \mathcal{N}_i$ for all t : as time marches on, no new private signals will ever be discovered, only what is known about the private signals in the neighborhood \mathcal{N}_i gets refined. In more tangible terms, the POSET requirement for an agent i would imply that in observing any of her neighbors $j \in \mathcal{N}_i$, she not only observes agent j but also observes anything that agent j observes (except for agent j 's private signal).¹

¹We can regard the directed neighborhood relationship as a binary relation on the set of vertices: i is in

Table 2.2: List of the variables for Bayesian calculations in POSET groups (BAYES-POSET).

$\mathcal{S}_{i,t}$	$\mathcal{S}_{i,t} \subset \mathcal{S}_i$ is the list of all private signals that are deemed feasible for agent i at time t , by an agent who has observed her actions in a POSET network structure up until time t .
$\mathbf{a}_{i,t}(s_i)$	$\mathbf{a}_{i,t}(s_i) \in \mathcal{A}_i$ is the optimal choice of agent i at time t , given her observations in the POSET up until time t conditioned on the event that her initial private is s_i .
$\mathcal{I}_{i,t}(s_i)$	$\mathcal{I}_{i,t}(s_i) = \{s_i\} \times \prod_{j \in \mathcal{N}_i} \mathcal{S}_{j,t}$ is the list of all signal profiles that are deemed feasible by agent i for the POSET of her neighbors, given her observations of their actions up until time t conditioned on own private signal being s_i .

The special structure of POSET networks mitigates the issue of hidden observations, and as a result, Bayesian inferences in a POSET structure are significantly less complex. In particular, the fact that the agent has access to all observations of her neighbors while observing their actions allows her to directly map an observed action to refined information about the private signals of the particular agent taking that action. We make this intuition precise in what follows by giving the exact description of the Bayesian calculations that an agent performs in a POSET structure.

Note from Table 2.2 that agent i needs only to keep track of $\mathcal{S}_{j,t}$ for all $j \in \mathcal{N}_i$, i.e. the private signals that she deems feasible for each of her neighbors individually. This is due to the fact that in a POSET structure all agents whose actions may influence (directly or indirectly) the recommendations of a decision maker are already directly observed by her; any other agent's private signals would be immaterial to her decisions, as she would never make any observations that might have been influenced by those other agent's private signals. At the t -th decision epoch, the information that is at the disposal of agent i constitutes the list of private signals that agent i deems feasible for each of her neighbors $j \in \mathcal{N}_i$ given her observations up to time t . The goal of the agent at time t is to update her list of feasible signal profiles from $\mathcal{I}_{i,t}$ to $\mathcal{I}_{i,t+1}$ by incorporating her observations of her neighboring actions at time t : $\mathbf{a}_{j,t}, j \in \mathcal{N}_i$. The POSET structure allows the list of feasible signal profiles at time t to be decomposed according to the signals that are feasible for each of the neighbors individually, i.e. $\mathcal{I}_{i,t} = \{s_i\} \times \prod_{j \in \mathcal{N}_i} \mathcal{S}_{j,t}$; the updating is thus achieved by incorporating the respective actions $\mathbf{a}_{j,t}$ for each $j \in \mathcal{N}_i$ individually and transforming the

relation \mathcal{R}_N with j iff $j \in \mathcal{N}_i$. Then the POSET property would ensure that \mathcal{R}_N is a transitive relation on the set of vertices. Note that the neighborhood relationship as defined does not specify a partial order on the set of vertices because it does not satisfy the antisymmetry property. To provide for the anti-symmetry property, one needs to identify all pairs of vertices with a bidirectional link between them; thus by identifying agents who have a bidirectional link between them we obtain the neighborhood partial order \preceq_N between the set of agents in a POSET group: $i \succeq_N j, \forall j \in \mathcal{N}_i$.

respective $\mathcal{S}_{j,t}$ into $\mathcal{S}_{j,t+1}$. Agent i could then refine her belief and come up with improved recommendations based on (2.1.1) and (2.1.2). After initializing $\mathcal{S}_{j,0} = \mathcal{S}_j$ for all $j \in \mathcal{N}_i$ and $\mathcal{I}_{i,0} = \{\mathbf{s}_i\} \times \prod_{j \in \mathcal{N}_i} \mathcal{S}_{j,0}$, at any time t the transformation from $\mathcal{I}_{i,t}$ into $\mathcal{I}_{i,t+1}$ in a POSET structure can be achieved as follows:

(I2: BAYES-POSET). The information available to agent i at time t :

- $\mathcal{S}_{j,t} \subset \mathcal{S}_j$ for all $j \in \mathcal{N}_i$ is the list of private signals that agent i deems feasible for her neighbor $j \in \mathcal{N}_i$ given her observations up to time t .

(A2: BAYES-POSET). Calculations of agent i at time t for deciding $\mathbf{a}_{i,t+1}$ in a POSET:

1. For all $j \in \mathcal{N}_i$ do:

- Initialize $\mathcal{S}_{j,t+1} = \mathcal{S}_{j,t}$, and for all $s_j \in \mathcal{S}_{j,t+1}$ do:
 - Calculate $\mathbf{a}_{j,t}(s_j)$ given $\mathcal{I}_{j,t}(s_j) = \{s_j\} \times \prod_{k \in \mathcal{N}_j} \mathcal{S}_{k,t}$.
 - If $\mathbf{a}_{j,t} \neq \mathbf{a}_{j,t}(s_j)$, then set $\mathcal{S}_{j,t+1} = \mathcal{S}_{j,t+1} \setminus \{s_j\}$.

2. Update $\mathcal{I}_{i,t+1} = \{\mathbf{s}_i\} \times \prod_{j \in \mathcal{N}_i} \mathcal{S}_{j,t+1}$.

In Appendix A.2, we determine the computational complexity of (A2:BAYES-POSET) as follows:

Theorem 2.3 (Efficient Bayesian group decisions in POSETs). *There exists an algorithm with running time $O(Amn^2M^2)$ which given the private signal of agent i and the previous actions of her neighbors $\{\mathbf{a}_{j,\tau} : j \in \mathcal{N}_i, \tau < t\}$ in any POSET, calculates $\mathbf{a}_{i,t}$, the updated action of agent i at time t .*

The highly connected structure of POSETs leads to the rich history of observations from the neighboring actions that allows for efficient computation of Bayesian decisions in POSETs. On the other hand, one can also design efficient algorithms that are tailored to the special symmetries of the signal or network structure; for example, if all agents observe i.i.d. binary signals and take their best guess of the underlying binary state (cf. Appendix D).

We end this section by highlighting the observation from Appendix C that in a path of length n , the $(n - t)$ -th agent gets fixed in her decisions after time t ; and in particular, no agents will change their recommendations after $t \geq n - 1$ (see the left graph in Fig. 2.3 for the case $n = 4$). The following proposition extends our above realization about




Figure 2.3: On the left, a directed path of length four. On the right, a directed graph is acyclic if and only if it has a topological ordering; a topological ordering of a DAG orders its vertices such that every edge goes from a lesser node (to the left) to a higher one (to the right).

the bounded convergence time of group decision process over paths to all directed acyclic graphs (DAGs), cf. e.g. [96]. Such ordered structures include many cases of interest in real-world applications with a conceivable hierarchy among players: each agent observe her inferiors and is observed by her superiors or vice versa.¹ A simple examination of the dynamics in the case of two communicating agents (with a bidirectional link between them) reveals how the conclusion of this proposition can be violated in loopy structures.

Proposition 2.1 (Bounded convergence time of group decision process over DAGs). *Let \mathcal{G} be a DAG on n nodes with a topological ordering \prec , and let the agents be labeled in accordance with this topological order as follows: $n \prec n-1 \prec \dots \prec 1$. Then every agent $n-t$ gets fixed in her decisions after time t ; and in particular, no agents will change their recommendations after $t \geq n-1$.*

2.2 The Case of Revealed Beliefs

Let us label $\theta_j \in \Theta := \{\theta_1, \dots, \theta_m\}, j \in \{1, \dots, m\}$ by $\bar{e}_j \in \mathbb{R}^m$ which is a column vector of all zeros except for its j -th element which is equal to one. Furthermore, we relax the requirement that the action spaces $\mathcal{A}_i, i \in [n]$ are finite sets; instead, for each agents $i \in [n]$ let \mathcal{A}_i be the m -dimensional probability simplex: $\mathcal{A}_i = \Delta\Theta = \{(x_1, \dots, x_m)^T \in \mathbb{R}^m : \sum_1^m x_i = 1 \text{ and } x_i \geq 0, \forall i\}$. If the utility assigned to each action $\bar{a} := (a_1, \dots, a_m)^T \in \mathcal{A}_i$ and at every state $\theta_j \in \Theta$, measures the quadratic squared distance between \bar{a} and \bar{e}_j , then

¹The key property of DAGs is their topological ordering [96, Proposition 2.1.3]: a topological ordering of a directed graph is an ordering of its vertices into a sequence, such that for every edge the start vertex of the edge occurs earlier in the sequence than the ending vertex of the edge, and DAGs can be equivalently characterized as the graphs that have topological orderings. This topological ordering property allows for the influences of other agents to be addressed and analyzed in an orderly fashion, starting with the closest agents and expanding to farther and farther agents as time proceeds (see the right graph in Fig. 2.3). This topological ordering can be obtained by removing a vertex with no neighbors (which is guaranteed to exist in any DAG) and by repeating this procedure in the resultant DAG. Using a depth-first search (DFS) one can devise an algorithm that is linear-time in the number of nodes and edges and determines a topological ordering of a given DAG.

it is optimal for each agent i at any given time t to reveal her belief about the unknown state as $\sum_{\hat{\theta} \in \Theta} u_i(a_i, \hat{\theta}) \mu_{i,t}(\hat{\theta})$ in (2.1.2) is uniquely maximized over $a_i \in \mathcal{A}_i$ by $\mathbf{a}_{i,t} = (\mu_{i,t}(\theta_1), \dots, \mu_{i,t}(\theta_m))^T$.

Thence under the prescribed quadratic utility and by taking actions over the probability simplex, agents announce their beliefs truthfully at every epoch. Such agents engage in the discussion by repeatedly exchanging their beliefs about an issue of common interest, which is modeled by the state θ . For example in the course of a political debate, Θ can be the set of all political parties and it would take a binary value in a bipartisan system (cf. Fig. 1.3, on the left). The value/identity of θ is not known to the agents but they each receive a private signal about the unknown θ and starting from a full-support prior belief $\nu(\cdot)$, at any time $t \in \{0, 1, 2, \dots\}$ each agent holds a belief $\mu_{i,t}$, which is her Bayesian posterior on Θ given her knowledge of the signal structure and priors as well as her history of observations, which include her initial private signal as well as the beliefs that she has observed in her neighbors throughout past times $\tau < t$.

Consider the finite state space $\Theta = \{\theta_1, \dots, \theta_m\}$ and for all $2 \leq k \leq m$, let:

$$\Phi_{i,t}(\theta_k) := \log \left(\frac{\mu_{i,t}(\theta_k)}{\mu_{i,t}(\theta_1)} \right), \quad \beta(\theta_k) := \log \left(\frac{\nu(\theta_k)}{\nu(\theta_1)} \right). \quad (2.2.1)$$

Moreover, for any $\theta_k \in \Theta$ let $\lambda_{\theta_k} : \cup_{i \in [n]} \mathcal{S}_i \rightarrow \mathbb{R}$ be the real valued function measuring log-likelihood ratio of the signal s_i under states θ_k and θ_1 , defined as $\lambda_{\theta_k}(s_i) := \log(\ell_i(s_i|\theta_k)/\ell_i(s_i|\theta_1))$. This is a measure of the information content that the signal s_i provides for distinguishing any state θ_k from θ_1 . Here and throughout Sections 2.2 and 2.2.3, we assume the agents have started from uniform prior beliefs and the size of the state space is $m = 2$, thence we enjoy a slightly simpler notation: with uniform priors $\beta(\theta_k) = \log(\nu(\theta_k)/\nu(\theta_1)) = 0$ for all i, k , whereas otherwise knowing the (common) priors the agents can always compensate for the effect of the priors as they observe each other's beliefs; with a binary state space $\Theta = \{\theta_1, \theta_2\}$, the agents need to only keep track of one set of belief and likelihood ratios corresponding to the pair (θ_1, θ_2) , whereas in general the agents should form and calculate $m - 1$ ratio terms for each of the pairs (θ_1, θ_k) , $k = 2, \dots, m$ to have a fully specified belief. For a binary state space with no danger of confusion we can use the simplified notation $\lambda_i = \lambda_{\theta_2}(s_i) := \log(\ell_i(s_i|\theta_2)/\ell_i(s_i|\theta_1))$, and $\Phi_{i,t} = \Phi_{i,t}(\theta_2) = \log(\mu_{i,t}(\theta_2)/\mu_{i,t}(\theta_1))$.

Problem 2.1 (GROUP-DECISION). *At any time t , given the graph structure \mathcal{G} , the private signal s_i and the history of observed neighboring beliefs $\mu_{j,\tau}$, $j \in \mathcal{N}_i$, $\tau \in [t]$ determine the Bayesian posterior belief $\mu_{i,t+1}$.*

In general GROUP-DECISION is a hard problem as we will describe in Subsection 2.2.3. Here, we introduce a special class of structures which play an important role in determining the type of calculations that agent i should undertake to determine her posterior

belief (recall that the t -radius ego-net of agent i , $\tilde{\mathcal{N}}_i^t$, is the set of all agents who are within distance t of or closer to agent i):

Definition 2.2 (Transparency). *The graph structure \mathcal{G} is transparent to agent i at time t , if for all $j \in \mathcal{N}_i$ and every $\tau \leq t - 1$ we have that: $\Phi_{j,\tau} = \sum_{k \in \tilde{\mathcal{N}}_j^\tau} \lambda_k$, for any choice of signal structures and all possible initial signals.*

The initial belief exchanges reveal the likelihoods of the private signals in the neighboring agents. Hence, from her observations of the beliefs of her neighbors at time zero $\{\mu_{j,0}, j \in \mathcal{N}_i\}$, agent i learns all that she ever needs to know regarding the private signals of her neighbors so far as they influence her beliefs about the unknown state θ . However, the future neighboring beliefs (at time one and beyond) are less “transparent” when it comes to reflecting the neighbors’ knowledge of other private signals that are received throughout the network. In particular, the time one beliefs of the neighbors $\Phi_{j,1}, j \in \mathcal{N}_i$ is given by $\Phi_{j,1} = \sum_{k \in \tilde{\mathcal{N}}_j^1} \lambda_k$; hence, from observing the time one belief of a neighbor, agent i would only get to know $\sum_{k \in \tilde{\mathcal{N}}_j^1} \lambda_k$, rather than the individual values of λ_k for each $k \in \mathcal{N}_j$, which her neighbor j had gotten to know before reporting the belief $\Phi_{j,1} = \sum_{k \in \tilde{\mathcal{N}}_j^1} \lambda_k$ to agent i . Indeed, this is a fundamental aspect of inference problems in observational learning (in learning from other actors): similar to responsiveness in [85], which is defined as a property of the utility functions to determine whether players’ beliefs can be inferred from their actions, *transparency* in our belief exchange setup is defined as a property of the graph structure (see Remark 2.4 on why transparency is a structural property) which determines to what extent other players’ private signals can be inferred from observing the neighboring beliefs. We also have the following simple consequence:

Corollary 2.1 (Transparency at time one). *All graphs are transparent to all their agents at time one.*

Remark 2.4 (Transparency, statistical efficiency, and impartial inference). Such agents j whose beliefs satisfy the equation in Definition 2.2 at some time τ are said to hold a *transparent* or *efficient* belief; the latter signifies the fact that the such a belief coincides with the Bayesian posterior if agent j were given direct access to the private signals of every agent in $\tilde{\mathcal{N}}_j^\tau$. This is indeed the best possible (or statistically efficient) belief that agent j can hope to form given the information available to her at time τ ; it specializes the perfect aggregation property of Appendix D to the case of revealed beliefs. The same connection to the statistically efficient beliefs arise in the work of Eyster and Rabin who formulate the closely related concept of “impartial inference” in a model of sequential decisions by different players in successive rounds [97]; accordingly, impartial inference ensures that the full informational content of all signals that influence a player’s beliefs can be extracted and players can fully (rather than partially) infer their predecessors’ signals. In other words,

under impartial inference, players' immediate predecessors provide "sufficient statistics" for earlier movers that are indirectly observed [97, Section 3]. Last but not least, it is worth noting that statistical efficiency or impartial inference are properties of the posterior beliefs, and as such the signal structures may be designed so that statistical efficiency or impartial inference hold true for a particular problem setting; on the other hand, transparency is a structural property of the network and would hold true for any choice of signal structures and all possible initial signals.

The following is a sufficient graphical condition for agent i to hold an efficient (transparent) belief at time t : there are no agents $k \in \tilde{\mathcal{N}}_i^t$ that has multiple paths to agent i , unless it is among her neighbors (agent k is directly observed by agent i).

Proposition 2.2 (Graphical condition for transparency). *Agent i will hold a transparent (efficient) belief at time t if there are no $k \in \tilde{\mathcal{N}}_i^t \setminus \mathcal{N}_i$ such that for $j \neq j'$, both j and j' belonging to \mathcal{N}_i , we have $k \in \tilde{\mathcal{N}}_j^{t-1}$ and $k \in \tilde{\mathcal{N}}_{j'}^{t-1}$.*

Proof. The proof follows by induction on t , i.e. by considering the agents whose information reach agent i for the first time at t . The claim is trivially true at time one, since agent i can always infer the likelihoods of the private signals of each of her neighbors by observing their beliefs at time one. Now consider the belief of agent i at time t , the induction hypothesis implies that $\boldsymbol{\phi}_{i,t-1} = \sum_{k \in \tilde{\mathcal{N}}_i^{t-1}} \lambda_k$, as well as $\boldsymbol{\phi}_{j,t-1} = \sum_{k \in \tilde{\mathcal{N}}_j^{t-1}} \lambda_k$ and $\boldsymbol{\phi}_{j,t-2} = \sum_{k \in \tilde{\mathcal{N}}_j^{t-2}} \lambda_k$ for all $j \in \mathcal{N}_i$. To form her belief at time t (or equivalently its log-ratio $\boldsymbol{\phi}_{i,t}$), agent i should consider her most recent information $\{\boldsymbol{\phi}_{j,t-1} = \sum_{k \in \tilde{\mathcal{N}}_j^{t-1}} \lambda_k, j \in \mathcal{N}_i\}$ and use that to update her current belief $\boldsymbol{\phi}_{i,t-1} = \sum_{k \in \tilde{\mathcal{N}}_i^{t-1}} \lambda_k$. To prove the induction claim, it suffices to show that agent i has enough information to calculate the sum of log-likelihood ratios of all signals in her t -radius ego-net, $\tilde{\mathcal{N}}_i^t$; i.e. to form $\boldsymbol{\phi}_{i,t} = \sum_{k \in \tilde{\mathcal{N}}_i^t} \lambda_k$. This is the best possible belief that she can hope to achieve at time t , and it is the same as her Bayesian posterior, had she direct access to the private signals of all agents in her t -radius ego-net. To this end, by using her knowledge of $\boldsymbol{\phi}_{j,t-1}$ and $\boldsymbol{\phi}_{j,t-2}$ she can form:

$$\hat{\boldsymbol{\phi}}_{j,t-1} = \boldsymbol{\phi}_{j,t-1} - \boldsymbol{\phi}_{j,t-2} = \sum_{k \in \tilde{\mathcal{N}}_j^{t-1} \setminus \tilde{\mathcal{N}}_j^{t-2}} \lambda_k,$$

for all $j \in \mathcal{N}_i$. Since, $\boldsymbol{\phi}_{i,t-1} = \sum_{k \in \tilde{\mathcal{N}}_i^{t-1}} \lambda_k$ by the induction hypothesis, the efficient belief $\boldsymbol{\phi}_{i,t} = \sum_{k \in \tilde{\mathcal{N}}_i^t} \lambda_k$ can be calculated if and only if,

$$\hat{\boldsymbol{\phi}}_{i,t} = \boldsymbol{\phi}_{i,t} - \boldsymbol{\phi}_{i,t-1} = \sum_{k \in \tilde{\mathcal{N}}_i^t \setminus \tilde{\mathcal{N}}_i^{t-1}} \lambda_k, \quad (2.2.2)$$

can be computed. In the above formulation $\hat{\boldsymbol{\phi}}_{i,t}$ is an innovation term, representing the information that agent i learns from her most recent observations at time t . We now show

that under the assumption that any agent with multiple paths to an agent i is directly observed by her, the innovation term in (2.2.2) can be constructed from the knowledge of $\Phi_{j,t-1} = \sum_{k \in \bar{\mathcal{N}}_j^{t-1}} \lambda_k$, and $\Phi_{j,t-2} = \sum_{k \in \bar{\mathcal{N}}_j^{t-2}} \lambda_k$ for all $j \in \mathcal{N}_i$; indeed, we show that:

$$\hat{\Phi}_{i,t} = \sum_{j \in \mathcal{N}_i} \left(\hat{\Phi}_{j,t-1} - \sum_{\substack{k \in \mathcal{N}_i: \\ \nabla(k,j)=t-1}} \Phi_{k,0} \right), \text{ for all } t > 1. \quad (2.2.3)$$

Consider any $k \in \bar{\mathcal{N}}_i^t \setminus \bar{\mathcal{N}}_i^{t-1}$, these are all agents which are at distance exactly t , $t > 1$, from agent i , and no closer to her. No such $k \in \bar{\mathcal{N}}_i^t \setminus \bar{\mathcal{N}}_i^{t-1}$ is a direct neighbor of agent i and the structural assumption therefore implies that there is a unique neighbor of agent i , call this unique neighbor $j_k \in \mathcal{N}_i$, satisfying $k \in \bar{\mathcal{N}}_{j_k}^{t-1} \setminus \bar{\mathcal{N}}_{j_k}^{t-2}$. On the other hand, consider any $j \in \mathcal{N}_i$ and some $k \in \bar{\mathcal{N}}_j^{t-1} \setminus \bar{\mathcal{N}}_j^{t-2}$, such an agent k is either a neighbor of i or else at distance exactly $t > 1$ from agent i and therefore $k \in \bar{\mathcal{N}}_i^t \setminus \bar{\mathcal{N}}_i^{t-1}$, and element j would be the unique neighbor $j_k \in \mathcal{N}_i$ satisfying $k \in \bar{\mathcal{N}}_{j_k}^{t-1} \setminus \bar{\mathcal{N}}_{j_k}^{t-2}$. Subsequently, we can partition

$$\bar{\mathcal{N}}_i^t \setminus \bar{\mathcal{N}}_i^{t-1} = \uplus_{j \in \mathcal{N}_i} \bar{\mathcal{N}}_j^{t-1} \setminus (\bar{\mathcal{N}}_j^{t-2} \cup \mathcal{N}_i),$$

and therefore we can rewrite the left-hand side of (2.2.2) as follows:

$$\begin{aligned} \hat{\Phi}_{i,t} &= \sum_{\substack{k \in \\ \bar{\mathcal{N}}_i^t \setminus \bar{\mathcal{N}}_i^{t-1}}} \lambda_k = \sum_{\substack{k \in \uplus_{j \in \mathcal{N}_i} \\ \bar{\mathcal{N}}_j^{t-1} \setminus (\bar{\mathcal{N}}_j^{t-2} \cup \mathcal{N}_i)}} \lambda_k = \sum_{j \in \mathcal{N}_i} \sum_{\substack{k \in \bar{\mathcal{N}}_j^{t-1} \setminus \\ (\bar{\mathcal{N}}_j^{t-2} \cup \mathcal{N}_i)}} \lambda_k \\ &= \sum_{j \in \mathcal{N}_i} \left(\sum_{\substack{k \in \bar{\mathcal{N}}_j^{t-1} \setminus \\ \bar{\mathcal{N}}_j^{t-2}}} \lambda_k - \sum_{\substack{k \in \mathcal{N}_i \cap \\ \bar{\mathcal{N}}_j^{t-1} \setminus \bar{\mathcal{N}}_j^{t-2}}} \lambda_k \right) = \sum_{j \in \mathcal{N}_i} \left(\hat{\Phi}_{j,t-1} - \sum_{\substack{k \in \mathcal{N}_i: \\ \nabla(k,j)=t-1}} \Phi_{k,0} \right), \end{aligned}$$

as claimed in (2.2.3), completing the proof. \square

Note that in the course of the proof of Proposition 2.2, for the structures that satisfy the sufficient condition for transparency, we obtain a simple algorithm for updating beliefs by setting the total innovation at every step equal to the sum of the most recent innovations observed at each of the neighbors, correcting for those neighbors who are being recounted:

1. Initialize: $\Phi_{i,0} = \lambda_i$, $\hat{\Phi}_{j,0} = \Phi_{j,0} = \lambda_j$, $\Phi_{i,1} = \sum_{j \in \bar{\mathcal{N}}_i^1} \Phi_{j,0}$.

2. For $t > 1$ set:

$$\begin{aligned}\hat{\Phi}_{j,t-1} &= \Phi_{j,t-1} - \Phi_{j,t-2}, \\ \hat{\Phi}_{i,t} &= \sum_{j \in \mathcal{N}_i} [\hat{\Phi}_{j,t-1} - \sum_{\substack{k \in \mathcal{N}_i: \\ \nabla(k,j)=t-1}} \Phi_{k,0}], \\ \Phi_{i,t} &= \Phi_{i,t-1} + \hat{\Phi}_{i,t}.\end{aligned}$$

Rooted (directed) trees are a special class of transparent structures, which also satisfy the sufficient structural condition of Proposition 2.2; indeed, in case of a rooted tree for any agent k that is indirectly observed by agent i , there is a unique path connecting k to i . As such the correction terms for the sum of innovations observed in the neighbors is always zero, and we have $\hat{\Phi}_{i,t} = \sum_{j \in \mathcal{N}_i} \hat{\Phi}_{j,t-1}$, i.e. the innovation at every time step is equal to the total innovations observed in all the neighbors.

Example 2.5 (Transparent structures). Fig. 2.4 illustrates cases of transparent and non-transparent structures. We refer to them as first, second, third, and fourth in their respective order from left to right. All structures except the first one are transparent. To see how the transparency is violated in the first structure, consider the beliefs of agent i : $\Phi_{i,0} = \lambda_i$, $\Phi_{i,1} = \lambda_i + \lambda_{j_1} + \lambda_{j_2}$; at time two, agent 1 observes $\Phi_{j_1,1} = \lambda_{j_1} + \lambda_{k_1} + \lambda_{k_2}$ and $\Phi_{j_2,1} = \lambda_{j_2} + \lambda_{k_2} + \lambda_{k_3}$. Knowing $\Phi_{j_1,0} = \lambda_{j_1}$ and $\Phi_{j_2,0} = \lambda_{j_2}$ she can infer the value of the two sub-sums $\lambda_{k_1} + \lambda_{k_2}$ and $\lambda_{k_1} + \lambda_{k_3}$, but there is no way for her to infer their total sum $\lambda_{j_1} + \lambda_{j_2} + \lambda_{k_1} + \lambda_{k_2} + \lambda_{k_3}$. Agent i cannot hold an efficient or transparent belief at time two. The issue is resolved in the second structure by adding a direct link so that agent k_2 is directly observed by agent i ; the sufficient structural condition of Proposition 2.2 is thus satisfied and we have $\Phi_{i,2} = \lambda_i + \lambda_{j_1} + \lambda_{j_2} + \lambda_{k_1} + \lambda_{k_2} + \lambda_{k_3}$. In structure three, we have $\Phi_{i,2} = \lambda_i + \lambda_{j_1} + \lambda_{j_2} + \lambda_{k_1} + \lambda_{k_2} = \lambda_i + \Phi_{j_1,1} + \Phi_{j_2,0}$. Structure four is also transparent and we have $\Phi_{i,2} = \lambda_i + \lambda_{j_1} + \lambda_{j_2} + \lambda_{k_1} + \lambda_{k_2} + \lambda_{k_3} + \lambda_{k_4} = \lambda_i + \Phi_{j_1,1} + \Phi_{j_2,1}$ and $\Phi_{i,3} = \lambda_i + \lambda_{j_1} + \lambda_{j_2} + \lambda_{k_1} + \lambda_{k_2} + \lambda_{k_3} + \lambda_{k_4} + \lambda_l = \lambda_i + \Phi_{j_1,1} + \Phi_{j_2,1} + (\Phi_{j_1,2} - \Phi_{j_1,1})$, where in the last equality we have used the fact that $\lambda_l = (\Phi_{j_1,2} - \Phi_{j_1,1})$. In particular, note that structures three and four violate the sufficient structural condition laid out in Proposition 2.2, despite both being transparent.




Figure 2.4: The last three structures are transparent but the first one is not.

When the transparency condition is violated, the neighboring agent’s beliefs is a complex non-linear function of the signal likelihoods of the upstream (indirectly observed) neighbors. Therefore, making inferences about the unobserved private signals from such “*nontransparent*” beliefs is a very complex task: it ultimately leads to agent i reasoning about feasible signal profiles that are consistent with her observations similar to the IEIS algorithm (A1:BAYES-GROUP). We elaborate on the belief calculations for the nontransparent case in Subsection 2.2.2, where we provide the version of IEIS algorithm that is tailored to belief communications and it can be used in the most general cases with non-transparent structures. When the transparency conditions are satisfied, the beliefs of the neighboring agents reveal the sum of log-likelihoods for the private signals of other agents within a distance t of agent i . Nevertheless, even when the network is transparent to agent i , cases arise where efficient algorithms for calculating Bayesian posterior beliefs for agent i are unavailable and indeed impossible (if $\mathcal{P} \neq \mathcal{NP}$). In Subsection 2.2.1, we describe the calculations of the Bayesian posterior belief when the transparency condition is satisfied. In Subsection 2.2.3, we show that well-known \mathcal{NP} -complete problems are special cases of the GROUP-DECISION problem and as such the latter is \mathcal{NP} -hard; there we also describe special cases where a more positive answer is available and provide an efficient algorithm accordingly.

2.2.1 Belief Calculations in Transparent Structures

Here we describe calculations of a Bayesian agent in a transparent structure. If the network is transparent to agent i , she has access to the following information from the beliefs that she has observed in her neighbors at times $\tau \leq t$, before deciding her belief for time $t + 1$:

- Her own signal \mathbf{s}_i and its log-likelihood λ_i .
- Her observations of the neighboring beliefs: $\{\mu_{j,\tau} : j \in \mathcal{N}_i, \tau \leq t\}$. Due to transparency, these beliefs reveal the following information about sums of log-likelihoods of private signals of subsets of other agents in the network: $\sum_{k \in \mathcal{N}_j^\tau} \lambda_k = \Phi_{i,\tau}$, for all $\tau \leq t$, and any $j \in \mathcal{N}_i$.

From the information available to her, agent i aims to learn as much as possible about the likelihoods of the private signals of others whom she does not observe; indeed, as she has already learned the likelihoods of the signals that her neighbors have observed from their reported beliefs at time one, at times $t > 1$ she is interested in learning about the agents that are further away from her up to the distance t . Her best hope for time $t + 1$ is to learn the sum of log-likelihoods of the signals of all agents that are within distance of at most $t + 1$ from her in the graph and to set her posterior belief accordingly; this however is not always possible as demonstrated for agent i in the leftmost graph of Fig. 2.4. To decide her belief, agent i constructs the following system of linear equations in $\text{card}(\bar{\mathcal{N}}_{t+1}) + 1$ unknowns: $\{\lambda_j : j \in \bar{\mathcal{N}}_{t+1}, \text{ and } \bar{\lambda}_{i,t+1}\}$, where $\bar{\lambda}_{i,t+1} = \sum_{j \in \bar{\mathcal{N}}_{t+1}} \lambda_j$ is the best possible (statistically efficient) belief for agent i at time $t + 1$:

$$\begin{cases} \sum_{k \in \bar{\mathcal{N}}_j^\tau} \lambda_k = \Phi_{j,\tau}, \text{ for all } \tau \leq t, \text{ and any } j \in \mathcal{N}_i, \\ \sum_{j \in \bar{\mathcal{N}}_i^{t+1}} \lambda_j - \bar{\lambda}_{i,t+1} = 0. \end{cases} \quad (2.2.4)$$

Agent i can apply the Gauss-Jordan method and convert the system of linear equations in $\text{card}(\bar{\mathcal{N}}_i^{t+1}) + 1$ variables to its reduced row echelon form. Next if in the reduced row echelon form $\bar{\lambda}_{i,t}$ is a basic variable with fixed value (its corresponding column has a unique non-zero element that is a one, and that one belongs to a row with all zero elements except itself), then she sets her belief optimally such that $\Phi_{i,t+1} = \bar{\lambda}_{i,t+1}$; this is the statistically efficient belief at time $t + 1$. Recall that in the case of a binary state space, log-belief ratio $\Phi_{i,t+1}$ uniquely determines the belief $\mu_{i,t+1}$.

Remark 2.6 (Statistical versus computational efficiency). Having $\Phi_{i,t+1} = \bar{\lambda}_{i,t+1}$ signifies the best achievable belief given the observations of the neighboring beliefs as it corresponds to the statistically efficient belief that the agent would have adopted, had she direct access to the private signals of every agent within distance $t + 1$ from her; notwithstanding the efficient case $\Phi_{i,t+1} = \bar{\lambda}_{i,t+1}$ does not necessarily imply that agent i learns the likelihoods of the signals of other agents in $\bar{\mathcal{N}}_i^{t+1}$; indeed, this was the case for agent i in the forth (transparent) structure of Example 2.5: agent i learns $\{\lambda_i, \lambda_{j_1}, \lambda_{j_2}, \lambda_{k_1} + \lambda_{k_2}, \lambda_{k_3} + \lambda_{k_4}, \lambda_l\}$ and in particular can determine the efficient beliefs $\bar{\lambda}_{i,2} = \lambda_i + \lambda_{j_1} + \lambda_{j_2} + \lambda_{k_1} + \lambda_{k_2} + \lambda_{k_3} + \lambda_{k_4}$ and $\bar{\lambda}_{i,3} = \lambda_i + \lambda_{j_1} + \lambda_{j_2} + \lambda_{k_1} + \lambda_{k_2} + \lambda_{k_3} + \lambda_{k_4} + \lambda_l$, but she never learns the actual values of the likelihoods $\{\lambda_{k_1}, \lambda_{k_2}, \lambda_{k_3}, \lambda_{k_4}\}$, individually. In other words, it is possible for agent i to determine the sum of log-likelihoods of signals of agents in her higher-order neighborhoods even though she does not learn about each signal likelihood individually. The case where $\bar{\lambda}_{i,t+1}$ can be determined uniquely so that $\Phi_{i,t+1} = \bar{\lambda}_{i,t+1}$, is not only statistically efficient but also computationally efficient as complexity of determining the Bayesian posterior belief at time $t + 1$ is the same as the complexity of performing Gauss-Jordan steps which is $O(n^3)$ for solving the $t \cdot \text{card}(\mathcal{N}_i)$ equations in $\text{card}(\bar{\mathcal{N}}_i^{t+1})$ unknowns. Note

that here we make no attempt to optimize these computations beyond the fact that their growth is polynomial in n . This is an interesting alignment that emerges between statistical and computational efficiency in group decision process, and it is in contrast with the trade-off between statistical and computational performance that is reported in other graphical inference problems such as sparse principal component analysis, planted partition and stochastic block models, as well as sub-matrix localization, where there is an “*information-computation gap*” between what is achievable in polynomial-time and what is statistically optimal (achieves the information theoretic limit); cf. [98, 99].

Next we consider the case where $\bar{\lambda}_{i,t+1}$ is not a basic variable in the reduced row echelon form of system (2.2.4) or it is a basic variable but its value is not fixed by the system and depends on how the free variables are set. In such cases agent i does not have access to the statistically efficient belief $\bar{\lambda}_{i,t+1}$. Instead she has to form her Bayesian posterior belief by inferring the set of all feasible signals for all agents in $\bar{\mathcal{N}}_i^{t+1}$ whose likelihoods are consistent with the system (2.2.4). To this end, she keeps track of the set of all signal profiles at any time t that are consistent with her information, system (2.2.4), at that time. Following the IEIS procedure of Section 2.1.1, let us denote the set of feasible signal profiles for agent i at time t by $\mathcal{I}_{i,t}$. The general strategy of agent i , would be to search over all elements of $\mathcal{I}_{i,t}$ and to eliminate (refute) any signal profile \bar{s} that is inconsistent with (i.e. does not satisfy) the \mathcal{N}_i new equations revealed to her from the transparent beliefs of her neighbors. For a signal profile $\bar{s} = (s_1, s_2, \dots, s_n)$, let $\lambda_i(\bar{s}) := \log(\ell_i(s_i|\theta_2)/\ell_i(s_i|\theta_1))$ denote the log-likelihood ratio of its i -th component private signal. Given the list of feasible signal profiles $\mathcal{I}_{i,t}$ for agent i at time t , we formalize the calculations of agent i , subject to observation of the transparent beliefs of her neighbors $\phi_{j,t}, j \in \mathcal{N}_i$, as follows:

(A3: BAYES-TRANSPARENT). Calculations of agent i at time t for deciding $\mu_{i,t+1}$ in a structure that is transparent to her:

1. Initialize $\mathcal{I}_{i,t+1} = \mathcal{I}_{i,t}$.
2. For all $\bar{s} \in \mathcal{I}_{i,t+1}$ and any $j \in \mathcal{N}_i$ do:
 - If $\phi_{j,t} \neq \sum_{k \in \bar{\mathcal{N}}_j^t} \lambda_k(\bar{s})$, then set $\mathcal{I}_{i,t+1} = \mathcal{I}_{i,t+1} \setminus \{\bar{s}\}$.
3. Given $\mathcal{I}_{i,t+1}$, calculate the updated belief $\mu_{i,t+1}$ according to (2.1.1).

Despite the relative simplification that is brought about by transparency, in general there is an exponential number of feasible signal profiles and verifying them for the new \mathcal{N}_i equations would take exponential time. The belief calculations may be optimized by inferring the largest subset of individual likelihood ratios whose summation is fixed by

system (2.2.4). The verification and refutation process can then be restricted to the remaining signals whose sum of log-likelihoods is not fixed by system (2.2.4). For example in leftmost structure of Fig. 2.4, agent i will not hold a transparent belief at time 2 but she can determine the sub-sum $\lambda_i + \lambda_{j_1} + \lambda_{j_2}$ and her belief would involve a search only over the profile of the signals of the remaining agents $(s_{k_1}, s_{k_2}, s_{k_3})$. At time two, she finds all $(s_{k_1}, s_{k_2}, s_{k_3})$ that agree with the additionally inferred sub-sums $\phi_{j_1,1} - \phi_{j_1,0} = \lambda_{k_1} + \lambda_{k_2}$ and $\phi_{j_2,1} - \phi_{j_2,0} = \lambda_{k_2} + \lambda_{k_3}$; indeed we can express $\phi_{i,2}$ as follows:

$$\phi_{i,2} = \lambda_i + \lambda_{j_1} + \lambda_{j_2} + \log \frac{\sum_{(s_{k_1}, s_{k_2}, s_{k_3}) \in \mathcal{I}_{i,2}} \ell_{k_1}(s_{k_1}|\theta_2) \ell_{k_2}(s_{k_2}|\theta_2) \ell_{k_3}(s_{k_3}|\theta_2)}{\sum_{(s_{k_1}, s_{k_2}, s_{k_3}) \in \mathcal{I}_{i,2}} \ell_{k_1}(s_{k_1}|\theta_1) \ell_{k_2}(s_{k_2}|\theta_1) \ell_{k_3}(s_{k_3}|\theta_1)},$$

where

$$\mathcal{I}_{i,2} = \{(s_{k_1}, s_{k_2}, s_{k_3}) : \log \frac{\ell_{k_1}(s_{k_1}|\theta_2)}{\ell_{k_1}(s_{k_1}|\theta_1)} + \log \frac{\ell_{k_2}(s_{k_2}|\theta_2)}{\ell_{k_2}(s_{k_2}|\theta_1)} = \lambda_{k_1} + \lambda_{k_2}, \text{ and} \\ \log \frac{\ell_{k_1}(s_{k_1}|\theta_2)}{\ell_{k_1}(s_{k_1}|\theta_1)} + \log \frac{\ell_{k_3}(s_{k_3}|\theta_2)}{\ell_{k_3}(s_{k_3}|\theta_1)} = \lambda_{k_2} + \lambda_{k_3}\}.$$

Here we make no attempt in optimizing the computations for the refutation process in transparent structures beyond pointing out that they can increase exponentially with the network size. In Subsection 2.2.3, we show that the GROUP-DECISION problem is \mathcal{NP} -hard; and as such there are no algorithms that will scale polynomial in network size for all network structures (unless $\mathcal{P} = \mathcal{NP}$).

When transparency is violated the neighboring beliefs are highly non-linear functions of the log-likelihoods and the forward reasoning approach of (A3: BAYES-TRANSPARENT) can no longer be applied; indeed, when transparency is violated then beliefs represent what signal profiles agents regard as feasible rather than what they know about the log-likelihoods of signals of others whom they have directly or indirectly observed. In particular, the agent cannot use the reported beliefs of the neighbors directly to make inferences about the original causes of those reports which are the private signals. Instead, to keep track of the feasible signal profiles that are consistent with her observations the agent employs a version of the IEIS algorithm of Section 2.1.1 that is tailored to the case of revealed beliefs. We describe these calculations of the Bayesian agents for nontransparent structures in Subsection 2.2.2.

2.2.2 Belief Calculations in Nontransparent Structures

In general nontransparent structures where one or more of the neighboring beliefs do not satisfy the transparency conditions in Definition 2.2, agent i would have to follow an IEIS strategy similar to (A1: BAYES-GROUP) to construct her Bayesian posterior belief given

her observations of her neighbors' nontransparent beliefs. Accordingly, as in Table 2.1, for every profile of initial signals $\bar{s} = (s_1, s_2, \dots, s_n)$ she constructs a list of all signal profiles that she believes are deemed feasible by another agent j , given what she believes agent j may have observed up until time t conditioned on the initial signals being prescribed by \bar{s} . Subsequently, the information available to her at time t is the same as that in (I1: BAYES-GROUP); and she uses this information to update her list of feasible signal profiles from $\mathcal{I}_{i,t}$ to $\mathcal{I}_{i,t+1}$. Before presenting the exact calculations for determining the Bayesian posterior of agent i , note that rather than the conditionally feasible actions for each agent j , given by $\alpha_{j,t}^{(i)}(\bar{s})$ in Table 2.1, agent i in the case of revealed beliefs would instead keep track of $\bar{\mu}_{j,t}^{(i)}(\bar{s}) = (\mu_{j,t}^{(i)}(\bar{s}; \theta_1), \dots, \mu_{j,t}^{(i)}(\bar{s}; \theta_m))$, i.e. the belief that she deems optimal for each agent j , given what she believes agent j could have observed up until time t conditioned on the event that the initial signals of all agents are prescribed according to \bar{s} . Note that following (2.1.4), we have:

$$\mu_{j,\tau}^{(i)}(\bar{s}; \theta_k) = \frac{\sum_{\bar{s}' \in \mathcal{I}_{j,\tau}^{(i)}(\bar{s})} \mathcal{P}_{\theta_k}(\bar{s}') \nu(\theta_k)}{\sum_{l=1}^m \sum_{\bar{s}' \in \mathcal{I}_{j,\tau}^{(i)}(\bar{s})} \mathcal{P}_{\theta_l}(\bar{s}') \nu(\theta_l)}. \quad (2.2.5)$$

Calculations of agent i at time t enables her to update her information at time t to incorporate her newly obtained data which constitute her observations of her neighbors' most recent beliefs $\bar{\mu}_{j,t}$ for all $j \in \mathcal{N}_i$; whence she refines $\mathcal{I}_{i,t}$ to $\mathcal{I}_{i,t+1}$ and updates her belief using (2.1.1). This is achieved as follows:

(A4: BAYES-NONTRANSPARENT). Calculations of agent i at time t for deciding her Bayesian posterior $\bar{\mu}_{i,t+1}$:

1. For all $\bar{s} := (s_1, \dots, s_n) \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ and all $j \in \mathcal{N}_i^{t+1}$ do:
 - If $\nabla(j, i) = t + 1$, initialize $\mathcal{I}_{j,0}^{(i)}(\bar{s}) = \{s_j\} \times \prod_{k \neq j} \mathcal{S}_k$.
 - Else initialize $\mathcal{I}_{j,t+1-\nabla(j,i)}^{(i)}(\bar{s}) = \mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$ and for all $\tilde{s} \in \mathcal{I}_{j,t+1-\nabla(j,i)}^{(i)}(\bar{s})$ do:
 - For all $k \in \mathcal{N}_j$ if $\bar{\mu}_{k,t-\tau}^{(i)}(\tilde{s}) \neq \bar{\mu}_{k,t-\tau}^{(i)}(\bar{s})$, then $\mathcal{I}_{j,t+1-\tau}^{(i)}(\bar{s}) = \mathcal{I}_{j,t+1-\tau}^{(i)}(\bar{s}) \setminus \{\tilde{s}\}$.
2. Initialize $\mathcal{I}_{i,t+1} = \mathcal{I}_{i,t}$ and for all $\bar{s} \in \mathcal{I}_{i,t+1}$ do:
 - For all $j \in \mathcal{N}_i$ if $\bar{\mu}_{j,t} \neq \bar{\mu}_{j,t}^{(i)}(\bar{s})$, then $\mathcal{I}_{i,t+1} = \mathcal{I}_{i,t+1} \setminus \{\bar{s}\}$.
3. Given $\mathcal{I}_{i,t+1}$, calculate the updated belief $\bar{\mu}_{i,t+1}$ according to (2.1.1).

Following Appendix A.1, we know the computational complexity of (A4: BAYES - NONTRANSPARENT) increases exponentially in n and can be bounded as $O(n^2 M^{2n-1} m)$. There is a key difference between the refutation process in step (2) of (A3: BAYES-TRANSPARENT), when the network is transparent to agent i , and the pruning that takes place in step (2) of the (A4 : BAYES- NONTRANSPARENT) for general (non-transparent) networks. In the latter case, the agent needs to consider the beliefs of other far way agents at any possible signal profile and to simulate the subsequent beliefs of her neighbors conditioned on the particular signal profile; cf. step (1) of (A4: BAYES-NONTRANSPARENT). Each signal profile will be rejected and removed from the feasible set if the simulated belief of a neighbor conditioned on that signal profile does not agree with the actual (observed) beliefs at that time. On the other hand, in a transparent structure, the agent does not need to simulate the beliefs of other agents conditioned on a signal profile to investigate its feasibility; compare step (1) of (A3: BAYES-TRANSPARENT) with step (1) of (A4: BAYES-NONTRANSPARENT). She can directly verify whether the individual signals likelihoods satisfy the most recent set of constraints that are revealed to the agent at time t from the transparent beliefs of her neighbors; and if any one of the new equations is violated, then that signal profile will be rejected and removed from the feasible set. This constitutes an interesting bridge between statistical and computational efficiency in group decision processes.

2.2.3 Hardness of GROUP-DECISION

In this subsection we prove:

Theorem 2.7 (Hardness of GROUP-DECISION). *The GROUP-DECISION (Problem 1) is \mathcal{NP} -hard.*

We provide two reductions for the proof of Theorem 2.7, one reduction is to the SUBSET-SUM problem and the other reduction is to the EXACT-COVER problem. In both reductions, we use binary signal spaces for all the agents; however, the first reduction requires agents to receive different signals with non-identical probabilities whose variety increases with the network size so that we can accommodate the increasing set size. Our second reduction, on the other hand, works with i.i.d. binary signals but still relies on a complex network structure (with unbounded degrees, i.e. node degrees increase with the increasing network size) to realize arbitrary instances of EXACT-COVER.

The particular structures in which the two problems are realized are depicted in Fig. 2.5. The graph on the left is used for the SUBSET-SUM reduction and the graph on the right is used for the EXACT-COVER problem. The SUBSET-SUM problem asks if given a set of n positive integers p_1, \dots, p_n and another positive integer q , there is a non-empty subset

of $\{p_1, \dots, p_n\}$ that sum to q . We encode the n parameters $\{p_1, \dots, p_n\}$ of the SUBSET-SUM problem using the log-likelihood ratios of binary signals that the n agents in the top layer of the left graph in Fig. 2.5 receive. The encoding is such that $p_h = \bar{p}_h - \underline{p}_h$ for all $h \in [n]$, where \bar{p}_h and \underline{p}_h are the log-likelihood ratios of the one and zero signals for each of the n agents $l_h, h \in [n]$. Throughout this section and when working with binary signals, we use the over and under bars to indicate the log-likelihood ratios of the one and zero signals, respectively. Similarly, we denote the log-likelihood ratios of the signals of the two agents k_1 and k_2 by \bar{p}^* and \underline{p}^* and set them such that $-q = \bar{p}^* - \underline{p}^*$. The crux of our reduction is in designing the aggregate beliefs of agents j_1 and j_2 at time one in such a way that agent i needs to decide whether the observed aggregates are caused by all of the indirectly observed agents l_1, \dots, l_n and k_1, k_2 having reported zero signals to j_1 and j_2 ; or else it is possible that the contributions from some of the one signals among l_1, \dots, l_n is canceled out in the aggregate by the one signals in k_1 and k_2 . In the latter case, those agents, l_h , who have received one signals, $s_{l_h} = 1$, constitute a feasibility certificate for the SUBSET-SUM problem, as their respective values of p_h sum to q . In Appendix A.3, we show that the decision problem of agent i in the designed scenario (after her observations of the beliefs of j_1 and j_2) simplifies to the feasibility of the SUBSET-SUM problem with parameters p_1, \dots, p_n and q .

We use the right-hand side structure of Fig. 2.5 for the EXACT-COVER reduction. Unlike the SUBSET-SUM reduction, in the EXACT-COVER reduction we do not rely on unboundedly many types of signals. Given a set of n elements $\{j_1, \dots, j_n\}$ and a family of m subsets $\{l_1, \dots, l_m\}, l_h \subset \{j_1, \dots, j_n\}$ for all $h \in [m]$, the EXACT-COVER problem asks if it is possible to construct a non-intersecting cover (partition) of $\{j_1, \dots, j_n\}$ using a (disjoint) subfamily of $\{l_1, \dots, l_m\}$. Given any instances of the EXACT-COVER problem, we can encode the inclusion relations between the n elements $\{j_1, \dots, j_n\}$ and m subsets $\{l_1, \dots, l_m\}$ using the bipartite graph in the first two layers of the right structure in Fig. 2.5. Here each node represents the respective entity (element j_r or subset $l_h, r \in [n]$ and $h \in [m]$) of the same name: an edge from a node l_h to a node j_r for some $h \in [m]$ and $r \in [n]$ indicates that element j_r is included in the subset l_h . Our strategy is again to take any instance of the EXACT-COVER problem and design the signal structures such that agent i 's belief in the corresponding instance of GROUP-DECISION problem (with the network structure given in the right hand side of Fig. 2.5) would indicate her knowledge of the feasibility of the (arbitrarily chosen) instance of the EXACT-COVER problem (that is encoded by the first two layers of the right hand side graph in Fig. 2.5). We use \bar{p} and \underline{p} for the log-likelihood ratios of the one and zero signals of the l_1, \dots, l_m nodes and set these parameters such that $\bar{p} - \underline{p} = 1$. Similarly, we denote the log-likelihood ratios of the one and zero signals in the node k by \bar{p}^* and \underline{p}^* , and set them such that $\bar{p}^* - \underline{p}^* = -1$. In Appendix A.4, we design a set of observations for agent i such that her belief at time 2 would require her to know whether her observations of the beliefs of j_1, \dots, j_n are caused by all agents

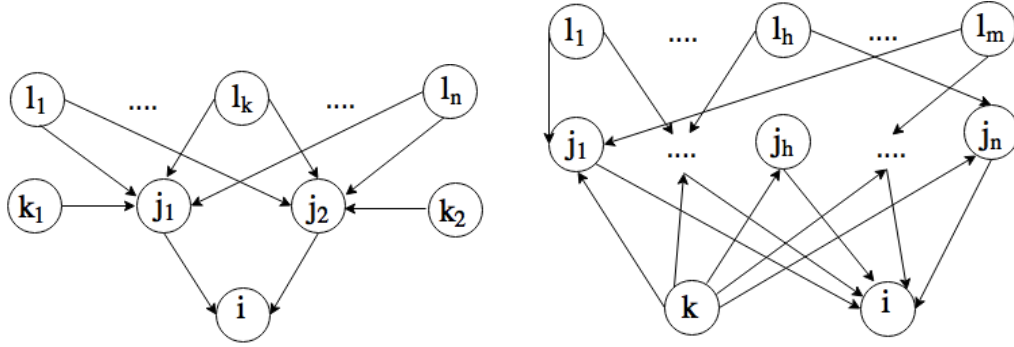


Figure 2.5: The graph structure on the left is used for the SUBSET-SUM reduction. The graph structure on the right is used for the EXACT-COVER reduction. In the particular instance of EXACT-COVER that is depicted on the right, we have that $j_1 \in l_1$ and $j_1 \in l_m$, as the links in the top two layers indicate the inclusion relations among subsets l_1, \dots, l_m and elements j_1, \dots, j_n .

l_1, \dots, l_m as well as agent k having received zero signals, or else whether it is possible that some of the agents among l_1, \dots, l_m have received one signals and their aggregate effects on the beliefs of j_1, \dots, j_n are canceled out by the one signal that agent k has received. The latter happens only when the corresponding instance of the EXACT-COVER problem (coded by the right hand graph of Fig. 2.5) is feasible. In such cases, those sets among l_1, \dots, l_m whose respective agents have receive one signals, $\{l_h : h \in [m], s_{l_h} = 1\}$, constitute a disjoint subfamily that covers $\{j_1, \dots, j_n\}$.

The detailed reductions are presented in Appendices A.3 and A.4. It is worth highlighting that our \mathcal{NP} -hardness reductions show that the GROUP-DECISION problem is hard to solve in the worst case. In other words, there exist network structures and particular profiles of private signals that lead to specific observations of the neighboring beliefs, such that making an inference about the observed beliefs and forming a Bayesian posterior belief conditioned on those observations is not possible in computation times that increase polynomially with the network size (unless $\mathcal{P} = \mathcal{NP}$). Alternatively, one may be interested in the complexity of computations in specific network structures with increasing size, such as trees, cycles, or complete graphs for which we know that beliefs can be computed efficiently by virtue of their transparency. Moreover, one may also be interested in the complexity of computations in an average sense (for “typical” network structures and “typical” private signals). Deriving complexity notions in these alternative settings is much more difficult; indeed, development of such alternative notions of complexity is an active area of research in theory of computation, cf. e.g. [100] for average-case complexity with respect to random inputs, and cf. e.g. [101] for the relevant complexity notions that apply to randomized algorithms.

Remark 2.8 (Beyond \mathcal{NP} -hardness). Both reductions are set up such that the feasibility of the corresponding \mathcal{NP} -complete problem (SUBSET-SUM or EXACT-COVER) is re-

flected in the time-two beliefs of agent i . However, the beliefs in both cases contain more information than the simple yes or no answer to the feasibility questions. Effectually, the information content of beliefs amounts to a weighted sum over all the feasibility certificates (each certificate is represented by a particular signal profile and is weighted by the likelihood of that particular signal profile). One possibility is to prove hardness in a class of functional problems such as $\#\mathcal{P}$. The class $\#\mathcal{P}$ is comprised of the counting problems associated with the problems in \mathcal{NP} . The latter is a class of decision problems for which the positive instances have an efficiently verifiable proof. While \mathcal{NP} captures the difficulty of finding any certificates, the class $\#\mathcal{P}$ captures the difficulty of counting the number of all valid certificates (if any). As such the problems in $\#\mathcal{P}$ are naturally harder than those in \mathcal{NP} (cf. e.g. [102, Chapter 17]).

The EXACT-COVER reduction relies critically on the fact that the number of directly observed neighbors (corresponding to the number of equations in the EXACT-COVER reduction) are allowed to increase. If the number of neighbors is fixed but different agents receive different signals with varying distributions, then our SUBSET-SUM reduction in Appendix A.3 again verifies that the hardness property holds. Our next example shows that either of the two structural features (increasing size of the neighborhood or infinitely many types of private signals among the indirectly observed agents) are needed to obtain a hard problem; indeed, an efficient calculation of beliefs may be possible when the neighborhood sizes are kept fixed and agents receive i.i.d. private signals.

For example in the left structure of Fig. 2.5 we can efficiently compute the beliefs if $\{l_1, \dots, l_n\}$ are receiving i.i.d. binary signals. To see how the belief of agent i at time two can be computed efficiently in the number of indirectly observed neighbors (n), suppose that the signal structures for agent i , her neighboring agents $\{j_1, j_2\}$, and the indirectly observed agents $\{k_1, k_2, l_1, \dots, l_n\}$ are the same as $\{j_h, h \in [n]\}$ and $\{l_h, h \in [m]\}$ in EXACT-COVER reduction of Appendix A.4: $\{i, j_1, j_2\}$ receiving non-informative signals and $\{k_1, k_2, l_1, \dots, l_n\}$ receiving i.i.d. binary signals, whose likelihoods satisfy $\lambda_r = s_r(\bar{p} - \underline{p}) + \underline{p}$ for all $r \in \{k_1, k_2, l_1, \dots, l_n\}$ as in (A.4.1) of Appendix A.4. Subsequently, $\Phi_{i,0} = \Phi_{j_1,0} = \Phi_{j_2,0} = \Phi_{i,1} = 0$, due to their initial noninformative signals. At time two, agent i has to incorporate the time one beliefs of her neighbors, which are themselves caused by the time zero beliefs of $k_1, k_2, l_1, \dots, l_n$: Given $\Phi_{j_r,1} = \lambda_{k_r} + \sum_{h=1}^n \lambda_{l_h}$, for $r = 1, 2$, agent i aims to determine her belief at time two (or equivalently $\Phi_{i,2}$). Using (A.4.1), we can write

$$\psi_{j_r} = s_{k_r} + \sum_{h=1}^n s_{l_h} = \frac{1}{\bar{p} - \underline{p}} (\Phi_{j_r,1} - \text{card}(\mathcal{N}_{j_r}) \underline{p}), r \in \{1, 2\},$$

where ψ_{j_r} are necessarily non-negative integers belonging to $[n+1]_0 = \{0\} \cup [n+1]$, due to their generation process; i.e. the fact that they count the number of one signals that

are received in the neighborhood \mathcal{N}_{j_r} of each of the neighbors $j_r, r = 1, 2$. To proceed, let $\eta \in [n]_0$ be the number of agents among $\{l_1, \dots, l_n\}$ who have received one signals. Depending on the relative values of ψ_{j_1} and ψ_{j_2} three cases may arise, and the subsequent beliefs $\Phi_{i,2}$ in each case are listed below:

1. If $\psi_{j_1} \neq \psi_{j_2}$, then exactly one of the two signals s_{k_1} and s_{k_2} is a one and the other one is zero, the latter corresponding to the lower of the two counts ψ_{j_1} and ψ_{j_2} . We further have that $\eta = \min\{\psi_{j_1}, \psi_{j_2}\}$ and $\Phi_{i,2} = (\eta + 1)\bar{p} + (n - \eta + 1)\underline{p}$.
2. If $\psi_{j_1} = \psi_{j_2} = 0$, then every body in the second order neighborhood of i has received a zero and we have $\Phi_{i,2} = (n + 2)\underline{p}$.
3. If $\psi_{j_1} = \psi_{j_2} \geq 1$, then either $s_{k_1} = s_{k_2} = 0$ and $\eta = \eta_0 = \psi_{j_1}$ or $s_{k_1} = s_{k_2} = 1$ and $\eta = \eta_1 = \psi_{j_1} - 1$. In this case, the belief of agent i at time two is given by: $\Phi_{i,2} = \log(f_{\theta_2}/f_{\theta_1})$, where $f_{\theta_r}, r = 1, 2$ is defined as: $f_{\theta_r} = \binom{n}{\eta_1} \ell(\mathbf{s} = 1|\theta_r)^{\eta_1+2} \ell(\mathbf{s} = 0|\theta_r)^{n-\eta_1} + \binom{m}{\eta_0} \ell(\mathbf{s} = 1|\theta_r)^{\eta_0} \ell(\mathbf{s} = 0|\theta_r)^{n-\eta_0+2}$.

In Appendix A.5, we devise a similar algorithm to calculate the time two belief of agent i in the left-hand-side structure of Fig. 2.5 (by counting the number of agents $h \in [m]$ for which $s_{l_h} = 1$), with time-complexity $O(n2^n m^{2^n} + m^{1+2^n} (2^{2^n})(3m + 2))$: increasing polynomially in m for a fixed neighborhood size (n).

Chapter 3

Heuristic Decision Making in Groups

In this chapter, we propose a no-recall 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 were highlighted in the previous chapter. 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.^{1,2}

¹The following papers cover the results of this chapter: [103–105]. Preliminary versions of the results were also presented in the following non-archival venues: 2015 Workshop on Social and Information Networks (in conjunction with the 16th ACM Conference on Economics and Computation), INFORMS 2015 Annual meeting, and 2016 NBER-NSF Seminar on Bayesian Inference in Econometrics and Statistics.

²I would like to thank Peter Krafft and Armand Makowski for insightful discussions, as well as the participants at the University of Pennsylvania, University of Maryland, Princeton and MIT local seminars for their comments and feedback.

In Section 3.1, we describe the mathematical details of the no-recall model; in particular, we explain the mathematical steps for deriving the so-called Bayesian or no-recall heuristics in a given decision scenario. In Section 3.2, 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 (no-recall) 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. Finally in Section 3.3, 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.

3.1 The No-Recall Model of Group Decision Making

We present the no-recall model of group decision making which explains the operations of system one (the fast/automatic system; cf. Subsection 1.1.2) in a group decision process. This model allows us 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 [106], 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.

Given the initial signal \mathbf{s}_i , agent i forms an initial Bayesian opinion $\mathcal{M}_{i,0}(\cdot)$ about the value of θ and chooses her action $\mathbf{a}_{i,0} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \int_{\Theta} u_i(\mathbf{a}_i, \theta') \mathcal{M}_{i,0}(d\theta')$, 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 $\{\mathbf{a}_{j,0}, j \in \bar{\mathcal{N}}_i\}$ 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:

$$\mathbf{a}_{i,1} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \int_{\Theta} u_i(\mathbf{a}_i, \theta) \mathcal{M}_{i,1}(d\theta), \quad (3.1.1)$$

maximizing her expected pay off conditional on everything that she has observed thus far; i.e. maximizing $\mathbb{E}_{i,1}\{u_i(\mathbf{a}_i, \theta)\} = \mathbb{E}_{\theta}\{u_i(\mathbf{a}_{i,1}, \theta) | \mathbf{s}_i, \mathbf{a}_{j,0} : j \in \bar{\mathcal{N}}_i\} = \int_{\Theta} u_i(\mathbf{a}_i, \theta) \mathcal{M}_{i,1}(d\theta)$. 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 $\mathbf{a}_{j,t-1} : j \in \bar{\mathcal{N}}_i$ 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:

$$f_i(\mathbf{a}_{j,0} : j \in \bar{\mathcal{N}}_i) := \mathbf{a}_{i,1} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{i,1}\{u_i(\mathbf{a}_i, \theta)\}. \quad (3.1.2)$$

derived in (3.1.1), 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: $\mathbf{a}_{i,t} = f_i(\mathbf{a}_{j,t-1} : j \in \bar{\mathcal{N}}_i), \forall t > 1$. We refer to the mapping $f_i : \prod_{j \in \bar{\mathcal{N}}_i} \mathcal{A}_j \rightarrow \mathcal{A}_i$ thus obtained, as a *Bayesian* or *no-recall heuristic*.

Remark 3.1 (“What you see is all there is!”). The heuristics thus obtained suffer from some 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 [9]. 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




Figure 3.1: Heuristic agents ignore the history of interactions (dotted arrows) in their inferences and always attribute actions to (imaginary) private signals; compare with the decision flow diagram for the inferences of two rational agents who communicate their actions repeatedly (Fig. 2.1).

of information; see Fig. 3.1. Such naivety in disregarding the history of interactions has been highlighted in our earlier works on Bayesian learning without recall [107], 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 3.2 (Naive inferences). Similar and related forms of naivety have been suggested in the literature. Eyster and Rabin [26, 97] 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 [108] 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 [109]; they analyze its implications to diffusion of information, and focus in particular, on the voting behavior.

3.2 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 σ -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:

$$\ell(s|\theta; \sigma_i, \delta_i) := \frac{d\mathcal{L}(\cdot|\theta; \sigma_i, \delta_i)}{d\mathcal{G}_s} = \sigma_i \left| \frac{\Lambda_k(\xi(ds))}{\mathcal{G}_s(ds)} \right| \tau(\sigma_i \xi(s), \delta_i) e^{\sigma_i \eta(\theta)^\top \xi(s) - \delta_i \gamma(\eta(\theta))} \quad (3.2.1)$$

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 Θ to \mathbb{R}^k , $\tau : \mathbb{R}^k \times (0, +\infty) \rightarrow (0, +\infty)$ is a positive weighting function, and

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

is a normalization factor that is constant when θ is fixed, even though $\delta_i > 0$ and $\sigma_i > 0$ vary. This normalization constant for each θ is uniquely determined by the functions $\eta(\cdot)$, $\xi(\cdot)$ and $\tau(\cdot)$. The parameter space Θ 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

$$\Omega_\eta := \left\{ \eta \in \mathbb{R}^k : \int_{s \in \mathcal{S}} \left| \frac{\Lambda_k(\xi(ds))}{\mathcal{G}_s(ds)} \right| \tau(\xi(s), 1) e^{\eta^\top \xi(s)} \mathcal{G}_s(ds) < \infty \right\}.$$

In (3.2.1), $\sigma_i > 0$ and $\delta_i > 0$ for each i are scaling factors that determine the quality or informativeness of the random sample s_i with regard to the unknown θ : fixing either one of the two factors σ_i or δ_i , the value of the other one increases with the increasing informativeness of the observed value $\xi(s_i)$. The following conjugate family of priors¹ are associated with the likelihood structure (3.2.1). 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 (α, β) , $\alpha \in \mathbb{R}^k$ and $\beta > 0$:

¹Consider a parameter space Θ , a sample space \mathcal{S} , and a sampling distribution $\mathcal{L}(\cdot|\theta) \in \Delta\mathcal{S}$, $\theta \in \Theta$. Suppose that s is a random variable which is distributed according to $\mathcal{L}(\cdot|\theta)$ for any θ . 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 $s = s$ belongs to \mathcal{F} .

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

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 θ) of the sufficient statistic $\xi(\mathbf{s}_i)$. Hence, we endow every agent $i \in [n]$ with the quadratic utility $u_i(\mathbf{a}, \theta) = -(\mathbf{a} - \mathbf{m}_{i,\theta})^\top (\mathbf{a} - \mathbf{m}_{i,\theta})$, $\forall \mathbf{a} \in \mathcal{A}_i = \mathbb{R}^k$, where $\mathbf{m}_{i,\theta} := \mathbb{E}_{i,\theta}\{\xi(\mathbf{s}_i)\} := \int_{s \in \mathcal{S}} \xi(s) \mathcal{L}(ds|\theta; \sigma_i, \delta_i) \in \mathbb{R}^k$.

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 $\mathbf{m}_{i,\theta}$. These settings are formalized under the following assumption that we term the *Exponential Family Signal-Utility Structure*.

Assumption 3.1 (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 (3.2.1) 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 $\mathbf{a} \in \mathcal{A}_i = \mathbb{R}^k$ and bears the quadratic utility $u_i(\mathbf{a}, \theta) = -(\mathbf{a} - \mathbf{m}_{i,\theta})^\top (\mathbf{a} - \mathbf{m}_{i,\theta})$, where $\mathbf{m}_{i,\theta} := \mathbb{E}_{i,\theta}\{\xi(\mathbf{s}_i)\} := \int_{s \in \mathcal{S}} \xi(s) \mathcal{L}(ds|\theta; \sigma_i, \delta_i) \in \mathbb{R}^k$.

The Bayesian heuristics $f_i(\cdot)$, $i \in [n]$ under the settings prescribed by the exponential family signal-utility structure (Assumption 3.1) are linear functions of the neighboring actions with specified coefficients that depend only on the likelihood structure parameters: n_i , σ_i and δ_i as well as the prior parameters: α_i and β_i , for all $i \in [n]$.

Theorem 3.3 (Affine action updates). *Under the exponential family signal-utility structure specified in Assumption 3.1, 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 δ_i are as follows:

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

The action profile at time t is the concatenation of all actions in a column vector: $\bar{\mathbf{a}}_t = (\mathbf{a}_{1,t}^\top, \dots, \mathbf{a}_{n,t}^\top)^\top$. The matrix T with entries T_{ij} , $i, j \in [n]$ given in Theorem 3.3 is called the social influence matrix. The constant terms ϵ_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 $\bar{\epsilon} = (\epsilon_1^\top, \dots, \epsilon_n^\top)^\top$ and refer to it as the vector of neighborhood biases. The evolution of action profiles under conditions of Theorem 3.3 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. Subsequently, the evolution of action profiles over time follows a non-homogeneous positive linear discrete-time dynamics, cf. [110]. 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 $\bar{\mathbf{a}}_e = ((I - T)^{-1} \otimes I_k) \bar{\epsilon}$ and $\lim_{t \rightarrow \infty} \bar{\mathbf{a}}_t = \bar{\mathbf{a}}_e$. 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. [111].

Example 3.4 (Gaussian Signals with Gaussian Beliefs). Mossel and Tamuz [55] 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 hold; 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 β_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 θ 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 3.3

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 $T_{ij} = \sigma_j / \sum_{p \in \tilde{\mathcal{N}}_i} \sigma_p > 0$ specify a convex combination: $\sum_{j \in \tilde{\mathcal{N}}_i} T_{ij} = 1$ for all i .

Example 3.5 (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 $\text{Gamma}(\alpha_i, \beta_i)$:

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

for all $\theta \in (0, \infty)$ and each $i \in [n]$. Note that here $\eta(\theta) = \log \theta$, $\gamma(\eta(\theta)) = \exp(\eta(\theta)) = \theta$, $\kappa(\alpha_i, \beta_i) = \Gamma(\alpha_i) \beta_i^{-\alpha_i}$, $\mathbf{m}_{i,\theta} = \delta_i \theta$, $\xi(\mathbf{s}_{i,p}) = \mathbf{s}_{i,p}$, $\sigma_i = 1$ and $\tau(\sigma_i \xi(\mathbf{s}_{i,p}), \delta_i) = \delta_i^{s_{i,p}} / (s_{i,p}!)$, for all i, p . This setting corresponds also to a case of Poisson observers with common rate θ and individual exposures $\delta_i, i \in [n]$, cf. [112, p. 54]. The posterior distribution over Θ 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 $n_i \delta_i + \beta_i$, [112, pp. 52–53]. Using a quadratic utility $-(\alpha - \delta_i \theta)^2$, the expected pay-off at time zero is maximized by the δ_i -scaled mean of the posterior Gamma belief distribution [112, 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: $\sum_{p=1}^{n_j} \mathbf{s}_{j,p} = (n_j + \beta_j \delta_j^{-1}) \mathbf{a}_{j,0} - \alpha_j, \forall j \in \tilde{\mathcal{N}}_i$, agent i can refine her belief into a Gamma distribution with parameters $\alpha_i + \sum_{j \in \tilde{\mathcal{N}}_i} [(n_j + \beta_j \delta_j^{-1}) \mathbf{a}_{j,0} - \alpha_j]$ and $\beta_i + \sum_{j \in \tilde{\mathcal{N}}_i} n_j \delta_j$. The subsequent optimal action at time 1 and the resultant Bayesian heuristics are as claimed in Theorem 3.3 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 $T_{ij} = n_j / \sum_{p \in \tilde{\mathcal{N}}_i} n_p > 0$ again specify a convex combination: $\sum_{j \in \tilde{\mathcal{N}}_i} T_{ij} = 1$ 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.

3.2.1 Linear Updating and Convergence

In general, the constant terms ϵ_i in Theorem 3.3 depend on the neighboring prior parameters $\alpha_j, j \in \tilde{\mathcal{N}}_i$ 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 3.3 reduce to linear update and $\epsilon_i = 0$. This requirement on the prior parameters is captured by our next assumption.

Assumption 3.2 (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.¹ 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 σ_i and δ_i . If we let $\alpha_i, \beta_i \rightarrow 0$ in the expressions of T_{ij} and ϵ_i from Theorem 3.3, then the affine action updates reduce to linear combinations and the succeeding corollary is immediate.

Corollary 3.1 (Linear updating). *Under the exponential family signal-utility structure (Assumption 3.1) with non-informative priors (Assumption 3.2); the Bayesian heuristics describe each updated action $\mathbf{a}_{i,t}$ as a linear combination of the neighboring actions $\mathbf{a}_{j,t-1}, j \in \tilde{\mathcal{N}}_i$: $\mathbf{a}_{i,t} = \sum_{j \in \tilde{\mathcal{N}}_i} T_{ij} \mathbf{a}_{j,t-1}$, where $T_{ij} = \delta_i \sigma_j n_j / (\sigma_i \sum_{p \in \tilde{\mathcal{N}}_i} n_p \delta_p)$.*

The action profiles under Corollary 3.1 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 [115, 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 3.1 (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.*

¹Conjugate 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 [113]. 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 [114, Section 3.5.3], while Laplace's classical principle of insufficient reason favors equiprobability leading to priors which are uniform over the parameter space.

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.

3.2.2 DeGroot Updates, Consensus, and Efficiency

In order for the linear action updates in Corollary 3.1 to constitute a convex combination as in the DeGroot model¹, we need to introduce some additional restrictions on the likelihood structure of the private signals.

Assumption 3.3 (Locally balanced likelihoods). *The likelihood structures given in (3.2.1) 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 3.3 signifies a local balance property for the two exponential family parameters σ_i and δ_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 σ_i and δ_i are both measures of accuracy or precision for private signals of agent i , the balance condition in Assumption 3.3 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 σ_i , when viewed as a function of their respective nodes (cf. [129, Section 2.1] for the definition and properties of harmonic functions):

$$\sigma_i = \sum_{j \in \mathcal{N}_i} \frac{n_j \sigma_j}{\sum_{k \in \mathcal{N}_i} n_k}, \delta_i = \delta, \forall i. \quad (3.2.2)$$

However, in a strongly conceded social network (3.2.2) cannot hold true unless σ_i is a constant: $\sigma_i = \sigma$ for all i . Similarly, when $\sigma_i = \sigma$ is a constant, then under Assumption 3.3 δ_i is spread as a harmonic function over the network nodes, and therefore can only take

¹The use of linear averaging rules for modeling opinion dynamics has a long history in mathematical sociology and social psychology [116]; their origins can be traced to French’s seminal work on “A Formal Theory of Social Power” [117]. 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 [118]. This model was later generalized to belief exchange dynamics and popularized by DeGroot’s seminal work [42] on linear opinion pools. In engineering literature, the possibility to achieve consensus in a distributed fashion (through local interactions and information exchanges between neighbors) is very desirable in a variety of applications such as load balancing [119], distributed detection and estimation [120–122], tracking [123], sensor networks and data fusion [124, 125], as well as distributed control and robotics networks [126, 127]. Early works on development of consensus algorithms originated in 1980s with the works of Tsitsiklis et.al [128] who propose a weighted average protocol based on a linear iterative approach for achieving consensus: each node repeatedly updates its value as a weighted linear combination of its own value and those received by its neighbors.

a constant value in a strongly connected network: $\delta_i = \delta$ for all i , cf. [129, Section 2.1, Maximum Principle]. In particular, fixing either of the parameters σ_i or δ_i for all agents and under the local balance condition in Assumption 3.3, it follows that the other parameter should be also fixed across the network; hence, the ratio σ_i/δ_i will be a constant for all i . Later when we consider the efficiency of consensus action we introduce a strengthening of Assumption 3.3, called globally balanced likelihood (cf. Assumption 3.4), where the ratio δ_i/σ_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 θ and common exposure δ) makes a case for likelihoods being locally balanced.

Theorem 3.6 (DeGroot updating and consensus). *Under the exponential family signal-utility structure (Assumption 3.1), with non-informative priors (Assumption 3.2) and locally balanced likelihoods (Assumption 3.3); the updated action $\mathbf{a}_{i,t}$ is a convex combination of the neighboring actions $\mathbf{a}_{j,t-1}, j \in \tilde{\mathcal{N}}_i$: $\mathbf{a}_{i,t} = \sum_{j \in \tilde{\mathcal{N}}_i} T_{ij} \mathbf{a}_{j,t-1}$, $\sum_{j \in \tilde{\mathcal{N}}_i} T_{ij} = 1$ 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 3.6, 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 3.3:

Assumption 3.4 (Globally balanced likelihoods). *The likelihood structures given in (3.2.1) 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 3.4, $\sigma_i \delta_j = \sigma_j \delta_i$ for all i, j , and it follows that the local balance of likelihoods is automatically satisfied. According to Definition 3.1, 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 $\tilde{\mathcal{N}}_j^{\text{out}} := \{i \in [n] : [I + A]_{ij} = 1\}$ and refer to them as the (self-inclusive)

out-neighborhood of agent j . This is in contrast to her (self-inclusive) neighborhood $\bar{\mathcal{N}}_j$, which is the set of all agents whom she listens to. Both sets $\bar{\mathcal{N}}_j$ and $\bar{\mathcal{N}}_j^{\text{out}}$ include agent j as a member.

Theorem 3.7 ((In-)Efficiency of consensus). *Under the exponential family signal-utility structure (Assumption 3.1) and with non-informative priors (Assumption 3.2); in a strongly connected social network, the agents achieve consensus at an efficient action if, and only if, the likelihoods are globally balanced and $\sum_{p \in \bar{\mathcal{N}}_j^{\text{out}}} n_p \delta_p = \sum_{p \in \bar{\mathcal{N}}_i} n_p \delta_p$, 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 3.3 and equation (3.2.2), we pointed out that if either of the two parameters σ_i and δ_i that characterize the exponential family distribution of (3.2.1) 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 3.5) 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 3.2 (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 3.1 with $\sigma_i = \sigma$ and $\delta_i = \delta$ for all i). If the agents have non-informative priors (Assumption 3.2) and the social network is strongly connected, then the (no-recall) heuristic agents update their action according to the linear combination: $\mathbf{a}_{i,t} = \sum_{j \in \bar{\mathcal{N}}_i} \mathbb{T}_{ij} \mathbf{a}_{j,t-1}$, where $\mathbb{T}_{ij} = n_j / \sum_{p \in \bar{\mathcal{N}}_i} n_p$, and reach a consensus. The consensus action is efficient if, and only if, $\sum_{p \in \bar{\mathcal{N}}_j^{\text{out}}} n_p = \sum_{p \in \bar{\mathcal{N}}_i} n_p$ 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 3.6 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 3.1 and 3.2) 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: $\mathbf{a}_{i,t} = (1/|\bar{\mathcal{N}}_i|) \sum_{j \in \bar{\mathcal{N}}_i} \mathbf{a}_{j,t-1}$ 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 $\sum_{i=1}^n (\deg(i)/|\mathcal{E}|) \mathbf{a}_{i,0}$, cf. [130, 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, $\deg(i) = d$ is fixed for all i , and $|\mathcal{E}| = n \cdot d$. In Appendix E, we consider the general problem of minimum variance estimation of a complete sufficient statistic from several i.i.d. samples that the networked agents collect. We rely on linear averaging to combine the observations of all agents; moreover, when the agents receive streams of data over time, we modify the update rule to accommodate the most recent observations and demonstrate the efficiency of our algorithm by proving convergence to the globally efficient estimator given the observations of all agents. We supplement these results by investigating the rate of convergence and providing finite-time performance guarantees when applicable.

Remark 3.8 (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, $\sum_{p \in \mathcal{N}_i^{\text{out}}} n_p = \sum_{p \in \tilde{\mathcal{N}}_i} n_p$, 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, [130]; 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 way agent; however, she fails to correct for these repetitions in the sources of her observations, leading to the co-called persuasion bias, [47].

3.3 Log-Linear Belief Updates

When the state space Θ 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, as in Section 2.2. The Bayesian heuristic in this case leads to a log-linear updating of beliefs.

We can follow the steps of Appendix A.13 to derive the Bayesian heuristic f_i in (3.1.2)

by replicating the time-one Bayesian belief update for all future time-steps:¹

$$\boldsymbol{\mu}_{i,t}(\theta') = \frac{\boldsymbol{\mu}_{i,t-1}(\theta') \left(\prod_{j \in \mathcal{N}_i} \frac{\boldsymbol{\mu}_{j,t-1}(\theta')}{\nu_j(\theta')} \right)}{\sum_{\tilde{\theta} \in \Theta} \boldsymbol{\mu}_{i,t-1}(\tilde{\theta}) \left(\prod_{j \in \mathcal{N}_i} \frac{\boldsymbol{\mu}_{j,t-1}(\tilde{\theta})}{\nu_j(\tilde{\theta})} \right)}, \text{ for all } \theta' \in \Theta \text{ and at any } t > 1. \quad (3.3.1)$$

Remark 3.9 (History-neglect and no-Recall). In writing (3.3.1), 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 3.10 (Generalization to one-step recall belief updates). In updating her belief at time t according to (3.3.1), agent i is implicitly assuming that each of her neighbors, $j \in \mathcal{N}_i$, 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 (3.3.1). This modifications requires only a single unit of memory and leads to the following update rule:

$$\boldsymbol{\mu}_{i,t}(\theta') = \frac{\boldsymbol{\mu}_{i,t-1}(\theta') \left(\prod_{j \in \mathcal{N}_i} \frac{\boldsymbol{\mu}_{j,t-1}(\theta')}{\boldsymbol{\mu}_{j,t-2}(\theta')} \right)}{\sum_{\tilde{\theta} \in \Theta} \boldsymbol{\mu}_{i,t-1}(\tilde{\theta}) \left(\prod_{j \in \mathcal{N}_i} \frac{\boldsymbol{\mu}_{j,t-1}(\tilde{\theta})}{\boldsymbol{\mu}_{j,t-2}(\tilde{\theta})} \right)}, \text{ for all } \theta' \in \Theta \text{ and at any } t > 1. \quad (3.3.2)$$

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 F, we analyze the evolution of beliefs when the agents interact and update their beliefs according

¹It is notable that the Bayesian heuristic in (3.3.1) has a log-linear structure. Geometric averaging and logarithmic opinion pools have a long history in Bayesian analysis and behavioral decision models [131, 132] and they can be also justified under specific behavioral assumptions [45]. They are also quite popular as a non-Bayesian update rule in engineering literature for addressing problems such as distributed detection and estimation [133–137]. In [137] the authors use a logarithmic opinion pool to combine the estimated posterior probability distributions in a Bayesian consensus filter; and show that as a result: the sum of KullbackLeibler divergences between the consensual probability distribution and the local posterior probability distributions is minimized. Minimizing the sum of KullbackLeibler divergences as a way to globally aggregate locally measured probability distributions is proposed in [138, 139] where the corresponding minimizer is dubbed the KullbackLeibler average. Similar interpretations of the log-linear update are offered in [140] as a gradient step for minimizing either the KullbackLeibler distance to the true distribution, or in [141] as a posterior incorporation of the most recent observations, such that the sum of KullbackLeibler distance to the local priors is minimized; indeed, the Bayes' rule itself has a product form and the Bayesian posterior can be characterized as the solution of an optimization problem involving the KullbackLeibler divergence to the prior distribution and subjected to the observed data [142].

to (3.3.2). In particular, we show that (3.3.2) 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 (3.3.2) 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 [57, Remark V.4].

Both the linear action updates studied in the previous section as well as the weighted majority update rules that arise in the binary case and are studied in [107, 143] have a familiar algebraic structure over the respective action spaces (the Euclidean space \mathbb{R}^k and the Galois field $\mathcal{GF}(2)$). In the next subsection, we develop similar structural properties for belief updates in (3.3.1) 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 Θ). In the resultant abelian group, the updated belief in (3.3.1) can be expressed as an addition of the self and neighboring beliefs subtracted by their priors, cf. (3.3.3).¹

3.3.1 An Algebra of Beliefs

Given two beliefs $\mu_1(\cdot)$ and $\mu_2(\cdot)$ over Θ 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 Θ 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

¹It 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 [144]. 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 [145].

given any such μ we can uniquely identify its inverse as follows:

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

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

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

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:

$$\boldsymbol{\mu}_{i,t} = \bigoplus_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t-1} \ominus \boldsymbol{\nu}_j. \quad (3.3.3)$$

The above belief update has a structure similar to the linear action updates studied in (3.3): 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 3.2 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 μ be denoted by $n \odot \mu := \mu^n$, defined as follows:¹

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

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

$$\boldsymbol{\mu}_{i,t+1} = \bigoplus_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t} \ominus \boldsymbol{\nu}_j = \bigoplus_{j \in [n]} ([\mathbf{I} + \mathbf{A}]_{ij} \odot \boldsymbol{\mu}_{j,t}) \ominus ([\mathbf{A}]_{ij} \odot \boldsymbol{\nu}_j). \quad (3.3.7)$$

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

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

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

With some abuse of notation, we can concatenate the network beliefs at every time t into a column vector $\boldsymbol{\mu}_t = (\boldsymbol{\mu}_{1,t}, \dots, \boldsymbol{\mu}_{n,t})^\top$ and similarly for the priors $\boldsymbol{v} = (v_1, \dots, v_n)^\top$; thus (3.3.7) can be written in the vectorized format by using the matrix notation as follows:

$$\boldsymbol{\mu}_t = \{(I + A) \odot \boldsymbol{\mu}_{t-1}\} \ominus \{A \odot \boldsymbol{v}\}. \quad (3.3.8)$$

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

$$\boldsymbol{\mu}_t = \left\{ (I + A)^t \odot \boldsymbol{\mu}_0 \right\} \ominus \left\{ \left(\sum_{\tau=0}^{t-1} (I + A)^\tau A \right) \odot \boldsymbol{v} \right\}. \quad (3.3.9)$$

The above is key to understanding the evolution of beliefs under the Bayesian heuristics in (3.3.1), as we will explore next. In particular, when all agents have uniform priors $v_j = \bar{\mu}$ for all j , then (3.3.8) and (3.3.9) 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 3.2.1, which paved the way for transition from affine action updates into linear ones. In the case of beliefs over a finite state space Θ , the uniform prior $\bar{\mu}$ is non-informative. If all agents start from common uniform priors, the belief update in (3.3.1) simplifies to:

$$\boldsymbol{\mu}_{i,t}(\theta') = \frac{\prod_{j \in \tilde{\mathcal{N}}_i} \boldsymbol{\mu}_{j,t-1}(\theta')}{\sum_{\tilde{\theta} \in \Theta} \prod_{j \in \tilde{\mathcal{N}}_i} \boldsymbol{\mu}_{j,t-1}(\tilde{\theta})}. \quad (3.3.10)$$

Our main focus in the Subsection 3.3.2 is to understand how the individual beliefs evolve under (3.3.1), or (3.3.10) which is a spacial case of (3.3.1). 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 [115, 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 $\bar{\mathbf{l}}_i, \bar{\mathbf{r}}_i$ as the left and right eigenvectors corresponding to the i th eigenvalue of $I + A$, satisfying: $\|\bar{\mathbf{l}}_i\|_2 = \|\bar{\mathbf{r}}_i\|_2 = 1$, $\bar{\mathbf{l}}_i^\top \bar{\mathbf{r}}_i = 1$ for all i and $\bar{\mathbf{l}}_i^\top \bar{\mathbf{r}}_j = 0$, $i \neq j$; in particular, the left eigenspace associated with ρ is one-dimensional with the corresponding eigenvector $\bar{\mathbf{l}}_1 = \bar{\boldsymbol{\alpha}} = (\alpha_1, \dots, \alpha_n)^\top$, uniquely satisfying $\sum_{i=1}^n \alpha_i = 1$,

$\alpha_i > 0, \forall i \in [n]$, and $\bar{\alpha}^\top \mathbf{A} = (\rho + 1)\bar{\alpha}^\top$. The entry α_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 \mathbf{A} to write [111, Section 6]:

$$(\mathbf{I} + \mathbf{A})^t = (1 + \rho)^t \left(\bar{r}_1 \bar{\alpha}^\top + \sum_{i=2}^n (\lambda_i(\mathbf{I} + \mathbf{A}) / (1 + \rho))^t \bar{r}_i \bar{l}_i^\top \right). \quad (3.3.11)$$

Asymptotically, we get that all eigenvalues other than the Perron-Frobenius eigenvalue $1 + \rho$ are subdominant; hence, $(\mathbf{I} + \mathbf{A})^t \rightarrow (1 + \rho)^t \bar{r}_1 \bar{\alpha}^\top$ and $\boldsymbol{\mu}_t = (1 + \rho)^t \bar{r}_1 \bar{\alpha}^\top \odot \boldsymbol{\mu}_0$ 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 3.11.

3.3.2 Becoming Certain about the Group Aggregate

We begin our investigation of the evolution of beliefs under (3.3.1) 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 [9, 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:

$$\boldsymbol{\mu}^*(\theta') = \frac{\prod_{j=1}^n \ell_j(\mathbf{s}_j | \theta')}{\sum_{\tilde{\theta} \in \Theta} \prod_{j=1}^n \ell_j(\mathbf{s}_j | \tilde{\theta})}. \quad (3.3.12)$$

Our next theorem describes the asymptotic outcome of the group decision process when the agents report their beliefs and follow the Bayesian heuristic (3.3.1) to aggregate them. The outcome indicated in Theorem 3.11 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 Θ^\diamond are weighted by the centralities (α_i) of their respective nodes. Secondly, the consensus belief is concentrated uniformly over Θ^\diamond , its support does not include the entire state space Θ 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 θ^\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 (3.3.12) 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 Θ^\diamond .

Theorem 3.11 (Certainty about the group aggregate). *Under the no-recall belief update (3.3.1), 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 $\lim_{t \rightarrow \infty} \mu_{i,t}(\theta') = 1/|\Theta^\diamond|$ for $\theta' \in \Theta^\diamond$ and $\lim_{t \rightarrow \infty} \mu_{i,t}(\theta') = 0$ for $\theta' \notin \Theta^\diamond$. In particular, if the sum of signal log-likelihoods weighted by node centralities is uniquely maximized by θ^\diamond , i.e. $\{\theta^\diamond\} = \Theta^\diamond$, then $\lim_{t \rightarrow \infty} \mu_{i,t}(\theta^\diamond) = 1$.*

Remark 3.12 (Relative efficiency of balanced regular networks). The fact that log-likelihoods in Θ^\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 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 θ , given the entire set of observations from all agents in the network.

In the next subsection, we use the insights from the no-recall belief exchange model to propose a coordination scheme for teaming agents across several time-steps without exposing them to redundant beliefs.

3.3.3 Teaming for Efficient Belief Exchange subject

Studies in committee mechanism design [146, 147], and Group Decision Support Systems (GDSS) [148], strive to construct mechanisms for information aggregation so that the group members arrive at optimal results. The overconfidence that Theorem 3.11 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 (3.3.1) 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 (3.3.1). 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. (F.0.4)). 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 . 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 [149, 150]. A classic study of group size in 1958 by Slater [149] 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 [150, 151], 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 [152]. 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 [153, 154]. A good rule is to have sufficiently many members, but not greater than that [155, 156]. The authors in [157] model the quality of organizational decisions as a function of the number of and time-commitment of individual decision-makers and show that a limit on the optimal size follows under natural convexity assumptions for the objective function and costs. Their results are especially relevant to decision-making organizations, such as legal or medical consulting firms and the public sector bureaus. Overall, team size is considered to be important in determining group performance by both theoretical and experimental studies [158, 159]. It continues to attract attention in social and organizational psychology for both business and management applications [160–163].

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 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 3.1 (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 3.1 (no-redundancy). To describe the procedure formally, let $\bar{n} := D^{\lceil \log_D(n) \rceil}$, $X_{n,D,t} = D^{\lceil \log_D(n) - t \rceil} = \bar{n}/D^{t+1}$, $Z_{n,D} = \bar{n}/D = D^{\lceil \log_D(n) - 1 \rceil} = D^t X_{n,D,t}$, where t indexes the rounds of communications and $\lceil \log_D(n) \rceil$ is as defined above.

(A5: NO-REDUNDANCY-COORDINATION). For n people to exchange beliefs in $\lceil \log_D(n) \rceil$ 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 $t = 0, \dots, \lceil \log_D(n) \rceil - 1$, organize $Z_{n,D} = D^t X_{n,D,t}$ meetings in parallel, where the participants in each meeting are indexed as follows:

$$\{\tau + \tau' D^t + \tau'' D^{t+1} : \tau' = 0, 1, 2, \dots, D - 1\}, \quad (3.3.13)$$

and the range of the constants τ and τ'' 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 (A5), 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 (A5): 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 (3.3.13), τ indexes the members of each subgroup of size D^t according to their rankings and τ'' 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 3.13 (Six people coordinated to meet in pairs). In Fig. 3.2, we provide an example implementation of (A5) with $D = 2$ and $n = 6$. The top diagram in Fig. 3.2 depicts the propagation of information with advancing time steps, as groups are merged and agents are regrouped in new pairs. The bottom diagrams of Fig. 3.2 depict the flow of information in the network and across time as agents are regrouped according to (A5). The no-redundancy requirement of Condition 3.1 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 (A5) 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. 3.2, 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




Figure 3.2: Implementation of (A5) 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.

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 (3.3.1) simplifies as follows:

$$\mu_{i,t}(\theta') = \frac{\mu_{i,t-1}(\theta') \prod_{j \in G_{i,t}} \mu_{j,t-1}(\theta')}{\sum_{\tilde{\theta} \in \Theta} \mu_{i,t-1}(\tilde{\theta}) \prod_{j \in G_{i,t}} \mu_{j,t-1}(\tilde{\theta})}. \quad (3.3.14)$$

where $G_{i,t} \subset [n]$ is the group of people with whom agent i is scheduled to meet at time t . Under (3.3.14) and after attending $T_{n,D}$ meetings all agents will hold the (common) Bayesian aggregate belief in (3.3.12). 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. (F.0.4) in Appendix F).

Chapter 4

Learning from Stream of Observations

In this chapter, we investigate the no-recall heuristic behaviors when the agents observe a stream of private signals in addition to each other's actions. As before, we present the implications of various choices for the structure of the action space and utility functions for such agents and investigate the properties of learning, convergence, and consensus in special cases. The chief question of interest 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. This learning framework in which agents have access to an stream of new observations is in contrast with the group decision model of the previous chapter; the difference being in the fact that there 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.^{1,2}

¹The following papers cover the results of this chapter: [74, 105, 107, 164–167].

²Our initial results on applying the no-recall idea to the case of two communicating agents with streaming observations (Example 4.3) was derived in collaboration with Pooya Molavi. Our results in Subsection 4.3.3 (Learning without recall by random walks on directed graphs) are derived in collaboration with Shahin Shahrampour.

In Section 4.2 we specialize the no-recall model to a binary state, signal, and action space and investigate the evolution of actions in the resultant Ising model. We show that the no-recall action updates in this case are given by *weighted majority and threshold rules* that linearly combine the observed binary actions and log-likelihoods of private signal. We show that under these update rules the action profiles evolve as a Markov chain on the Boolean cube and that the properties of consensus and learning are subsequently determined by the equilibria of this Markov chain. Next in Section 4.3 we analyze the case where the network agents are announcing their beliefs at every epoch and the no-recall updates become log-linear, similar to Section 3.3 with the added complexity of having streaming observations. Naivety of agents when they have access to a stream of observations impedes their ability to learn; except in simple social structures such as cycles or rooted trees. We explain how the circular no-recall updates generalize to any strongly connected topology if agents choose a neighbor randomly at every round and restrict their belief update to the selected neighbor each time.

4.1 The No-Recall Model Applied to Streaming Data

In this section, we consider how the no-recall framework may be applied in a setting, where the agents are observing a stream of private signals that arrive sequentially at every time period, in addition to observing each other's actions. The following examples help highlight the complexities of inference from social data when it is coupled with streaming data.

Example 4.1 (Decisions of a Single Agent in a Binary World). For clarity of exposition we consider a single agent, i , who observes binary signals $\mathbf{s}_{i,t} \in \{\pm 1\}$ and takes actions $\mathbf{a}_{i,t} \in \{\pm 1\}$ at every instant $t \in \mathbb{N}_0$. She lives in a binary world where the truth θ can take one of the two values ± 1 with equal probability (cf. Fig. 1.3); and her actions are rewarded by $+1$ if $\mathbf{a}_{i,t} = \theta$ and are penalized by -1 otherwise. Suppose further that her probability of receiving the signal θ is $p > 0.5$; so that

$$\mathcal{P}_\theta\{\mathbf{s}_{i,t} = \theta\} = 1 - \mathcal{P}_\theta\{\mathbf{s}_{i,t} = -\theta\} = p > 0.5 .$$

We assume that the agent is myopic so that at every time t she is only concerned about her immediate reward at that decision epoch. Consider the decision problems of agent i at every time instant $t \in \mathbb{N}_0$; to model them as Markov decision processes we define her state as the collection of all private signals that she would ever observe: $\bar{\mathbf{s}}_\infty := (\mathbf{s}_{i,0}, \dots, \mathbf{s}_{i,t}, \dots) \in \mathcal{S}_i^{\mathbb{N}_0}$. Subsequently, at any time t she has only partial knowledge of her state. Let $\bar{\mathbf{s}}_t := (\mathbf{s}_{i,0}, \dots, \mathbf{s}_{i,t})$ be the collection of all private signals that is revealed to her up until time t . Her expected reward from taking an action $\mathbf{a}_{i,t}$ is then given by

$$r(\mathbf{a}_{i,t}, \bar{\mathbf{s}}_t) = \mathcal{P}_\theta\{\mathbf{a}_{i,t} = \theta \mid \bar{\mathbf{s}}_t\} - \mathcal{P}_\theta\{\mathbf{a}_{i,t} \neq \theta \mid \bar{\mathbf{s}}_t\} = 2\mathcal{P}_\theta\{\mathbf{a}_{i,t} = \theta \mid \bar{\mathbf{s}}_t\} - 1.$$

Since both $\{\theta = +1\}$ and $\{\theta = -1\}$ are initially equally likely, the agents' optimal decisions is given by:

$$\mathbf{a}_{i,t}^* = \begin{cases} +1, & \text{if } \mathcal{P}_{+1}\{\bar{\mathbf{s}}_t\} \geq \mathcal{P}_{-1}\{\bar{\mathbf{s}}_t\}, \\ -1, & \text{otherwise.} \end{cases} \quad (4.1.1)$$

Using the likelihood ratio statistics

$$\frac{\mathcal{P}_{+1}\{\bar{\mathbf{s}}_t\}}{\mathcal{P}_{-1}\{\bar{\mathbf{s}}_t\}} = \left(\frac{p}{1-p} \right)^{\sum_{\tau=0}^t \mathbf{s}_{i,\tau}},$$

we can rewrite (4.1.1) as a threshold rule in terms of the sufficient statistic $\mathbf{S}_{i,t} = \sum_{\tau=0}^t \mathbf{s}_{i,\tau}$ that is the running total of the observed signals:

$$\mathbf{a}_{i,t}^* = \begin{cases} +1, & \text{if } \mathbf{S}_{i,t} \geq 0, \\ -1, & \text{otherwise.} \end{cases} \quad (4.1.2)$$

Example 4.2 (Two Communicating Agents). The authors in [56] have considered more sophisticated scenarios involving two agents, one of whom (called j) observes the other's (called i) actions (unidirectionally) in addition to her own sequence of private signals. They distinguish two cases: in case one the more informed agent j only observes her neighbor's penultimate action $\mathbf{a}_{i,t-1}$; in case two she observes the whole sequence of actions taken by her neighbor, $(\mathbf{a}_{i,0}, \dots, \mathbf{a}_{i,t-1})$, and use them in her decisions. The optimal (Bayesian) decision for agent j in both cases can be derived as threshold rules on the sum $\mathbf{S}_{j,t}$; however, the respective thresholds are time-varying and more complex. In particular, if agent j in addition to her private signals also observes the last action of agent i , who only observes private signals, then the optimal action of agent j involves a time-varying threshold expressed below, while the optimal action of agent i is the same as (4.1.2):

$$\mathbf{a}_{j,t}^* = \begin{cases} \text{sign}(\mathbf{S}_{j,t}), & \text{if } |\mathbf{S}_{j,t}| \geq \eta_t^*, \\ \mathbf{a}_{i,t-1}, & \text{otherwise,} \end{cases}$$

where

$$\eta_t^* = \left(\log \frac{\mathcal{P}_{\theta}\{\mathbf{a}_{i,t-1} = \theta\}}{\mathcal{P}_{\theta}\{\mathbf{a}_{i,t-1} \neq \theta\}} \right) / \left(\frac{p}{1-p} \right),$$

is set to optimally balance the probability of a mistake by agent i with the strength of private signals of agent j ; thus enabling the latter to decide whether to imitate her neighbor or else disregard her neighbors' action and act based on her own signals. If j observes the entire sequence of actions taken by i the optimal threshold depend also on the length

of last run that is the number of time periods in which agent i has taken the same action as her last choice, cf. [56, Proposition 14]; the authors further highlight the difficulties in the case where both agents i and j observe each other's actions. In particular, while the optimal decisions in the bidirectional case are not known, the increased interaction and the more information that is available as a result of bidirectional communication do not lead to a faster rate of learning. The slower rate of learning is caused by a process of ‘‘Bayesian groupthink’’, when the agents' mistakes reinforce each other, thus preventing them from taking a correct action despite the ample evidence presented by their private signals. \square

In general, when a rational agent observes her neighbors in a network, she should compensate for repetitions in the sources of her information: the same neighbors' actions are repeatedly observed and neighboring actions may be affected by the past actions of the agent herself; thence major challenges of Bayesian inference for social learning are due to the private signals and third party interactions that are hidden from the agent. Moreover, existence of loops in the network cause dependencies and correlations in the information received from different neighbors, which further complicates the inference task. we now present an application of the no-recall scheme to this streaming observations setting that avoids the complexities of Bayesian inference in such circumstances.

For clarity of exposition in the rest of the chapter, we restrict attention to finite signal and action spaces. Given $\mathbf{s}_{i,0}$, agent i forms an initial Bayesian opinion $\mu_{i,0}(\cdot)$ about the value of θ , which is given by

$$\mu_{i,0}(\hat{\theta}) = \frac{\nu_i(\hat{\theta})l_i(\mathbf{s}_{i,0} | \hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \nu_i(\tilde{\theta})l_i(\mathbf{s}_{i,0} | \tilde{\theta})}, \forall \hat{\theta} \in \Theta. \quad (4.1.3)$$

She then chooses the action: $\mathbf{a}_{i,0} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \sum_{\hat{\theta} \in \Theta} u_i(\mathbf{a}_i, \hat{\theta}) \mu_{i,0}(\hat{\theta})$, maximizing her expected reward: $\mathbb{E}_{i,0}\{u_i(\mathbf{a}_{i,0}, \theta)\}$. 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: $\mathbf{a}_{j,0}$, $j \in \mathcal{N}_i$. Given her extended set of observations at time $t = 1$, she makes a second and possibly different move $\mathbf{a}_{i,1}$ according to

$$\mathbf{a}_{i,1} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \sum_{\hat{\theta} \in \Theta} u_i(\mathbf{a}_i, \hat{\theta}) \mu_{i,1}(\hat{\theta}), \quad (4.1.4)$$

maximizing her expected pay off conditional on everything that she has observed thus far: $\mathbb{E}_{i,1}\{u_i(\mathbf{a}_{i,1}, \hat{\theta})\} = \mathbb{E}\{u_i(\mathbf{a}_{i,1}, \hat{\theta}) | \mathbf{s}_{i,0}, \mathbf{a}_{j,0} : j \in \mathcal{N}_i\}$.

Of particular significance in our description of the behavior of agents in the succeeding time periods $t > 1$, is the relation

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

derived in (4.1.4), which given the observations of agent i from time $t = 0$, specifies her (Bayesian) pay-off maximizing action for time $t = 1$. Note that in writing (4.1.4), we assumed that the agents do not receive any private signals at $t = 1$ and there is therefore no $s_{i,1}$ appearing in the updates of any agent i ; and this convention is exactly to facilitate the derivation of mapping $f_i : \mathcal{S}_i \times \prod_{j \in \mathcal{N}_i} \mathcal{A}_j \rightarrow \mathcal{A}_i$, from the private signal space and action spaces of the neighbors to succeeding actions of each agent. In every following instance we aim to model the inferences of agents about their observations as being rational but memoryless: as of those who come to know their immediate observations which include the actions of their neighbors and their last private signals, but cannot trace these observations to their roots and has no ability to reason about why their neighbors may be behaving the way they do. In particular, such agents have no incentives for experimenting with false reports, as their lack of memory prevents them from reaping the benefits of their experiment, including any possible revelations that a truthful report may not reveal. Subsequently, we argue on normative grounds that such rational but memoryless agents would replicate the behavior of a Bayesian (fully-rational) agent between times zero and one; whence by regarding their observations as being direct consequences of inferences that are made based on the initial priors, they reject any possibility of a past history beyond their immediate observations:¹

$$\mathbf{a}_{i,t} = f_i(\mathbf{s}_{i,t}, \mathbf{a}_{j,t-1} : j \in \mathcal{N}_i), \forall t > 1.$$

On the other hand, note that rationality of agents constrains their beliefs $\mu_{i,t}(\cdot)$ given their immediate observations; hence, we can also write

$$\mathbf{a}_{i,t} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{i,t}\{u_i(\mathbf{a}_i, \hat{\theta})\}, \quad (4.1.5)$$

¹Memory constraints have been also looked at in the context of social learning [168, Chapter 5]. In recent results, [169] considers the model of a decision maker who chooses between two actions with pay-offs that depend on the true state of the world. Furthermore, the decision maker must always summarize her information into one of finitely many states, leading to optimal decision rules that specify the transfers between states. The problem of learning with finite memory in the context of hypothesis testing was originally formulated by [170, 171] under memory constraints for storing the test statistics. Accordingly, while sufficient statistics are very useful computational tools their utility for memory reduction is not clear. Subsequent results provide sophisticated algorithms using automata to perform the task of hypothesis testing using test statistics that take only finitely many values and to guarantee an asymptotically vanishing error probability [172–175]. More recently, [176] have considered this problem in a setting where agents each receive an independent private signal and make decisions sequentially. Memory in this context refers to the number of immediate predecessors whose decisions are observable by any given agent at the time of making her decision. Accordingly, while the almost sure convergence of the sequence of individual decisions to the correct state is not possible in this finite memory setting, the authors construct decision rules that achieve convergence and learning in probability. They next go on to consider the behavior of rational (pay-off maximizing) agents in this context and show that in no equilibrium of the associated Bayesian game learning can occur.

or equivalently $\mathbf{a}_{i,t} \leftarrow \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \sum_{\hat{\theta} \in \Theta} \mathbf{u}_i(\mathbf{a}_i, \hat{\theta}) \mu_{i,t}(\hat{\theta})$.

We next demonstrate this procedure through an example with two agents observing private and communicating their beliefs repeatedly.

Example 4.3 (The case of Two Agents Exchanging Beliefs). Consider first the case of two agents, labeled one and two, which communicate with each other using a two-way communication link. Following the formation of initial beliefs according to (4.1.3) for $j \in \{1, 2\}$, the agents communicate their initial beliefs to each other, that is $\mu_{1,0}(\cdot)$ is made known to agent two and $\mu_{2,0}(\cdot)$ is made known to agent one. Each agent now updates its belief according to the Bayes' rule. This, in the case of agent two, leads to the refined belief $\mu_{2,1}(\cdot)$ given by:

$$\begin{aligned} \mu_{2,1}(\hat{\theta}) &= \mathcal{P}(\theta = \hat{\theta} \mid \mathbf{s}_{2,0}, \mu_{1,0}(\cdot)) = \frac{\mathcal{P}(\theta = \hat{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot))}{\mathcal{P}(\mathbf{s}_{2,0}, \mu_{1,0}(\cdot))} \\ &= \frac{\mathcal{P}(\theta = \hat{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot))}{\sum_{\tilde{\theta} \in \Theta} \mathcal{P}(\tilde{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot))}, \forall \hat{\theta} \in \Theta. \end{aligned} \quad (4.1.6)$$

To calculate $\mathcal{P}(\tilde{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot))$, $\forall \tilde{\theta} \in \Theta$, proceed as follows. Conditioning on $\mathbf{s}_{1,0}$ yields:

$$\mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot)) = \sum_{s \in \mathcal{S}_1} \mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot) \mid \mathbf{s}_{1,0} = s) \mathcal{P}(\mathbf{s}_{1,0} = s). \quad (4.1.7)$$

Next, note that if $\mathbf{s}_{1,0} = s$, then $\mu_{1,0}(\cdot)$ is known deterministically, and is given by (4.1.3) with $\mathbf{s}_{1,0}$ replaced by s . Whence, if $s \in \mathcal{I}_1(\mu_{1,0}(\cdot))$, then $\mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot) \mid \mathbf{s}_{1,0} = s) = \mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0} \mid \mathbf{s}_{1,0} = s)$, and else $\mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot) \mid \mathbf{s}_{1,0} = s) = 0$. Therefore, (4.1.7) can be rewritten as:

$$\begin{aligned} \mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot)) &= \sum_{s \in \mathcal{I}_1(\mu_{1,0}(\cdot))} \mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0} \mid \mathbf{s}_{1,0} = s) \mathcal{P}(\mathbf{s}_{1,0} = s) \\ &= \sum_{s \in \mathcal{I}_1(\mu_{1,0}(\cdot))} \mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0}, \mathbf{s}_{1,0} = s). \end{aligned}$$

In the latter, $\mathcal{P}(\cdot)$ can be expressed in terms of the signal structures and common priors as $\mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0}, \mathbf{s}_{1,0} = s) = \nu(\tilde{\theta}) \mathbf{l}_2(\mathbf{s}_{2,0} \mid \tilde{\theta}) \mathbf{l}_1(s \mid \tilde{\theta})$ leading to $\mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot)) = \sum_{s \in \mathcal{I}_1(\mu_{1,0}(\cdot))} \nu(\tilde{\theta}) \mathbf{l}_2(\mathbf{s}_{2,0} \mid \tilde{\theta}) \mathbf{l}_1(s \mid \tilde{\theta})$. Next note that for all $s \in \mathcal{I}_1(\mu_{1,0}(\cdot))$, it is inferred from (4.1.3) that $\nu(\tilde{\theta}) \mathbf{l}_1(s \mid \tilde{\theta}) = \mu_{1,0}(\tilde{\theta}) \sum_{\bar{\theta} \in \Theta} \nu(\bar{\theta}) \mathbf{l}_1(s \mid \bar{\theta})$, which along with (A.13.2) yields:

$$\begin{aligned} \mathcal{P}(\theta = \tilde{\theta}, \mathbf{s}_{2,0}, \mu_{1,0}(\cdot)) &= \mathbf{l}_2(\mathbf{s}_{2,0} \mid \tilde{\theta}) \mu_{1,0}(\tilde{\theta}) \sum_{s \in \mathcal{I}_1(\mu_{1,0}(\cdot))} \sum_{\bar{\theta} \in \Theta} \nu(\bar{\theta}) \mathbf{l}_1(s \mid \bar{\theta}) \\ &= \mathbf{l}_2(\mathbf{s}_{2,0} \mid \tilde{\theta}) \mu_{1,0}(\tilde{\theta}) \mathcal{K}_1(\mu_{1,0}(\cdot)). \end{aligned} \quad (4.1.8)$$

Substituting (4.1.8) for $\mathcal{P}(\tilde{\theta}, \mathbf{s}_{2,0}, \boldsymbol{\mu}_{1,0}(\cdot))$ in (4.1.6) and canceling out $\mathcal{K}_1(\boldsymbol{\mu}_{1,0}(\cdot))$ from the the numerator and denominator leads to:

$$\boldsymbol{\mu}_{2,1}(\hat{\theta}) = \frac{\mathfrak{l}_2(\mathbf{s}_{2,0} | \hat{\theta}) \boldsymbol{\mu}_{1,0}(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mathfrak{l}_2(\mathbf{s}_{2,0} | \tilde{\theta}) \boldsymbol{\mu}_{1,0}(\tilde{\theta})}, \forall \hat{\theta} \in \Theta, \quad (4.1.9)$$

which is the *Bayesian* inferred belief that agent two holds about the true state of the world θ , following its observation of the private signal $\mathbf{s}_{2,0}$ as well as agent one's initial belief $\boldsymbol{\mu}_{1,0}(\cdot)$. Replacing $\boldsymbol{\mu}_{2,1}(\cdot)$, $\boldsymbol{\mu}_{1,0}(\cdot)$, and $\mathbf{s}_{2,0}$ with $\boldsymbol{\mu}_{2,t}(\cdot)$, $\boldsymbol{\mu}_{1,t-1}(\cdot)$, and $\mathbf{s}_{2,t}$ in (4.1.9) leads to a *Non-Bayesian* update rule for the succeeding steps $t > 1$, as follows:

$$\boldsymbol{\mu}_{2,t}(\hat{\theta}) = \frac{\mathfrak{l}_2(\mathbf{s}_{2,t} | \hat{\theta}) \boldsymbol{\mu}_{1,t-1}(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mathfrak{l}_2(\mathbf{s}_{2,t} | \tilde{\theta}) \boldsymbol{\mu}_{1,t-1}(\tilde{\theta})}, \forall \hat{\theta} \in \Theta. \quad (4.1.10)$$

Swapping the roles of indices 1 and 2 in (4.1.10) yields the corresponding rule for agent one. To see the distinction, consider the Bayesian opinion of agent two at time $t = 2$, that is when agent two has observed two private signals $\mathbf{s}_{2,0}$ and $\mathbf{s}_{2,2}$, as well as agent one's beliefs $\boldsymbol{\mu}_{1,1}(\cdot)$ and $\boldsymbol{\mu}_{1,0}(\cdot)$. This Bayesian opinion is given for all $\hat{\theta} \in \Theta$ by:

$$\boldsymbol{\mu}_{2,2}(\hat{\theta}) = \frac{\mathfrak{l}_2(\mathbf{s}_{2,2} | \hat{\theta}) \mathfrak{l}_2(\mathbf{s}_{2,0} | \hat{\theta}) \boldsymbol{\mu}_{1,0}(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mathfrak{l}_2(\mathbf{s}_{2,2} | \tilde{\theta}) \mathfrak{l}_2(\mathbf{s}_{2,0} | \tilde{\theta}) \boldsymbol{\mu}_{1,0}(\tilde{\theta})}. \quad (4.1.11)$$

Note that no private signals are received at time $t = 1$, and therefore $\mathbf{s}_{2,1}$ does not appear in the preceding formulation. Moreover, given the knowledge of $\mathbf{s}_{2,0}$ and conditioned on the value of $\mathbf{s}_{1,0}$ both $\boldsymbol{\mu}_{1,0}(\cdot)$ and $\boldsymbol{\mu}_{1,1}(\cdot)$ would be known deterministically and that is why the latter does not appear in (4.1.11). In forming the Bayesian opinion in (4.1.11), the agent is using the entire history of its available information up to time $t = 2$. This is in contrast with the opinion suggested by (4.1.10), which takes into account only the immediately available private signal and the other agent's last opinion. An interesting feature of (4.1.10) is that the agent's own prior beliefs do not appear in their update rules. This, however, is seen to be not true in the general case of a social network unless the agent in question has only a single neighbor.

As a numeric example, let $\Theta = \{1, 2, 3\}$ and suppose that each agent $i \in \{1, 2\}$ receives binary private signals $\mathbf{s}_{i,t}$, $t > 1$, with the probability of zero at each state of the world given as follows:

likelihoods	$\hat{\theta} = 1$	$\hat{\theta} = 2$	$\hat{\theta} = 3$
$\mathfrak{l}_1(\mathbf{s}_{1,t} = 0 \hat{\theta})$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{4}$
$\mathfrak{l}_2(\mathbf{s}_{2,t} = 0 \hat{\theta})$	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{1}{3}$

At time $t = 0$ the state $\theta = 2$ is realized, at which point each agent receives an

initial signal $s_{i,0}$ and forms an opinion $\mu_{i,0}(\cdot)$ using the Bayes rule (4.1.3). The agents then communicate and update their beliefs according to (4.1.10) at every following step of time $t \geq 1$. Note that at the realized state of the world $\Theta = 2$ neither of the agents is able to correctly infer the realized state based only on their private signals. It is because the signal structures are such that agent one cannot distinguish between the first two states $\hat{\theta} \in \{1, 2\}$ and agent two is unable to distinguish between the states $\hat{\theta} \in \{2, 3\}$. Nonetheless, the evolution of the belief dynamics and the difference between the two agents beliefs in Figs. 4.1 and 4.5 indicate that both agents have indeed learned the true state of the world. In other words, by communicating in a social network the two agents are able to effectively combine each other's observations so that overall they both learn the true state of the world, while alone neither would have been able to.

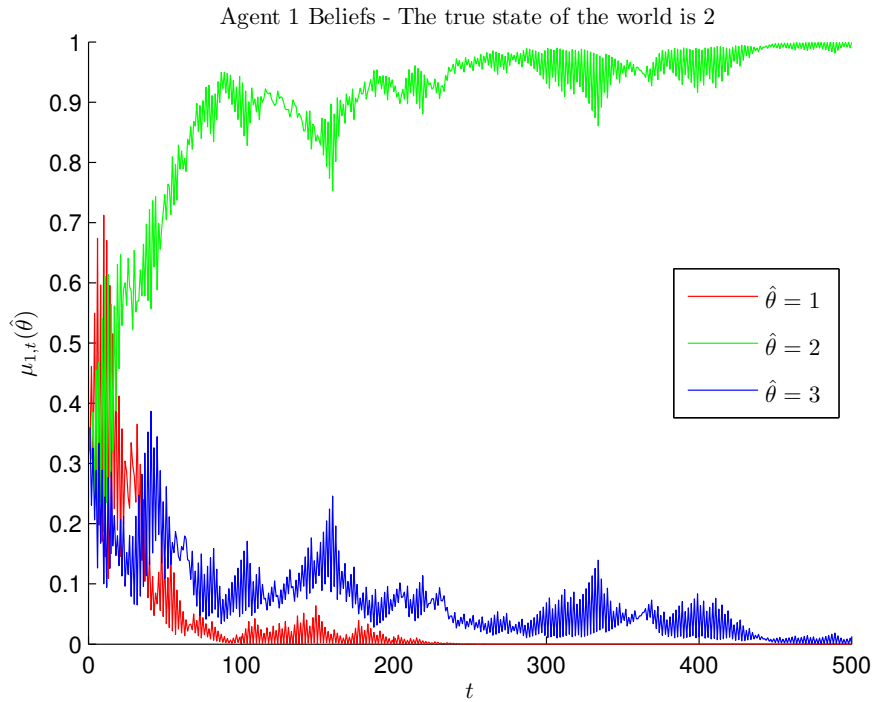


Figure 4.1: Evolution of the first agent's beliefs over time

We can indeed formalize the convergence result in case of two communicating agents as a proposition:

Proposition 4.1. Suppose that $\theta \in \Theta$ is such that for any $\hat{\theta} \in \Theta$ and $\hat{\theta} \neq \theta$, there exist an agent $i \in \{1, 2\}$ with $\mathcal{L}_i(\cdot | \theta) \neq \mathcal{L}_i(\cdot | \hat{\theta})$, i.e. $\mathcal{L}_i(s | \theta) \neq \mathcal{L}_i(s | \hat{\theta})$ for some $s \in \mathcal{S}_i$. If $\theta = \theta$ is realized by the nature and the two agent follow the update rule in (4.1.10), then both agents learn the true state θ of the world, \mathcal{P} -almost surely.

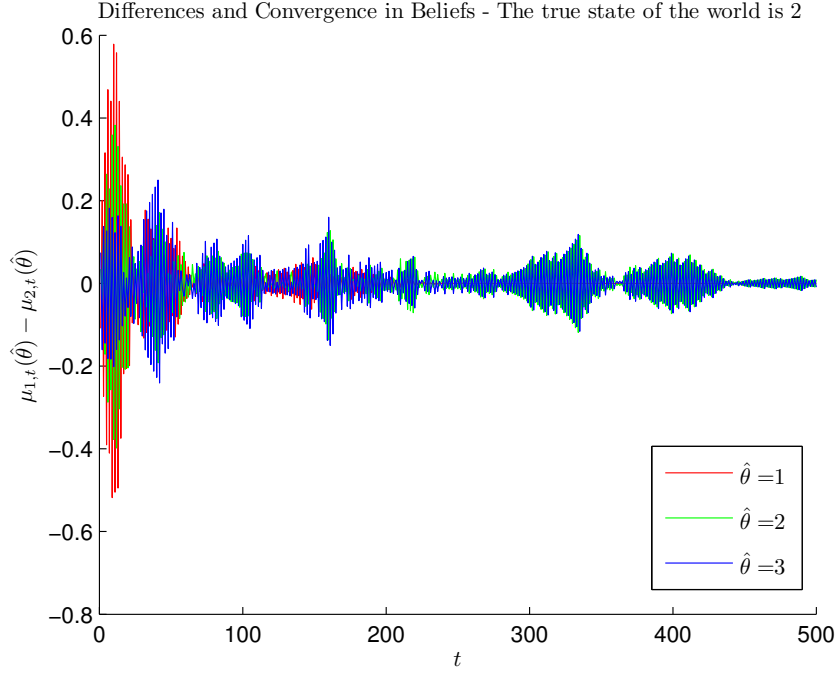


Figure 4.2: The difference in agents' beliefs over time

Proof. For all $\omega \in \Omega$, let $\mathbb{1}_{\theta=\theta}(\omega) = 1$ if $\theta(\omega) = \theta$ and $\mathbb{1}_{\theta=\theta}(\omega) = 0$, otherwise. Let $\{i, j\} = \{1, 2\}$ and define $\mathcal{F}_{i,t} = \sigma(\mathbf{s}_{i,t}, \mathbf{s}_{j,t-t}, \mathbf{s}_{i,t-2}, \dots, \mathbf{s}_{i,1}, \mathbf{s}_{i,0})$ for t odd and $\mathcal{F}_{i,t} = \sigma(\mathbf{s}_{i,t}, \mathbf{s}_{j,t-t}, \dots, \mathbf{s}_{i,2}, \mathbf{s}_{j,1}, \mathbf{s}_{i,0})$ for t even, as the sigma fields generated by the private signals of the agents in alternating order. Note in particular that $\mathcal{F}_{0,t} = \sigma(\{\mathbf{s}_{i,0}\})$ and $\{\mathcal{F}_{i,t}, t \in \mathbb{W}\}$ is a filtration on the measure space (Ω, \mathcal{F}) . Moreover, from the Bayes rule in (4.1.3), it follows that:

$$\mu_{i,0}(\theta) = \mathcal{P}\{\theta = \theta \mid \mathbf{s}_{i,0}\} = \mathbb{E}\{\mathbb{1}_{\theta=\theta} \mid \mathcal{F}_{i,0}\}. \quad (4.1.12)$$

Now two successive applications of (4.1.10) at times t and $t-1$, yields $\forall t > 1$:

$$\mu_{i,t}(\theta) = \frac{l_i(\mathbf{s}_{i,t} \mid \theta) \mu_{j,t-1}(\theta)}{\sum_{\tilde{\theta} \in \Theta} l_i(\mathbf{s}_{i,t} \mid \tilde{\theta}) \mu_{j,t-1}(\tilde{\theta})}, \quad \mu_{j,t-1}(\theta) = \frac{l_j(\mathbf{s}_{j,t-1} \mid \theta) \mu_{i,t-2}(\theta)}{\sum_{\tilde{\theta} \in \Theta} l_j(\mathbf{s}_{j,-1} \mid \tilde{\theta}) \mu_{i,t-2}(\tilde{\theta})},$$

and replacing for $\mu_{j,t-1}(\theta)$ in $\mu_{i,t}(\theta)$ yields:

$$\mu_{i,t}(\theta) = \frac{l_i(\mathbf{s}_{i,t} \mid \theta) l_j(\mathbf{s}_{j,t-1} \mid \theta) \mu_{i,t-2}(\theta)}{\sum_{\tilde{\theta} \in \Theta} l_i(\mathbf{s}_{i,t} \mid \tilde{\theta}) l_j(\mathbf{s}_{j,t-1} \mid \tilde{\theta}) \mu_{i,t-2}(\tilde{\theta})}, \quad (4.1.13)$$

for all $t > 1$. However, starting from the Bayesian opinion in (4.1.12) the above is exactly the Bayesian update of agent i 's belief from time $t-2$ to time t , given that at time $t-1$

agent j has observed the signal $\mathbf{s}_{j,t-1}$ and at time t agent i has observed the signal $\mathbf{s}_{i,t}$. Whence, (4.1.13) can be combined with (4.1.12) to get $\mu_{i,t}(\theta) = \mathbb{E}\{\mathbb{1}_{\theta=\theta} \mid \mathcal{F}_{i,t}\}, \forall t \geq 0$. The beliefs form a bounded martingale with respect to the filtration introduced above, and stage is set for the martingale convergence theorem. Indeed, the proof is now immediate, since as a consequence of the supposition $\{\theta = \theta\} \in \mathcal{F}_{i,\infty}$ and by Levy's zero-one law [177], $\lim_{t \rightarrow \infty} \mu_{i,t}(\theta) = \mathbb{E}\{\mathbb{1}_{\theta=\theta} \mid \mathcal{F}_{i,\infty}\} = \mathbb{1}_{\theta=\theta}$, \mathcal{P} -almost surely. ■

The proof verifies the intuition that by following (4.1.10), the agents can enjoy the full benefits of each other's observations. Indeed, (4.1.13) implies that the belief of an agent (call it i) at time t is the Bayesian update of its belief at time $t - 2$, given that at time $t - 1$ the other agent (call it j) has observed the signal $\mathbf{s}_{j,t-1}$ and at time t , agent i has observed the signal $\mathbf{s}_{i,t}$. This further indicates that the exponential convergence rate that holds true in the case of Bayesian learning can be applied here as well [168]. In Appendix A.13, a set of steps that parallel (4.1.6) to (4.1.10) are followed to derive an update rule for the general case of agents exchanging beliefs in a social network, which we discuss in Section 4.3.

In the sequel, we explore various structures for the action space and the resultant update rules f_i . In Section 4.2, we show how a common heuristic such as weighted majority can be explained as a rational but memoryless behavior with actions taken from a binary set. In Section 4.3, we shift focus to a finite state space and the probability simplex as the action space. There agents exchange beliefs and the belief updates are log-linear.

4.2 Weighted Majority and Threshold Rules

In this section, we consider a binary state space $\Theta = \{+1, -1\}$, and suppose that the agents have a common binary action space $\mathcal{A}_i = \{-1, 1\}$, for all i . Let their utilities be given by $u_i(\alpha, \theta) = 2\mathbb{1}_\alpha(\theta) - 1$, for any agent i and all $\theta, \alpha \in \{-1, 1\}$; here, $\mathbb{1}_\alpha(\theta)$ is equal to one only if $\theta = \alpha$ and is equal to zero otherwise. Subsequently, the agent is rewarded by $+1$ every time she correctly determines the value of θ and is penalized by -1 otherwise.

We can now calculate $\sum_{\hat{\theta} \in \Theta} u_i(\alpha_i, \hat{\theta}) \mu_{i,t}(\hat{\theta}) = \alpha_i(\mu_{i,t}(+1) - \mu_{i,t}(-1)) = \alpha_i(2\mu_{i,t}(+1) - 1)$, $\forall \alpha_i \in \{-1, 1\}$; and from (4.1.5), we get¹

$$\mathbf{a}_{i,t} = \begin{cases} 1 & \text{if } \mu_{i,t}(+1) \geq \mu_{i,t}(-1), \\ -1 & \text{if } \mu_{i,t}(+1) < \mu_{i,t}(-1), \end{cases} \quad (4.2.1)$$

We can thus proceed to derive the memoryless update rule f_i under the above prescribed settings. This is achieved by the following expression of the action update of agent i at

¹In writing (4.2.1) we follow the convention that agents choose $+1$ when they are indifferent between their two options. Similarly, the sign function is assumed to take the value $+1$ when its argument is zero. This assumption is consistently followed everywhere throughout this paper, except in Proposition 4.2 and its proof in Appendix A.12, see the footnote therein for further details.

time 1. Throughout this section and without any loss of generality, we assume that $\theta = -1$. Recall that $\lambda_1(s_i) := \log(\ell_i(s_i|+1)/\ell_i(s_i|-1))$ is the log-likelihood ratio of signal s_i .

Lemma 4.1 (Time-One Bayesian Actions). *The Bayesian action of agent i at time one following her observations of actions of her neighbors at time zero and her own private signal at time zero is given by $\mathbf{a}_{i,1} = \text{sign}(\sum_{j \in \mathcal{N}_i} w_j \mathbf{a}_{j,0} + \eta_i + \lambda_1(\mathbf{s}_{i,0}))$, where w_i and η_i are constants for each i and they are completely determined by the initial prior and signal structures of agent i and her neighbors.*

The exact expressions of the constants w_i , η_i and their derivations can be found in Appendix A.10. Indeed, making the necessary substitutions we derive the following memory-less update f_i for all $t > 1$: $\mathbf{a}_{i,t} = \text{sign}\left(\sum_{j \in \mathcal{N}_i} w_j \mathbf{a}_{j,t-1} + \eta_i + \lambda_1(\mathbf{s}_{i,t})\right)$. This update rule has a familiar format as a weighted majority and threshold function¹ with the weights and threshold given by w_i and $\mathbf{t}_{i,t} := -\lambda_1(\mathbf{s}_{i,t}) - \eta_i$, the latter being random and time-varying.

Following this model, every agent $i \in [n]$ chooses her action $\mathbf{a}_{i,t} \in \{\pm 1\}$ as the sign of $\sum_{j \in \mathcal{N}_i} w_j \mathbf{a}_{j,t-1} + \eta_i + \lambda_1(\mathbf{s}_{i,t})$. Subsequently, in processing her available data and choosing her action $\mathbf{a}_{i,t}$, every agent seeks to maximize $(\sum_{j \in \mathcal{N}_i} w_j \mathbf{a}_{j,t-1} \mathbf{a}_{i,t} + \mathbf{a}_{i,t}(\eta_i + \lambda_1(\mathbf{s}_{i,t})))$. Hence, we can interpret each of the terms appearing as the argument of the sign function, in accordance with how they influence agent i 's choice of action. In particular, the term $\eta_i + \lambda_1(\mathbf{s}_{i,t})$ represents the propensity of agent i in choosing the false action $\theta_1 := 1$ at time t , and it is determined by the log-likelihood ratio of private signal $\lambda_1(\mathbf{s}_{i,t})$, as well as her innate tendency towards $+1$ irrespective of any observations. The latter is reflected in the constant $\eta_i := \log(v_i(\theta_1)/v_i(\theta_2)) + \log V_i$ based on the log-ratio of her initial prior belief and her knowledge of her neighbor's signal structures, as captured by the constant V_i in (A.10.2) of Appendix A.10. The latter is increasing in $\ell_j(s_j | \theta_1)$ and decreasing in $\ell_j(s_j | \theta_2)$ for any fixed signal $s_j \in \mathcal{S}_j$, $j \in \mathcal{N}_i$; cf. Lemma A.3 of Appendix A.10.

By the same token, we can also interpret the interaction terms $w_j \mathbf{a}_{j,t-1} \mathbf{a}_{i,t}$. Lemma A.4 of Appendix A.10 establishes that constants w_j are non-negative for every agent $j \in [n]$. Hence, in maximizing $\sum_{j \in \mathcal{N}_i} w_j \mathbf{a}_{j,t-1} \mathbf{a}_{i,t} + \mathbf{a}_{i,t}(\eta_i + \lambda_1(\mathbf{s}_{i,t}))$ through her choice of $\mathbf{a}_{i,t} \in \pm 1$ at every time t , agent i aspires to align her choice with as many of her neighbors $j \in \mathcal{N}_i$ as possible. However, in doing so she weighs more the actions of those of her neighbors

¹Majority and threshold functions are studied in the analysis of Boolean functions [178, Chapter 5] and several properties of them including their noise stability are of particular interest [179–181]. This update rule also appears as the McCulloch-Pitts model of an artificial neuron [182], with important applications in neural networks and computing [183]. This update rule is also important in the study of the Glauber Dynamics in the Ising model, where the ± 1 states represent atomic spins. The spins are arranged in a graph and each spin configuration has a probability associated with it depending on the temperature and the interaction structure [184, Chapter 15], [185]. The Ising model provides a natural setting for the study of cooperative behavior in social networks. Recent studies have explored the applications of Ising model for analysis of social and economic phenomena such as rumor spreading [186], study of market equilibria [187], and opinion dynamics [188].

$j \in \mathcal{N}_i$ who have larger constants w_j . The constant $w_j := \log W_j$ with W_j given in (A.10.3) of Appendix A.10 is a measure of observational ability of agent j as relates to our model: agents with large constants w_j are those who hold expert opinions in the social network and they play a major role in shaping the actions of their neighboring agents. Positivity of w_i for any $i \in [n]$, per Lemma A.4 of Appendix A.10, also signifies a case of positive externalities: an agent is more likely to choose an action if her neighbors make the same decision.

4.2.1 Analysis of Convergence and Learning in Ising Networks

To begin with the analysis of the binary action update dynamics derived above, we introduce some useful notation. For all $t \in \mathbb{N}_0$, let $\bar{\mathbf{a}}_t := (\mathbf{a}_{1,t}, \dots, \mathbf{a}_{n,t})^\top$ be the profile of actions taken by all agents at time t . Subsequently, we are interested in the probabilistic evolution of the action profiles $\bar{\mathbf{a}}_t$, $t \in \mathbb{N}_0$ under the following dynamics

$$\mathbf{a}_{i,0} = \text{sign} \left(\log \frac{\nu_i(\theta_1)}{\nu_i(\theta_2)} + \lambda_1(\mathbf{s}_{i,0}) \right), \quad (4.2.2)$$

$$\mathbf{a}_{i,t} = \text{sign} \left(\sum_{j \in \mathcal{N}_i} w_j \mathbf{a}_{j,t-1} + \eta_i + \lambda_1(\mathbf{s}_{i,t}) \right), \quad t \geq 1, \quad (4.2.3)$$

for all $i \in [n]$. The two constants w_i and η_i for each agent i are specified in Appendix A.10 and they depend only on the signal structure and initial prior of that agent and her neighbors. The evolution of action profiles $\bar{\mathbf{a}}_t$ in (4.2.3) specifies a finite Markov chain that jumps between the vertices of the Boolean hyper cube, $\{\pm 1\}^n$. The analysis of the time-evolution of action profiles is facilitated by the classical results from the theory of finite Markov chains with the details spelled out in Appendix A.11.

If the signal structures are rich enough to allow for sufficiently strong signals (having large absolute log-likelihood ratios), or if the initial priors are sufficiently balanced (dividing the probability mass almost equally between θ_1 and θ_2), then any action profiles belonging to $\{\pm 1\}^n$ is realizable as $\bar{\mathbf{a}}_0$ with positive probability under (4.2.2). In particular, any recurrent state of the finite Markov chain over the Boolean cube is reachable with positive probability and the asymptotic behavior can be only determined up to a distribution over the first set of communicating recurrent states that is reached by $\bar{\mathbf{a}}_t$, cf. Proposition A.1 of Appendix A.11. However, if a recurrent class constitutes a singleton, then our model makes sharper predictions: $\lim_{t \rightarrow \infty} \bar{\mathbf{a}}_t$ almost surely exists and is identified as an absorbing state of the finite Markov chain. This special case is treated next due to its interesting implications.

4.2.2 Equilibrium, Consensus, and (Mis-)Learning

We begin by noting that the absorbing states of the Markov chain of action profiles specify the equilibria under the action update dynamics in (4.2.3). Formally, an equilibrium $\bar{\mathbf{a}}^* \in \{\pm 1\}^n$ is such that if the dynamics in (4.2.3) is initialized by $\bar{\mathbf{a}}_0 = \bar{\mathbf{a}}^*$, then with probability one it satisfies $\bar{\mathbf{a}}_t = \bar{\mathbf{a}}^*$ for all $t \geq 1$. Subsequently, the set of all equilibria is completely characterized as the set of all absorbing states, i.e. any action profiles $\bar{\mathbf{a}}^* \in \{\pm 1\}^n$ satisfying $P(\bar{\mathbf{a}}^*, \bar{\mathbf{a}}^*) = 1$, where $P : \{\pm 1\}^n \times \{\pm 1\}^n \rightarrow [0, 1]$ specifies the transition probabilities in the Markov chain of action profiles, as defined in (A.11.1) of Appendix A.11. It is useful to express this condition in terms of the model parameters as follows. The proof is included in Appendix A.12 and with a caveat explained in its footnote.

Proposition 4.2 (Characterization of the Equilibria). *An action profile $(\mathbf{a}_1^*, \dots, \mathbf{a}_n^*) \in \{\pm 1\}^n$ is an equilibrium of (4.2.3) if, and only if, $-\min_{s_i \in \mathcal{S}_i} \mathbf{a}_i^* (\lambda_1(s_i) + \eta_i) \leq \sum_{j \in \mathcal{N}_i} w_j \mathbf{a}_j^* \mathbf{a}_i^*, \forall i \in [n]$.*

Of particular interest are the two action profiles $(1, \dots, 1)^\top$ and $(-1, \dots, -1)^\top$ which specify a consensus among the agents in their chosen actions. The preceding characterization of equilibria is specialized next to highlight the necessary and sufficient conditions for the agents to be at equilibrium whenever they are in consensus.

Corollary 4.1 (Equilibrium at Consensus). *The agents will be in equilibrium at consensus if, and only if, $\max_{s_i \in \mathcal{S}_i} |\lambda_1(s_i) + \eta_i| < \sum_{j \in \mathcal{N}_i} w_j, \forall i \in [n]$.*

The requirement of learning under our model is for the agents to reach a consensus on truth. That is for the action profiles $\bar{\mathbf{a}}_t$ to converge to (θ, \dots, θ) as $t \rightarrow \infty$. In particular, as in Corollary 4.1, we need agents to be at equilibrium when in consensus; hence, there would always be a positive probability for the agents to reach consensus on an untruth: with a positive probability, the agents (mis-)learn.

Next in Section 4.3 we show that when the action space is rich enough to reveal the beliefs of the agents, then the rational but memoryless behavior culminates in a log-linear updating of the beliefs with the observations. The analysis of convergence and learning under these log-linear updates consumes the bulk of that section.

4.3 Log-Linear Learning

In Appendix A.13, for a finite state space, a quadratic utility function, and with agents taking actions over the m -dimensional probability simplex, we calculate the following Bayesian belief at time one, in terms of the observed neighboring beliefs and private signal at time zero. The steps are very similar to Example 4.3.

Lemma 4.2 (Time-One Bayesian Beliefs). *The Bayesian belief of agent i at time one following her observations of beliefs of her neighbors at time zero and her own private signal at time zero is given by*

$$\mu_{i,1}(\hat{\theta}) = \frac{\nu_i(\hat{\theta})l_i(\mathbf{s}_{i,0} | \hat{\theta}) \left(\prod_{j \in \mathcal{N}(i)} \frac{\mu_{j,0}(\hat{\theta})}{\nu_j(\hat{\theta})} \right)}{\sum_{\tilde{\theta} \in \Theta} \nu_i(\tilde{\theta})l_i(\mathbf{s}_{i,0} | \tilde{\theta}) \left(\prod_{j \in \mathcal{N}(i)} \frac{\mu_{j,0}(\tilde{\theta})}{\nu_j(\tilde{\theta})} \right)}, \forall \hat{\theta} \in \Theta. \quad (4.3.1)$$

Subsequently, at any time step $t > 1$, each agent i observes the realized values of $\mathbf{s}_{i,t}$ as well as the current beliefs of her neighbors $\mu_{j,t-1}(\cdot)$, $\forall j \in \mathcal{N}_i$ and forms a refined opinion $\mu_{i,t}(\cdot)$, using the following rule:

$$\mu_{i,t}(\hat{\theta}) = \frac{\nu_i(\hat{\theta})l_i(\mathbf{s}_{i,t} | \hat{\theta}) \left(\prod_{j \in \mathcal{N}_i} \frac{\mu_{j,t-1}(\hat{\theta})}{\nu_j(\hat{\theta})} \right)}{\sum_{\tilde{\theta} \in \Theta} \nu_i(\tilde{\theta})l_i(\mathbf{s}_{i,t} | \tilde{\theta}) \left(\prod_{j \in \mathcal{N}_i} \frac{\mu_{j,t-1}(\tilde{\theta})}{\nu_j(\tilde{\theta})} \right)}, \quad (4.3.2)$$

for all $\hat{\theta} \in \Theta$ and at any $t > 1$. In writing (4.3.2), every time agent i regards each of her neighbors $j \in \mathcal{N}_i$ as having started from prior belief $\nu_j(\cdot)$ and arrived at their currently reported belief $\mu_{j,t-1}(\cdot)$ directly, hence rejecting any possibility of a past history. This is equivalent to the assumption that the reported beliefs of every neighbor are formed from a private observation and a fixed prior, and not through repeated communications.

Such a rule is of course not the optimum Bayesian update of agent i 's belief at any step $t > 1$, because the agent is not taking into account the complete observed history of her private signals and neighbors' beliefs and is instead, basing her inference entirely on the immediately observed signal and neighboring beliefs; hence, the name *memoryless*. Here, the status of a *Rational but Memoryless* agent is akin to a person who is possessed of a knowledge but cannot see how she has come to be possessed of that knowledge. Likewise, it is by the requirement of rationality in such a predicament that we impose a fixed prior $\nu_i(\cdot)$ on every agent i and carry it through for all times t . Indeed, it is the grand tradition of Bayesian statistics, as advocated in the prominent and influential works of [189], [190], [191], [192] and many others, to argue on normative grounds that rational behavior in a decision theoretic framework forces individuals to employ Bayes rule and appropriate it to their personal priors.

4.3.1 Analysis of Convergence and Log-Linear Learning

A main question of interest is whether the agents can learn the true realized value θ :

Definition 4.1 (Learning). *An agent i is said to learn the truth, if $\lim_{t \rightarrow \infty} \mu_{i,t}(\theta) = 1$, \mathcal{P}_θ -almost surely.*

We begin our analysis of convergence and learning under the update rule in (4.3.2) by considering the case of a single agent i , who starts from a prior belief $\nu_i(\cdot)$ and sequentially updates her beliefs according to Bayes rule:

$$\mu_{i,t}(\hat{\theta}) = \frac{\mu_{i,t-1}(\hat{\theta}) \ell_i(\mathbf{s}_{i,t} | \hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mu_{i,t-1}(\tilde{\theta}) \ell_i(\mathbf{s}_{i,t} | \tilde{\theta})}, \forall \hat{\theta} \in \Theta. \quad (4.3.3)$$

The Bayesian belief update in (4.3.3) linearizes in terms of the log-ratio of beliefs and signal likelihoods, $\Phi_{i,t}(\cdot)$ and $\lambda_{\check{\theta}}(\cdot)$, leading to

$$\Phi_{i,t}(\check{\theta}) = \log \left(\frac{\nu_i(\check{\theta})}{\nu_i(\theta)} \right) + \sum_{\tau=0}^t \lambda_{\check{\theta}}(\mathbf{s}_{i,\tau}) \rightarrow \log \left(\frac{\nu_i(\check{\theta})}{\nu_i(\theta)} \right) + (t+1) \mathbb{E}_\theta \{ \lambda_{\check{\theta}}(\mathbf{s}_{i,0}) \}, \quad (4.3.4)$$

\mathcal{P}_θ -almost surely, as $t \rightarrow \infty$; by the strong law of large numbers [177, Theorem 22.1] applied to the sequence of \mathbb{E}_θ -integrable, independent and identically distributed variables: $\lambda_{\check{\theta}}(\mathbf{s}_{i,t})$, $t \in \mathbb{N}_0$. In particular, if $D_{\text{KL}}(\ell_i(\cdot|\theta) \parallel \ell_i(\cdot|\check{\theta})) := -\mathbb{E}_\theta \{ \lambda_{\check{\theta}}(\mathbf{s}_{i,t}) \} > 0$, then $\Phi_{i,t}(\check{\theta}) \rightarrow -\infty$ almost surely and agent i asymptotically rejects the false state $\check{\theta}$ in favor of the true state θ , putting a vanishing belief on the former relative to the latter. Therefore, the single Bayesian agent following (4.3.3) learns the truth if and only if $D_{\text{KL}}(\ell_i(\cdot|\theta) \parallel \ell_i(\cdot|\check{\theta})) > 0$ for all $\check{\theta} \neq \theta$ and the learning is asymptotically exponentially fast at the rate $\min_{\check{\theta} \in \Theta \setminus \{\theta\}} D_{\text{KL}}(\ell_i(\cdot|\theta) \parallel \ell_i(\cdot|\check{\theta}))$ as shown in [165].¹

The preceding result is also applicable to the case of a Bayesian agents with direct (centralized) access to all observations across the network: consider an outside Bayesian agent $\hat{\delta}$ who shares the same common knowledge of the prior and signal structures with the networked agents; in particular, $\hat{\delta}$ knows the signal structures $\ell_i(\cdot|\hat{\theta})$, for all $\hat{\theta} \in \Theta$ and $i \in [n]$; thence, making the same inferences as any other agent when given access to the same observations. Consider next a Gedanken experiment where $\hat{\delta}$ is granted direct access to all the signals of every agent at all times. The analysis leading to (4.3.4) can be applied to the evolution of log belief ratios for $\hat{\delta}$, whose observations at every time $t \in \mathbb{N}_0$

¹Note from the information inequality for the Kullback-Leibler divergence that $D_{\text{KL}}(\cdot|\cdot) \geq 0$ and the inequality is strict whenever $\ell_i(\cdot|\check{\theta}) \neq \ell_i(\cdot|\theta)$, i.e. $\exists s \in \mathcal{S}_i$ such that $\ell_i(s|\check{\theta}) \neq \ell_i(s|\theta)$ [193, Theorem 2.6.3]. Further note that whenever $\ell_i(\cdot|\check{\theta}) \equiv \ell_i(\cdot|\theta)$ or equivalently $D_{\text{KL}}(\ell_i(\cdot|\theta) \parallel \ell_i(\cdot|\check{\theta})) = 0$, then the two states $\check{\theta}$ and θ are statically indistinguishable to agent i : there is no way for agent i to distinguish between $\check{\theta}$ and θ , based only on her received signals. This is because both θ and $\check{\theta}$ induce the same probability distribution on her sequence of observed i.i.d. signals. Since different states $\hat{\theta} \in \Theta$ are distinguished through their different likelihood functions $\ell_i(\cdot|\hat{\theta})$; the more refined such differences are, the better the states are distinguished. Hence, the proposed asymptotic rate is one measure of resolution for the likelihood structure of agent i .

is an element of the product space $\prod_{i \in [n]} \mathcal{S}_i$. Subsequently, the centralized Bayesian beliefs concentrate on the true state at the asymptotically exponentially fast rate of

$$R_n := \min_{\check{\theta} \in \Theta \setminus \{\theta\}} D_{\text{KL}} \left(\prod_{i \in [n]} \ell_i(\cdot | \theta) \parallel \prod_{i \in [n]} \ell_i(\cdot | \check{\theta}) \right) = \min_{\check{\theta} \in \Theta \setminus \{\theta\}} \sum_{i \in [n]} D_{\text{KL}}(\ell_i(\cdot | \theta) \parallel \ell_i(\cdot | \check{\theta})). \quad (4.3.5)$$

Next to understand the evolution of beliefs under the log-linear updates in (4.3.2), consider the network graph structure as encoded by its adjacency matrix A defined as $[A]_{ij} = 1 \iff (j, i) \in \mathcal{E}$, and $[A]_{ij} = 0$ otherwise. For a strongly connected \mathcal{G} the Perron-Frobenius theory [115, Theorem 1.5] implies that A has a simple positive real eigenvalue, denoted by $\rho > 0$, which is equal to its spectral radius. Moreover, the left eigenspace associated with ρ is one-dimensional with the corresponding eigenvector $\bar{\alpha} = (\alpha_1, \dots, \alpha_n)^\top$, uniquely satisfying $\sum_{i=1}^n \alpha_i = 1$, $\alpha_i > 0$, $\forall i \in [n]$, and $\bar{\alpha}^\top A = \rho \bar{\alpha}^\top$. The entry α_i is also called the centrality of agent i and as the name suggests, it is a measure of how central is the location of agent in the network. Our main result state that almost sure learning cannot be realized in a strongly connected network unless it has unit spectral radius which is the case only of a directed circle.

Theorem 4.4 (No Learning when Spectral Radius $\rho > 1$). *In a strongly connected social network and under the memoryless belief updates in (4.3.2), no agents can learn the truth unless the spectral radius $\rho = 1$.*

Proof outline: A complete proof is included in Appendix A.14, but here we provide a description of the mechanism and the interplay between the belief aggregation and information propagation. To facilitate the exposition of the underlying logic we introduce some notation. We define a global (network-wide) random variable $\Phi_t(\check{\theta}) := \sum_{i=1}^n \alpha_i \Phi_{i,t}(\check{\theta})$, where α_i is the centrality of agent i and $\Phi_{i,t}(\check{\theta})$ characterizes how biased (away from the truth and towards $\check{\theta}$) the network beliefs and priors are at each point in time. In particular, if any agent is to learn the truth, then $\Phi_t(\check{\theta}) \rightarrow -\infty$ as $t \rightarrow \infty$ for all the false states $\check{\theta} \in \Theta \setminus \{\theta\}$. To proceed, we define another network-wide random variable $\Lambda_t(\check{\theta}) := \sum_{i=1}^n \alpha_i \lambda_{\check{\theta}}(s_{i,t})$ which characterizes the information content of the observed signals (received information) for the entire network, at each time t . Moreover, since the received signal vectors $\{(s_{1,t}, \dots, s_{n,t}), t \in \mathbb{N}_0\}$ are i.i.d. over time, $\forall \check{\theta} \neq \theta$, $\{\Lambda_t(\check{\theta}), t \in \mathbb{N}_0\}$ constitutes a sequence of i.i.d. random variables satisfying $\mathbb{E}\{\Lambda_t(\check{\theta})\} = -\sum_{i=1}^n \alpha_i D_{\text{KL}}(\ell_i(\cdot | \theta) \parallel \ell_i(\cdot | \check{\theta})) \leq 0$. In order for the agents to learn the true state of the world based on their observations, it is necessary that at each false state $\check{\theta} \neq \theta$ some agent be able to distinguish $\check{\theta}$ from the truth θ , in which case $\mathbb{E}\{\Lambda_t(\check{\theta})\} < 0$, and we can refer to this criterion as *global identifiability* for the true state θ . This global identifiability condition can be also viewed in the following sense: consider a gedanken experiment where an external fully rational observer \hat{o} is granted direct access to all the signals of all agents in

the network and assume further that she shares the same common knowledge of the prior and signal structures with the network agents. Then $\hat{\theta}$ learns the truth if, and only if, it is globally identifiable.

In Appendix A.14 we argue that under the update rules in (4.3.2) the global belief ratio statistics $\Phi_t(\check{\theta})$ evolves as a sum of weighted i.i.d. variables $\rho^\tau \Lambda_{t-\tau}(\check{\theta})$:

$$\Phi_t(\check{\theta}) = \sum_{\tau=0}^t \rho^\tau (\Lambda_{t-\tau}(\check{\theta}) + (1 - \rho)\beta(\check{\theta})), \quad (4.3.6)$$

where $\beta(\check{\theta}) := \sum_{i=1}^n \alpha_i \log(v_i(\check{\theta})/v_i(\theta))$ is a measure of bias in the initial prior beliefs. The weights in (4.3.6) form a geometric progression in ρ ; hence, the variables increase unbounded in their variance and convergence cannot hold true in a strongly connected social network, unless $\rho = 1$. This is due to the fact that ρ upper bounds the average degree of the graph [194, Chapter 2], and every node in a strongly connected graph has degree greater than or equal to one, subsequently $\rho \geq 1$ for all strongly connected graphs. \square

Remark 4.5 (Polarization, data incest and unlearning). The unlearning in the case of $\rho > 1$ in Theorem 4.4, which applies to all strongly connected topologies except directed circles (where $\rho = 1$, see Subsection 4.3.2 below), is related to the inefficiencies associated with social learning and 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 [195]. In particular, when $\rho > 1$ the noise or randomness in the agents' observations is amplified at every stage of network interactions; since the agents fail to correct for the repetitions in the sources of their observations as in the case of persuasion bias argued by DeMarzo, Vayanos and Zwiebel [47], or data incest argued by Krishnamurthy and Hoiles [27]. When $\rho > 1$ the effect of the agents' priors is also amplified through the network interactions and those states $\hat{\theta}$ for which $\beta(\hat{\theta}) > 0$ in (4.3.6), will be asymptotically rejected as $\sum_{\tau=0}^t \rho^\tau (1 - \rho)\beta(\check{\theta}) \rightarrow -\infty$, irrespectively of the observed data $\Lambda_\tau(\check{\theta})$, $\tau \in \mathbb{N}_0$. This phenomenon arises as agents engage in excessive anti-imitative behavior, compensating for the neighboring priors at every period [97]. It is justified as a case of choice shift toward more extreme opinions [5, 6] or group polarization [7, 8], when like-minded people after interacting with each other and under the influence of their mutually positive feedback become more extreme in their opinions, and less receptive of opposing beliefs.

4.3.2 Learning in Circles and General Connected Topologies

For a strongly connected digraph \mathcal{G} , if $\rho = 1$, then it has to be the case that all nodes have degree 1 and the graph is a directed circle. Subsequently, the progression for $\Phi_t(\check{\theta})$ in (4.3.6) reduces to sum of i.i.d. variables in \mathcal{L}^1 and by the strong law of large numbers [177,

Theorem 22.1], it converges almost surely to the mean value


$$\Phi_t(\check{\theta}) = \beta(\check{\theta}) + \sum_{\tau=0}^t \Lambda_\tau(\check{\theta}) \rightarrow \beta(\check{\theta}) + (t+1)\mathbb{E}\{\Lambda_0(\check{\theta})\} \rightarrow -\infty,$$

as $t \rightarrow \infty$, provided that $\mathbb{E}\{\Lambda_0(\check{\theta})\} < 0$, i.e. if the truth is globally identifiable. Note also the analogy with (4.3.4), where $\Lambda_t(\check{\theta})$ is replaced by $\lambda_{\check{\theta}}(\mathbf{s}_{i,t})$ as both represent the observed signal(s) or received information at time t . Indeed, if we further assume that $\nu_i(\cdot) \equiv \nu(\cdot)$ for all i , i.e. all agents share the same common prior, then (4.3.2) for a circular network becomes

$$\mu_{i,t}(\hat{\theta}) = \frac{\mu_{j,t-1}(\hat{\theta})\ell_i(\mathbf{s}_{i,t} | \hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mu_{j,t-1}(\tilde{\theta})\ell_i(\mathbf{s}_{i,t} | \tilde{\theta})}, \forall \hat{\theta} \in \Theta, \quad (4.3.7)$$

where $j \in [\mathfrak{n}]$ is the unique vertex $j \in \mathcal{N}(i)$. Update (4.3.7) replicates the Bayesian update of a single agents in (4.3.3) but the self belief $\mu_{i,t-1}(\cdot)$ on the right-hand side being is replaced by the belief $\mu_{j,t-1}(\cdot)$ of the unique neighbor $\{j\} = \mathcal{N}_i$. Indeed, the learning in this case is asymptotically exponentially fast at the rate $(1/\mathfrak{n}) \min_{\check{\theta} \in \Theta \setminus \{\theta\}} \sum_{j=1}^{\mathfrak{n}} \mathbb{D}_{\text{KL}}(\ell_j(\cdot | \theta) || \ell_j(\cdot | \check{\theta})) = 1/3R_3$; hence, the same exponential rate as that of a central Bayesian can be achieved through the BWR update rule, except for a $1/\mathfrak{n}$ factor that decreases with the increasing cycle length, cf. [165].

Example 4.6 (Eight Agents with Binary Signals in a Tri-State World.) Consider the network of agents in Fig. 4.3 with the true state of the world being 1, the first of the tree possible states $\Theta = \{1, 2, 3\}$. The agents receive binary signals about the true state θ according to the likelihoods listed in the table.



likelihoods	$\hat{\theta} = 1$	$\hat{\theta} = 2$	$\hat{\theta} = 3$
$l_1(\mathbf{s}_{1,t} = 0 \mid \hat{\theta})$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{5}$
$l_2(\mathbf{s}_{2,t} = 0 \mid \hat{\theta})$	$\frac{1}{2}$	$\frac{2}{3}$	$\frac{1}{2}$
$l_3(\mathbf{s}_{3,t} = 0 \mid \hat{\theta})$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$

Figure 4.3: A hybrid structure

We begin by the observation that this network can be thought of as a rooted directed tree, in which the root node is replaced with a directed circle (the root circle).¹ Next note that the root circle is comprised of three agents and none of them can learn the truth on their own. Indeed, agent 3 does not receive any informative signals; therefore, in isolation i.e. using (4.3.3), her beliefs shall never depart from their initial priors. We further set $l_j(\cdot \mid \cdot) \equiv l_3(\cdot \mid \cdot)$ for all $j \in [8] \setminus [3]$, so that all the peripheral follower agents are also unable to infer anything about the true state of the world from their own private signals.

Starting from a uniform common prior and following the proposed rules (4.3.7), all agents asymptotically learn the true state, even though none of them can learn the true state on their own. The plots in Figs. 4.4 and 4.5 depict the evolutions of the beliefs for the third agent as well as the difference between the beliefs for the first and eighth agents. We can further show that all agents learn the true state at the same exponentially fast asymptotic rate. In fact, the three nodes belonging to the directed circle learn the true state of the world at the exponentially fast asymptotic rate of $(1/3)\mathbb{R}_3$ noted above, irrespectively of the peripheral nodes. The remaining peripheral nodes then follow up with the beliefs of root circle nodes, except for a vanishing difference that increases with the increasing distance of a peripheral node from the root circle: following (4.3.7), the first three agents form a circle of leaders where they combine their observations and reach a consensus; every other agent in the network then follows whatever state that the leaders have collectively agreed upon.

¹Any weakly connected digraph \mathcal{G} which has only degree zero or degree one nodes can be drawn as a rooted tree, whose root is replaced by a directed circle, a so-called root circle. This is true since any such digraph can have at most one directed circle and all other nodes that are connected to this circle should be directed away from it, otherwise \mathcal{G} would have to include a node of degree two or higher.

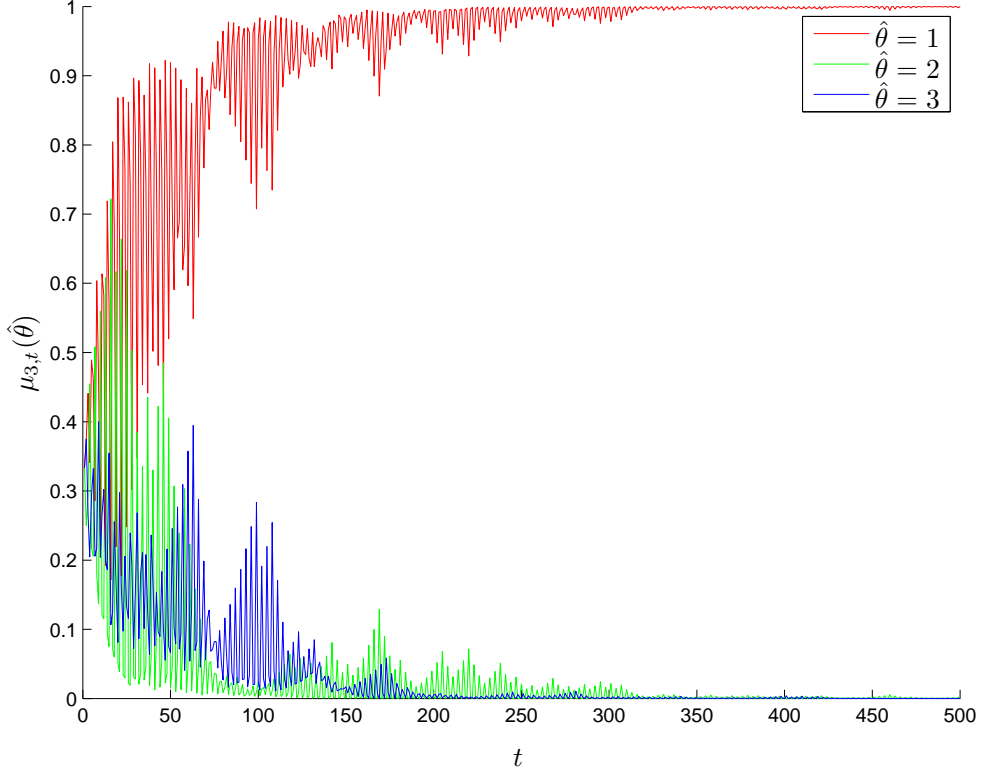


Figure 4.4: Evolution of the third agent's beliefs over time

In the next subsection, we show the application of the update rule in (4.3.7) to general strongly connected topologies where agents have more than just a single neighbor in their neighborhoods. It is proposed to choose a neighbor $j \in \mathcal{N}_i$ independently at random every time and then apply (4.3.7) with the reported belief from that neighbor.

4.3.3 Learning by Random Walks on Directed Graphs

Here we propose the application of the Learning without Recall updates that we described in the previous section to general networks, by requiring that at every time step t , node i make a random choice from her set of neighbors \mathcal{N}_i and uses that choice for the unique j in (4.3.7). To this end, let $\sigma_t \in \prod_{i \in [n]} \mathcal{N}_i$, $t \in \mathbb{N}$ be a sequence of independent and identically distributed random vectors such that $\forall t \in \mathbb{N}$, $\sigma_{t,i} \in \mathcal{N}_i$ is that neighbor of i which she chooses to communicate with at time t . Hence, for all t and any i , (4.3.7) becomes

$$\mu_{i,t}(\hat{\theta}) = \frac{\mu_{\sigma_{t,i},t-1}(\hat{\theta}) \ell_i(\mathbf{s}_{i,t} | \hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mu_{\sigma_{t,i},t-1}(\tilde{\theta}) \ell_i(\mathbf{s}_{i,t} | \tilde{\theta})}, \forall \hat{\theta} \in \Theta. \quad (4.3.8)$$

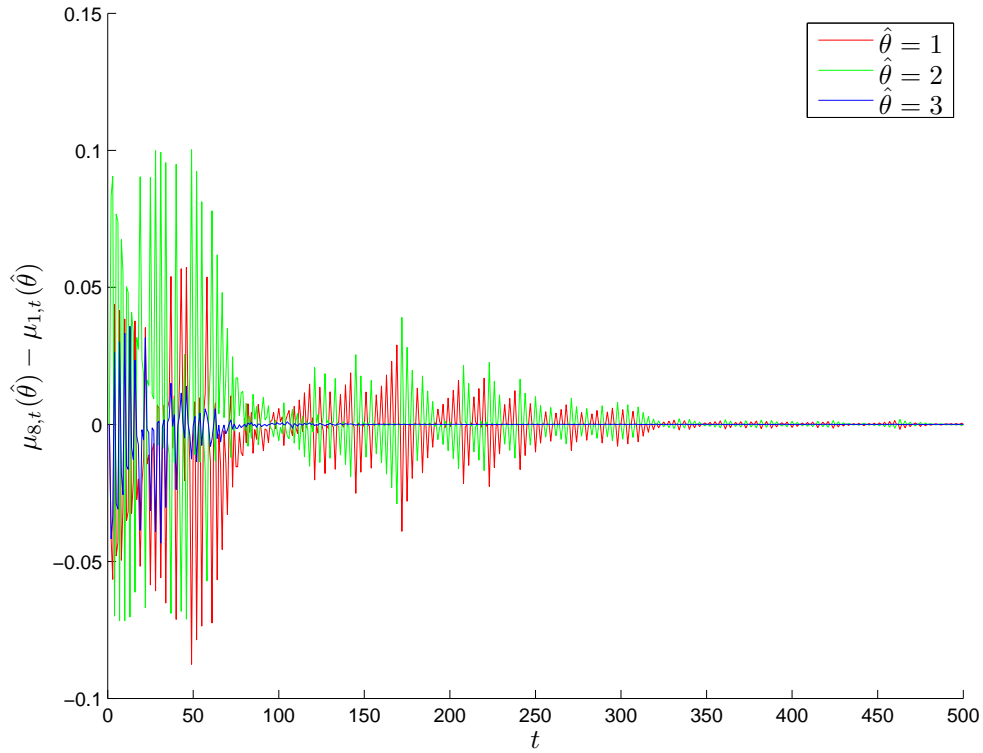


Figure 4.5: The difference between the first and eighth agents' beliefs over time

To proceed, annex the random choice of neighbors for every node $i \in [n]$ and all times $t \in \mathbb{N}$ to the original probability space; and for $t \in \mathbb{N}$ arbitrary, let $\mathcal{P}\{\sigma_{t,i} = j\} = p_{i,j} > 0$. Wherefore, $\sum_{j \in \mathcal{N}_i} p_{i,j} = 1 - p_{i,i} \leq 1$, and $p_{i,j} = 0$ whenever $j \notin \mathcal{N}_i$. Let P be the row stochastic matrix whose (i, j) -th entry is equal to $p_{i,j}$. Let $\mathbb{1}_{\{\sigma_{t,i}=j\}} = 1$ if $\sigma_{t,i} = j$ and $\mathbb{1}_{\{\sigma_{t,i}=j\}} = 0$ otherwise. Then (4.3.8) can be written as

$$\begin{aligned} \mu_{i,t}(\hat{\theta}) &= \sum_{j=1}^n \mathbb{1}_{\{\sigma_{t,i}=j\}} \frac{\mu_{j,t-1}(\hat{\theta}) \ell_i(\mathbf{s}_{i,t} | \hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mu_{j,t-1}(\tilde{\theta}) \ell_i(\mathbf{s}_{i,t} | \tilde{\theta})} \\ &= \ell_i(\mathbf{s}_{i,t} | \hat{\theta}) \prod_{j=1}^n \left(\frac{\mu_{j,t-1}(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mu_{j,t-1}(\tilde{\theta}) \ell_i(\mathbf{s}_{i,t} | \tilde{\theta})} \right)^{\mathbb{1}_{\{\sigma_{t,i}=j\}}} \end{aligned}$$

We dub this procedure “*gossips without recall*” and analyze its properties in Appendix A.15. In particular, consider the following global identifiability condition:

Definition 4.2 (Global Identifiability). *In a strongly connected topology, the true state θ is globally identifiable, if for all $\check{\theta} \neq \theta$ there exists some agent $m \in [n]$ such that*

$D_{\text{KL}}(\ell_m(\cdot|\theta)\|\ell_m(\cdot|\check{\theta})) < 0$, i.e. m can distinguish between $\check{\theta}$ and θ based only on her private signals.

Here again if truth is globally identifiable, all agents learn the truth at an asymptotically exponentially fast rate given by $\min_{\check{\theta} \in \Theta \setminus \{\theta\}} \sum_{j=1}^n \pi_j D_{\text{KL}}(\ell_j(\cdot|\theta)\|\ell_j(\cdot|\check{\theta}))$, where π_j are the probabilities in the stationary distribution of the Markov chain whose transition probabilities are the same as the probabilities for the random choice of neighbors at every point in time.¹

It is notable that the asymptotic rate here is a weighted average of the KL distances $D_{\text{KL}}(\ell_i(\cdot|\theta)\|\ell_i(\cdot|\check{\theta}))$, in contrast with the arithmetic (unweighted) mean $(1/n)R_n$ that arise in the circular case. Both rates are upper bounded by the centralized Bayesian learning rate of R_n calculated in (4.3.5). Finally, we point out that the rate of distributed learning upper bounds the (weighted) average of individual learning rates. It is due to the fact that observations of different agents complement each other, and while one agent may be good at distinguishing one false state from the truth, she can rely on observational abilities of other agents for distinguishing the remaining false states: consider agents 1 and 2 in Example 4.6, the former can distinguish $\hat{\theta} = 2$ from $\theta = 1$, while the latter is good at distinguishing $\hat{\theta} = 3$ from $\theta = 1$; together they can distinguish all states. Hence, the overall rate of distributed learning upper bounds the average of individual learning rates, and is itself upper bounded by the learning rate of a central Bayesian agent:

$$\frac{1}{n} \sum_{i=1}^n \min_{\check{\theta} \in \Theta \setminus \{\theta\}} D_{\text{KL}}(\ell_i(\cdot|\theta)\|\ell_i(\cdot|\check{\theta})) < \min_{\check{\theta} \in \Theta \setminus \{\theta\}} \frac{1}{n} \sum_{i=1}^n D_{\text{KL}}(\ell_i(\cdot|\theta)\|\ell_i(\cdot|\check{\theta})) = \frac{1}{n} R_n < R_n.$$

In Appendix A.15 we establish the following conditions for learning under the without recall updates in (4.3.3) and (4.3.7), where the neighbor j is chosen randomly with strictly positive probabilities specified in transition matrix P :

Theorem 4.7 (Almost-Sure Learning). *Under the gossip without recall updates in a strongly connected network where the truth is globally identifiable, all agents learn the truth asymptotically almost surely. The learning is asymptotically exponentially fast with the rate $\sum_{m=1}^n \pi_m D_{\text{KL}}(\ell_m(\cdot|\theta)\|\ell_m(\cdot|\check{\theta}))$, where (π_1, \dots, π_n) is the stationary distribution of the transition matrix P .*

¹In many distributed learning models over random and switching networks, agents must have positive self-reliant at any time; as for instance in gossip algorithms [196] and ergodic stationary processes [197]. This condition however is relaxed under (4.3.7), as our agents rely entirely on the beliefs of their neighbors every time that they select a neighbor to communicate with. Moreover, unlike the majority of results that rely on the convergence properties of products of stochastic matrices and are applicable only to irreducible and aperiodic communication matrices, cf. [130, Proposition 1]; the convergence results in [166] do not require the transition probability matrix to be aperiodic, as it relies on properties of ergodic Markov chains and holds true for any irreducible, finite-state chain [198, Theorems 1.5.6 and 1.7.7].

Example 4.8 (Eight Agents with Binary Signals in a Tri-State World). As an illustration consider the network of agents in Fig. 4.6 with the true state of the world being 1, the first of the tree possible states $\Theta = \{1, 2, 3\}$. The likelihood structure for the first three agents is given in the table and note that none of them can learn the truth on their own; indeed, agent 3 does not receive any informative signals and her beliefs shall never depart from their initial priors following (4.3.3). We further set $l_j(\cdot | \cdot) \equiv l_3(\cdot | \cdot)$ for all $j \in [8] \setminus [3]$, so that all the remaining agents are also unable to infer anything about the true state of the world from their own private signals.




Figure 4.6: Network Structure for Example 1

likelihoods	$\hat{\theta} = 1$	$\hat{\theta} = 2$	$\hat{\theta} = 3$
$l_1(\mathbf{s}_{1,t} = 0 \hat{\theta})$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{5}$
$l_2(\mathbf{s}_{2,t} = 0 \hat{\theta})$	$\frac{1}{2}$	$\frac{2}{3}$	$\frac{1}{2}$
$l_3(\mathbf{s}_{3,t} = 0 \hat{\theta})$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$

Starting from a uniform common prior and following the proposed gossip without recall scheme with neighbors chosen uniformly at random, all agents asymptotically learn the true state, even though none of them can learn the true state on their own. The plots in Figs. 4.7 and 4.8 depict the belief evolution for the second agent, as well as the difference between the beliefs for the third and eighth agents. It is further observable that all agents learn the true state at the same exponentially fast asymptotic rate of learning.

Fixing the priors over time will not result in convergence of beliefs, except in very specific cases as discussed above. In the next subsection, we investigate the properties of convergence and learning under the update rules in (4.3.2), where the priors $\nu_j(\cdot)$ are replaced by time-varying distributions $\xi_{i,j}(\cdot, t)$ that parametrize the log-linear updating of the agents' beliefs over time.

4.3.4 Log-Linear Learning with Time-Varying Priors

In this subsection, we consider the performance of no-recall belief updates (4.3.2), where the priors $\nu_j(\cdot), \forall j$ are replaced with time-varying distributions $\xi_{i,j}(\cdot, t)$, and argue that




Figure 4.7: Evolution of the second agents beliefs over time

a *Rational but Memoryless* agent i would make her rational inference about the opinion $\mu_{j,t-1}(\cdot)$ that agent j reports to her at time t according to some time-varying prior $\xi_{i,j}(\cdot, t), j \in [n]$. Here, any choice of distributions $\xi_{i,j}(\cdot, t), j \in [n]$ should satisfy the information constraints of a *Rational but Memoryless* agent, so long as such a choice does not require an agent to recall any information other than what she has just observed $s_{i,t}$ and what her neighbors have just reported to her $\mu_{j,t-1}(\cdot), j \in \mathcal{N}_i$. A memoryless yet rational agent of this type can process the beliefs of her neighbors, but cannot recall how these beliefs were formed. Accordingly, (4.3.2) becomes

$$\mu_{i,t}(\check{\theta}) = \frac{\xi_{i,i}(\check{\theta}, t) l_i(s_{i,t} | \check{\theta}) \left(\prod_{j \in \mathcal{N}_i} \frac{\mu_{j,t-1}(\check{\theta})}{\xi_{i,j}(\check{\theta}, t)} \right)}{\sum_{\hat{\theta} \in \Theta} \xi_{i,i}(\hat{\theta}, t) l_i(s_{i,t} | \hat{\theta}) \left(\prod_{j \in \mathcal{N}_i} \frac{\mu_{j,t-1}(\hat{\theta})}{\xi_{i,j}(\hat{\theta}, t)} \right)}, \quad (4.3.9)$$

for all $\check{\theta} \in \Theta$ and at any $t > 1$. In writing (3.3.1), the time one update is regarded as a function that maps the priors, the private signal, and the neighbors' beliefs to the agent's posterior belief; and in using the time one update in the subsequent steps as in (3.3.1), every time agent i regards each of her neighbors $j \in \mathcal{N}_i$ as having started from some prior belief $\xi_{i,j}(\cdot, t)$ and arrived at their currently reported belief $\mu_{j,t-1}(\cdot)$ directly after observing a private signal, hence rejecting any possibility of a past history. Such a rule is of course not the optimum Bayesian update of agent i 's belief at any step $t > 1$, because the agent is not




Figure 4.8: The difference between the third and eighth agents beliefs over time

taking into account the complete observed history of her private signals and neighboring beliefs and is instead, basing her inference entirely on the immediately observed signal and neighboring beliefs; hence, the name *memoryless*. In the next section, we address the choice of random and time-varying priors $\xi_{i,j}(\cdot, t)$, $j \in \mathcal{N}_i$ while examining the properties of convergence and learning under the update rules in (3.3.1).

We begin our analysis of (4.3.9) by forming the log-ratios of beliefs, signals, and priors under the true and false states as:

$$\begin{aligned}\Phi_{i,t}(\check{\theta}) &:= \log(\mu_{i,t}(\check{\theta})/\mu_{i,t}(\theta)), \\ \lambda_{i,t}(\check{\theta}) &:= \log(\ell_i(\mathbf{s}_{i,t}|\check{\theta})/\ell_i(\mathbf{s}_{i,t}|\theta)), \\ \gamma_{i,j}(\check{\theta}, t) &:= \log(\xi_{i,j}(\check{\theta}, t)/\xi_{i,j}(\theta, t)),\end{aligned}$$

for all i, j and t . Consequently, (3.3.1) can be linearized as follows:

$$\Phi_{i,t}(\check{\theta}) = \gamma_{i,i}(\check{\theta}, t) + \lambda_{i,t}(\check{\theta}) + \sum_{j \in \mathcal{N}_i} (\Phi_{j,t-1}(\check{\theta}) - \gamma_{i,j}(\check{\theta}, t)). \quad (4.3.10)$$

The network graph structure is encoded by its adjacency matrix A defined as $[A]_{ij} = 1 \iff (j, i) \in \mathcal{E}$, and $[A]_{ij} = 0$ otherwise. For a strongly connected \mathcal{G} the Perron-

Frobenius theory, cf. [115, Theorem 1.5], implies that A has a simple positive real eigenvalue, denoted by $\rho > 0$, which is equal to its spectral radius. Moreover, the left eigenspace associated with ρ is one-dimensional with the corresponding eigenvector $\bar{\alpha} = (\alpha_1, \dots, \alpha_n)^\top$, uniquely satisfying $\sum_{i=1}^n \alpha_i = 1$, $\alpha_i > 0$, $\forall i \in [n]$, and $\bar{\alpha}^\top A = \rho \bar{\alpha}^\top$. The quantity α_i is also known as the eigenvector centrality of vertex i in the network, cf. [199, Section 7.2]. Multiplying both sides of (4.3.10) by α_i and summing over all i we obtain that

$$\begin{aligned} \Phi_t(\check{\theta}) &= \text{tr}\{\Xi_t(\check{\theta})\} + \Lambda_t(\check{\theta}) + \rho \Phi_{t-1}(\check{\theta}) - \text{tr}\{\Xi_t(\check{\theta})A^\top\} \\ &= \sum_{\tau=0}^t \rho^\tau (\Lambda_{t-\tau} + \text{tr}\{(I - A^\top)\Xi_{t-\tau}(\check{\theta})\}), \end{aligned} \quad (4.3.11)$$

where $\Phi_t(\check{\theta}) := \sum_{i=1}^n \alpha_i \phi_{i,t}(\check{\theta})$ and $\Lambda_t(\check{\theta}) := \sum_{i=1}^n \alpha_i \lambda_{i,t}(\check{\theta})$ are global (network-wide) random variables, and $\Xi_t(\check{\theta})$ is a random $n \times n$ matrix whose i, j -th entry is given by $[\Xi_t(\check{\theta})]_{i,j} = \alpha_i \gamma_{i,j}(\check{\theta}, t)$. At each epoch of time, $\Phi_t(\check{\theta})$ characterizes how biased (away from the truth and towards $\check{\theta}$) the network beliefs are, and $\Lambda_t(\check{\theta})$ measures the information content of the received signal across all the agents in the network. Note that in writing (4.3.11), we use the fact that

$$\sum_{i=1}^n \alpha_i \sum_{j \in \mathcal{N}_i} \phi_{j,t-1}(\check{\theta}) = \bar{\alpha}^\top A \bar{\Phi}_{t-1}(\check{\theta}) = \rho \bar{\alpha}^\top \bar{\Phi}_{t-1}(\check{\theta}) = \rho \Phi_{t-1}(\check{\theta}),$$

where $\bar{\Phi}_t(\check{\theta}) := (\phi_{1,t}(\check{\theta}), \dots, \phi_{n,t}(\check{\theta}))^\top$. On the other hand, since the received signal vectors $\{\mathbf{s}_{i,t}, i \in [n], t \in \mathbb{N}_0\}$ are i.i.d. over time, $\{\Lambda_t(\check{\theta}), t \in \mathbb{N}_0\}$ constitutes a sequence of i.i.d. random variables satisfying

$$\mathbb{E} \{\Lambda_t(\check{\theta})\} = \sum_{i=1}^n \alpha_i \mathbb{E}\{\lambda_{i,t}(\check{\theta})\} = \sum_{i=1}^n \alpha_i \lambda_i(\check{\theta}) \leq 0, \quad (4.3.12)$$

where $\bar{\lambda}(\check{\theta}) := (\lambda_1(\check{\theta}), \dots, \lambda_n(\check{\theta}))^\top :=$

$$- (\text{D}_{\text{KL}}(\ell_1(\cdot|\check{\theta})\|\ell_1(\cdot|\theta)), \dots, \text{D}_{\text{KL}}(\ell_n(\cdot|\check{\theta})\|\ell_n(\cdot|\theta)))^\top, \quad (4.3.13)$$

and the non-positivity of (4.3.12) follows from the information inequality for the Kullback-Leibler divergence: $\text{D}_{\text{KL}}(\cdot|\cdot) \geq 0$, and is strict whenever $\ell_i(\cdot|\check{\theta}) \neq \ell_i(\cdot|\theta)$ for some i , i.e. $\exists s \in \mathcal{S}_i, i \in [n]$ such that $\ell_i(s|\check{\theta}) \neq \ell_i(s|\theta)$, cf. [193, Theorem 2.6.3]. In particular, if for all $\check{\theta} \neq \theta$ there exists an agent i with $\lambda_i(\check{\theta}) < 0$, then $\mathbb{E} \{\Lambda_t(\check{\theta})\} < 0$ and we say that the truth θ is *globally identifiable*. Indeed, if any agent is to learn the truth, then we need that $\Phi_t(\check{\theta}) \rightarrow -\infty$ as $t \rightarrow \infty$ for all the false states $\check{\theta} \neq \theta$.

Fixing the priors over time will not result in convergence of beliefs, except in very spe-

cific cases such as directed circles or with randomly chosen neighbors that were discussed in preceding subsections. In the sequel, we investigate the properties of convergence and learning under the update rules in (3.3.1), where the parameterizing priors $\xi_{i,j}(\cdot, t)$ are chosen to be random and time-varying variables, leading to the log-linear updating of the agents' beliefs over time. We distinguish two cases depending on whether the agents do not recall their own self-beliefs or they do, leading respectively to time-invariant or time-varying log-linear update rules.

Priors Set to a Geometric Average:

We propose setting the time-varying priors $\xi_{i,j}(\cdot, t), j \in \mathcal{N}_i$ of each agent i and at every time t , proportionally to the geometric average of the beliefs reported to her by all her neighbors at every time t : $\prod_{j \in \mathcal{N}_i} \mu_{j,t-1}(\cdot)^{1/\deg(i)}$. Therefore, (3.3.1) becomes

$$\mu_{i,t}(\check{\theta}) = \frac{l_i(\mathbf{s}_{i,t} | \check{\theta}) \left(\prod_{j \in \mathcal{N}_i} \mu_{j,t-1}(\check{\theta}) \right)^{1/\deg(i)}}{\sum_{\hat{\theta} \in \Theta} l_i(\mathbf{s}_{i,t} | \hat{\theta}) \left(\prod_{j \in \mathcal{N}_i} \mu_{j,t-1}(\hat{\theta}) \right)^{1/\deg(i)}}. \quad (4.3.14)$$

To analyze the evolution of beliefs with this choice of priors, let

$$\bar{\lambda}_t(\check{\theta}) := (\lambda_{1,t}(\check{\theta}), \dots, \lambda_{n,t}(\check{\theta}))^T,$$

be the stacked vector of log-likelihood ratios of received signals for all agents at time t . Hence, we can write the vectorized update $\bar{\Phi}_t(\check{\theta}) = T\bar{\Phi}_{t-1}(\check{\theta}) + \bar{\lambda}_t(\check{\theta})$, where T is the normalized adjacency of the graph defined by $[T]_{ij} = 1/\deg(i)[A]_{ij}$ for all i and j . We can now iterate the vectorized update to get $\bar{\Phi}_t(\check{\theta}) = \sum_{\tau=0}^{t-1} T^\tau \bar{\lambda}_{t-\tau}(\check{\theta}) + T^t \bar{\Psi}(\check{\theta})$. Next note from the analysis of convergence for DeGroot model, cf. [130, 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), then $\lim_{\tau \rightarrow \infty} T^\tau = \mathbb{1} \bar{s}^T$, where $\bar{s} := (s_1, \dots, s_n)^T$ 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, the Cesàro mean together with the strong law implies that $\lim_{t \rightarrow \infty} \frac{1}{t} \bar{\Phi}_{i,t}(\check{\theta}) = -\sum_{i=1}^n s_i \lambda_i(\check{\theta})$, almost surely for all $\check{\theta} \neq \theta$, and the agents learn the truth asymptotically exponentially fast, at the rate $\min_{\check{\theta} \neq \theta} \sum_{i=1}^n -s_i \lambda_i(\check{\theta})$; $\lambda_i(\check{\theta})$ as defined in (4.3.13) measures the ability of agent i to distinguish between $\check{\theta}$ and θ .

If the agents $i \in [n]$ recall their self-beliefs $\mu_{i,t-1}(\cdot), i \in [n]$ when making decisions or performing inferences at time t , then we set $\xi_{i,i}(\cdot, t) \equiv \mu_{i,t-1}(\cdot)$ for all i and t . Furthermore, we set $\xi_{i,j}(\cdot, t) \equiv \mu_{j,t-1}(\cdot)^{\eta_t} / \zeta_j(t)$ for all $i, j \in \mathcal{N}_i$ and t , where $\zeta_j(t) := \sum_{\hat{\theta} \in \Theta} \mu_{j,t-1}(\hat{\theta})^{\eta_t}$ is the normalization constant to make the exponentiated probabilities sum to one. The choice of $0 < \eta_t < 1$ as time-varying exponents to be determined shortly, is motivated by the requirements of convergence under (3.3.1). Subsequently, in Appendix

A.16 we investigate the properties of convergence and learning for the following log-linear update rule with time-varying coefficients

$$\boldsymbol{\phi}_{i,t}(\check{\theta}) = \boldsymbol{\phi}_{i,t-1}(\check{\theta}) + \lambda_{i,t}(\check{\theta}) + (1 - \eta_t) \sum_{j \in \mathcal{N}_i} \boldsymbol{\phi}_{j,t-1}(\check{\theta}),$$

where $1 - \eta_t$ is the weight that the agent puts on her neighboring beliefs (relative to her own) at any time t . Let $\alpha_t := \rho(1 - \eta_t)$. In Appendix A.16, we identify $\sum_{u=1}^{\infty} \alpha_u < \infty$ as a necessary condition for convergence of beliefs with log-linear time-varying updates, in the case of agents who have no recollection of the past, excepting their own immediate beliefs. With α_t representing the relative weight on the neighboring beliefs, such agents facing individual identification problems can still learn the truth in a strongly connected and aperiodic network by relying on each other's observations; provided that as time evolves they put less and less weight on the neighboring beliefs and rely more on their private observations: $\alpha_t \rightarrow 0$ as $t \rightarrow \infty$.

For such agents who recall their immediate self-beliefs $\boldsymbol{\mu}_{i,t}(\cdot)$, if we relax the requirement that $\xi_{i,i}(\cdot, t) \equiv \boldsymbol{\mu}_{i,t-1}(\cdot)$ then it is possible to achieve asymptotic almost sure exponentially fast learning using time-invariant updates just as in (4.3.14). In particular, for any $0 < \eta < 1$ fixed, we can set $\xi_{i,i}(\cdot, t)$ proportional to $\boldsymbol{\mu}_{i,t-1}(\cdot)^\eta$ for all i , and we can further set $\xi_{i,j}(\cdot, t)$ at every time t , for any i , and all $j \in \mathcal{N}_i$ proportional to $\boldsymbol{\mu}_{j,t-1}(\cdot)^{1-(1-\eta)/\deg(i)}$. Subsequently, (3.3.1) becomes

$$\boldsymbol{\mu}_{i,t}(\check{\theta}) = \frac{l_i(\mathbf{s}_{i,t} \mid \check{\theta}) \boldsymbol{\mu}_{i,t-1}(\check{\theta})^\eta \left(\prod_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}(\check{\theta}) \right)^{\frac{1-\eta}{\deg(i)}}}{\sum_{\hat{\theta} \in \Theta} l_i(\mathbf{s}_{i,t} \mid \hat{\theta}) \boldsymbol{\mu}_{i,t-1}(\hat{\theta})^\eta \left(\prod_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}(\hat{\theta}) \right)^{\frac{1-\eta}{\deg(i)}}}. \quad (4.3.15)$$

To analyze (4.3.15), we form the log-belief and likelihood ratios and set $B = (\eta I + (1 - \eta)T)$, where T is the same normalized adjacency as in Subsection 4.3.4. Hence, we recover the vectorized iterations $\bar{\boldsymbol{\phi}}_t(\check{\theta}) = B \bar{\boldsymbol{\phi}}_{t-1}(\check{\theta}) + \bar{\lambda}_t(\check{\theta}) = \sum_{\tau=0}^t B^\tau \bar{\lambda}_{t-\tau}(\check{\theta}) + B^t \bar{\psi}(\check{\theta})$ and it follows from [200, Theorems 5.1.1 and 5.1.2] that for a strongly connected network \mathcal{G} , $\lim_{\tau \rightarrow \infty} B^\tau = \mathbf{1} \bar{\mathbf{s}}^\top$, where $\bar{\mathbf{s}} := (s_1, \dots, s_n)^\top$ is the unique stationary distribution associated with the Markov chain whose probability transition matrix is T (or equivalently B); whence it follows from the Cesàro mean and the strong law that the agents learn the truth asymptotically exponentially fast, at the rate $\min_{\check{\theta} \neq \theta} \sum_{i=1}^n -s_i \lambda_i(\check{\theta})$, similar to Subsection 4.3.4. Note that here, unlike Subsection 4.3.4 but similarly to [166], we only use properties of ergodic chains and existence and uniqueness of their stationary distributions; hence, relaxing the requirement for the social network to be aperiodic.

An interesting extension of the streaming data model is when the agents' access to new private observations is intermittent. In Appendix G, we consider the case where nodes

(agents) have an initial data set and they communicate their beliefs in order to combine their local data and come up with the model parameter that best describes all their data collectively; i.e. the global maximum likelihood estimator. We then shift attention to a learning framework where agents instead of starting with an initial data set receive new observations at every round. The data received at every point can provide differing and possibly complementary information about the unknown parameter and the number of data points that is observed at every round varies randomly.

Chapter 5

Conclusions & Future Directions

In this final chapter, we summarize our contributions from the two parts of the thesis. We also provide some reflections and comparative remarks to highlight the intuition behind several of the results and relations between them. We end by giving some directions for future work on each of the two parts.

5.1 Summary & Discussion of Part one

In **Chapter 2**, we analyzed recommendations of rational agents in a group decision process, as they each observe an exogenous initial private signal and are exposed to the recommendations of (a subset) of other group members in the ensuing decision epochs. Such agents in a group decision process (GDP) have purely informational externalities, but they still need to interpret (and learn from) the actions of others subject to the fact that they are acting rationally. Other members' actions reveal additional information about the state of the world, which can be then used to make better future recommendations. Indeed, the actions of neighbors are informative only to the extent that they reveal information about their private signals; and as time marches on, more information is revealed about the private signals of neighbors, and neighbors of neighbors, etc. Hence, after a long enough time, all players would be (partially) informed about the private signals of all other players if the graph is connected. We analyzed the complexity of decision-making in this information structure. Iterated elimination of infeasible signals (IEIS) curbs some of the complexities of inference in group decision-making, although its running time is exponential. These computations simplify and become efficient in a POSET structure where the agent has direct access to all observations of her neighbors (except their private signals). The computations also simplify in special symmetric settings, for example with i.i.d. binary signals over a directed path, or a rooted (directed) tree (cf. Appendix D).

In the special case that agents reveal their beliefs to each other, we introduce and analyze a structural property of the graph, referred to as transparency, which plays a critical role in characterizing the complexity of the computations when forming Bayesian posterior beliefs. Bayesian beliefs in transparent structures are both easy to compute and statically efficient; in the sense that they coincide with the Bayesian posterior of the agent, had she

direct access to the private signals of all agents whom she has observed, either directly, or indirectly through their influences on her neighbors, neighbors of neighbors, etc.

We proved the \mathcal{NP} -hardness of the Bayesian belief exchange problem by providing reductions that show well-known \mathcal{NP} -complete problems such as SUBSET-SUM and EXACT-COVER are special cases of GDP. The former relies on the increasing variety of signal types and the latter relies on the increasing neighborhood size.

Transparency of the network structure to agent i allows her to trace the reported beliefs of her neighbors directly to their root causes which are the private signals of other agents. When transparency is violated, the neighboring beliefs are complicated highly non-linear functions of the signal likelihoods and the forward reasoning approach can no longer be applied to search for possible signals that lead to observed beliefs; indeed, if transparency is violated, then the observed beliefs only represent what signal profiles are regarded as feasible by the neighbors. This is quite different from the transparent case where the beliefs of neighbors directly reflect their knowledge about the likelihoods of signals that occur in the higher-order neighborhoods. In other words, in a nontransparent structure, agent i cannot use the reported beliefs of her neighbors to make direct inferences about the original causes of those reports which are the private signals in the higher-order neighborhoods. Instead, to keep track of the feasible signal profiles that are consistent with her observations agent i should consider what beliefs other agents would hold under each of the possible signal profiles and to prune the infeasible ones following an IEIS procedure. A similar observation can be made in the case of POSETs and actions: as compared with general graphs, POSETs remove the need to simulate the network at a given signal profile to reject or approve it. Instead, we can directly verify if each individual private signal agrees with the observed action of its respective agent and if it does not, then it is rejected and removed from the list of feasible private signals.

Although determining the posterior beliefs during a GDP is, in general, \mathcal{NP} -hard, for transparent structures the posterior belief at each step can be computed efficiently using the reported beliefs of the neighbors. Furthermore, the optimality of belief exchange over transparent structures is a unique structural feature of the inference set up in GDP. It provides an interesting and distinct addition to known optimality conditions for inference problems over graphs. In particular, the transparent structures over which efficient and optimal Bayesian belief exchange is achievable include many loopy structures in addition to trees.¹

¹It is well known that if a Bayesian network has a tree (singly connected or polytree) structure, then efficient inference can be achieved using belief propagation (message passing or sum-product algorithms), cf. [201]. However, in general loopy structures, belief propagation only gives a (potentially useful) approximation of the desired posteriors [202]. Notwithstanding, our Bayesian belief exchange set up also greatly simplifies in the case of tree structures, admitting a trivial sum of innovations algorithm. The authors in [203] study the complexity landscape of inference problems over graphical models in terms of their treewidth. For bounded treewidth structures the junction-tree method (performing belief propagation on the tree decomposi-

5.2 Summary & Discussion of Part Two

In **Chapter 3**, we proposed 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 [45].

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.

tion of the graph) works efficiently [204] but there is no class of graphical models with unbounded treewidth in which inference can be performed in time polynomial in treewidth.

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.

In **Chapter 4**, we extended the no-recall model of inference and belief formation to a social and observational learning scenario, where agents attempt to learn some unknown state of the world which belongs to a finite state space. Conditioned on the true state, a sequence of i.i.d. private signals are generated and observed by each agent of the network. The private signals do not provide the agents with adequate information to identify the truth on their own. Hence, agents interact with their neighbors to augment their imperfect observations with those of their neighbors.

Following the no-recall approach, the complexities of a fully rational inference at the forthcoming epochs are avoided, while some essential features of Bayesian inference are preserved. We analyzed the specific form of no-recall updates in two cases of binary state and action space, as well as a finite state space with actions taken over the probability simplex. In the case of binary actions the no-recall updates take the form of a linear majority rule, whereas if the action spaces are rich enough for the agents to reveal their beliefs, then belief updates take a log-linear format. In each case we investigate the properties of convergence, consensus and learning. The latter is particularly interesting when the truth is identifiable through the aggregate private observations of all individuals in a strongly connected social network, but not individually.

On the one hand, the specific forms of the no-recall update rules in each case help us

better understand the mechanisms of naive inference, when rational agents are devoid of their ability to make recollections. On the other hand, our results also highlight the consequences of such naivety in shaping the mass behavior; by comparing our predictions with the rational learning outcomes. In particular, we saw in Subsection 4.2.2 that there is a positive probability for rational but memoryless agents in an Ising model to mis-learn by reaching consensus on an untruth. However Bayesian (fully rational) beliefs constitute a bounded martingale; hence, when truth is identifiable and number of observations increases, the beliefs of rational agents converge almost surely to a point mass centered at the true state [83, 84]. Similarly Theorem 4.4 states the impossibility of asymptotic learning under the no-recall belief updates, whenever the spectral radius of the interconnection graph adjacency is greater than one.

Finally, based on our results in Appendices E and G, we can point out some key differences between the linear and log-linear update rules in the way they accommodate streaming observations. The requirement of averaging over time for linear updates necessitates that new observations be discounted as $1/t$ with increasing time, which avoids fluctuations with new observations in the limit. The same principle governs the discounting or diminishing step sizes in the case of consensus+innovation algorithm [205], as well as other online learning methods [206]. However, in case of log-linear update rules no such discounting is necessary. Because the product-nature of such rules imply that as beliefs approach a point mass their multiplication with the product of likelihoods of new observations will have less and less effect. This in turn allows us to effectively accommodate the varying sizes of data sets at every time-period using log-linear update rules with fixed coefficients.

Another key difference between the linear and log-linear updates is that the weights in the former need to be adjusted for the initial sample sizes, whereas the latter require no such adjustment of weights. Accordingly, in the linear case an agent weighs each neighbor's report differently and in accordance with the quality and quantity of their samples. On the other hand, when communicating their beliefs for log-linear updating the quality of each neighbor's signal is already internalized in their reported beliefs; hence, when incorporating its neighboring beliefs, an agent regards the reported beliefs of all its neighbors equally, and irrespective of the quality of their sample points. These observations lead to the conclusion that: *log-linear aggregation schemes (as opposed to linear ones) are very effective design tools for dealing with various types of heterogeneities that arise in networked systems.*

5.3 Directions for Future Work

In extensions of part one, it would be particularly interesting if one can provide a tight graphical characterization for transparency or provide other useful sufficient conditions that ensure transparency and complement our Proposition 2.2. Furthermore, the nature of the two reductions in Appendices A.3 and A.4 leave space for strengthening the complexity class of the GROUP-DECISION problem beyond \mathcal{NP} , cf. Remark 2.8 in Section 2.2.3. Another possibility is to prove that the beliefs are even hard to approximate, by exploiting the gap that exists between the log-ratio of beliefs, depending on whether the underlying instance of the decision problem is feasible or not. More importantly, one would look for other characterizations of the complexity landscape and find other notions of simplicity that are different from transparency. Another promising venue is to further investigate the algorithmic and complexity theoretic foundation for the actions. Our hardness results are specific to the case of belief exchange and it would be valuable to develop parallel results for the action exchange case. An open problem is to investigate configurations and structures for which the computation of Bayesian actions is achievable in polynomial-time. It is also of interest to know the quality of information aggregation; i.e. under what conditions on the signal structure and network topology, Bayesian actions coincide with the best action given the aggregate information of all agents. In Appendix D, we take some preliminary steps in this direction for i.i.d. binary signals and actions.

For part two, it is valuable to extend the no-recall model of group decision-making to agents with bounded memory. We take some initial steps along this direction in Appendix F with one-step recall belief update rules. Nonetheless, the analysis of Bayesian update even in the simplest cases become increasingly complex. As an example, consider a rational agent who recalls only the last two epochs of her past. In order for such an agent to interpret her observations in the penultimate and ultimate steps, she needs not only a prior to interpret her neighbor's beliefs at the penultimate step, but also a prior to interpret her neighbor's inferences about what she reported to them at the penultimate step leading to their ultimate beliefs. In other words, she needs a prior on what her neighbor's regard as her prior when they interpret what she reports to them as her penultimate belief. Indeed, such belief hierarchies are commonplace in game-theoretic analysis of incomplete information and are captured by the formalism of type space [207, 208].

On a broader scale, our Bayesian model of group decision-making in part one, extends the project selection model of Sah and Stiglitz [92–94] to an iterative setup with managers iterating on their decisions before committing to a final choice. Such connections to organization science can be deepened to obtain useful insights about the operations of teams in medical, legal and other industrial decision-making organizations. We have explored these connections to some extent in Chapter 2 (Remark 2.2) by pointing out how addition of new members in an organization can be done in a way that curbs the increased com-

plexity of organizational decision-making. Similarly, in Subsection 3.3.3 we give a method for organizing teams of decision-makers to avoid redundancy and increase the efficiency of the aggregate group-decision outcome. However, there is much space for improvement in this regard. For example, the transparent structures of Subsection 2.2 can be developed into a useful architecture for organizational decision-making. Overall with the exception of the hand full of works that we explained in the Introduction (Chapter 1), the algorithmic, computational, and optimality aspects of decision-making in organizations are vastly unexplored, leaving many untapped potentials for applying our techniques in organization science.

The no-recall model of part two, 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 future research, one can 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. Examples for such as approach can be found in [209, 210] where the effects of confirmatory and hindsight biases are modeled and analyzed.

Appendix A

Proofs & Mathematical Derivations

A.1 Complexity of Bayesian Decisions Using (A1: BAYES-GROUP)

Suppose that agent i has reached her t -th decision epoch in a general network structure. Given her information (II) at time t , for all $\bar{s} = (s_1, \dots, s_n) \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ and any $j \in \bar{\mathcal{N}}_i^{t+1}$ she has to update $\mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$ into $\mathcal{I}_{j,t+1-\nabla(j,i)}^{(i)}(\bar{s}) \subset \mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$. If $\nabla(j,i) = t+1$ then agent j is being considered for the first time at the t -th decision epoch and $\mathcal{I}_{j,0}^{(i)}(\bar{s}) = \{s_j\} \times \prod_{k \neq j} \mathcal{S}_k$ is initialized without any calculations. However if $\nabla(j,i) \leq t$, then $\mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$ can be updated into $\mathcal{I}_{j,t+1-\nabla(j,i)}^{(i)}(\bar{s}) \subset \mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$ only by verifying the condition $\alpha_{k,t-\nabla(j,i)}^{(i)}(\tilde{s}) = \alpha_{k,t-\nabla(j,i)}^{(i)}(\bar{s})$ for every $\tilde{s} \in \mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$ and $k \in \mathcal{N}_j$: any $\tilde{s} \in \mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$ that violates this condition for some $k \in \mathcal{N}_j$ is eliminated and $\mathcal{I}_{j,t+1-\nabla(j,i)}^{(i)}(\bar{s})$ is thus obtained by pruning $\mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$.

Verification of $\alpha_{k,t-\nabla(j,i)}^{(i)}(\tilde{s}) = \alpha_{k,t-\nabla(j,i)}^{(i)}(\bar{s})$ involves calculations of $\alpha_{k,t-\nabla(j,i)}^{(i)}(\tilde{s})$ and $\alpha_{k,t-\nabla(j,i)}^{(i)}(\bar{s})$ according to (2.1.4). The latter requires the addition of $\text{card}(\mathcal{I}_{k,t-\nabla(j,i)}^{(i)}(\bar{s}))$ product terms $u_k(\alpha_k, \hat{\theta}) \mathcal{L}(\bar{s}' | \hat{\theta}) \nu(\hat{\theta}) = u_k(\alpha_k, \hat{\theta}) \ell_1(s'_1 | \hat{\theta}) \dots \ell_n(s'_n | \hat{\theta}) \nu(\hat{\theta})$ for each $\bar{s}' \in \mathcal{I}_{k,t-\nabla(j,i)}^{(i)}(\bar{s})$, $\hat{\theta} \in \Theta$, and $\alpha_k \in \mathcal{A}_k$ to evaluate the left hand-side of (2.1.4). Hence, we can estimate the total number of additions and multiplications required for calculation of each (conditionally) feasible action $\alpha_{k,t-\nabla(j,i)}^{(i)}(\bar{s})$ as $A \cdot (n+2) \cdot m \cdot \text{card}(\mathcal{I}_{k,t-\nabla(j,i)}^{(i)}(\bar{s}))$, where $m := \text{card}(\Theta)$ and $A = \max_k \in [n] \text{card}(\mathcal{A}_k)$. Hence the total number of additions and multiplications undertaken by agent i at time t for determining actions $\alpha_{k,t-\nabla(j,i)}^{(i)}(\bar{s})$ can be estimated as follows:

$$A \cdot (n+2) \cdot \text{card}(\Theta) \cdot \sum_{j: \nabla(j,i) \leq t} \sum_{k \in \mathcal{N}_j} \text{card}(\mathcal{I}_{k,t-\nabla(j,i)}^{(i)}(\bar{s})) \leq A \cdot (n+2) \cdot n \cdot M^{n-1} \cdot m, \quad (\text{A.1.1})$$

where we upper-bound the cardinality of the union of the higher-order neighborhoods of agent i by the total number of agents: $\text{card}(\bar{\mathcal{N}}_i^{t+1}) \leq n$ and use the inclusion relationship $\mathcal{I}_{k,t-\nabla(j,i)}^{(i)}(\bar{s}) \subset \mathcal{I}_{k,0}^{(i)}(\bar{s}) = \{s_k\} \times \prod_{j \neq k} \mathcal{S}_j$ to upper-bound $\text{card}(\mathcal{I}_{k,t-\nabla(j,i)}^{(i)}(\bar{s}))$ by M^{n-1} where M is the largest cardinality of finite signal spaces, \mathcal{S}_j , $j \in [n]$. As the above calculations

are performed at every signal profile $\bar{s} \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ the total number of calculations (additions and multiplications) required for the Bayesian decision at time t can be bounded as follows:

$$A \cdot M^n \leq A \cdot C_1 \leq (n + 2) \cdot n \cdot M^{2n-1} \cdot m, \quad (\text{A.1.2})$$

where we apply (A.1.1) for the right-hand side. In particular, the calculations grow exponential in the number of agents n . Once agent i calculates the action sets $\alpha_{k,t-\nabla(j,i)}^{(i)}(\bar{s})$ for all $k \in \mathcal{N}_j$ with $\nabla(j,i) \leq t$ she can then update the feasible signal profiles $\mathcal{I}_{j,t-\nabla(j,i)}^{(i)}(\bar{s})$, following step 1 of (A1), to obtain $\mathcal{I}_{j,t+1-\nabla(j,i)}^{(i)}(\bar{s})$ for all $j : \nabla(j,i) \leq t + 1$ and any $\bar{s} \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$. This in turn enables her to calculate the conditional actions of her neighbors $\alpha_{j,t}^{(i)}(\bar{s})$ at every signal profile and to eliminate any \bar{s} for which the conditionally feasible action set $\alpha_{j,t}^{(i)}(\bar{s})$ does not agree with the observed action $\mathbf{a}_{j,t}$ for some $j \in \mathcal{N}_i$. She can thus update her list of feasible signal profiles from $\mathcal{I}_{i,t}$ to $\mathcal{I}_{i,t+1}$ and adopt the corresponding Bayesian belief $\mu_{i,t+1}$ and action $\mathbf{a}_{i,t+1}$. The latter involves an additional $(n + 2)m\text{Acad}(\mathcal{I}_{i,t+1})$ additions and multiplication which are nonetheless dominated by the number calculations required in (A.1.2) for the simulation of other agents' actions at every signal profile. \square

A.2 Computational Complexity of (A2:BAYES-POSET)

According to (I2), in a POSET structure at time t agent i has access to the list of feasible private signals for each of her neighbors: $\mathcal{S}_{j,t}$, $j \in \mathcal{N}_i$ given their observations up until that point in time. The feasible signal set for each agent $j \in \mathcal{N}_i$ is calculated based on the actions taken by others and observed by agent j until time $t - 1$ together with possible private signals that can explain her history of choices: $\mathbf{a}_{j,0}$, $\mathbf{a}_{j,1}$, and so on up until her most recent choice which is $\mathbf{a}_{j,t}$. At time t , agent i will have access to all the observations of every agent in her neighborhood and can vet their most recent choices $\mathbf{a}_{j,t}$ against their observations to eliminate the incompatible private signals from the feasible set $\mathcal{S}_{j,t}$ and obtain an updated list of feasible signals $\mathcal{S}_{j,t+1}$ for each of her neighbors $j \in \mathcal{N}_i$. This pruning is achieved by calculating $\mathbf{a}_{j,t}(s_j)$ given $\mathcal{I}_{j,t}(s_j) = \{s_j\} \times \prod_{k \in \mathcal{N}_j} \mathcal{S}_{k,t}$ for each $s_j \in \mathcal{S}_{j,t}$ and removing any incompatible s_j that violates the condition $\mathbf{a}_{j,t} = \mathbf{a}_{j,t}(s_j)$; thus obtaining the pruned set $\mathcal{S}_{j,t+1}$. The calculation of $\mathbf{a}_{j,t}(s_j)$ given $\mathcal{I}_{j,t}(s_j) = \{s_j\} \times \prod_{k \in \mathcal{N}_j} \mathcal{S}_{k,t}$ is performed according to (2.1.4) but the decomposition of the feasible signal profiles based on the relation $\mathcal{I}_{j,t}(s_j) = \{s_j\} \times \prod_{k \in \mathcal{N}_j} \mathcal{S}_{k,t}$ together with the independence of private signals across different agents help reduce the number of additions and multiplications involved as follows:

$$\begin{aligned}
\mathbf{a}_{j,t}(s_j) &\leftrightarrow \arg \max_{\mathbf{a}_j \in \mathcal{A}_j} \sum_{\hat{\theta} \in \Theta} u_j(\mathbf{a}_j, \hat{\theta}) \frac{\sum_{\bar{s}' \in \mathcal{I}_{i,t}(s_j)} \mathcal{L}(\bar{s}' | \hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \sum_{\bar{s}' \in \mathcal{I}_{i,t}(s_j)} \mathcal{L}(\bar{s}' | \tilde{\theta}) \nu(\tilde{\theta})} \\
&= \arg \max_{\mathbf{a}_j \in \mathcal{A}_j} \sum_{\hat{\theta} \in \Theta} u_j(\mathbf{a}_j, \hat{\theta}) \frac{\mathcal{L}(s_j | \hat{\theta}) \prod_{k \in \mathcal{N}_j} \sum_{s_k \in \mathcal{S}_{k,t}} \mathcal{L}(s_k | \hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mathcal{L}(s_j | \tilde{\theta}) \prod_{k \in \mathcal{N}_j} \sum_{s_k \in \mathcal{S}_{k,t}} \mathcal{L}(s_k | \tilde{\theta}) \nu(\tilde{\theta})}.
\end{aligned}$$

Hence, the calculation of the conditionally feasible action $\mathbf{a}_{j,t}(s_j)$ for each $s_j \in \mathcal{S}_{j,t}$ can be achieved through $\text{card}(\Theta)A \sum_{k \in \mathcal{N}_j} \text{card}(\mathcal{S}_{k,t})$ additions and $\text{card}(\Theta) (\text{card}(\mathcal{N}_j) + 2) A$ multiplications; subsequently, the total number of additions and multiplications required for agent i to update the feasible private signals of each of her neighbor can be estimated as follows:

$$A \sum_{j \in \mathcal{N}_i} \text{card}(\Theta) \text{card}(\mathcal{S}_{j,t}) \left[\sum_{k \in \mathcal{N}_j} \text{card}(\mathcal{S}_{k,t}) + \text{card}(\mathcal{N}_j) + 2 \right] \leq An^2M^2m + An^2Mm + 2nMmA, \quad (\text{A.2.1})$$

where M , n , m and A are as in (A.1.2). After updating her lists for the feasible signal profiles of all her agents the agent can refine her list of feasible signal profiles $\mathcal{I}_{i,t+1} = \{\mathbf{s}_i\} \times \prod_{j \in \mathcal{N}_i} \mathcal{S}_{j,t+1}$ and determine her belief $\mu_{i,t+1}$ and refined choice $\mathbf{a}_{i,t+1}$. The latter is achieved through an extra $\text{card}(\Theta)A \sum_{j \in \mathcal{N}_i} \text{card}(\mathcal{S}_{j,t+1})$ additions and $\text{card}(\Theta)A (\text{card}(\mathcal{N}_i) + 2)$ multiplications, which are dominated by the required calculations in (A.2.1). Most notably, the computations required of the agent for determining her Bayesian choices in a POSET increase polynomially in the number of agents n , whereas in a general network structure using (A1) these computations increase exponentially fast in the number of agents n . \square

A.3 Proof of Theorem 2.7: A SUBSET-SUM Reduction

The SUBSET-SUM problem can be described as follows and it is known to be \mathcal{NP} -complete (cf. [211, A3.2, SP13]):

Problem A.1 (SUBSET-SUM). *Given a set of n positive integers $\{p_1, \dots, p_n\}$ and a positive integer q , determine if any non-empty subset of $\{p_1, \dots, p_n\}$ sum to q .*

We now describe the reduction to an arbitrary instance of SUBSET-SUM from a particular instance of GROUP-DECISION. Consider the problem of determining the belief of agent i at time 2, $\mu_{i,2}$, in the graph \mathcal{G} with $n + 5$ nodes and $2n + 4$ edges as in the left graph in Fig. 2.5: agent i has two neighbors j_1 and j_2 , who themselves have n neighbors in common l_1, \dots, l_n . Furthermore, j_1 and j_2 each has one additional neighbor, k_1 and

k_2 respectively, whom they do not share. We take the signal spaces of i , j_1 , and j_2 to be singletons $\mathcal{S}_i = \mathcal{S}_{j_1} = \mathcal{S}_{j_2} = \{\hat{s}\}$, so that their private signals reveal no information and as such $\Phi_{i,0} = \Phi_{j_1,0} = \Phi_{j_2,0} = \Phi_{i,1} = 0$, following the simplifying assumptions of the binary state space with common uniform priors. We assume that each of the remaining agents l_1, l_2, \dots, l_n have a binary signal space $\mathbf{s}_{l_h} \in \{0, 1\}$, with the probabilities that are set such that

$$\underline{p}_h := \log \left(\frac{\ell(\mathbf{s}_{l_h} = 0|\theta_2)}{\ell(\mathbf{s}_{l_h} = 0|\theta_1)} \right), \bar{p}_h := \log \left(\frac{\ell(\mathbf{s}_{l_h} = 0|\theta_2)}{\ell(\mathbf{s}_{l_h} = 1|\theta_1)} \right), p_h = \bar{p}_h - \underline{p}_h, \text{ for all } h \in [n]. \quad (\text{A.3.1})$$

As for the agents k_1 and k_2 , they also receive binary signals but with probabilities that are set such that for $r = 1, 2$:

$$\underline{p}^* := \log \left(\frac{\ell(\mathbf{s}_{k_r} = 0|\theta_2)}{\ell(\mathbf{s}_{k_r} = 0|\theta_1)} \right), \bar{p}^* := \log \left(\frac{\ell(\mathbf{s}_{k_r} = 1|\theta_2)}{\ell(\mathbf{s}_{k_r} = 1|\theta_1)} \right), -q = \bar{p}^* - \underline{p}^*. \quad (\text{A.3.2})$$

Suppose further that at time 2, agent i observes the beliefs of both agents j_1 and j_2 to be as follows: $\Phi_{j_1,1} = \Phi_{j_2,1} = \sum_{h=1}^m \underline{p}_h + \underline{p}^*$. Note that in the above notation we have

$$\lambda_{k_r} = \mathbf{s}_{k_r} \frac{(\bar{p}^* - \underline{p}^*)}{-q} + \underline{p}^*, \text{ and } \lambda_{l_h} = \mathbf{s}_{l_h} \frac{(\bar{p}_{l_h} - \underline{p}_{l_h})}{p_h} + \underline{p}_{l_h}, r = 1, 2, h \in [n]. \quad (\text{A.3.3})$$

These quantities are important as they determine the beliefs of agents j_1 and j_2 at time one, which are reported to agent i for processing her belief update at time 2. In particular, at time 2, and from the fact that $\Phi_{j_1,1} = \Phi_{j_2,1} = \sum_{h=1}^m \underline{p}_h + \underline{p}^*$ agent i infers the following information:

$$\Phi_{j_1,1} = \sum_{h=1}^n \lambda_{l_h} + \lambda_{k_1} = \sum_{h=1}^m \underline{p}_h + \underline{p}^*, \text{ and } \Phi_{j_2,1} = \sum_{h=1}^n \lambda_{l_h} + \lambda_{k_2} = \sum_{h=1}^m \underline{p}_h + \underline{p}^*.$$

Replacing from (A.3.1), (A.3.2) and (A.3.3), the preceding relations can be written in terms of the private signals \mathbf{s}_{l_h} , $h \in [n]$ and $\mathbf{s}_{k_1}, \mathbf{s}_{k_2}$ as follows:

$$\sum_{h=1}^n \mathbf{s}_{l_h} p_h - \mathbf{s}_{k_1} q = 0, \text{ and } \sum_{h=1}^n \mathbf{s}_{l_h} p_h - \mathbf{s}_{k_2} q = 0. \quad (\text{A.3.4})$$

Note that the constant term $\sum_{h=1}^m \underline{p}_h + \underline{p}^*$ is canceled out from both sides of the two equations leading to the homogeneous system in (A.3.4). To compute her Bayesian posterior belief $\mu_{i,2}$ or equivalently $\Phi_{i,2}$, agent i should first solve the arbitrary instance of SUBSET-SUM for the given parameters: SUBSET-SUM($p_1, \dots, p_n; q$). If she determines that the answer to SUBSET-SUM($p_1, \dots, p_n; q$) is negative then she concludes that

all agents must have received zero signals and she sets her belief accordingly: $\Phi_{i,2} = \sum_{h=1}^n s_{t_h} p_h - s_{k_1} q - s_{k_1} q + \sum_{h=1}^m \underline{p}_h + 2\underline{p}^* = \sum_{h=1}^m \underline{p}_h + 2\underline{p}^*$; in particular, we have:

$$\text{If } \text{SUBSET-SUM}(p_1, \dots, p_n; q) = \text{FALSE}, \text{ then } \Phi_{i,2} = \sum_{h=1}^m \underline{p}_h + 2\underline{p}^*. \quad (\text{I})$$

It is also worth highlighting that when $\text{SUBSET-SUM}(p_1, \dots, p_n; q) = \text{FALSE}$ the belief of agent i at time two is in fact an efficient belief but that does not imply the transparency of the graph structure because the latter is a structural property that should hold true for all choices of the signal structure parameters; in fact, the graph structure in Fig. 2.5, on the left, is not transparent. On the other hand, if the answer to $\text{SUBSET-SUM}(p_1, \dots, p_n; q)$ is positive, then agent i concludes that in addition to the case of all zero signals, there are additional cases (i.e. feasible signal profiles) in which some agents receive a one signal. In any such cases, we should necessarily have that $s_{k_1} = s_{k_1} = 1$, in order for (A.3.4) to remain satisfied. Subsequently, for all such nontrivial signal profiles we have that:

$$\begin{aligned} (-qs_{k_1} + \underline{p}^*) + (-qs_{k_2} + \underline{p}^*) + \sum_{h=1}^n (p_h s_{t_h} + \underline{p}_h) &= \underline{p}^* + \frac{(-q + \underline{p}^*)}{\bar{p}^*} + \sum_{h=1}^n \underline{p}_h \\ &= \sum_{h=1}^n \underline{p}_h + \underline{p}^* + \bar{p}^* < \sum_{h=1}^n \underline{p}_h + 2\underline{p}^*, \end{aligned}$$

where in the first equality we use (A.3.4) to cancel out the indicated terms and in the last inequality we use the fact that $\bar{p}^* = (-q) + \underline{p}^* < \underline{p}^*$. Agent i thus needs to find all these feasible signal profiles and set her belief at time two based on the the set of all feasible signal profiles. In particular, since in all the non-trivial cases (feasible signal profiles that are not all zero signals), $\sum_{k=1}^n \lambda_k + \lambda_{k_1} + \lambda_{k_2} = \sum_{h=1}^n \underline{p}_h + \underline{p}^* + \bar{p}^* < \sum_{h=1}^n \underline{p}_h + 2\underline{p}^*$ we have that:

$$\text{If } \text{SUBSET-SUM}(p_1, \dots, p_n; q) = \text{TRUE}, \text{ then } \Phi_{i,2} < \sum_{h=1}^n \underline{p}_h + 2\underline{p}^*. \quad (\text{II})$$

This concludes the reduction because if an algorithm is available that solves any instances of GROUP-DECISION in polynomial time then by inspecting the output of that algorithm according to (I) and (II) for the particular instance of GROUP-DECISION described above, agent i can decide the feasibility of any instance of the SUBSET-SUM problem in polynomial time. \square

A.4 Proof of Theorem 2.7: An EXACT-COVER Reduction

EXACT-COVER is the fourteenth on Karp's list of 21 \mathcal{NP} -complete problems. It is described as follows [212]:

Problem A.2 (EXACT-COVER). A set of n items $\{j_1, \dots, j_n\} = \{j_h : h \in [n]\}$ and a family $\{l_1, \dots, l_m\}$ of m subsets: $l_h \subset \{j_1, \dots, j_n\}$ for all $h \in [m]$, are given. Determine if there is a subfamily of disjoint subsets belonging to $\{l_1, \dots, l_m\}$ such that their union is $\{j_1, \dots, j_n\}$: $\{l_{h_1}, \dots, l_{h_p}\} \subset \{l_1, \dots, l_m\}$, $l_{h_q} \cap l_{h_{q'}} = \emptyset$ for all $q, q' \in [p]$, and $\cup_{q=1}^p l_{h_q} = \{j_1, \dots, j_n\}$.

The input to EXACT-COVER can be represented by a graph $\hat{\mathcal{G}}_{m,n}$ on the $m + n$ nodes $\{l_1, \dots, l_m ; j_1, \dots, j_n\}$ which is bipartite between $\{l_1, \dots, l_m\}$ and $\{j_1, \dots, j_n\}$ and the only edges are from nodes l_h to $j_{h'}$ whenever subset l_h contains the element $j_{h'}$ for some $h \in [m]$ and $h' \in [n]$ in the description of EXACT-COVER. Henceforth, we use the notation $\text{EXACT-COVER}(\hat{\mathcal{G}}_{m,n})$ to denote the output of EXACT-COVER for an arbitrary input $\mathcal{G}_{m,n}$: $\text{EXACT-COVER}(\hat{\mathcal{G}}_{m,n}) \in \{\text{TRUE}, \text{FALSE}\}$. If there is a subset l_h , $h \in [m]$ that alone covers all the items $\{j_1, \dots, j_n\}$, then the answer to $\text{EXACT-COVER}(\hat{\mathcal{G}}_{m,n})$ is (trivially) true, and we can thus check for and remove this case in our polynomial reduction.

To construct the reduction from an arbitrary instance of EXACT-COVER to a particular instance of GROUP-DECISION, we consider the decision problem of agent i in a graph \mathcal{G} that is derived from $\hat{\mathcal{G}}_{m,n}$ by adding two additional nodes i and k and $2n$ additional edges: n edges that are directed from node k to each of $\{j_1, \dots, j_n\}$ and another n edges from each of $\{j_1, \dots, j_n\}$ to node i (cf. the right graph in Fig. 2.5). We assume that agents i and j_1, \dots, j_n can only receive the non-informative signal \hat{s} : $\mathcal{S}_i = \mathcal{S}_{j_1} = \mathcal{S}_{j_2} = \dots = \mathcal{S}_{j_n} = \{\hat{s}\}$; hence, $\Phi_{i,0} = \Phi_{j_1,0} = \Phi_{j_2,0} = \dots = \Phi_{j_n,0} = \Phi_{i,1} = 0$.

We assume that agents l_1, \dots, l_m observe initial i.i.d. binary signals: $\mathbf{s}_{l_h} \in \{0, 1\}$ with the probabilities set such that for all $h \in [m]$:

$$\log \left(\frac{\ell(\mathbf{s}_{l_h} = 1 | \theta_2)}{\ell(\mathbf{s}_{l_h} = 1 | \theta_1)} \right) = \bar{p}, \log \left(\frac{\ell(\mathbf{s}_{l_h} = 0 | \theta_2)}{\ell(\mathbf{s}_{l_h} = 0 | \theta_1)} \right) = \underline{p}, \bar{p} - \underline{p} = 1.$$

Similarly, agent k receives a binary signal but with probabilities such that

$$\log \left(\frac{\ell(\mathbf{s}_k = 1 | \theta_2)}{\ell(\mathbf{s}_k = 1 | \theta_1)} \right) = \bar{p}^*, \log \left(\frac{\ell(\mathbf{s}_k = 0 | \theta_2)}{\ell(\mathbf{s}_k = 0 | \theta_1)} \right) = \underline{p}^*, \bar{p}^* - \underline{p}^* = -1.$$

Note that with the above setting,

$$\lambda_k = \mathbf{s}_k \left(\frac{\bar{p}^* - \underline{p}^*}{-1} \right) + \underline{p}^*, \text{ and } \lambda_{l_h} = \mathbf{s}_{l_h} \left(\frac{\bar{p} - \underline{p}}{1} \right) + \underline{p}, h \in [m]. \quad (\text{A.4.1})$$

At time two, agent i observes that each of her neighbors j_1, \dots, j_n have changed their beliefs from their initial uniform priors, such that

$$\Phi_{j_r,1} = \text{card}(\{h \in [m] : j_r \in l_h\}) \underline{p} + \underline{p}^* = (\text{card}(\mathcal{N}_{j_r}) - 1) \underline{p} + \underline{p}^*. \quad (\text{A.4.2})$$

Note that $\mathcal{N}_{j_r} = \{k\} \cup \{l_h : h \in [m], j_r \in l_h\}$, and $\text{card}(\mathcal{N}_{j_r}) - 1 = \text{card}(\{h \in [m] : j_r \in l_h\})$

counts the number of subsets $l_h, h \in [m]$ that cover item j_r in the original description of EXACT-COVER (Problem A.2). To make a Bayesian inference about the reported beliefs in (A.4.2) and to decide her time two belief $\mu_{i,2}$ (or equivalently $\Phi_{i,2}$), agent i should first consider the following construction of the reported beliefs, $\Phi_{j_r,1}$ for all $r \in [n]$:

$$\begin{aligned}\Phi_{j_r,1} &= \lambda_k + \sum_{l_h \in \mathcal{N}_{j_r} \setminus \{k\}} \lambda_{l_h} = (-\mathbf{s}_k + \underline{\mathbf{p}}^*) + \sum_{l_h \in \mathcal{N}_{j_r} \setminus \{k\}} (\mathbf{s}_{l_h} + \underline{\mathbf{p}}) \\ &= \left(\sum_{l_h \in \mathcal{N}_{j_r} \setminus \{k\}} \mathbf{s}_{l_h} - \mathbf{s}_k \right) + (\text{card}(\mathcal{N}_{j_r}) - 1) \underline{\mathbf{p}} + \underline{\mathbf{p}}^*,\end{aligned}\quad (\text{A.4.3})$$

Combining her observations in (A.4.2) with the construction of the reported beliefs in (A.4.3), agent i should consider the solutions of the resultant system of n equations in the following $m + 1$ binary variables: $\mathbf{s}_{l_1}, \dots, \mathbf{s}_{l_m}$ and \mathbf{s}_k . In particular, she has to decide whether her observations in (A.4.2) are the result of k and l_1, \dots, l_m having all received zero signals, or else if it is possible that agent k has received a one signal ($\Phi_{k,0} = \lambda_k = -1 + \underline{\mathbf{p}}^*$) and a specific subset of the agents l_1, \dots, l_m have also received one signals ($\Phi_{l_h,0} = \lambda_{l_h} = 1 + \underline{\mathbf{p}}$, for all l_h who see $\mathbf{s}_{l_h} = 1$) enough to exactly balance the net effect, leading to (A.4.2). The latter is possible only if there is a non-trivial solution to the following system:

$$\sum_{l_h \in \mathcal{N}_{j_r} \setminus \{k\}} \mathbf{s}_{l_h} - \mathbf{s}_k = 0, \text{ for all } r \in [n]; (\mathbf{s}_{l_1}, \dots, \mathbf{s}_{l_m}, \mathbf{s}_k) \in \{0, 1\}^{m+1}. \quad (\text{A.4.4})$$

This is equivalent to the feasibility of EXACT-COVER($\hat{\mathcal{G}}_{m,n}$) since the latter can be formulated as the following 0/1-integer program:

$$\sum_{h \in [m]: j_r \in l_h} \mathbf{s}_{l_h} = 1, \text{ for all } j_r \in \{j_1, \dots, j_n\}; (\mathbf{s}_{l_1}, \dots, \mathbf{s}_{l_m}) \in \{0, 1\}^m. \quad (\text{A.4.5})$$

Note that a variable \mathbf{s}_{l_h} in System (A.4.5) will be one only if the corresponding set l_h is chosen in the solution of the feasible EXACT-COVER; moreover, the constraints in (A.4.5) express the requirement that the chosen sets do not intersect at any of the elements $\{j_1, \dots, j_n\}$. In other words, each of the n items are contained in one and exactly one subset: for each $p \in [n]$, there is a unique $h \in [m]$ such that $j_p \in l_h$ and $\mathbf{s}_{l_h} = 1$. System (A.4.4) having a non trivial solution is equivalent to the feasibility of System (A.4.5), because in any non trivial solution of (A.4.4) we should necessarily have $\mathbf{s}_k = 1$; and furthermore, from our construction of the graph \mathcal{G} based on the EXACT-COVER input $\hat{\mathcal{G}}_{m,n}$ we have that $\mathcal{N}_{j_p} = \{l_h : h \in [m], j_p \in l_h\} \cup \{k\}$ for all $j_p \in \{j_1, \dots, j_n\}$.

Note that since in our polynomial reduction we have removed the case where all of the items $\{j_1, \dots, j_n\}$ are covered by one subset l_h for some $h \in [m]$, in any nontrivial solution

of exact cover, we have that $s_{l_h} = 1$ for at least two distinct values of $h \in [m]$: at least two subsets are needed for all the elements to be covered in a feasible EXACT-COVER. Subsequently, if agent i determines that $\text{EXACT-COVER}(\hat{\mathcal{G}}_{m,n})$ is FALSE, then she concludes that all agents must have received zero signals and she sets her belief accordingly: $\Phi_{i,2} = \lambda_k + \sum_{h=1}^m \lambda_{l_h} + \sum_{r=1}^n \lambda_{j_r} + \lambda_i = \underline{p}^* + m\underline{p}$, where we use the facts that $\lambda_{j_r} = \lambda_i = 0$ for all $r \in [n]$, as well as $\lambda_k = \underline{p}^*$ and $\lambda_{l_h} = \underline{p}$ for all l_1, \dots, l_m with zero signals. Put succinctly,

$$\text{If } \text{EXACT-COVER}(\hat{\mathcal{G}}_{m,n}) = \text{FALSE}, \text{ then } \Phi_{i,2} = \underline{p}^* + m\underline{p}. \quad (\text{III})$$

However, if the answer to $\text{EXACT-COVER}(\hat{\mathcal{G}}_{m,n})$ is TRUE, then for any additional feasible signal profile that agent i identifies and determines to satisfy (A.4.4), it is necessarily true that $s_k = 1$ and $s_{l_h} = 1$, for at list two distinct agents among $\{l_1, \dots, l_m\}$; hence, for any such additionally feasible signal profiles it is always true that

$$\lambda_k + \sum_{h=1}^m \lambda_{l_h} = -s_k + \underline{p}^* + \sum_{h=1}^m s_{l_h} + m\underline{p} \geq 1 + \underline{p}^* + m\underline{p},$$

where in the latter lower-bound we use the facts that s_k as well as at least two of s_{l_h} are one in any non-trivially feasible signal profile, i.e. $-s_k + \sum_{h=1}^m s_{l_h} \geq 1$. In particular, we conclude that

$$\text{If } \text{EXACT-COVER}(\hat{\mathcal{G}}_{m,n}) = \text{TRUE}, \text{ then } \Phi_{i,2} > \underline{p}^* + m\underline{p}. \quad (\text{IV})$$

Hence we conclude the \mathcal{NP} -hardness of GROUP-DECISION by its reduction to EXACT-COVER. Because if the polynomial time computation of beliefs in GROUP-DECISION was possible, then by inspecting the computed beliefs according to (III) and (IV) for the particular instance of GROUP-DECISION (with i.i.d. binary signals) described above, agent i can decide the feasibility of any instance of the EXACT-COVER problem in polynomial time. \square

A.5 Belief Calculations in Bounded Neighborhoods with i.i.d. Signals

In this example, we consider a variation of the right-hand-side structure in Fig. 2.5 in which agent k is removed and also n , the number of directly observed neighbors of agent i , is fixed. We show that the belief of agent i at time two can be computed efficiently in the number of indirectly observed neighbors (m). We suppose that the signal structures for agent i , her neighboring agents j_1, \dots, j_n , and the indirectly observed agents l_1, \dots, l_m are as in Appendix A.4; subsequently, $\Phi_{i,0} = \Phi_{j_1,0} = \Phi_{j_2,0} = \dots = \Phi_{j_n,0} = \Phi_{i,1} = 0$.

At time two, agent i has to incorporate the time one beliefs of her neighbors, which are themselves caused by the time zero beliefs of l_1, \dots, l_m : Given $\Phi_{j_r,1} = \sum_{l_h \in \mathcal{N}_{j_r}} \lambda_{l_h}$, for $r = 1, \dots, n$, agent i aims to determine her belief at time two (or equivalently $\Phi_{i,2}$). Using (A.4.1), we can write $\Psi_{j_r} = \sum_{l_h \in \mathcal{N}_{j_r}} \mathbf{s}_{l_h}$, where

$$\Psi_{j_r} = \frac{1}{\bar{p} - \underline{p}} (\Phi_{j_r,1} - \text{card}(\mathcal{N}_{j_r}) \underline{p}), r \in [n],$$

are necessarily non-negative integers belonging in to $[m]_0 = \{0\} \cup [m]$, due to their generation process, i.e. the fact that they count the number of one signals that are received in the neighborhood \mathcal{N}_{j_r} of each of the neighbors $j_r, r \in [n]$. For all $r \in [n]$ and $r' \in [m]$, let $a_{j_r, l_{r'}} = 1$ if $l_{r'} \in \mathcal{N}_{j_r}$ and $a_{j_r, l_{r'}} = 0$ otherwise. Denoting $\bar{a}_{j_r} = (a_{j_r, l_1}, \dots, a_{j_r, l_m})$ and using the transpose notation T , we can rewrite Ψ_{j_r} as an inner product $\Psi_{j_r} = \bar{a}_{j_r} \bar{\mathbf{s}}^T$, where $\bar{\mathbf{s}} = (\mathbf{s}_{l_1}, \dots, \mathbf{s}_{l_m})$. To proceed for each $r \in [m]$, let $\bar{a}_{l_r} = (a_{j_1, l_r}, \dots, a_{j_n, l_r})$. To determine her belief, agent i acts as follows:

1. For each $\bar{\kappa} = (\kappa_1, \dots, \kappa_n) \in \{0, 1\}^n$, let $\Psi_{\bar{\kappa}} = \{l_r : \bar{a}_{l_r} = \bar{\kappa}\}$, note that $\Psi_{\bar{\kappa}}$ are non-intersecting, possibly empty sets, whose union is equal to $\{l_1, \dots, l_m\}$. Also let $\eta_{\bar{\kappa}}$ be the number of agents belonging to $\Psi_{\bar{\kappa}}$ who have received one signals; the rest having received zero signals, the variables $\eta_{\bar{\kappa}}, \bar{\kappa} \in \{0, 1\}^n$ should satisfy:

$$\sum_{\bar{\kappa} \in \Xi_r} \eta_{\bar{\kappa}} = \Psi_{j_r}, \text{ for all } r \in [n], \text{ where } \Xi_r = \{\bar{\kappa} : \bar{\kappa} = (\kappa_1, \dots, \kappa_n), \kappa_r = 1\}. \quad (\text{A.5.1})$$

2. Note that $\eta_{\bar{\kappa}} \in [\text{card}(\Psi_{\bar{\kappa}})]_0$ for each $\bar{\kappa} \in \{0, 1\}^n$, and to determine her belief, agent i needs to find the set Γ_i of all such non-negative integer solutions of (A.5.1):

- Initialize $\Gamma_i = \emptyset$.
- For each $\bar{\eta} := (\eta_{\bar{\kappa}}, \bar{\kappa} \in \{0, 1\}^n) \in \prod_{\bar{\kappa} \in \{0, 1\}^n} [\text{card}(\Psi_{\bar{\kappa}})]_0$, if all $\eta_{\bar{\kappa}}, \bar{\kappa} \in \{0, 1\}^n$ satisfy (A.5.1) for each $r \in [n]$, then set $\Gamma = \Gamma \cup \{\bar{\eta}\}$.

3. Having thus found Γ_i , agent i sets her belief (or equivalently its log-ratio) as follows:

$$\Phi_{i,2} = \log \frac{\sum_{\bar{\eta} \in \Gamma_i} \left[\prod_{\bar{\kappa} \in \{0, 1\}^n} \binom{\text{card}(\Psi_{\bar{\kappa}})}{\eta_{\bar{\kappa}}} \right] \left[\ell(\mathbf{s} = 1 | \theta_2)^{\sum_{\bar{\kappa} \in \{0, 1\}^n} \eta_{\bar{\kappa}}} \right] \left[\ell(\mathbf{s} = 0 | \theta_2)^{m - \sum_{\bar{\kappa} \in \{0, 1\}^n} \eta_{\bar{\kappa}}} \right]}{\sum_{\bar{\eta} \in \Gamma_i} \left[\prod_{\bar{\kappa} \in \{0, 1\}^n} \binom{\text{card}(\Psi_{\bar{\kappa}})}{\eta_{\bar{\kappa}}} \right] \left[\ell(\mathbf{s} = 1 | \theta_1)^{\sum_{\bar{\kappa} \in \{0, 1\}^n} \eta_{\bar{\kappa}}} \right] \left[\ell(\mathbf{s} = 0 | \theta_1)^{m - \sum_{\bar{\kappa} \in \{0, 1\}^n} \eta_{\bar{\kappa}}} \right]}. \quad (\text{A.5.2})$$

Note that with private signals restricted to be i.i.d. binary signals, the set Γ_i in fact represents

the set of all private signals profiles that are deemed feasible by agent i at time two, as with $\mathcal{I}_{i,t}$ in (2.1.1). The symmetry of the binary structure allows for the summation over the feasible signal profiles to be simplified as in (A.5.2) by counting the number of ways in which the agents would receive one signals within each of subsets $\Psi_{\bar{\kappa}}$, $\bar{\kappa} \in \{0, 1\}^n$; this is achieved by the product of the binomial coefficients in (A.5.2). The overall complexity of computing the Bayesian posterior belief in (A.5.2) can now be bounded by a total of $O(n2^n m^{2^n})$ additions and multiplications for computing the set Γ_i and another $O(m^{1+2^n}(2^{2^n})(3m+2))$ for computing the beliefs (or their ratio) in (A.5.2). Note that we made no effort in optimizing these computations beyond the fact that they increase polynomially in m for a fixed neighborhood size (n). \square

A.6 Proof of Theorem 3.3

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. [114, 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 (3.2.1) is $\mathcal{V}(\cdot; \alpha_i + \sigma_i \sum_{p=1}^{n_i} \xi(\mathbf{s}_{i,p}), \beta_i + n_i \delta_i)$, 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 $u_i(a, \theta) = -(a - m_{i,\theta})^\top (a - m_{i,\theta})$, the expected pay-off at any time time t is maximized is by choosing [213, 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.¹

Lemma A.1 (Proposition 3.3.14 of [114]). *Let $\zeta \in \mathbb{R}^k$ be a parameter and suppose that the parameter space Ω_ζ 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 $\kappa'(\alpha, \beta) e^{\zeta^\top \alpha - \beta \gamma(\zeta)}$ w.r.t. Λ_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 $\tau'(\mathbf{s}) e^{\zeta^\top \mathbf{s} - \gamma'(\zeta)}$ w.r.t. Λ_k , then*

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

¹In 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 [214].

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

$$\begin{aligned}
\int_{\theta \in \Theta} m_{i, \theta} \mathcal{V}(d\theta; \alpha, \beta) &= \int_{\theta \in \Theta} \int_{s \in \mathcal{S}} \xi(s) \mathcal{L}(ds | \theta; \sigma, \delta) \mathcal{V}(d\theta; \alpha, \beta) \tag{A.6.1} \\
&= \int_{\theta \in \Theta} \left| \frac{\Lambda_k(\eta(d\theta))}{\mathcal{G}_\theta(d\theta)} \right| \frac{e^{\eta(\theta)^T \alpha - \beta \gamma(\eta(\theta))}}{\kappa(\alpha, \beta)} \mathcal{G}_\theta(d\theta) \times \dots \\
&\quad \dots \int_{s \in \mathcal{S}} \xi(s) \sigma \left| \frac{\Lambda_k(\xi(ds))}{\mathcal{G}_s(ds)} \right| \tau(\sigma \xi(s), \delta) e^{\sigma \eta(\theta)^T \xi(s) - \delta \gamma(\eta(\theta))} \mathcal{G}_s(ds) \\
&= \int_{\zeta \in \Omega_\theta} \left| \frac{\Lambda_k(\eta(d\theta))}{\mathcal{G}_\theta(d\theta)} \right| \frac{e^{\eta(\theta)^T \alpha - \beta \gamma(\eta(\theta))}}{\kappa(\alpha, \beta)} \mathcal{G}_\theta(d\theta) \times \dots \\
&\quad \dots \int_{s \in \mathcal{S}} \xi(s) \sigma \left| \frac{\Lambda_k(\xi(ds))}{\mathcal{G}_s(ds)} \right| \tau(\sigma \xi(s), \delta) e^{\sigma \eta(\theta)^T \xi(s) - \delta \gamma(\eta(\theta))} \mathcal{G}_s(ds) \\
&= \int_{\zeta \in \Omega_\eta} \frac{e^{\zeta^T \alpha - \frac{\beta}{\sigma} \gamma'(\zeta)}}{\kappa(\alpha, \beta)} \Lambda_k(d\zeta) \int_{s' \in \mathcal{S}'} \frac{s' \tau'(s')}{\sigma} e^{\zeta^T s' - \gamma'(s')} \Lambda_k(ds') = \frac{\alpha \delta}{\sigma \beta},
\end{aligned}$$

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 A.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:

$$\mathbf{a}_{i,0} = \frac{\sum_{p=1}^{n_i} \xi(\mathbf{s}_{i,p}) + \sigma_i^{-1} \alpha_i}{n_i + \delta_i^{-1} \beta_i}. \tag{A.6.2}$$

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

$$\sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}) = (n_j + \delta_j^{-1} \beta_j) \mathbf{a}_{j,0} - \sigma_j^{-1} \alpha_j, \forall j \in \mathcal{N}_i. \tag{A.6.3}$$

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 (A.6.1), to get that the expected pay-off maximizing action at time 1 is given by:

$$\mathbf{a}_{i,1} = \frac{\delta_i \left(\alpha_i + \sum_{j \in \mathcal{N}_i} \sigma_j \sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}) \right)}{\sigma_i \left(\beta_i + \sum_{j \in \mathcal{N}_i} n_j \delta_j \right)}. \quad (\text{A.6.4})$$

Finally, we can use (A.6.3) 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 3.3. \square

A.7 Proof of Theorem 3.6

The balancedness of likelihoods (Assumption 3.3) ensures that the coefficients of the linear combination from Corollary 3.1 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 (A.6.2), and at every $t > 1$ they update their actions according to $\bar{\mathbf{a}}_t = (\mathbb{T} \otimes \mathbb{I}_k) \bar{\mathbf{a}}_{t-1} = (\mathbb{T}^t \otimes \mathbb{I}_k) \bar{\mathbf{a}}_0$, where $\bar{\mathbf{a}}_t = (\mathbf{a}_{1,t}^\top, \dots, \mathbf{a}_{n,t}^\top)^\top$ and \mathbb{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. [130, 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 \mathbb{T} are all non-zero), then $\lim_{\tau \rightarrow \infty} \mathbb{T}^\tau = \mathbf{1} \bar{\mathbf{s}}^\top$, where $\bar{\mathbf{s}} := (s_1, \dots, s_n)^\top$ is the unique left eigenvector associated with the unit eigenvalue of \mathbb{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 $\sum_{i=1}^n s_i \mathbf{a}_{i,0} = \sum_{i=1}^n s_i \left(\sum_{p=1}^{n_i} \xi(\mathbf{s}_{i,p})/n_i \right)$. \square

A.8 Proof of Theorem 3.7

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 A.2 (Global MVUE). *Under the exponential family signal-utility structure (Assumption 3.1), the (global) MVUE of $m_{i,\theta}$ given the entire set of observations of all the*

agents across the network is given by:

$$\mathbf{a}_i^* = \frac{\delta_i \left(\alpha_i + \sum_{j=1}^n \sigma_j \sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}) \right)}{\sigma_i \left(\beta_i + \sum_{j=1}^n n_j \delta_j \right)}. \quad (\text{A.8.1})$$

If we further impose non-informative priors (Assumption 3.2), then the global MVUE for each i can be rewritten as

$$\mathbf{a}_i^* = \frac{\delta_i \left(\sum_{j=1}^n \sigma_j \sum_{p=1}^{n_j} \xi(\mathbf{s}_{j,p}) \right)}{\sigma_i \left(\sum_{j=1}^n n_j \delta_j \right)} = \frac{\delta_i}{\sigma_i} \sum_{j=1}^n \frac{\sigma_j n_j}{\sum_{p=1}^n n_p \delta_p} \mathbf{a}_{j,0}. \quad (\text{A.8.2})$$

This lemma can be proved easily. Following the same steps that lead to (A.6.4), yields (A.8.1). Next by making the necessary substitutions under Assumption 3.2, (A.8.2) is obtained. From (A.8.2), 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 3.3) is automatically satisfied and given non-informative priors Theorem 3.6 guarantees convergence to consensus in a strongly connected social network. Moreover, we can rewrite (A.8.2) as $\mathbf{a}_i^* = \mathbf{a}^* = \left(\sum_{j=1}^n \delta_j n_j \mathbf{a}_{j,0} \right) / \sum_{p=1}^n n_p \delta_p$, for all i . Hence, if the consensus action $\left(\sum_{i=1}^n s_i \mathbf{a}_{i,0} \right)$ in the proof of Theorem 3.6, Appendix A.7) 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 A.7. Using $\delta_i \sigma_j = \sigma_i \delta_j$, we can also rewrite the coefficients T_{ij} of the DeGroot update in Theorem 3.6 as $T_{ij} = \delta_j n_j / \left(\sum_{p \in \mathcal{N}_i} n_p \delta_p \right)$.

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

$$\sum_{i=1}^n s_i T_{ij} = \sum_{i=1}^n \left(\frac{\delta_i n_i}{\sum_{j=1}^n \delta_j n_j} \right) \frac{\delta_j n_j [I + A]_{ij}}{\sum_{p \in \mathcal{N}_i} n_p \delta_p} = s_j = \frac{\delta_j n_j}{\sum_{j=1}^n \delta_j n_j}, \quad (\text{A.8.3})$$

or equivalently,

$$\sum_{i,j \in \mathcal{N}_i} \frac{\delta_i n_i}{\sum_{p \in \mathcal{N}_i} n_p \delta_p} = \sum_{i \in \mathcal{N}_j^{\text{out}}} \frac{\delta_i n_i}{\sum_{p \in \mathcal{N}_i} n_p \delta_p} = 1, \quad (\text{A.8.4})$$

for all j . Under the global balance condition (Assumption 3.4), $\delta_i \sigma_j = \delta_j \sigma_i$, the weights $T_{ij} = \delta_j n_j / \left(\sum_{p \in \mathcal{N}_i} n_p \delta_p \right)$ as given above, correspond to transition probabilities of a node-weighted random walk on the social network graph, cf. [215, 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 $(s_1, \dots, s_n)^\top$ can be calculated in closed form as follows [216, Section 3.2]:

$$s_i = \frac{\sum_{j \in \tilde{\mathcal{N}}_i} w_{i,j}}{\sum_{i=1}^n \sum_{j \in \tilde{\mathcal{N}}_i} w_{i,j}}. \quad (\text{A.8.5})$$

In a node-weighted random walk we can replace $w_{i,j} = w_i w_j$ for all $j \in \tilde{\mathcal{N}}_i$ and (A.8.5) simplifies into

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

Similarly to (A.8.3), the consensus action will be efficient if and only if

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

$$\left(\sum_{k=1}^n w_k \right) \sum_{j \in \tilde{\mathcal{N}}_i} w_j = \sum_{i=1}^n \left(w_i \sum_{j \in \tilde{\mathcal{N}}_i} w_j \right), \forall i,$$

which holds true only if $\sum_{j \in \tilde{\mathcal{N}}_i} w_j$ is a common constant that is the same for all agents, i.e. $\sum_{j \in \tilde{\mathcal{N}}_i} w_j = \sum_{j \in \tilde{\mathcal{N}}_i} \delta_j n_j = C' > 0$ for all $i \in [n]$. Next replacing in (A.8.4) yields that, in fact, $C' = \sum_{i \in \tilde{\mathcal{N}}_j^{\text{out}}} \delta_i n_i$ for all j , completing the proof for the conditions of efficiency. \square

A.9 Proof of Theorem 3.11

Using the log-ratio variables and their concatenations defined at the beginning of Appendix F, we can rewrite the log-linear belief updates of (3.3.1) in a linearized vector format as

shown below:

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

Next we use the spectral decomposition in (3.3.11) to obtain: ¹

$$\begin{aligned}
\bar{\Phi}_{t+1}(\hat{\theta}, \check{\theta}) &= (1 + \rho)^{t+1} \bar{r}_1 \mathbf{L}(\hat{\theta}, \check{\theta}) + ((1 + \rho)^{t+1} - \sum_{\tau=0}^t (1 + \rho)^\tau \rho) \bar{r}_1 \beta(\hat{\theta}, \check{\theta}) + o((1 + \rho)^{t+1}) \\
&= (1 + \rho)^{t+1} \left(\bar{r}_1 \mathbf{L}(\hat{\theta}, \check{\theta}) + (1 - \sum_{\tau=0}^t (1 + \rho)^{\tau-t-1} \rho) \bar{r}_1 \beta(\hat{\theta}, \check{\theta}) + o(1) \right) \\
&\rightarrow (1 + \rho)^{t+1} \bar{r}_1 \mathbf{L}(\hat{\theta}, \check{\theta}), \tag{A.9.1}
\end{aligned}$$

where we adopt the following notations for the global log likelihood and prior ratio statistics: $\beta(\hat{\theta}, \check{\theta}) := \bar{\alpha}^\top \bar{\gamma}(\hat{\theta}, \check{\theta})$ and $\mathbf{L}(\hat{\theta}, \check{\theta}) := \bar{\alpha}^\top \bar{\lambda}(\hat{\theta}, \check{\theta})$; furthermore, in calculation of the limit in the last step of (A.9.1) we use the geometric summation identity $\sum_{\tau=0}^{\infty} \rho(1+\rho)^{\tau-1} = 1$.

To proceed denote $\mathbf{L}(\hat{\theta}) := \sum_{i=1}^n \alpha_i \ell_i(\mathbf{s}_i | \hat{\theta})$ so that $\mathbf{L}(\hat{\theta}, \check{\theta}) = \mathbf{L}(\hat{\theta}) - \mathbf{L}(\check{\theta})$. Since Θ^\diamond consists of the set of all maximizers of $\mathbf{L}(\hat{\theta})$, we have that $\mathbf{L}(\hat{\theta}, \check{\theta}) < 0$ whenever $\check{\theta} \in \Theta^\diamond$ and $\hat{\theta} \notin \Theta^\diamond$. Next recall from (A.9.1) that $\bar{\Phi}_{t+1}(\hat{\theta}, \check{\theta}) \rightarrow (1 + \rho)^{t+1} \bar{r}_1 \mathbf{L}(\hat{\theta}, \check{\theta})$ where \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 $\lim_{t \rightarrow \infty} \mu_{i,t}(\check{\theta}) = 0, \forall \check{\theta} \notin \Theta^\diamond$ as claimed in the Theorem. In the special case that Θ^\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 $i \in [n]$.

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

A.10 Proof of Lemma 4.1: Time-One Bayesian Actions

Note that given her observation of private signal $\mathbf{s}_{i,0}$, the posterior probability assigned by agent i to the set θ_1 is given by (4.1.3) with $\hat{\theta} = \theta_1$. We form a dichotomy of the signal space \mathcal{S}_i of each agent into \mathcal{S}_i^1 and \mathcal{S}_i^{-1} ; by setting $\mathcal{S}_i^1 := \{s \in \mathcal{S}_i : \ell_i(s|\theta_1)\nu_i(\theta_1) \geq \ell_i(s|\theta_2)\nu_i(\theta_2)\}$ and $\mathcal{S}_i^{-1} := \mathcal{S}_i \setminus \mathcal{S}_i^1$. It thus follows from (4.2.2) that for any $j \in \mathcal{N}_i$ the observation that $\mathbf{a}_{j,0} = 1$ is equivalent to the information that $\{\mathbf{s}_{j,0} \in \mathcal{S}_j^1\}$ and $\mathbf{a}_{j,0} = -1$ is equivalent to the information that $\{\mathbf{s}_{j,0} \in \mathcal{S}_j^{-1}\}$. Thereby, the belief of agent i at time $t = 1$ given her observation of the actions of her neighbors and the private signal $\mathbf{s}_{j,0}$ is given by

$$\mu_{i,1}(\theta_1) = \frac{\ell_i(\mathbf{s}_{i,0}|\theta_1) \prod_{j \in \mathcal{N}_i} \left(\sum_{\mathbf{s}_j \in \mathcal{S}_j^{\mathbf{a}_{j,0}}} \ell_j(\mathbf{s}_j|\theta_1) \right) \nu_i(\theta_1)}{\sum_{\hat{\theta} \in \Theta} \ell_i(\mathbf{s}_{i,0}|\hat{\theta}) \prod_{j \in \mathcal{N}_i} \left(\sum_{\mathbf{s}_j \in \mathcal{S}_j^{\mathbf{a}_{j,0}}} \ell_j(\mathbf{s}_j|\hat{\theta}) \right) \nu_i(\hat{\theta})},$$

and we can thus form the ratio

$$\begin{aligned} \frac{\mu_{i,1}(\theta_1)}{\mu_{i,1}(\theta_2)} &= \frac{\ell_i(\mathbf{s}_{i,0}|\theta_1)\nu_i(\theta_1)}{\ell_i(\mathbf{s}_{i,0}|\theta_2)\nu_i(\theta_2)} \prod_{j \in \mathcal{N}_i} \left(\frac{\sum_{\mathbf{s}_j \in \mathcal{S}_j^{\mathbf{a}_{j,0}}} \ell_j(\mathbf{s}_j|\theta_1)}{\sum_{\mathbf{s}_j \in \mathcal{S}_j^{\mathbf{a}_{j,0}}} \ell_j(\mathbf{s}_j|\theta_2)} \right) \\ &= \frac{\ell_i(\mathbf{s}_{i,0}|\theta_1)\nu_i(\theta_1)}{\ell_i(\mathbf{s}_{i,0}|\theta_2)\nu_i(\theta_2)} V_i \prod_{j \in \mathcal{N}_i} W_j^{\mathbf{a}_{j,0}}, \end{aligned} \quad (\text{A.10.1})$$

where for all $i \in [n]$ we have defined

$$V_i = \prod_{j \in \mathcal{N}_i} \left(\frac{\sum_{\mathbf{s}_j \in \mathcal{S}_j^1} \ell_j(\mathbf{s}_j|\theta_1)}{\sum_{\mathbf{s}_j \in \mathcal{S}_j^1} \ell_j(\mathbf{s}_j|\theta_2)} \times \frac{\sum_{\mathbf{s}_j \in \mathcal{S}_j^{-1}} \ell_j(\mathbf{s}_j|\theta_1)}{\sum_{\mathbf{s}_j \in \mathcal{S}_j^{-1}} \ell_j(\mathbf{s}_j|\theta_2)} \right)^{1/2}, \quad (\text{A.10.2})$$

$$W_i = \left(\frac{\sum_{\mathbf{s}_i \in \mathcal{S}_i^1} \ell_i(\mathbf{s}_i|\theta_1)}{\sum_{\mathbf{s}_i \in \mathcal{S}_i^1} \ell_i(\mathbf{s}_i|\theta_2)} \times \frac{\sum_{\mathbf{s}_i \in \mathcal{S}_i^{-1}} \ell_i(\mathbf{s}_i|\theta_2)}{\sum_{\mathbf{s}_i \in \mathcal{S}_i^{-1}} \ell_i(\mathbf{s}_i|\theta_1)} \right)^{1/2}. \quad (\text{A.10.3})$$

Furthermore let $w_i := \log W_i$ and $\eta_i := \log(V_i \nu_i(\theta_1) / \nu_i(\theta_2))$ be constants that are determined completely by the initial prior and signal structures of each agent and her neighbors. Subsequently, taking logarithms of both sides in (A.10.1) yields the following update rule for the log-ratio of the beliefs at time one,

$$\log \left(\frac{\mu_{i,1}(\theta_1)}{\mu_{i,1}(\theta_2)} \right) = \sum_{j \in \mathcal{N}_i} w_j \mathbf{a}_{j,0} + \eta_i + \lambda_1(\mathbf{s}_{i,0}). \quad (\text{A.10.4})$$

Finally, we can apply (4.2.1) to derive the claimed expression in Lemma 4.1 for the updated Bayesian action of agent i following her observations of her neighbors' actions $\mathbf{a}_{j,0}$, $j \in \mathcal{N}_i$ and her own private signal $\mathbf{s}_{i,0}$. We end our derivation by pointing out some facts concerning constants η_i and w_i which appear in (A.10.4).

Lemma A.3 (Monotonicity of η_i). *Consider any $i \in [n]$ and fix a signal $s_j \in \mathcal{S}_j$ for some $j \in \mathcal{N}_i$. It holds true that the constant η_i is increasing in $\ell_j(s_j | \theta_1)$ and decreasing in $\ell_j(s_j | \theta_2)$.*

Proof. The claim follows directly from the defining relation $\eta_i = \log(\nu_i(\theta_1) / \nu_i(\theta_2)) + \log V_i$, as replacing from (A.10.2) yields

$$\begin{aligned} \log V_i = & + \frac{1}{2} \sum_{j \in \mathcal{N}_i} \log \sum_{s_j \in \mathcal{S}_j^1} \ell_j(s_j | \theta_1) + \frac{1}{2} \sum_{j \in \mathcal{N}_i} \log \sum_{s_j \in \mathcal{S}_j^{-1}} \ell_j(s_j | \theta_1) \\ & - \frac{1}{2} \sum_{j \in \mathcal{N}_i} \log \sum_{s_j \in \mathcal{S}_j^1} \ell_j(s_j | \theta_2) - \frac{1}{2} \sum_{j \in \mathcal{N}_i} \log \sum_{s_j \in \mathcal{S}_j^{-1}} \ell_j(s_j | \theta_2). \end{aligned} \quad (\text{A.10.5})$$

The proof now follows upon the realization that for any fixed $s_j \in \mathcal{S}_j$, $j \in \mathcal{N}_i$ the term $\ell_j(s_j | \theta_1)$ appears in one of the first two terms appearing with a plus sign in (A.10.5), and the term $\ell_j(s_j | \theta_2)$ appears in one of the last first two terms appearing with a minus sign in (A.10.5). Hence, when all else kept constant, $\log V_i$ and subsequently η_i is increasing in $\ell_j(s_j | \theta_1)$ and decreasing in $\ell_j(s_j | \theta_2)$. \square

Lemma A.4 (Positivity of w_i). *It holds true for any $i \in [n]$ that $w_i \geq 0$.*

Proof. First note from the definitions of the sets \mathcal{S}_i^1 and \mathcal{S}_i^{-1} that $\forall s \in \mathcal{S}_i^1$,

$$\frac{\ell_i(s | \theta_1)}{\ell_i(s | \theta_2)} \geq \frac{\nu_i(\theta_2)}{\nu_i(\theta_1)}, \text{ and } \forall s \in \mathcal{S}_i^{-1}, \frac{\ell_i(s | \theta_2)}{\ell_i(s | \theta_1)} > \frac{\nu_i(\theta_1)}{\nu_i(\theta_2)}.$$

Next we sum the numerators and denominators of the likelihood ratios of the signals in each of sets \mathcal{S}_i^1 and \mathcal{S}_i^{-1} ; invoking basic algebraic properties from the resultant fractions

yields

$$\frac{\sum_{s \in \mathcal{S}_i^1} \ell_i(s|\theta_1)}{\sum_{s \in \mathcal{S}_i^1} \ell_i(s|\theta_2)} \geq \frac{\nu_i(\theta_2)}{\nu_i(\theta_1)}, \text{ and } \frac{\sum_{s \in \mathcal{S}_i^{-1}} \ell_i(s|\theta_2)}{\sum_{s \in \mathcal{S}_i^{-1}} \ell_i(s|\theta_1)} > \frac{\nu_i(\theta_1)}{\nu_i(\theta_2)}.$$

Subsequently, replacing form (A.10.3) yields that

$$W_i := \frac{\sum_{s \in \mathcal{S}_i^1} \ell_i(s|\theta_1)}{\sum_{s \in \mathcal{S}_i^1} \ell_i(s|\theta_2)} \times \frac{\sum_{s \in \mathcal{S}_i^{-1}} \ell_i(s|\theta_2)}{\sum_{s \in \mathcal{S}_i^{-1}} \ell_i(s|\theta_1)} \geq \frac{\nu_i(\theta_2)}{\nu_i(\theta_1)} \times \frac{\nu_i(\theta_1)}{\nu_i(\theta_2)} = 1,$$

and proof follows from the defining relation $w_i := \log W_i \geq 0$. \square

A.11 A Markov Chain on the Boolean Cube

To begin, for any vertex of the Boolean hypercube $\bar{\mathbf{a}} := (a_1, \dots, a_n)^\top \in \{\pm 1\}^n$ and each agent i , define the function $\pi_i : \{\pm 1\}^n \rightarrow [0, 1]$ as $\pi_i(\bar{\mathbf{a}}) := \mathcal{P}\{\bar{\mathbf{a}}_{i,t+1} = +1 \mid \bar{\mathbf{a}}_t = \bar{\mathbf{a}}\} = \mathcal{P}_\theta\{-\lambda_1(\mathbf{s}_{i,t+1}) \leq \sum_{j \in \mathcal{N}_i} w_j a_j + \eta_i\}$. The transition probabilities for the Markov chain of action profiles on the Boolean hypercube are given by

$$\mathcal{P}(\bar{\mathbf{a}}', \bar{\mathbf{a}}) := \mathcal{P}\{\bar{\mathbf{a}}_{t+1} = \bar{\mathbf{a}}' \mid \bar{\mathbf{a}}_t = \bar{\mathbf{a}}\} = \prod_{i: a'_i = +1} \pi_i(\bar{\mathbf{a}}) \prod_{i: a'_i = -1} (1 - \pi_i(\bar{\mathbf{a}})), \quad (\text{A.11.1})$$

for all $t \in \mathbb{N}_0$ and any pair of vertices $\bar{\mathbf{a}}' := (a'_1, \dots, a'_n)^\top \in \{\pm 1\}^n$ and $\bar{\mathbf{a}} \in \{\pm 1\}^n$.

It follows from the classification of states and chains in [200, Section 2.4] that $\{\pm 1\}^n$ can be partitioned into sets of transient communication classes: $\mathcal{C}'_1, \dots, \mathcal{C}'_{r'}$, and recurrent (ergodic) communication classes: $\mathcal{C}_1, \dots, \mathcal{C}_r$. Moreover, as $t \rightarrow \infty$, $\bar{\mathbf{a}}_t$ almost surely belongs to $\cup_{i \in [r]} \mathcal{C}_i$. It is further true that if $\bar{\mathbf{a}}_{t_0} \in \mathcal{C}_i$ for some $i \in [r]$ and $t_0 \in \mathbb{N}$, then $\bar{\mathbf{a}}_t \in \mathcal{C}_i$ almost surely for all $t \geq t_0$: the process will almost surely leave any set of transient action profiles, i.e. $\cup_{i \in [r']} \mathcal{C}'_i$, and will almost surely remain in the first recurrent set that it reaches before any other. Let $\mathbf{r}^* := \arg \min_{\rho \in [r]} \{t : \bar{\mathbf{a}}_t \in \mathcal{C}_\rho\}$ be the random variable that determines the first ergodic set of action profiles that is reached by the Markov chain process $\{\bar{\mathbf{a}}_t, t \in \mathbb{N}_0\}$; suppose $\tau := \text{card}(\mathcal{C}_{r^*})$ and further denote $\mathcal{C}_{r^*} := \{\mathbf{a}_1^*, \dots, \mathbf{a}_\tau^*\}$. The asymptotic behavior of the process can now be characterized as follows.

Proposition A.1 (Asymptotic Distribution of Action Profiles). *Let $\bar{\mathbf{p}} := (\mathbf{p}_1, \dots, \mathbf{p}_\tau)^\top$ be the stationary distribution over \mathcal{C}_{r^*} which uniquely satisfies $\mathbf{p}_k \sum_{j=1}^\tau \mathcal{P}(\mathbf{a}_k^*, \mathbf{a}_j^*) = \sum_{j=1}^\tau \mathcal{P}(\mathbf{a}_k^*, \mathbf{a}_j^*) \mathbf{p}_j$, for all $k \in [\tau]$. Then $\mathcal{P}\{\lim_{t \rightarrow \infty} \bar{\mathbf{a}}_t = \mathbf{a}_k^*\} = \mathbf{p}_k$, for all $k \in [\tau]$. \square*

A.12 Proof of Proposition 4.2: Equilibrium Action Profiles

Any equilibrium $\bar{a}^* := (a_1^*, \dots, a_n^*)$ of (4.2.3) should satisfy $a_i^* = \text{sign}(\sum_{j \in \mathcal{N}_i} w_j a_j^* + \eta_i + \lambda_1(s_{i,t}))$, with probability one for all i and t . Hence, $\bar{a}^* \in \{\pm 1\}^n$ is an equilibrium of (4.2.3) if, and only if, $-\lambda_1(s_i) \leq \sum_{j \in \mathcal{N}_i} w_j a_j^* + \eta_i, \forall s_i \in \mathcal{S}_i$ whenever $a_i^* = 1$, and $-\lambda_1(s_i) \geq \sum_{j \in \mathcal{N}_i} w_j a_j^* + \eta_i, \forall s_i \in \mathcal{S}_i$ whenever $a_i^* = -1$. By multiplying both sides of the inequalities by a_i^* in each case and reordering the terms we derive the claimed characterization of the equilibria or the absorbing states under the action update dynamics in (4.2.3).¹ \square

A.13 Proof of Lemma 4.2: Time-One Bayesian Beliefs

We begin by applying the Bayes rule to the observation of agent i at time 1 which include her neighbors' initial beliefs $\{\mu_{j,0}(\cdot); j \in \mathcal{N}_i\}$ as well as her private signal $s_{i,0}$. Accordingly, for any $\hat{\theta} \in \Theta$:

$$\begin{aligned} \mu_{i,1}(\hat{\theta}) &= \mathcal{P}_{i,0}(\hat{\theta} \mid s_{i,0}, \{\mu_{j,0}(\cdot); j \in \mathcal{N}_i\}) = \frac{\mathcal{P}_{i,0}(\hat{\theta}, s_{i,0}, \{\mu_{j,0}(\cdot); j \in \mathcal{N}_i\})}{\mathcal{P}_{i,0}(s_{i,0}, \{\mu_{j,0}(\cdot); j \in \mathcal{N}_i\})} \\ &= \frac{\mathcal{P}_{i,0}(\hat{\theta}, s_{i,0}, \{\mu_{j,0}(\cdot); j \in \mathcal{N}_i\})}{\sum_{\tilde{\theta} \in \Theta} \mathcal{P}_{i,0}(\tilde{\theta}, s_{i,0}, \{\mu_{j,0}(\cdot); j \in \mathcal{N}_i\})}. \end{aligned} \quad (\text{A.13.1})$$

The succeeding steps follow those in [74] for the case of two communicating agents. For any $j \in [n]$ and all $\pi(\cdot) \in \Delta\Theta$, define the correspondence $\mathcal{I}_j : \Delta\Theta \rightarrow \mathcal{P}(\mathcal{S}_j)$ and function $\mathcal{K}_j : \Delta\Theta \rightarrow \mathbb{R}$, given by:

$$\begin{aligned} \mathcal{I}_j(\pi(\cdot)) &= \{s \in \mathcal{S}_j : \pi(\hat{\theta}) = \frac{\nu_j(\hat{\theta})l_j(s \mid \hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \nu_j(\tilde{\theta})l_j(s \mid \tilde{\theta})}, \forall \hat{\theta} \in \Theta\}, \\ \mathcal{K}_j(\pi(\cdot)) &= \sum_{s \in \mathcal{I}_j(\pi(\cdot))} \sum_{\tilde{\theta} \in \Theta} \nu_j(\tilde{\theta})l_j(s \mid \tilde{\theta}). \end{aligned} \quad (\text{A.13.2})$$

In (A.13.2), $\mathcal{I}_j(\pi(\cdot))$ signifies the set of private signals for agent j , which are consistent with the observation of belief $\pi(\cdot)$ in that agent. By the same token, $\mathcal{K}_j(\pi(\cdot))$ in (A.13.2) is the ex-ante probability for the event that the private signal of agent j belongs to the set

¹Here, and in writing the conditions for the case of $a_i^* = -1$ as non-strict inequalities we have violated our earlier convention that agents choose $+1$ when they are indifferent between $+1$ and -1 . Instead, we are assuming that ties are broken in favor of the equilibrium action profile. This assumption facilitates compact expression of the characterizing conditions for the equilibrium action profiles, and it will have no effect unless with some pathological settings of the signal structure and priors leading to $\sum_{j \in \mathcal{N}_i} w_j a_j^* + \eta_i = 0$ for some $s_i \in \mathcal{S}_i, i \in [n]$.

$\mathcal{I}_j(\pi(\cdot))$.

The terms $\mathcal{P}_{i,0}(\tilde{\theta}, \mathbf{s}_{i,0}, \{\boldsymbol{\mu}_{j,0}(\cdot); j \in \mathcal{N}_i\})$ for $\tilde{\theta} \in \Theta$, which appear in the both numerator and denominator of (A.13.1) can be simplified by conditioning on the neighbors' observed signals $\{s_j, j \in \mathcal{N}_i\}$ as follows in (A.13.3).

$$\begin{aligned} \mathcal{P}_{i,0}(\theta, \mathbf{s}_{i,0}, \{\boldsymbol{\mu}_{j,0}(\cdot); j \in \mathcal{N}_i\}) &= \quad (\text{A.13.3}) \\ &\sum_{\substack{s_j \in \mathcal{S}_j, \\ j \in \mathcal{N}_i}} \mathcal{P}(\theta, \mathbf{s}_{i,0}, \{\boldsymbol{\mu}_{j,0}(\cdot); j \in \mathcal{N}_i\} \mid \{s_j, j \in \mathcal{N}_i\}) \times \mathcal{P}_{i,0}(\{s_j, j \in \mathcal{N}_i\}). \end{aligned}$$

We next express $\mathcal{P}_{i,0}(\cdot)$ in terms of the priors and signal structures leading to:

$$\begin{aligned} \mathcal{P}_{i,0}(\tilde{\theta}, \mathbf{s}_{i,0}, \{\boldsymbol{\mu}_{j,0}(\cdot); j \in \mathcal{N}_i\}) &= \sum_{\{s_j \in \mathcal{I}_j(\boldsymbol{\mu}_{j,0}(\cdot)), j \in \mathcal{N}_i\}} \nu_i(\tilde{\theta}) l_i(\mathbf{s}_{i,0} \mid \tilde{\theta}) \prod_{j \in \mathcal{N}_i} l_j(s_j \mid \tilde{\theta}) \quad (\text{A.13.4}) \\ &= \frac{\nu_i(\tilde{\theta}) l_i(\mathbf{s}_{i,0} \mid \tilde{\theta})}{\prod_{j \in \mathcal{N}_i} \nu_j(\tilde{\theta})} \prod_{j \in \mathcal{N}_i} \left(\sum_{s_j \in \mathcal{I}_j(\boldsymbol{\mu}_{j,0}(\cdot))} \nu_j(\tilde{\theta}) l_j(s_j \mid \tilde{\theta}) \right). \end{aligned}$$

Bayes rule in (4.1.3), together with the functions defined in (A.13.2), can now be used to eliminate the product terms involving s_j from (A.13.4) and get:

$$\begin{aligned} \mathcal{P}_{i,0}(\tilde{\theta}, \mathbf{s}_{i,0}, \{\boldsymbol{\mu}_{j,0}(\cdot); j \in \mathcal{N}_i\}) &= \frac{\nu_i(\tilde{\theta}) l_i(\mathbf{s}_{i,0} \mid \tilde{\theta})}{\prod_{j \in \mathcal{N}_i} \nu_j(\tilde{\theta})} \prod_{j \in \mathcal{N}_i} \left(\boldsymbol{\mu}_{j,0}(\tilde{\theta}) \sum_{s_j \in \mathcal{I}_j(\boldsymbol{\mu}_{j,0}(\cdot))} \sum_{\tilde{\theta} \in \Theta} \nu_j(\tilde{\theta}) l_j(s_j \mid \tilde{\theta}) \right) \\ &= \nu_i(\tilde{\theta}) l_i(\mathbf{s}_{i,0} \mid \tilde{\theta}) \left(\prod_{j \in \mathcal{N}_i} \frac{\boldsymbol{\mu}_{j,0}(\tilde{\theta})}{\nu_j(\tilde{\theta})} \right) \prod_{j \in \mathcal{N}_i} \mathcal{K}_j(\boldsymbol{\mu}_{j,0}(\cdot)). \quad (\text{A.13.5}) \end{aligned}$$

Upon replacing (A.13.5) in (A.13.1), the product terms involving $\mathcal{K}_j(\boldsymbol{\mu}_{j,0}(\cdot))$ cancel out and (4.3.1) follows. \square

A.14 Proof of Theorem 4.4: No Learning when $\rho > 1$

We begin the analysis of the beliefs propagation under (4.3.2) by forming the ratio

$$\frac{\boldsymbol{\mu}_{i,t}(\check{\theta})}{\boldsymbol{\mu}_{i,t}(\theta)} = \frac{\nu_i(\check{\theta})}{\nu_i(\theta)} \times \frac{l_i(\mathbf{s}_{i,t} \mid \check{\theta})}{l_i(\mathbf{s}_{i,t} \mid \theta)} \times \prod_{j \in \mathcal{N}_i} \frac{\boldsymbol{\mu}_{j,t-1}(\check{\theta})}{\boldsymbol{\mu}_{j,t-1}(\theta)} \times \frac{\nu_j(\theta)}{\nu_j(\check{\theta})},$$

for any false state $\check{\theta} \in \Theta \setminus \{\theta\}$ and each agent $i \in [n]$ at all times $t \in \mathbb{N}$. The above has the advantage of removing the normalization factor in the dominator out of the picture; thence, focusing instead on the evolution of belief ratios, which has a log-linear format. The latter motivates definitions of log-likelihood ratios for signals, beliefs, and priors as follows. Similarly to $\lambda_{\check{\theta}}(\mathbf{s}_{i,t})$ and $\Phi_{i,t}(\check{\theta})$, define the log-ratios of prior beliefs as $\gamma_i(\check{\theta}) := \log(\nu_i(\check{\theta})/\nu_i(\theta))$. Starting from the above iterations for the belief ratio and taking the logarithms of both sides yields

$$\Phi_{i,t}(\check{\theta}) = \gamma_i(\check{\theta}) + \lambda_{\check{\theta}}(\mathbf{s}_{i,t}) + \sum_{j \in \mathcal{N}_i} \Phi_{j,t-1}(\check{\theta}) - \gamma_j(\check{\theta}).$$

Multiplying both sides of (4.3.10) by α_i , which is the centrality of agent i , and summing over all $i \in [n]$ yields that

$$\Phi_t(\check{\theta}) = \sum_{i=1}^n \alpha_i \gamma_i(\check{\theta}) + \sum_{i=1}^n \alpha_i \lambda_{\check{\theta}}(\mathbf{s}_{i,t}) + \sum_{i=1}^n \alpha_i \sum_{j \in \mathcal{N}_i} (\Phi_{j,t-1}(\check{\theta}) - \gamma_j(\check{\theta})).$$

First note that we can write

$$\sum_{i=1}^n \alpha_i \gamma_i(\check{\theta}) - \sum_{i=1}^n \alpha_i \sum_{j \in \mathcal{N}_i} \gamma_j(\check{\theta}) = \text{tr} \left\{ (I - A^T) \bar{\alpha} \bar{\gamma}(\check{\theta})^T \right\} = (1 - \rho) \beta(\check{\theta}), \quad (\text{A.14.1})$$

where $\bar{\gamma}(\check{\theta}) := (\gamma_1(\check{\theta}), \dots, \gamma_n(\check{\theta}))^T$.

Next note that by the choice of $\bar{\alpha}$ as the eigenvector corresponding to the ρ eigenvalue of matrix A we get

$$\sum_{i=1}^n \alpha_i \sum_{j \in \mathcal{N}_i} \Phi_{j,t-1}(\check{\theta}) = \bar{\alpha}^T A \bar{\Phi}_{t-1}(\check{\theta}) = \rho \bar{\alpha}^T \bar{\Phi}_{t-1}(\check{\theta}) = \rho \Phi_{t-1}(\check{\theta}). \quad (\text{A.14.2})$$

where $\bar{\Phi}_t(\check{\theta}) := (\Phi_{1,t}(\check{\theta}), \dots, \Phi_{n,t}(\check{\theta}))^T$. Now replacing (A.14.1) and (A.14.2) in (4.3.11) yields the following recursion for $\Phi_t(\check{\theta})$:

$$\Phi_t(\check{\theta}) = \Lambda_t(\check{\theta}) + \rho \Phi_{t-1}(\check{\theta}) + (1 - \rho) \beta(\check{\theta}), \quad (\text{A.14.3})$$

initialized by $\Phi_0(\check{\theta}) = \beta(\check{\theta}) + \Lambda_0(\check{\theta})$, where $\beta(\check{\theta}) := \sum_{i=1}^n \alpha_i \log(\nu_i(\check{\theta})/\nu_i(\theta))$ is a constant that is determined by the initial prior beliefs, and it measures the total bias in the network relative between the two states $\check{\theta}$ and θ . In particular, if the agents are unbiased starting from uniform priors on Θ , then $\beta(\check{\theta}) = 0$, $\forall \check{\theta} \in \Theta$. Note also that the assumption of full support priors implies that $|\beta(\check{\theta})|$ is finite. By iterating (A.14.3) for $t \in \mathbb{N}$ we obtain (4.3.6). Next note that in a strongly connected graph every node has a degree greater than

or equal to one so that $\rho \geq 1$, [194, Chapter 2]. If $\rho > 1$, then the term $\rho^t \Lambda_0(\check{\theta})$ increases in variance as $t \rightarrow \infty$, and unless $\Lambda_0(\check{\theta}) < \epsilon$ with \mathcal{P}_θ -probability one for some $\epsilon < 0$, almost sure convergence to $-\infty$ for $\Phi_t(\check{\theta})$ in (4.3.6) cannot hold true. \square

A.15 Proof of Theorem 4.7: Learning by Random Walks on Directed Graphs

To analyze the propagation of beliefs under (4.3.8) we form the belief ratio

$$\frac{\mu_{i,t}(\check{\theta})}{\mu_{i,t}(\theta)} = \frac{\ell_i(\mathbf{s}_{i,t} | \check{\theta})}{\ell_i(\mathbf{s}_{i,t} | \theta)} \prod_{j=1}^n \left(\frac{\mu_{j,t-1}(\check{\theta})}{\mu_{j,t-1}(\theta)} \right)^{\mathbb{1}_{\{\sigma_{t,i}=j\}}}, \quad (\text{A.15.1})$$

for any false state $\check{\theta} \in \Theta \setminus \{\theta\}$ and each agent $i \in [n]$ at all times $t \in \mathbb{N}$. The above has the advantage of removing the normalization factor in the dominator out of the picture; thence, focusing instead on the evolution of belief ratios. To proceed, we take the logarithms of both sides in (A.15.1) to obtain

$$\log \left(\frac{\mu_{i,t}(\check{\theta})}{\mu_{i,t}(\theta)} \right) = \log \left(\frac{\ell_i(\mathbf{s}_{i,t} | \check{\theta})}{\ell_i(\mathbf{s}_{i,t} | \theta)} \right) + \sum_{j=1}^n \mathbb{1}_{\{\sigma_{t,i}=j\}} \log \left(\frac{\mu_{j,t-1}(\check{\theta})}{\mu_{j,t-1}(\theta)} \right). \quad (\text{A.15.2})$$

Next we can iterate (A.15.2) to replace for $(\mu_{j,t-1}(\check{\theta})/\mu_{j,t-1}(\theta))$ and so on, from which we get (A.15.3) at the top of next page. Also note,

$$\sum_{i_1=1}^n \cdots \sum_{i_t=1}^n \mathbb{1}_{\{\sigma_{t,i}=i_1\}} \cdots \mathbb{1}_{\{\sigma_{1,i_{t-1}}=i_t\}} = 1,$$

almost surely, and in fact every where on Ω , so that the initial prior belief ratio $\log(\nu(\check{\theta})/\nu(\theta))$ always appears in the summation (A.15.3), and it simplifies as in (A.15.4).

$$\begin{aligned}
\log \left(\frac{\boldsymbol{\mu}_{i,t}(\check{\theta})}{\boldsymbol{\mu}_{i,t}(\theta)} \right) &= \log \left(\frac{\ell_i(\mathbf{s}_{i,t} | \check{\theta})}{\ell_i(\mathbf{s}_{i,t} | \theta)} \right) + \sum_{i_1=1}^n \mathbb{1}_{\{\sigma_{t,i}=i_1\}} \log \left(\frac{\ell_{i_1}(\mathbf{s}_{i_1,t-1} | \check{\theta})}{\ell_{i_1}(\mathbf{s}_{i_1,t-1} | \theta)} \right) \\
&+ \mathbb{1}_{\{\sigma_{t,i}=i_1\}} \sum_{i_2=1}^n \mathbb{1}_{\{\sigma_{t-1,i_1}=i_2\}} \log \left(\frac{\ell_{i_2}(\mathbf{s}_{i_2,t-2} | \check{\theta})}{\ell_{i_2}(\mathbf{s}_{i_2,t-2} | \theta)} \right) \\
&+ \mathbb{1}_{\{\sigma_{t,i}=i_1\}} \mathbb{1}_{\{\sigma_{t-1,i_1}=i_2\}} \sum_{i_3=1}^n \mathbb{1}_{\{\sigma_{t-2,i_2}=i_3\}} \log \left(\frac{\ell_{i_3}(\mathbf{s}_{i_3,t-3} | \check{\theta})}{\ell_{i_3}(\mathbf{s}_{i_3,t-3} | \theta)} \right) + \dots \\
&+ \mathbb{1}_{\{\sigma_{t,i}=i_1\}} \mathbb{1}_{\{\sigma_{t-1,i_1}=i_2\}} \dots \mathbb{1}_{\{\sigma_{1,i_{t-2}}=i_{t-1}\}} \sum_{i_t=1}^n \mathbb{1}_{\{\sigma_{1,i_{t-1}}=i_t\}} \left\{ \log \left(\frac{\ell_{i_t}(\mathbf{s}_{i_t,0} | \check{\theta})}{\ell_{i_t}(\mathbf{s}_{i_t,0} | \theta)} \right) + \log \left(\frac{\nu(\check{\theta})}{\nu(\theta)} \right) \right\}
\end{aligned} \tag{A.15.3}$$

$$\begin{aligned}
\log \left(\frac{\boldsymbol{\mu}_{i,t}(\check{\theta})}{\boldsymbol{\mu}_{i,t}(\theta)} \right) &= \log \left(\frac{\ell_i(\mathbf{s}_{i,t} | \check{\theta})}{\ell_i(\mathbf{s}_{i,t} | \theta)} \right) + \log \left(\frac{\nu(\check{\theta})}{\nu(\theta)} \right) + \sum_{i_1=1}^n \mathbb{1}_{\{\sigma_{t,i}=i_1\}} \left\{ \log \left(\frac{\ell_{i_1}(\mathbf{s}_{i_1,t-1} | \check{\theta})}{\ell_{i_1}(\mathbf{s}_{i_1,t-1} | \theta)} \right) \right. \\
&+ \sum_{i_2=1}^n \mathbb{1}_{\{\sigma_{t,i_1}=i_2\}} \left\{ \log \left(\frac{\ell_{i_2}(\mathbf{s}_{i_2,t-2} | \check{\theta})}{\ell_{i_2}(\mathbf{s}_{i_2,t-2} | \theta)} \right) + \dots + \sum_{i_\tau=1}^n \mathbb{1}_{\{\sigma_{t-\tau+1,i_{\tau-1}}=i_\tau\}} \left\{ \log \left(\frac{\ell_{i_\tau}(\mathbf{s}_{i_\tau,t-\tau} | \check{\theta})}{\ell_{i_\tau}(\mathbf{s}_{i_\tau,t-\tau} | \theta)} \right) \right. \right. \\
&+ \dots + \sum_{i_{t-1}=1}^n \mathbb{1}_{\{\sigma_{1,i_{t-2}}=i_{t-1}\}} \left\{ \log \left(\frac{\ell_{i_{t-1}}(\mathbf{s}_{i_{t-1},1} | \check{\theta})}{\ell_{i_{t-1}}(\mathbf{s}_{i_{t-1},1} | \theta)} \right) + \sum_{i_t=1}^n \mathbb{1}_{\{\sigma_{1,i_{t-1}}=i_t\}} \log \left(\frac{\ell_{i_t}(\mathbf{s}_{i_t,0} | \check{\theta})}{\ell_{i_t}(\mathbf{s}_{i_t,0} | \theta)} \right) \right. \\
&\left. \left. \left. \right\} \dots \right\}.
\end{aligned} \tag{A.15.4}$$

We now claim that whenever $t \rightarrow \infty$ and the network graph \mathcal{G} is strongly concerted, with \mathcal{P} -probability one the likelihood ratios of private signals from any node $m \in [n]$ appears in the summation (A.15.4) as $\ell_m(\mathbf{s}_{m,t-\tau} | \check{\theta})/\ell_m(\mathbf{s}_{m,t-\tau} | \theta)$ for infinitely many values of τ . The gist of the proof is in realizing the correspondence between the summation (A.15.3) and a random walk on the directed graph \mathcal{G} that starts at time t on node i , proceeds in the reversed time direction, and terminates at time zero. The jumps in this random walk are made from each node i to one of her in-neighbors $j \in \mathcal{N}_i$ and in accordance with the probabilities $p_{i,j}$ specified by matrix $P = [p_{i,j}]$. Indeed, we can denote the random sequence of nodes that are hit by this random walk as $(i, \mathbf{i}_1, \dots, \mathbf{i}_t)$ where the random variables $\mathbf{i}_\tau \in [n]$, $\tau \in [t]$ are defined recursively by $\mathbf{i}_1 := \sigma_{t,i}$, $\mathbf{i}_2 := \sigma_{t-1,\sigma_{t,i_1}}$, $\mathbf{i}_3 := \sigma_{t-2,\sigma_{t-1,i_2}}$, \dots ,

$\mathbf{i}_t := \sigma_{1, \sigma_{2, i_{t-1}}}$. Whence (A.15.4) is written succinctly as

$$\log \left(\frac{\boldsymbol{\mu}_{i,t}(\check{\theta})}{\boldsymbol{\mu}_{i,t}(\theta)} \right) = \log \left(\frac{\ell_i(\mathbf{s}_{i,t} | \check{\theta})}{\ell_i(\mathbf{s}_{i,t} | \theta)} \right) + \log \left(\frac{\nu(\check{\theta})}{\nu(\theta)} \right) + \sum_{\tau=1}^t \log \left(\frac{\ell_{i_\tau}(\mathbf{s}_{i_\tau, t-\tau} | \check{\theta})}{\ell_{i_\tau}(\mathbf{s}_{i_\tau, t-\tau} | \theta)} \right). \quad (\text{A.15.5})$$

As $t \rightarrow \infty$, the sequence $\mathbf{i}_\tau, \tau \in \mathbb{N}$ forms a Markov process with transition matrix P . Given (A.15.5), our claim can be restated as that for every $m \in [n]$ and as $t \rightarrow \infty$ there are infinitely many values of $\tau \in \mathbb{N}$ for which $\mathbf{i}_\tau = m$, and it is true because in a finite state Markov chain with transition matrix P every state is persistent (recurrent) and will be hit infinitely many times provided that the directed graph \mathcal{G} is strongly connected [198, Theorem 1.5.6], i.e. we have that $\forall m \in [n]$,

$$\mathcal{P}\{\mathbf{i}_\tau = m, \text{ for infinitely many } \tau\} = 1.$$

For any agent $m \in [n]$ let $\mathcal{T}_m := \{\tau_{m,j}, j \in \mathbb{N}\}$ be the sequence of stopping times that record the first, second and so on passage times of node m by the process $\mathbf{i}_\tau, \tau \in \mathbb{N}$. That is we have $\tau_{m,1} = \inf\{\tau \in \mathbb{N} : \mathbf{i}_\tau = m\}$ and for $j > 1$, $\tau_{m,j} = \inf\{\tau > \tau_{m,j-1} : \mathbf{i}_\tau = m\}$. Using the above notation, (A.15.5) can be rewritten as

$$\begin{aligned} \log \left(\frac{\boldsymbol{\mu}_{i,t}(\check{\theta})}{\boldsymbol{\mu}_{i,t}(\theta)} \right) &= \log \left(\frac{\ell_i(\mathbf{s}_{i,t} | \check{\theta})}{\ell_i(\mathbf{s}_{i,t} | \theta)} \right) + \log \left(\frac{\nu(\check{\theta})}{\nu(\theta)} \right) \\ &\quad + \sum_{m=1}^n \sum_{\substack{\tau \in \mathcal{T}_m, \\ \tau \leq t}} \log \left(\frac{\ell_m(\mathbf{s}_{m,t-\tau} | \check{\theta})}{\ell_m(\mathbf{s}_{m,t-\tau} | \theta)} \right). \end{aligned} \quad (\text{A.15.6})$$

On the other hand, note that $\log(\ell_m(\mathbf{s}_{m,t-\tau_{m,j}} | \check{\theta}) / \ell_m(\mathbf{s}_{m,t-\tau_{m,j}} | \theta)), j \in \mathbb{N}$ is a sequence of independent and identically distributed signals, so that by the strong of large numbers we obtain that with \mathcal{P} -probability one,

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^n \log \left(\frac{\ell_m(\mathbf{s}_{m,t-\tau_{m,j}} | \check{\theta})}{\ell_m(\mathbf{s}_{m,t-\tau_{m,j}} | \theta)} \right) &= \mathbb{E} \log \left(\frac{\ell_m(\mathbf{s}_{m,0} | \check{\theta})}{\ell_m(\mathbf{s}_{m,0} | \theta)} \right) \\ &:= -D_{\text{KL}}(\ell_m(\cdot | \check{\theta}) \| \ell_m(\cdot | \theta)) \leq 0, \end{aligned} \quad (\text{A.15.7})$$

where the non-positivity follows from the information inequality for the Kullback-Leibler divergence $D_{\text{KL}}(\cdot \| \cdot)$ and is strict whenever $\ell_m(\cdot | \check{\theta}) \not\equiv \ell_m(\cdot | \theta)$, i.e. $\exists s \in \mathcal{S}_i$ such that $\ell_i(s | \check{\theta}) \neq \ell_i(s | \theta)$ [193, Theorem 2.6.3]. Note that whenever $\ell_i(\cdot | \hat{\theta}) \equiv \ell_i(\cdot | \theta)$ or equivalently $D_{\text{KL}}(\ell_i(\cdot | \hat{\theta}) \| \ell_i(\cdot | \theta)) = 0$, then the two states $\hat{\theta}$ and θ are statically indistinguishable to agent i . In other words, there is no way for agent i to differentiate $\hat{\theta}$ from θ based

only on her private signals. This follows from the fact that both θ and $\hat{\theta}$ induce the same probability distribution on her sequence of observed i.i.d. signals. On the other hand, having $D_{\text{KL}}(\ell_m(\cdot|\theta)\|\ell_m(\cdot|\check{\theta})) < 0$ for some agent $m \in [n]$ would ensure per (A.15.7) and persistence of state m that with \mathcal{P} -probability one,

$$\sum_{\substack{\tau \in \mathcal{T}_m, \\ \tau \leq t}} \log \left(\frac{\ell_m(\mathbf{s}_{m,t-\tau} | \check{\theta})}{\ell_m(\mathbf{s}_{m,t-\tau} | \theta)} \right) \rightarrow -\infty$$

as $t \rightarrow \infty$ in (A.15.6); consequently, $\log(\mu_{i,t}(\check{\theta})/\mu_{i,t}(\theta)) \rightarrow -\infty$ for all agent $i \in [n]$ and any such $\check{\theta} \in \Theta$, $\check{\theta} \neq \theta$. Indeed, having $\log(\mu_{i,t}(\check{\theta})/\mu_{i,t}(\theta)) \rightarrow -\infty$ for all $\check{\theta} \neq \theta$ is necessary and sufficient for learning, and Definition 4.2 provides the required characterization, as claimed in Theorem 4.7, for learning under the without recall updates in (4.3.3) and (4.3.7), where the neighbor j is chosen randomly with strictly positive probabilities specified in transition matrix \mathcal{P} .

We can now extend the above analysis to derive an asymptotic rate of learning for the agents that is exponentially fast and is expressed as $\sum_{m=1}^n \pi_m D_{\text{KL}}(\ell_m(\cdot|\theta)\|\ell_m(\cdot|\check{\theta})) < 0$, where $\bar{\pi} := (\pi_1, \dots, \pi_n)$ is the stationary distribution of the transition matrix \mathcal{P} , which for a strongly connected \mathcal{G} is the unique probability distribution on $[n]$ satisfying $\bar{\pi}\mathcal{P} = \bar{\pi}$. To see how, for each agent $m \in [n]$ and all time t , define $\mathcal{T}_m(t) := \{\tau_{m,j}, j \in \mathbb{N} : \tau_{m,j} \leq t\}$ and divide both sides of (A.15.6) by t to obtain

$$\begin{aligned} \frac{1}{t} \log \left(\frac{\mu_{i,t}(\check{\theta})}{\mu_{i,t}(\theta)} \right) &= \frac{1}{t} \log \left(\frac{\ell_i(\mathbf{s}_{i,t} | \check{\theta})}{\ell_i(\mathbf{s}_{i,t} | \theta)} \right) + \frac{1}{t} \log \left(\frac{\nu(\check{\theta})}{\nu(\theta)} \right) \\ &\quad + \frac{1}{t} \sum_{m=1}^n \sum_{\tau \in \mathcal{T}_m(t)} \log \left(\frac{\ell_m(\mathbf{s}_{m,t-\tau} | \check{\theta})}{\ell_m(\mathbf{s}_{m,t-\tau} | \theta)} \right). \end{aligned}$$

Upon invoking (A.15.7) we obtain

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \left(\frac{\mu_{i,t}(\check{\theta})}{\mu_{i,t}(\theta)} \right) = - \sum_{m=1}^n \lim_{t \rightarrow \infty} \frac{|\mathcal{T}_m(t)|}{t} D_{\text{KL}}(\ell_m(\cdot|\theta)\|\ell_m(\cdot|\check{\theta})). \quad (\text{A.15.8})$$

Finally the ergodic theorem ensures that the average time spent in any state $m \in [n]$ converges almost surely to its stationary probability π_m , i.e. with probability one [198, Theorem 1.10.2]:

$$\lim_{t \rightarrow \infty} \frac{|\mathcal{T}_m(t)|}{t} = \pi_m.$$

Hence, (A.15.8) becomes

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \left(\frac{\boldsymbol{\mu}_{i,t}(\check{\theta})}{\boldsymbol{\mu}_{i,t}(\theta)} \right) = - \sum_{m=1}^n \pi_m D_{\text{KL}} (\ell_m(\cdot|\theta) \parallel \ell_m(\cdot|\check{\theta})),$$

completing the proof for the claimed asymptotically exponentially fast rate.

A.16 Analysis of Convergence when Agents Recall their Self Beliefs

Using $\chi_t := \rho(1 - \eta_t)$ and $B := (1/\rho)A$, the previous equation can be written in vectorized format as follows

$$\bar{\boldsymbol{\Phi}}_t(\check{\theta}) = (I + \chi_t B) \bar{\boldsymbol{\Phi}}_{t-1}(\check{\theta}) + \bar{\boldsymbol{\lambda}}_t(\check{\theta}) = \sum_{\tau=0}^t P^{(t,\tau)} \bar{\boldsymbol{\lambda}}_{\tau}(\check{\theta}) + P^{(t,0)} \bar{\boldsymbol{\Psi}}(\check{\theta}), \quad (\text{A.16.1})$$

where $P^{(t,t)} := I$, and $P^{(t,\tau)} := \prod_{u=\tau+1}^t (I + \chi_u B)$ for $\tau < t$. Next note that $P^{(t,\tau)}$ can be expanded as follows

$$P^{(t,\tau)} = \sum_{j=0}^{t-\tau} M_j^{(t,\tau)} B^j,$$

where $M_0^{(t,\tau)} = 1$, $\tau \leq t$ and

$$M_j^{(t,\tau)} = \sum_{u_1=\tau+1}^{t-j+1} \sum_{u_2=u_1+1}^{t-j+2} \cdots \sum_{u_j=u_{j-1}+1}^t \chi_{u_1} \chi_{u_2} \cdots \chi_{u_j}.$$

Consequently, (A.16.1) can be rewritten as

$$\bar{\boldsymbol{\Phi}}_t(\check{\theta}) = \sum_{\tau=0}^t \sum_{j=0}^{t-\tau} M_j^{(t,\tau)} B^j \bar{\boldsymbol{\lambda}}_{\tau}(\check{\theta}) + \sum_{j=0}^t M_j^{(t,0)} B^j \bar{\boldsymbol{\Psi}}(\check{\theta}), \quad (\text{A.16.2})$$

To proceed, for fixed τ and j let $M_j^{(\tau)} := \lim_{t \rightarrow \infty} M_j^{(t,\tau)}$. Next consider the summands in (A.16.2) for each τ , $0 \leq \tau \leq t$. Note that for j fixed, $\{B^j \bar{\boldsymbol{\lambda}}_{\tau}(\check{\theta}), \tau \in \mathbb{N}_0\}$ is a sequence of independent and identically distributed random vectors. On the other hand, since B has unit spectral radius and given that $\chi_u > 0$, having $M_1^{(0)} := \sum_{u=1}^{\infty} \chi_u < \infty$ is sufficient to ensure that the random vectors $M_j^{(t,\tau)} B^j \bar{\boldsymbol{\lambda}}_{\tau}(\check{\theta})$ are all in \mathcal{L}^2 and have variances that are bounded

uniformly in the choice of t , τ . This is because for any t , τ , and j we have that

$$M_j^{(t,\tau)} \leq M_j^{(\tau)} \leq M_j^{(0)} \leq \frac{1}{j!} \left(M_1^{(0)} \right)^j,$$

all as a consequence of positivity, $\kappa_u > 0$. In particular, with $M_1^{(0)} < \infty$ we can bound

$$\left\| \sum_{j=0}^{\infty} M_j^{(0)} B^j \bar{\psi}(\check{\theta}) \right\| \leq \sum_{j=0}^{\infty} M_j^{(0)} \|B^j \bar{\psi}(\check{\theta})\| \leq \|\bar{\psi}(\check{\theta})\| \exp(M_1^0), \quad (\text{A.16.3})$$

so that the contribution made by the initial bias of the network is asymptotically bounded and therefore sub-dominant when

$$\lim_{t \rightarrow \infty} \sum_{\tau=0}^t \sum_{j=0}^{t-\tau} M_j^{(t,\tau)} B^j \bar{\lambda}_{\tau}(\check{\theta}) = (-\infty)_n,$$

almost surely; here, by $(-\infty)_n$ we mean the entry-wise convergence of the column vector to $-\infty$ for the each of the n entries, corresponding to the n agents. In the sequel we investigate conditions under which this almost sure convergence would hold true.

We begin by noting that the condition $M_1^{(0)} < \infty$ is indeed necessary for convergence, because if $M_1^{(0)} = \infty$, then the term $M_1^{(t,0)} B \bar{\lambda}_0(\check{\theta})$ appearing in (A.16.2) for $j = 1$ and $\tau = 0$ increases unbounded in its variance as $t \rightarrow \infty$, so that (A.16.2) cannot converge in an almost sure sense.

Next note that with $M_1^{(0)} < \infty$ we can invoke Kolmogorov's criterion, [217, Section X.7], to get that as $t \rightarrow \infty$ the summation in (A.16.2) converges almost surely to its expected value, i.e. the following almost sure limit holds true

$$\lim_{t \rightarrow \infty} \Phi_{i,t}(\check{\theta}) = \sum_{\tau=0}^{\infty} \sum_{j=0}^{\infty} M_j^{(\tau)} [B^j \bar{\lambda}(\check{\theta})]_i + \sum_{j=0}^{\infty} M_j^{(0)} [B^j \bar{\psi}(\check{\theta})]_i, \quad (\text{A.16.4})$$

the second term being bounded per (A.16.3). Moreover, for a strongly connected social network \mathcal{G} , if it is aperiodic, then the matrix B is a primitive matrix; and in particular for all $j \geq d := \text{diam}(\mathcal{G}) + 1$, every entry of B^j is strictly greater than zero, and if the truth is globally identifiable then one can take an absolute constant $\epsilon > 0$ such that $[B^j \bar{\lambda}(\check{\theta})]_i < -\epsilon$ whenever $j \geq d$, while $[B^j \bar{\lambda}(\check{\theta})]_i \leq 0$ for any j . Subsequently, we get that

$$\sum_{\tau=0}^{\infty} \sum_{j=0}^{\infty} M_j^{(\tau)} [B^j \bar{\lambda}(\check{\theta})]_i \leq -\epsilon \sum_{\tau=0}^{\infty} \sum_{j=d}^{\infty} M_j^{(\tau)}. \quad (\text{A.16.5})$$

Combining the results of (A.16.3), (A.16.4) and (A.16.5) leads to the following characterization: all agents will learn the truth (that is $\lim_{t \rightarrow \infty} \Phi_{i,t}(\check{\theta}) = -\infty$, almost surely for all i

and any $\check{\theta} \neq \theta$), if $M_1^{(0)} < \infty$ and $\sum_{\tau=0}^{\infty} \sum_{j=d}^{\infty} M_j^{(\tau)} = \infty$. Notice the preceding conditions are indeed not far from necessity. Firstly, we need $M_1^{(0)} < \infty$ to bound the growth of variance for convergence, as noted above. Moreover, with B having a unit spectral radius we can bound

$$|[B^j \bar{\lambda}(\check{\theta})]_i| \leq \|B^j \bar{\lambda}(\check{\theta})\| \leq \|\bar{\lambda}(\check{\theta})\|,$$

so that we can lower-bound $\lim_{t \rightarrow \infty} \Phi_{i,t}(\check{\theta})$ in (A.16.4) as follows

$$-\|\bar{\lambda}(\check{\theta})\| \sum_{\tau=0}^{\infty} \sum_{j=0}^{\infty} M_j^{(\tau)} - \|\bar{\psi}(\check{\theta})\| \exp(M_1^0) \leq \lim_{t \rightarrow \infty} \Phi_{i,t}(\check{\theta}).$$

Consequently, if $\sum_{\tau=0}^{\infty} \sum_{j=0}^{\infty} M_j^{(\tau)} < \infty$, then $\lim_{t \rightarrow \infty} \Phi_{i,t}(\check{\theta})$ is almost surely bounded away from $-\infty$ and agents do not learn the truth.

For a strongly connected and aperiodic social network \mathcal{G} , matrix B has a single eigenvalue at one (corresponding to the largest eigenvalue of the adjacency A) and all of the other eigenvalues of B have magnitudes strictly less than one. Therefore, by the iterations of the power method, [199, Section 11.1], we know that $[B^j \bar{\lambda}(\check{\theta})]_i$ converges to $\sum_{k=1}^n \alpha_k \lambda_k(\check{\theta})$ for every i , and the convergence is geometrically fast in the magnitude-ratio of the first and second largest eigenvalues of the adjacency matrix A . Subsequently, for a strongly connected and aperiodic social network \mathcal{G} , we can replace $-\epsilon$ and d in (A.16.5) by $\epsilon + \sum_{i=1}^n \alpha_i \lambda_i(\check{\theta}) < 0$ and some constant $D(\epsilon)$. Here, $\epsilon > 0$ is a small but arbitrary and $D(\epsilon)$ is chosen large enough in accordance with the geometric rate of $[B^j \bar{\lambda}(\check{\theta})]_i \rightarrow \sum_{k=1}^n \alpha_k \lambda_k(\check{\theta})$, such that $|[B^j \bar{\lambda}(\check{\theta})]_i - \sum_{k=1}^n \alpha_k \lambda_k(\check{\theta})| < \epsilon$ for all $j \geq D(\epsilon)$, and the analysis of convergence and its rate can thus be refined.

Furthermore, having

$$K_1 < \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^t \sum_{j=d}^{t-\tau} M_j^{(t,\tau)} \leq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^t \sum_{j=d}^{t-\tau} M_j^{(t,\tau)} < K_2,$$

for some positive constants $0 < K_1 < K_2$ implies that the learning rate is asymptotically exponentially fast, as was the case for all the other update rules that we discussed in this paper. However, depending on how slow or fast (compared to t) is the convergence $\sum_{\tau=0}^t \sum_{j=d}^{t-\tau} M_j^{(t,\tau)} \rightarrow \infty$ as $t \rightarrow \infty$, the almost sure asymptotic rate at which for some $\check{\theta} \neq \theta$ and $i \in [n]$, $\mu_{i,t}(\check{\theta}) \rightarrow 0$ as $t \rightarrow \infty$ could be slower or faster than an exponential.

Appendix B

Rational Equilibrium and Consensus in Symmetric Groups

Following [218] and [15], the asymptotic outcome of Bayesian group decision can be characterized as a Markov perfect Bayesian equilibrium in a repeated game of incomplete information that is played by successive generations of short-lived players. Short-lived agents inherit the beliefs of the player playing in the previous stage in their role while also observing the last stage actions of the players in their social neighborhood. Rational myopia arise by nature of short-lived agents, and the equilibrium concept can be used to study the rational myopic decisions, subject to the assumption that other players are also playing their myopic best responses given their own information. Markov perfect Bayesian equilibrium is the appropriate solution concept for the study of Bayesian group decision process because a Markovian strategy for agent i can depend on the information available to her, $\mathbf{h}_{i,t}$, in her role as agent i at time t , only to extent that $\mathbf{h}_{i,t}$ is informative about θ , the pay-off relevant state of the world. In [218] the authors provide the following recursive construction of Markovian strategy profiles: consider the probability triplet $(\Omega, \mathcal{B}, \mathcal{P})$, where $\Omega = \Theta \times \prod_{i \in [n]} \mathcal{S}_i$, \mathcal{B} is the Borel sigma algebra, and \mathcal{P} assigns probabilities to the events in \mathcal{B} consistently with the common prior ν and the product of the likelihoods $\mathcal{L}(\cdot | \cdot)$; for each i , let $\sigma_{i,0} : \Omega \rightarrow \mathcal{A}_i$ be a measurable map defined on $(\Omega, \mathcal{B}, \mathcal{P})$ that specifies the time zero action of agent i as a function of her private signal, and let $\mathcal{H}_{i,0}$ denote the information available to agent i at time zero which is the smallest sub-sigma algebra of \mathcal{B} that makes \mathbf{s}_i measurable. Then for any time t , we can define a Markovian strategy $\sigma_{i,t}$, recursively, as a random variable which is measurable with respect to $\mathcal{H}_{i,t}^{\bar{\sigma}_{t-1}}$, where $\bar{\sigma}_{t-1} = (\bar{\sigma}_{1,t-1}, \dots, \bar{\sigma}_{n,t-1})$, $\bar{\sigma}_{i,t-1} = (\sigma_{i,0}, \dots, \sigma_{i,t-1})$ for all i , and $\mathcal{H}_{i,t}^{\bar{\sigma}_{t-1}}$ is the smallest sub-sigma algebra of \mathcal{B} that makes \mathbf{s}_i and $\bar{\sigma}_{j,t-1}, j \in \mathcal{N}_i$ measurable. The contributions of [218] and [15] consist of proving convergence to an equilibrium profile $\bar{\sigma}_\infty$ and showing consensus properties for the equilibrium profile, the former (convergence result) relies on the compactness of the action space, while the latter (asymptotic consensus result) relies on an imitation principle argument that works for common (symmetric among the agents) utility and action structures.¹ Both results rely on some analytical properties

¹In symmetric groups all agents have the same action space $\mathcal{A}_i = \mathcal{A}_j$ for all i, j and identical utility functions $u_i(\mathbf{a}, \theta) = u_j(\mathbf{a}, \theta)$ for all $\mathbf{a} \in \mathcal{A}$ and any $\theta \in \Theta$. Symmetric settings arise very naturally in

of the utility function as well, such as supermodularity¹ in [218] or continuity (where the action spaces are metric compact spaces) and boundedness between L_2 integrable functions in [15]. Other works have looked at different asymptotics; in particular, information aggregation and learning as the number of agents grows [53, 54]. In this work we are interested in the computations that are required of a Bayesian agent in order for her to achieve her optimal recommendations at every (finite) step during a group decision process, rather than the asymptotic and equilibrium properties of such recommendations.

group-decision scenarios where people have similar preferences about the group-decision outcome and seek the same truth or a common goal. In such scenarios, the question of consensus or unanimity is of particular importance, as it gives a sharp prediction about the group decision outcome and emergence of agreement among individual decision makers.

¹In general, the supermodularity of the utilities signifies strategic complementarity between the actions of the players, as is the case for [218]; however, in the absence of strategic externalities (as is the case for GDP) supermodularity implies a case of diminishing returns: $u_i(\cdot, \cdot)$ is strictly supermodular iff $u_i(\min\{\alpha, \alpha'\}, \theta) + u_i(\max\{\alpha, \alpha'\}, \theta) > u_i(\alpha, \theta) + u_i(\alpha', \theta)$, for all $\alpha \neq \alpha' \in \mathcal{A}_i$ and each $\theta \in \Theta$.

Appendix C

Bayesian Calculations for Forward Reasoning

In this Appendix, we explain the Bayesian calculations that are involved in a forward reasoning implementation of a decision process, and in particular focus on the case of a directed path. Taking up after (2.1.3), at time one having observed her neighbor's actions from time zero agent i learns something about the private signals of each of her neighbors $\{s_j, j \in \mathcal{N}_i\}$. In particular, from her observation of $\{a_{j,0}, j \in \mathcal{N}_i\}$ she infers that s_j for each $j \in \mathcal{N}_i$ should necessarily satisfy:

$$s_j \in \left\{ s_j \in \mathcal{S}_j : a_{j,0} \in \arg \max_{a_j \in \mathcal{A}_j} \frac{\sum_{\hat{\theta} \in \Theta} u_j(a_j, \hat{\theta}) \ell_j(s_j | \hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \ell_j(s_j | \tilde{\theta}) \nu(\tilde{\theta})} \right\}.$$

Subsequently, she crosses out any signal profile $\bar{s} \in \mathcal{I}_{i,0}$ for which s_j does not satisfy

$$a_{j,0} \in \arg \max_{a_j \in \mathcal{A}_j} \frac{\sum_{\hat{\theta} \in \Theta} u_j(a_j, \hat{\theta}) \ell_j(s_j | \hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \ell_i(s_j | \tilde{\theta}) \nu(\tilde{\theta})},$$

thus pruning $\mathcal{I}_{i,0}$ into the smaller set $\mathcal{I}_{i,1}$. Agent i then updates her Bayesian posterior $\mu_{i,1}$ and changes her recommendation $a_{i,1}$ according to (2.1.1) and (2.1.2), respectively. At time 2 the agent observes her neighbors recommendations $\{a_{j,1}, j \in \mathcal{N}_i\}$ for a second time. The second interaction informs her about what actions her neighbor's neighbors could have taken at time zero and in turn what private signals they could have observed at time zero; in addition, she also refines what she has already learned about private observations of her neighbors in \mathcal{N}_i based on their actions at time one, $\{a_{j,1}, j \in \mathcal{N}_i\}$, cf. Fig. C.1.

Considering her neighbors' neighbors for the first time at $t = 2$, agent i calculates the time-one actions of all of the agents in \mathcal{N}_i^2 for each of the signal profiles belonging to $\mathcal{I}_{i,1}$ and uses the result to calculate the time two actions of all her neighbors for each $\bar{s} \in \mathcal{I}_{i,1}$. Any \bar{s} for which the calculated time 2 action of some neighbor $j \in \mathcal{N}_i$ does not agree with the observed action $a_{j,2}$ is subsequently removed from $\mathcal{I}_{i,1}$ and the updated list $\mathcal{I}_{i,2}$ is thus obtained. A similar set of calculations is repeated at time three: for every signal profile in $\mathcal{I}_{i,2}$ that have survived the pruning process up until $t = 3$, the agent starts by calculating the actions of agents in \mathcal{N}_i^3 at time zero (as determined by their private signals fixed in

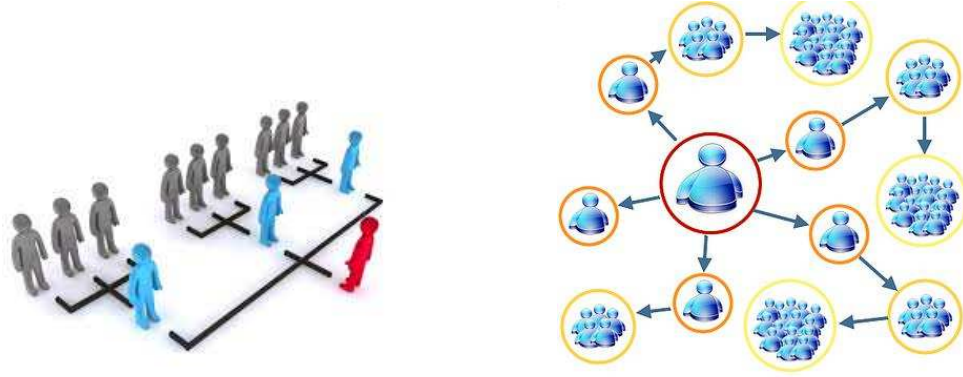


Figure C.1: Learning about the private signals of neighbors, neighbors of neighbors, and so on.

3). These actions along with the rest of the private signal in turn determine the choices of the agents in \mathcal{N}_i^2 at time one as well those of \mathcal{N}_i at time two; subsequently, the agent can compare the latter calculated actions with her most recent observation of the actions of her neighbors and eliminate the signal profiles for which there is a mismatch.

For clarity consider a path of length n , and label them in the order of their location in the path agent n being the leaf node with no neighbors and agent 1 being the agent who has (indirect) access to everybody's decisions (see Fig. 2.3 for the case $n = 4$). We begin by the observation that all agents become fixed in their decisions after $n-1$ steps; in particular, agent n never changes her decisions:

$$\mathbf{a}_{n,0} = \mathbf{a}_{n,1} = \dots = \mathbf{a}_{n,t} \leftarrow \arg \max_{\mathbf{a} \in \mathcal{A}_n} \frac{\sum_{\hat{\theta} \in \Theta} u_n(\mathbf{a}, \hat{\theta}) l_n(\mathbf{s}_n | \hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} l_n(\mathbf{s}_n | \tilde{\theta}) \nu(\tilde{\theta})}, \forall t.$$

Agent $n-1$, will get fixed in her decision after two steps (after learning her private signal and the decision of agent n):

$$\mathbf{a}_{n-1,0} \leftarrow \arg \max_{\mathbf{a} \in \mathcal{A}_{n-1}} \frac{\sum_{\hat{\theta} \in \Theta} u_{n-1}(\mathbf{a}, \hat{\theta}) l_{n-1}(\mathbf{s}_{n-1} | \hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} l_{n-1}(\mathbf{s}_{n-1} | \tilde{\theta}) \nu(\tilde{\theta})}, \quad (\text{C.0.1})$$

$$\mathbf{a}_{n-1,1} = \mathbf{a}_{n-1,2} = \dots = \mathbf{a}_{n-1,t} \quad (\text{C.0.2})$$

$$\leftarrow \arg \max_{\mathbf{a} \in \mathcal{A}_{n-1}} \frac{\sum_{\hat{\theta} \in \Theta} u_{n-1}(\mathbf{a}, \hat{\theta}) l_{n-1}(\mathbf{s}_{n-1} | \hat{\theta}) \left(\sum_{s \in \mathcal{S}_{n,1}^{(n-1)}} l_n(\mathbf{s} | \hat{\theta}) \right) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} l_{n-1}(\mathbf{s}_{n-1} | \tilde{\theta}) \left(\sum_{s \in \mathcal{S}_{n,1}^{(n-1)}} l_n(\mathbf{s} | \tilde{\theta}) \right) \nu(\tilde{\theta})}, \forall t,$$

where

$$\mathcal{S}_{n,1}^{(n-1)} = \left\{ s \in \mathcal{S}_n : \mathbf{a}_{n,0} \in \arg \max_{\mathbf{a} \in \mathcal{A}_n} \frac{\sum_{\hat{\theta} \in \Theta} u_n(\mathbf{a}, \hat{\theta}) \ell_n(s|\hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \ell_n(s|\tilde{\theta}) \nu(\tilde{\theta})} \right\},$$

is the information that $n - 1$ infers about the private signal of n from observing her action $\mathbf{a}_{n,0}$, i.e. the list of private signals that $n - 1$ deems feasible for n given her observation of $\mathbf{a}_{n,0}$. Note that in this notation $\mathcal{S}_{n,0}^{(n-1)} = \mathcal{S}_n$, since at time zero agent $n - 1$ knows nothing about agent n and any value of s_n is deemed feasible.

Agent $n - 2$, will become fixed in decisions after three rounds: once she observes the final decision of agent $n - 1$ at time one, i.e. $\mathbf{a}_{n-1,1}$, she updates her decisions to $\mathbf{a}_{n-2,2}$ and then she remains fixed at this decision for the ensuing decision epochs. Her final decision is influenced not only by her knowledge of what private signals agent $n - 1$ might have observed but also by those of agent n . In particular, her decisions $\mathbf{a}_{n-2,0}$ and $\mathbf{a}_{n-2,1}$ are obtained based on her private signal s_{n-2} and the initial action of agent $n - 1$, identical to (C.0.1) and (C.0.2) with $n - 1$ and n replaced for $n - 2$ and $n - 1$. A time $t = 2$, agent $n - 2$ already knows that the private signal of agent $n - 1$ should necessarily belong to $\mathcal{S}_{n-1,1}^{(n-2)}$ given by:

$$\mathcal{S}_{n-1,1}^{(n-2)} = \left\{ s \in \mathcal{S}_{n-1} : \mathbf{a}_{n-1,0} \in \arg \max_{\mathbf{a} \in \mathcal{A}_{n-1}} \frac{\sum_{\hat{\theta} \in \Theta} u_{n-1}(\mathbf{a}, \hat{\theta}) \ell_{n-1}(s|\hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \ell_{n-1}(s|\tilde{\theta}) \nu(\tilde{\theta})} \right\}.$$

To decide about $\mathbf{a}_{n-2,2}$, agent $n - 2$ should reason about the initial signal and action of agent n , and also update her knowledge of the private signal of agent $n - 1$. This can be achieved as follows: for each $\mathbf{a} \in \mathcal{A}_n$, let $\mathcal{S}_n^{\mathbf{a}}$ be the list of initial signals that can cause agent n to take action \mathbf{a} at time zero:

$$\mathcal{S}_n^{\mathbf{a}} = \left\{ s \in \mathcal{S}_n : \mathbf{a} \in \arg \max_{\mathbf{a} \in \mathcal{A}_n} \frac{\sum_{\hat{\theta} \in \Theta} u_n(\mathbf{a}, \hat{\theta}) \ell_n(s|\hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \ell_n(s|\tilde{\theta}) \nu(\tilde{\theta})} \right\}, \forall \mathbf{a} \in \mathcal{A}_n. \quad (\text{C.0.3})$$

To analyze the decisions of agent $n - 2$ at times 2 and onward, we should determine all pairs of signals that she deems feasible for agents n and $n - 1$ combined, given her observations of $\mathbf{a}_{n-1,0}$ and $\mathbf{a}_{n-1,1}$; in particular, as the initial action of agent n directly influence her observation of $\mathbf{a}_{n-1,1}$, agent $n - 2$ should reject any initial action $\mathbf{a} \in \mathcal{A}_n$ for agent n that cannot be supported by at least one private signal $s_{n-1} \in \mathcal{S}_{n-1,1}^{(n-2)}$. In the due process she should also remove any private signal $s_{n-1} \in \mathcal{S}_{n-1,1}^{(n-2)}$ that cannot be combined by any of the feasible initial actions of agent n to justify the decision of agent $n - 1$ to choose $\mathbf{a}_{n-1,1}$. Both these goal are achieved by the construction of $\mathcal{I}_{n-2,2}$ the list of all feasible pairs $(s_n, s_{n-1}) \in \mathcal{S}_n \times \mathcal{S}_{n-1}$ that are not refuted by her observation of the two actions $\mathbf{a}_{n-1,0}$ and $\mathbf{a}_{n-1,1}$. Furthermore, since agent 2 will not obtain any additional

information about private signals of any of the other remaining agents, this list is never again updated; neither is her decisions in the ensuing epochs: $\mathbf{a}_{n-2,2} = \mathbf{a}_{n-2,3} = \dots = \mathbf{a}_{n-2,t}$ for all t . These calculations and the subsequent decision $\mathbf{a}_{n-2,2}$ are summarized below.

Calculations of agent $n - 2$ at time 2 for deciding $\mathbf{a}_{n-2,2} = \mathbf{a}_{n-2,3} = \dots = \mathbf{a}_{n-2,t}$ for all t , in a path of length n :

1. Initialize $\mathcal{I}_{n-2,2} = \emptyset$.
2. For each $\mathbf{a} \in \mathcal{A}_n$ do:
 - Set $\mathcal{S}_n^{\mathbf{a}}$ according to (C.0.3).
 - For each $s \in \mathcal{S}_{n-1,1}^{(n-2)}$, and every $\mathbf{a} \in \mathcal{A}_n$ check if

$$\mathbf{a}_{n-1,1} \in \arg \max_{\mathbf{a} \in \mathcal{A}_{n-1}} \frac{\sum_{\hat{\theta} \in \Theta} \mathbf{u}_{n-1}(\mathbf{a}, \hat{\theta}) \ell_{n-1}(s|\hat{\theta}) \left(\sum_{s \in \mathcal{S}_n^{\mathbf{a}}} \ell_n(s|\hat{\theta}) \right) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \ell_{n-1}(s_{n-1}|\tilde{\theta}) \left(\sum_{s \in \mathcal{S}_{n,1}^{(n-1)}} \ell_n(s|\tilde{\theta}) \right) \nu(\tilde{\theta})},$$

is satisfied, then set $\mathcal{I}_{n-2,2} = \mathcal{I}_{n-2,2} \cup \mathcal{S}_n^{\mathbf{a}} \times \{s\}$.

3. Set $\mathbf{a}_{n-2,2} = \mathbf{a}_{n-2,3} = \dots = \mathbf{a}_{n-2,t}$ as follows:

$$\mathbf{a}_{n-2,2} \stackrel{\text{(C.0.4)}}{\leftarrow} \arg \max_{\mathbf{a} \in \mathcal{A}_{n-2}} \sum_{\hat{\theta} \in \Theta} \mathbf{u}_{n-2}(\mathbf{a}, \hat{\theta}) \frac{\ell_{n-2}(\mathbf{s}_{n-2}|\hat{\theta}) \sum_{(s,s') \in \mathcal{I}_{n-2,2}} \ell_n(s|\hat{\theta}) \ell_{n-1}(s'|\hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \ell_{n-2}(\mathbf{s}_{n-2}|\tilde{\theta}) \sum_{(s,s') \in \mathcal{I}_{n-2,2}} \ell_n(s|\tilde{\theta}) \ell_{n-1}(s'|\tilde{\theta}) \nu(\tilde{\theta})}.$$

Next, to consider choices of agent $n - 3$ notice that her calculations and inferences at times $t = 0$, $t = 1$, and $t = 2$ are identical to those of agent $n - 2$, except for the shift in the indices of the agents whose private signals (and initial actions) are being inferred. However, at time $t = 3$, in order to interpret $\mathbf{a}_{n-2,2}$, which is set according to (C.0.4), agent $n - 3$ must simulate $\mathcal{I}_{n-2,2}$ for all possible pairs and combine them with what she has learned so far about private signals of agent $n - 2$ to reject any triple of signals for the three upstream agents that are not supported by her observation of $\mathbf{a}_{n-2,2}$. Such a need for

simulation of the feasible signal profiles of the neighboring agents contingent on all of their possible “observations”, arise as a result of forward (causal) reasoning on the part of the agents to try and combine the incoming information with what they have already learned: agents isolate all the new information about the private signals of a far away agent that has reached them for the first time and use that to refine what they already know about previously discovered agents. Such a forward implementation is very cumbersome for complex network structures; in other words, forward simulation of other people’s inferences at all their possible observations is extremely inefficient and scales very poorly when applied to general group structures. Hence, in Section 2.1.1 we propose a recursive implementation of these calculations that relies on iterative elimination of infeasible signal profiles for all the agents whose decisions directly or indirectly influence the decisions of agent i . Although the calculations scale exponentially with the network size, they are much more amenable for application to complex structures as they circumvent the need to simulate the inferences of the neighboring agents at their possible observations (which the agent do not always observe directly). This is true with the exception of some very well-connected agents who have indeed direct access to all the observations of their neighbors and can thus analyze their decisions using forward reasoning, we expand on this special case (called POSETs) in Section 2.1.2 and explain how the Bayesian calculations simplify as a result.

We end this Section by expanding on the above realization that in a path of length n , every agent $n-t$ gets fixed in decisions after time t ; and in particular, no agents will change their recommendations after $t \geq n-1$. There is an easy inductive proof upon noting that indeed agent n , who is a leaf node with no access to the recommendations of anybody else, will never change initial action. Moreover, if agent $n-t+1$ fixes her decision at time $t-1$, then agent $n-t$ would necessarily fix her decision at time t as she receives no new information following her observation of $\mathbf{a}_{n-t+1,t-1}$. This finite time convergence property of paths can be extended to more general structures where a “strict” partial order can be imposed on the set of agents, and in such a way that this order respects the neighborhood relationships among the agents. The strictness property restricts our method to structures without loops or bidirectional links, which are widely known as directed acyclic graphs (DAGs), cf. e.g. [96]. Proposition 2.1 in Section 2.1.2 extends our above realization about the bounded convergence time of group decision process over paths to all DAGs.

Proposition C.1 (Bounded convergence time of group decision process over DAGs). *Let \mathcal{G} be a DAG on n nodes with a topological ordering \prec , and let the agents be labeled in accordance with this topological order as follows: $n \prec n-1 \prec \dots \prec 1$. Then every agent $n-t$ gets fixed in her decisions after time t ; and in particular, no agents will change their recommendations after $t \geq n-1$.*

Appendix D

Efficient Rational Choice in Symmetric Binary Environments

Suppose that the agents are in a binary environment with two states: $\theta_1 = 0$ and $\theta_2 = 1$ and uniform priors $\nu(\theta_1) = \nu(\theta_2) = 1/2$. They receive i.i.d. binary initial signals $\mathbf{s}_i \in \{0, 1\}$, $i \in [n]$, such that for some $p > 1/2$ fixed we have $\ell_i(\mathbf{s}_i|\theta) = p$ if $\mathbf{s}_i = \theta$ and $\ell_i(\mathbf{s}_i|\theta) = 1 - p$, otherwise. Since, $p > 1/2$, at time zero all agents will act based on their signals by simply choosing $\mathbf{a}_{i,0} = \mathbf{s}_i$. At time two, agent i gets to learn the private signals of her neighbors from their time zero actions and therefore takes the action that indicates the majority over the signals observed in by her and everybody in her immediate neighborhood. Since the signals are i.i.d. the agent could be indifferent between her actions, thus in the sequel, we assume that the agent sticks with own signal whenever her information makes her indifferent between her actions. This assumption may seem natural and harmless but in fact, it leads to drastically different behaviors in the case of a directed path. Consider, the directed path of length four in Fig. 2.3, on the left. At time one, if any agent observes an action that contradicts own, then she will be indifferent between the two actions. If we assume that agents resolve their indifference by reverting to their private signals then no agent will ever shift her initial action in a directed path and we have $\mathbf{a}_{i,t} = \mathbf{s}_i$ for all t and i . However, a starkly different outcome will emerge if we instead assume that the agents will always shift their actions whenever they are indifferent. In particular, for a directed path we get that at any time the agents will take the optimal action given the initial signals of everybody in their t -radius ego-net (i.e. *perfect information aggregation*): $\mathbf{a}_{i,t} \in \arg \max_{x \in \{0,1\}} \sum_{j \in \tilde{\mathcal{N}}_i^t} \mathbb{1}\{\mathbf{s}_j = x\}$, where we use the indicator function notation: $\mathbb{1}\{\mathcal{P}\}$ is one if \mathcal{P} is true, and it is zero otherwise. At time t , $\tilde{\mathcal{N}}_i^t$ is the set of all agents who directly or indirectly influence the decisions of agent i , and perfect aggregation ensures that the action of agent i at time t coincides with her optimal action if she had given direct access to all the signals of all agents who have influenced her decision directly or indirectly; hence, the name “*perfect aggregation*”. We can verify that perfect aggregation holds true in any directed path by induction; in particular, consider agent one in the left graph of Fig. 2.3: she gets to know about the private signal of agent two at time one, after observing $\mathbf{a}_{2,0} = \mathbf{s}_2$, next at time two, she observes $\mathbf{a}_{2,1} = \mathbf{a}_{3,0} = \mathbf{s}_3$ (note that if agents switch actions when indifferent then, at time one in a directed path all agents will replicate their neighbor’s time

zero actions); hence, agent one learns the private signal of agent three at time two leading her to take the majority over all three signals $\{s_1, s_2, s_3\}$. By the same token at time three agent one observes $\mathbf{a}_{2,2}$, which is the majority over $\{s_2, s_3, s_4\}$ and having learned s_2 and s_3 from her observation at previous time steps, she can now infer the value of s_4 and thus take the majority over all the private signals $\{s_1, s_2, s_3, s_4\}$, achieving the global optimum at time three. More generally, in any directed path, at time t agent i learns the values of the private signals of the agent at distance t from her. She achieves this by combining two pieces of information: (i) her knowledge at time $t - 1$, which constitutes the values of the signals of all agent at distance $t - 1$ from her, (ii) her observation at time t , which constitutes her neighbor's action at time $t - 1$ and is equal to the majority over all signals within distance t of agent i , excluding herself. Put more succinctly, suppose $i > t$ and let the agents be labeled in accordance with the topological ordering of the directed path, then knowing the values of the all $t - 1$ preceding signals and also the majority over all t preceding signals, agent i can learn the value of the t -th signal at time t .

Therefore, switching actions or staying with own past actions when indifferent, makes the difference between no aggregation at all ($\mathbf{a}_{i,t} = \mathbf{s}_i$ for all t, i) and perfect aggregation in the case of a directed path; indeed, by switching their actions at time one, after observing their neighbor's time zero action (or equivalently private signal) the agent can pass along her information about her neighbor's signal to the person who is observing her. In the case of a directed path, this indirect signaling is enough to ensure perfect aggregation. The exact same argument can be applied to the case of a rooted ordered tree, cf. Fig. D.1 on the left; in such a structure the set of all agents who influence the actions of some agent i always constitute a directed path that starts with the root node and ends at the particular agent i . As such when computing her Bayesian actions in a rooted ordered tree agent i need only consider the unique path that connects her to the root node; thus reducing her computations to those of an agent in a directed path.

In a general structure (beyond rooted trees and symmetric binary environments), perfect aggregation can be defined as:

$$\mathbf{a}_{i,t} \in \arg \max_{\mathbf{a}_i \in \mathcal{A}_i} \frac{\sum_{\hat{\theta} \in \Theta} u_i(\mathbf{a}_i, \hat{\theta}) \prod_{j \in \mathcal{N}_i^t} \ell_j(\mathbf{s}_j | \hat{\theta}) \nu(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \prod_{j \in \mathcal{N}_i^t} \ell_j(\mathbf{s}_j | \tilde{\theta}) \nu(\tilde{\theta})}, \text{ for all } i, t.$$

While many asymptotic characterizations are available for the efficiency of the GDP equilibrium outcomes (cf. Appendix B), deriving tight conditions that ensures perfect aggregation for GDP in general structures is a significant open problem. Our focus in this paper is on the computations of the Bayesian agent; hence we address the efficiency of information aggregation only to the extent that it relates to the computations of the Bayesian agent. In particular, when investigating the complexity of Bayesian belief exchange in Section 2.2, we introduce and study a graph property, called ‘‘transparency’’, that ensures perfect

aggregation for beliefs.

We end our discussion of the symmetric binary environment by considering an oriented tree of depth three and focusing on the actions of the root node, call her agent i (cf. Fig. D.1, on the right). At time zero all agents report their initial signals as their actions $\mathbf{a}_{i,0} = \mathbf{s}_i$; having learned her neighbors' private signals, at time one each agent takes a majority over all the signals in her immediate neighborhood (including own signal). Indeed, this is true for any graph structure in a symmetric binary environment that $\mathbf{a}_{i,0} = \mathbf{s}_i$ and $\mathbf{a}_{i,1} \in \arg \max_{x \in \{0,1\}} \sum_{j \in \tilde{\mathcal{N}}_i^1} \mathbb{1}\{\mathbf{s}_j = x\}$ for all i . At time two, agent i is informed about the time-one actions of her neighbors which gives her the majority values over each of their respective local neighborhoods $\tilde{\mathcal{N}}_j^1, j \in \mathcal{N}_i$. In a (singly connected) tree structure these neighborhoods are non-intersecting; hence, agent i can form a refined belief ratio at time two, by summing over all (mutually exclusive) signal profiles that lead to each of the observed majority values in each local neighborhood $\tilde{\mathcal{N}}_j^1, j \in \mathcal{N}_i$ and then form their product, using the fact that signals are generated independently across the non-intersecting neighborhoods:

$$\frac{\mu_{i,2}(0)}{\mu_{i,2}(1)} = \frac{p^{1-s_i}(1-p)^{s_i} \prod_{j \in \mathcal{N}_i} p^{1-a_{j,0}}(1-p)^{a_{j,0}} f_p^{a_{j,1}}(\lfloor d_j/2 \rfloor + 1, d_j)^{|a_{j,1}-a_{j,0}|} f_p^{a_{j,1}}(\lceil d_j/2 \rceil, d_j)^{1-|a_{j,1}-a_{j,0}|}}{(1-p)^{1-s_i} p^{s_i} \prod_{j \in \mathcal{N}_i} (1-p)^{1-a_{j,0}} p^{a_{j,0}} f_{1-p}^{a_{j,1}}(\lfloor d_j/2 \rfloor + 1, d_j)^{|a_{j,1}-a_{j,0}|} f_{1-p}^{a_{j,1}}(\lceil d_j/2 \rceil, d_j)^{1-|a_{j,1}-a_{j,0}|}},$$

where we use $d_j := \text{card}(\mathcal{N}_j)$ and for non-negative integers x, y and $0 < p < 1$ we define:

$$f_p^a(x, y) = \sum_{\eta_j=x}^y \binom{y}{\eta_j} p^{\eta_j(1-a)+(d_j-\eta_j)a} (1-p)^{\eta_j a+(d_j-\eta_j)(1-a)}, \quad (\text{D.0.1})$$

where $\lfloor d_j/2 \rfloor$ and $\lceil d_j/2 \rceil$ are respectively, the greatest integer less than or equal to $d_j/2$, and the least integer greater than or equal to $d_j/2$. Note that the summations in (D.0.1) are over the set of signal profiles that agent i deems feasible for each of the disjoint neighborhoods $\mathcal{N}_j, j \in \mathcal{N}_i$. Computation of these summations and their use in the belief ratio $\mu_{i,2}(0)/\mu_{i,2}(1)$ are simplified by fixing the majority population η_j in each neighborhood \mathcal{N}_j : $\lfloor d_j/2 \rfloor + 1 \leq \eta_j \leq d_j$ if $\mathbf{a}_{j,1} \neq \mathbf{a}_{j,0}$ and $\lceil d_j/2 \rceil \leq \eta_j \leq d_j$ if $\mathbf{a}_{j,1} = \mathbf{a}_{j,0}$; then using the binomial coefficients to count the number of choices to form the fixed majority population η_j out of the total neighborhood size $d_j = \text{card}(\mathcal{N}_j)$. Given $\mu_{i,2}(0)/\mu_{i,2}(1)$, agent i can take actions as follows: $\mathbf{a}_{i,2} = 1$ if $\mu_{i,2}(0)/\mu_{i,2}(1) < 1$, $\mathbf{a}_{i,2} = 0$ if $\mu_{i,2}(0)/\mu_{i,2}(1) > 1$, and $\mathbf{a}_{i,2} = 1 - \mathbf{a}_{i,1}$ if $\mu_{i,2}(0)/\mu_{i,2}(1) = 1$.




Figure D.1: On the left, the Bayesian computations of agent i in a rooted ordered tree reduces to those in the unique path connecting her to the root (the leftmost node); On the right, an oriented (singly connected) tree with depth three.

Appendix E

Minimum Variance Unbiased Estimation and Online Learning

There is a large body of literature on decentralized detection with the notable examples of [120, 121, 219]; and recently there is a renewed interest in this topic due to its applications to sensor and robotic networks [220–224] and the emergence of new literature considering network of sensor and computational units [135, 141, 225]. Other relevant results investigate the formation and evolution of beliefs in social networks and subsequent shaping of the individual and mass behavior through social learning [136, 226, 227]. Obtaining a global consensus by combining noisy and unreliable locally sensed data is a key step in many wireless sensor network applications; subsequently, many sensor fusion schemes offer reasonable recipes to address this requirement [124, 125]. In many such applications, each sensor forms an estimate of the field using its local measurements and then the sensors initiate distributed optimization to fuse their local estimates. If all the data from every sensor in the network can be collected in a fusion center, then a jointly optimal decision is readily available by solving the global optimization problem given all the data. However, many practical considerations limit the applicability of such a centralized solution. This gives rise to the distributed sensing problems that include distributed network consensus or agreement [12, 13, 120], and distributed averaging [228]; with close relations to the consensus and coordination problems that are studied in the distributed control theory [126, 127, 229].

In this appendix, we allow the parameter space Θ to be any measurable set, and in particular not necessarily finite. Consider again the network of n agents and suppose that each agent $i \in [n]$ observes an i.i.d. samples \mathbf{s}_i from a common distribution $\ell(\cdot|\theta)$ over a measurable sample space \mathcal{S} . We consider an undirected network graph and let the symmetric network graph structure be encoded by its modified adjacency matrix $\mathbf{A} = [\mathbf{a}_{ij}]_{i,j=1}^n$, defined according to the Metropolis-Hastings weights [230]: $\mathbf{a}_{ij} = 1/\max\{d_i, d_j\}$ if $(j, i) \in \mathcal{E}$, and $[\mathbf{A}]_{ij} = 0$ otherwise for $i \neq j$; furthermore, $\mathbf{a}_{ii} = 1 - \sum_{j \neq i} \mathbf{a}_{ij}$.

We assume that $\ell(\cdot|\theta)$ belongs to a one-parameter exponential family so that it admits a probability density or mass function that can be expressed as

$$\ell(\mathbf{s}|\theta) = \tau(\mathbf{s})e^{\alpha(\theta)^T \xi(\mathbf{s}) - \beta(\alpha(\theta))}, \quad (\text{E.0.1})$$

where $\xi(s) \in \mathbb{R}$ is a measurable function acting as a complete sufficient statistic for the i.i.d. random samples \mathbf{s}_i , and $\alpha : \Theta \rightarrow \mathbb{R}$ is a mapping from the parameter space Θ to the real line \mathbb{R} , $\tau(s) > 0$ is a positive weighting function, and

$$\beta(\alpha) := \ln \int_{s \in \mathcal{S}} \tau(s) e^{\alpha \xi(s)} ds, \quad (\text{E.0.2})$$

is a normalization factor known as the log-partition function. In (E.0.1), $\xi(\cdot)$ is a complete sufficient statistic for θ . It is further true that $\sum_{i=1}^n \xi(\mathbf{s}_i)$ is a complete sufficient statistic for the n i.i.d. signals that the agents have received [213, Section 1.6.1]. The agents aim to estimate the expected value of $\xi(\cdot)$: $m_\theta = \mathbb{E}\{\xi(\mathbf{s}_i)\}$, with as little variance as possible. The Lehmann-Scheffé theory (cf. [231, Theorem 7.5.1]) implies that any function of the complete sufficient statistic that is unbiased for m_θ is the almost surely unique minimum variance unbiased estimator of m_θ . In particular, the minimum variance unbiased estimator of m_θ given the initial data sets of all nodes in the network is given by: $\mathbf{m}_n = (1/n) \sum_{i=1}^n \xi(\mathbf{s}_i)$. The agents can compute this value using any average consensus algorithm [232]; guaranteeing convergence to average of the initial values asymptotically.

The agents initialize with: $\boldsymbol{\mu}_{i,0} = \xi(\mathbf{s}_i)$, and in any future time period the agents communicate their values and update them according to the following rule:

$$\boldsymbol{\mu}_{i,t} = \alpha_{ii} \boldsymbol{\mu}_{i,t-1} + \sum_{j \in \mathcal{N}_i} \alpha_{ij} \boldsymbol{\mu}_{j,t-1}. \quad (\text{III})$$

The mechanisms for convergence in this case rely on the product of stochastic matrices, similar to mixing of Markov chains (cf. [184, 225]); hence, many available results on mixing rates of Markov chains can be employed to provide finite time guarantees after T iteration of the average consensus algorithm for fixed T . Such results often rely on the eigenstructure (eigenvalues/eigenvectors) of the communication matrix A , and the facts that it is a primitive matrix and its ordered eigenvalues satisfy $-1 < \lambda_n(A) \leq \lambda_{n-1}(A) \leq \dots \leq \lambda_1(A) = 1$, as a consequence of the Perron-Frobenius theory [115, Theorems 1.5 and 1.7].

Theorem E.1 (Minimum Variance Unbiased Estimation). *Under (III), $\lim_{t \rightarrow \infty} \boldsymbol{\mu}_{i,t} = \mathbf{m}_n$ almost surely, for all i . Furthermore, $|\boldsymbol{\mu}_{i,t} - \mathbf{m}_n| \leq \epsilon$, whenever*

$$t > \left(\log(\epsilon) - \log(\mathbf{M}_n \sqrt{n-1}) \right) / \log \beta^*,$$

where $\mathbf{M}_n = \max_{i \in [n]} |\xi(\mathbf{s}_i)|$ and $\beta^* = \max\{\lambda_2(A), |\lambda_n(A)|\}$.

Proof. Define the concatenated variables $\bar{\boldsymbol{\mu}}_t = (\boldsymbol{\mu}_{1,t}, \dots, \boldsymbol{\mu}_{n,t})^\top$, $\bar{\boldsymbol{\lambda}} = (\xi(\mathbf{s}_1), \dots, \xi(\mathbf{s}_n))^\top$ and note that $\mathbf{m}_n = (1/n) \sum_{i=1}^n \xi(\mathbf{s}_i) = (1/n) \mathbf{1} \mathbf{1}^\top \bar{\boldsymbol{\lambda}}$. Initialized by $\bar{\boldsymbol{\mu}}_0 = \bar{\boldsymbol{\lambda}}$, the evolution of beliefs under (III) can be written in the following vectorized form: $\bar{\boldsymbol{\mu}}_t = \mathbf{A} \bar{\boldsymbol{\mu}}_{t-1} = \mathbf{A}^t \bar{\boldsymbol{\lambda}}$, and as in Appendix A.9 for a connected network \mathcal{G} , we have that $\lim_{t \rightarrow \infty} \mathbf{A}^t = (1/n) \mathbf{1} \mathbf{1}^\top$; and subsequently, $\lim_{t \rightarrow \infty} \bar{\boldsymbol{\mu}}_t = (1/n) \mathbf{1} \mathbf{1}^\top \bar{\boldsymbol{\lambda}} = \mathbf{1} \mathbf{m}_n$. Hence, the claim about the almost sure limits of every agents' beliefs is verified. To investigate the rate of convergence of $\bar{\boldsymbol{\mu}}_t$ to $\mathbf{1} \mathbf{m}_n$ we can write: $\bar{\boldsymbol{\mu}}_t - \mathbf{1} \mathbf{m}_n = (\mathbf{A}^t - (1/n) \mathbf{1} \mathbf{1}^\top) \bar{\boldsymbol{\lambda}}$. Hence for each node i we have:

$$\begin{aligned} |\boldsymbol{\mu}_{i,t} - \mathbf{m}_n| &= \left| \sum_{j=1}^n \left([\mathbf{A}^t]_{ij} - \frac{1}{n} \right) \xi(\mathbf{s}_j) \right| \\ &\leq \sum_{j=1}^n \left| [\mathbf{A}^t]_{ij} - \frac{1}{n} \right| |\xi(\mathbf{s}_j)| \leq \mathbf{M}_n \sum_{j=1}^n \left| [\mathbf{A}^t]_{ij} - \frac{1}{n} \right|. \end{aligned} \quad (\text{E.0.3})$$

We next use the fact that \mathbf{A} can specify the transition probabilities of an aperiodic irreducible Markov chain with uniform stationary distribution. In particular, it is a time-reversible Markov chain and [233, Proposition 3] implies that

$$\sum_{j=1}^n \left| [\mathbf{A}^t]_{ij} - \frac{1}{n} \right| \leq \sqrt{(n-1)} (\beta^*)^t, \quad (\text{E.0.4})$$

where $\beta^* = \max\{\lambda_2(\mathbf{A}), |\lambda_n(\mathbf{A})|\}$ and $0 \leq \beta^* < 1$ as a consequence of Perron-Frobenius theory [115, Theorems 1.5 and 1.7] applied to the primitive matrix \mathbf{A} . Replacing (E.0.4) in (E.0.3) yields that for all i the distance to the limiting values \mathbf{m}_n decrease at least exponential fast and can be bounded as follows: $|\boldsymbol{\mu}_{i,t} - \mathbf{m}_n| \leq \mathbf{M}_n \sqrt{(n-1)} (\beta^*)^t$. The claimed finite time guarantee now follows upon setting $\mathbf{M}_n \sqrt{(n-1)} (\beta^*)^t < \epsilon$ or equivalently: $t > (\log(\epsilon) - \log(\mathbf{M}_n \sqrt{n-1})) / \log \beta^*$. \square

We now take a brief look at the case where the initial data sets are of different sizes n_i , so that each agent has access to a set of n_i initial data points $\mathbf{s}_i^1, \dots, \mathbf{s}_i^{n_i}$, each of which is identically distributed according to the common exponential family distribution $\ell(\cdot|\theta)$. We explain how (III) should be modified to accommodate the varying sample sizes. In this case, the globally efficient (minimum variance) estimator of the mean sufficient statistic \mathbf{m}_θ given all the initial data sets is as follows: $\mathbf{m}_n^* = \left(1 / \sum_{p=1}^n n_p \right) \sum_{i=1}^n \sum_{j=1}^{n_i} \xi(\mathbf{s}_i^j)$. The agents can initialize with: $\boldsymbol{\mu}_{i,0} = (1/n_i) \sum_{j=1}^{n_i} \xi(\mathbf{s}_i^j)$ for all i ; however, to ensure convergence to the right limit the coefficients of the linear update rule in (III) should be modified in accordance with the initial sample sizes. Let $\delta_i = n_i / \sum_{p=1}^n n_p$; the following modification of the Metropolis - Hastings weights explained in [230, 234], incorporates the sample sizes and ensures convergence of the linear iterations in (III) to the right limit \mathbf{m}_n^* :

$$\alpha_{ij} = \begin{cases} \frac{1}{d_i} \min\{1, \frac{\delta_j d_i}{d_j \delta_i}\} & \text{if } (j, i) \in \mathcal{E}, \\ 1 - \sum_{j \neq i} \alpha_{ij} & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{E.0.5})$$

In (E.0.5), we use the notation $d_i = \deg(i) - 1$ for the cardinality of \mathcal{N}_i , $i \in [n]$.

Next, suppose that every time $t \in \mathbb{N}$, each agent $i \in [n]$ receives an i.i.d. sample $\mathbf{s}_{i,t}$, in addition to communicating their current estimates $\boldsymbol{\mu}_{i,t}$. All signals $\{\mathbf{s}_{i,t} : i \in [n], t \in \mathbb{N}\}$ are distributed according to the same distribution $\ell(\cdot|\theta)$, and as before the agents aim to estimate the expected value of the complete sufficient statistic $\xi(\cdot)$ with as little variance as possible. Here we propose a $1/t$ discounting of new samples with increasing time t . This would enable the agents to learn the true value \mathbf{m}_θ asymptotically almost surely; and in such a way that the variance of their estimates decreases as $1/t$: linearly in time. The exact upper bound for $\text{Var}\{\boldsymbol{\mu}_{i,t}\}$ is derived in the proof of Theorem E.2 as follows:

$$\text{Var}\{\boldsymbol{\mu}_{i,t}\} \leq \frac{n(n-1)\mathbb{E}\{\xi(\mathbf{s}_{j,\tau})^2\}}{t(1-\beta^{*2})} + \frac{\text{Var}\{\xi(\mathbf{s}_{1,1})\}}{nt}. \quad (\text{E.0.6})$$

We can further use the properties of the exponential family to express the expectation and variance of the complete sufficient statistic $\xi(\cdot)$ in terms of the first and second derivatives of the log-partition function given in (E.0.2), cf. [213, Theorem 1.6.2]: $\mathbb{E}\{\xi(\mathbf{s}_{j,\tau})\} = \beta'(\alpha(\theta))$ and $\text{Var}\{\xi(\mathbf{s}_{1,1})\} = \beta''(\alpha(\theta))$. Hence, (E.0.6) becomes:

$$\text{Var}\{\boldsymbol{\mu}_{i,t}\} \leq \frac{n(n-1)(\beta'(\alpha(\theta))^2 + \beta''(\alpha(\theta)))}{t(1-\beta^{*2})} + \frac{\beta''(\alpha(\theta))}{nt}, \quad (\text{E.0.7})$$

The preceding upper-bound can be used to provide finite-time guarantees for the quality of the estimate $\boldsymbol{\mu}_{i,T}$ at any node i and after a finite termination time T . These bounds are comprised of two additive terms: the first terms on right-hand sides of (E.0.6) and (E.0.7) capture the rate at which the powers of Metropolis-Hastings weight matrix A approach their limit: $A^t \rightarrow \frac{1}{n}\mathbb{1}\mathbb{1}^\top$ as $t \rightarrow \infty$; the second term captures the diminishing variance of the estimates with the increasing number of samples, as gathered by all the agents in the network. The latter is a simple consequence of the Chebyshev inequality applied to the entire set of nt samples that are gathered by all the n agents up to time t . On the other hand, the first term on the right-hand side of the bounds is governed by the mixing rate of the Metropolis-Hastings weights; in particular, it is influenced by the structure of the network through β^* : the second largest magnitude of the eigenvalues of matrix A . The same structural effect appears through β^* in the bound claimed in Theorem E.1. A similar effect can be observed through α^* from the expression of the finite-time \mathbf{T} in the proof of Theorem G.1, (Appendix A.9).

Initializing $\mu_{i,0}$ arbitrarily, in any future time period $t \geq 1$ the agents observe a signal $\mathbf{s}_{i,t}$, communicate their current values $\mu_{i,t-1}$, and update their beliefs to $\mu_{i,t}$, according to the following rule:

$$\mu_{i,t} = \frac{t-1}{t} \left(\alpha_{ii} \mu_{i,t-1} + \sum_{j \in \mathcal{N}_i} \alpha_{ij} \mu_{j,t-1} \right) + \frac{1}{t} \xi(\mathbf{s}_{i,t}). \quad (\text{IV})$$

Theorem E.2 (Online Learning of Expected Values). *Under (IV), $\lim_{t \rightarrow \infty} \mu_{i,t} = m_\theta$ almost surely, for all i . Furthermore, $\text{Var}\{\mu_{i,t}\} = O(1/t)$ and $\mathbb{E}\{\mu_{i,t}\} = m_\theta$ for all t .*

Proof. Let $\bar{\mu}_t$ be as in the proof of Theorem E.1, $\bar{\lambda}_t = (\xi(\mathbf{s}_{1,t}), \dots, \xi(\mathbf{s}_{n,t}))^\top$, and $A_t = \frac{t-1}{t} A$. Under (IV) the beliefs evolve as follows:

$$\begin{aligned} \bar{\mu}_t &= A_t \bar{\mu}_{t-1} + \frac{1}{t} \bar{\lambda}_t = \frac{1}{t} \bar{\lambda}_t + \sum_{\tau=1}^{t-1} \left(\prod_{u=\tau+1}^t A_u \right) \frac{1}{\tau} \bar{\lambda}_\tau \\ &= \frac{1}{t} \bar{\lambda}_t + \sum_{\tau=1}^{t-1} \left(\frac{t-1}{t} \times \frac{t-2}{t-1} \times \dots \times \frac{\tau}{\tau+1} A^{t-\tau} \right) \frac{1}{\tau} \bar{\lambda}_\tau \\ &= \frac{1}{t} \sum_{\tau=1}^t A^{t-\tau} \bar{\lambda}_\tau. \end{aligned} \quad (\text{E.0.8})$$

As in proof of Theorem E.1, we have that $\lim_{\tau \rightarrow \infty} A^\tau = (1/n) \mathbb{1} \mathbb{1}^\top$, and we can invoke the Cesàro mean together with the strong law to conclude that

$$\lim_{t \rightarrow \infty} \bar{\mu}_t = \mathbb{1} \left(\lim_{t \rightarrow \infty} \frac{1}{nt} \sum_{\tau=1}^t \sum_{i=1}^n \xi(\mathbf{s}_{i,\tau}) \right) = \mathbb{1} \mathbb{E}\{\xi(\mathbf{s}_{i,1})\},$$

so that $\mu_{i,t} \rightarrow m_\theta$ with probability one for all agents $i \in [n]$; in particular, $\mu_{i,t}$ for each i is a strongly consistent estimator of $\bar{\theta}$. We can further bound the rate of decrease in $\text{Var}(\mu_{i,t})$ as t increases. Taking expectation of both sides in (E.0.8) yields that $\mathbb{E}\{\bar{\mu}_t\} = (1/t) \sum_{\tau=1}^t A^{t-\tau} \mathbb{1} m_\theta = \mathbb{1} m_\theta$. Hence, we can subtract $\mathbb{1} m_\theta$ from both sides of (E.0.8) and bound the variance of $\bar{\mu}_t$ in terms of the variance of i.i.d. random variable $\bar{\lambda}_t$ and rate of convergence (mixing) for $A^t \rightarrow (1/n) \mathbb{1} \mathbb{1}^\top$. Indeed, using (E.0.8) we can write

$$|\mu_{i,t} - m_\theta| = \left| \frac{1}{t} \sum_{\tau=1}^t \sum_{j=1}^n [A^{t-\tau}]_{ij} \xi(\mathbf{s}_{j,\tau}) - m_\theta \right|.$$

Next by adding and subtracting $\frac{1}{nt} \sum_{\tau=1}^t \sum_{j=1}^n \xi(\mathbf{s}_{j,\tau})$, which is the average of all signals

across all times and agents; and then applying the triangle inequality we obtain:

$$|\boldsymbol{\mu}_{i,t} - m_\theta| \leq \frac{1}{t} \left| \sum_{\tau=1}^t \sum_{j=1}^n \left([A^{t-\tau}]_{ij} - \frac{1}{n} \right) \xi(\mathbf{s}_{j,\tau}) \right| + \frac{1}{nt} \left| \sum_{\tau=1}^t \sum_{j=1}^n (\xi(\mathbf{s}_{j,\tau}) - m_\theta) \right|.$$

Taking squares of both sides and using the Cauchy-Schwartz inequality yields that

$$\begin{aligned} (\boldsymbol{\mu}_{i,t} - m_\theta)^2 &\leq \tag{E.0.9} \\ \frac{1}{t^2} &\left\{ \sum_{\tau=1}^t \sum_{j=1}^n \left([A^{t-\tau}]_{ij} - \frac{1}{n} \right)^2 \right\} \left\{ \sum_{\tau=1}^t \sum_{j=1}^n \xi(\mathbf{s}_{j,\tau})^2 \right\} \\ &+ \frac{1}{n^2 t^2} \left(\sum_{\tau=1}^t \sum_{j=1}^n (\xi(\mathbf{s}_{j,\tau}) - m_\theta) \right)^2. \end{aligned}$$

We next apply the Markov chain mixing time inequality (E.0.4) from the proof of Theorem E.1 to bound

$$\begin{aligned} \sum_{\tau=1}^t \sum_{j=1}^n \left([A^{t-\tau}]_{ij} - \frac{1}{n} \right)^2 &\leq \sum_{\tau=1}^t \left(\sum_{j=1}^n \left| [A^{t-\tau}]_{ij} - \frac{1}{n} \right| \right)^2 \\ &\leq \sum_{\tau=1}^t (n-1) (\beta^*)^{2(t-\tau)} \leq \frac{n-1}{1-\beta^{*2}}, \end{aligned} \tag{E.0.10}$$

where $\beta^* = \max\{\lambda_2(A), |\lambda_n(A)|\}$ and $0 \leq \beta^* < 1$. Furthermore, since $\{\xi(\mathbf{s}_{j,\tau}), j \in [n], t \in \mathbb{N}\}$ form a sequence of i.i.d. random variables with mean m_θ , we have:

$$\mathbb{E} \left\{ \left(\sum_{\tau=1}^t \sum_{j=1}^n (\xi(\mathbf{s}_{j,\tau}) - m_\theta) \right)^2 \right\} = nt \text{Var}\{\xi(\mathbf{s}_{j,\tau})\}.$$

We can now bound $\text{Var}\{\boldsymbol{\mu}_{i,t}\}$ by taking expectations of both sides in (E.0.9) and using (E.0.10) to get (E.0.6) and subsequently (E.0.7); whence, $\text{Var}\{\boldsymbol{\mu}_{i,t}\} = O(1/t)$ as claimed. \square

Unlike the log-linear update rules which could be easily modified to accommodate intermittent data streams with varying sizes (compare (I) and (II)), the linear update rules are not amenable to heterogeneities in the network. It is due the requirement to discount the new observations with increasing time and the need to adapt the linearity coefficients to the varying sample sizes (see (E.0.5)). These factors make the linear update rules unnameable to the case of intermittent observations (compare (III) and (IV)).

Appendix F

Generalization to Belief Updates with One-Step Recall

Consider the log-ratio statistics defined in (2.2.1) 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:

$$\begin{aligned}\bar{\Phi}_t(\hat{\theta}, \check{\theta}) &:= (\Phi_{1,t}(\hat{\theta}, \check{\theta}), \dots, \Phi_{n,t}(\hat{\theta}, \check{\theta}))^\top, \\ \bar{\Phi}_{\Delta,t} &:= (\hat{\Phi}_{1,t}(\hat{\theta}, \check{\theta}), \dots, \hat{\Phi}_{n,t}(\hat{\theta}, \check{\theta}))^\top, \\ \bar{\lambda}(\hat{\theta}, \check{\theta}) &:= (\lambda_1(\hat{\theta}, \check{\theta}), \dots, \lambda_n(\hat{\theta}, \check{\theta}))^\top, \\ \bar{\gamma}(\hat{\theta}, \check{\theta}) &:= (\gamma_1(\hat{\theta}, \check{\theta}), \dots, \gamma_n(\hat{\theta}, \check{\theta}))^\top.\end{aligned}$$

Under (3.3.2) the log-belief ratios evolve as follows:

$$\Phi_{i,t}(\hat{\theta}, \check{\theta}) = \Phi_{i,t-1}(\hat{\theta}, \check{\theta}) + \hat{\Phi}_{i,t}(\hat{\theta}, \check{\theta}), \text{ where } \hat{\Phi}_{i,t}(\hat{\theta}, \check{\theta}) := \sum_{j \in \mathcal{N}_i} \hat{\Phi}_{j,t-1}(\hat{\theta}, \check{\theta}), \quad (\text{F.0.1})$$

initialized by: $\Phi_{i,0}(\hat{\theta}, \check{\theta}) = \gamma_i(\hat{\theta}, \check{\theta}) + \lambda_i(\hat{\theta}, \check{\theta})$ and $\hat{\Phi}_{i,0}(\hat{\theta}, \check{\theta}) = \lambda_i(\hat{\theta}, \check{\theta})$, for all i . In [29, Proposition 4.3 and its following paragraph] we point out that (F.0.1) 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 $\hat{\Phi}_{i,t}(\hat{\theta}, \check{\theta})$ terms as innovations. Following (F.0.1), 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 (3.3.2) (or equivalently (F.0.1)) can be analyzed as follows:

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

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 τ (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) [199, 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:

$$\Phi_{i,t}(\hat{\theta}, \check{\theta}) = \gamma_i(\hat{\theta}, \check{\theta}) + \sum_{j=1}^n C_t(i, j)\lambda_j(\hat{\theta}, \check{\theta}). \quad (\text{F.0.2})$$

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 (F.0.2), 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:

$$\mu_{i,t}(\theta') = \frac{\nu_i(\theta') \prod_{j=1}^n \ell_j(\mathbf{s}_j|\theta')^{C_n(i,j)}}{\sum_{\theta'' \in \Theta} \nu_i(\theta'') \prod_{j=1}^n \ell_j(\mathbf{s}_j|\theta'')^{C_n(i,j)}}, \forall t \geq n. \quad (\text{F.0.3})$$

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:

$$\mu_{i,t}(\theta') = \frac{\nu_i(\theta') \prod_{j=1}^n \ell_j(\mathbf{s}_j|\theta')}{\sum_{\theta'' \in \Theta} \nu_i(\theta'') \prod_{j=1}^n \ell_j(\mathbf{s}_j|\theta'')}, \forall t \geq n. \quad (\text{F.0.4})$$

The final belief (F.0.4) 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 (F.0.3). 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 3.3.2. In Subsection 3.3.3, 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 G

Distributed Estimation and Learning from Intermittent Data

In sensor networks, due to the diverse sensing capabilities and other unpredictable physical factors, usually the quality and availability of local observations varies among the different sensors and over time. A main focus of this Appendix is to demonstrate how log-linear aggregation schemes can be modified to accommodate the heterogeneity of the sensed data both over time and across different sensors. Suppose that the set of n agents aim to collectively distinguish the true state θ from a set of finitely many possibilities Θ . Each agent $i \in [n]$ has access to a set of n_i initial data points $\mathbf{s}_i^1, \dots, \mathbf{s}_i^{n_i}$, each of which is identically distributed according to a common distribution $\ell_i(\cdot|\theta)$. In this section we give a procedure so that by forming a belief over the set Θ and iteratively updating these beliefs, the agents can determine the maximum likelihood estimator of θ given all the initial data sets: $\{\mathbf{s}_i^1, \dots, \mathbf{s}_i^{n_i}\}$, $i \in [n]$. Similar to Appendix E, suppose that the symmetric network graph structure is encoded by the weighted adjacency matrix $\mathbf{A} = [a_{ij}]_{i,j=1}^n$, and the weights are set according to the Metropolis-Hastings [230]: $a_{ij} = 1/\max\{d_i, d_j\}$ if $(j, i) \in \mathcal{E}$, and $[\mathbf{A}]_{ij} = 0$ otherwise for $i \neq j$; furthermore, $a_{ii} = 1 - \sum_{j \neq i} a_{ij}$.

The agents begin by forming: $\gamma_i(\tilde{\theta}) = \prod_{j=1}^{n_i} \ell_i(\mathbf{s}_i^j|\tilde{\theta})$, and initializing their beliefs to $\mu_{i,0}(\hat{\theta}) = \gamma_i(\hat{\theta}) / \sum_{\tilde{\theta} \in \Theta} \gamma_i(\tilde{\theta})$. In any future time period the agents update their belief after communication with their neighboring agents, and according to the following update rule for any $\hat{\theta}$:

$$\mu_{i,t}(\hat{\theta}) = \frac{\mu_{i,t-1}^{1+a_{ii}}(\hat{\theta}) \prod_{j \in \mathcal{N}_i} \mu_{j,t-1}^{a_{ij}}(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \mu_{i,t-1}^{1+a_{ii}}(\tilde{\theta}) \prod_{j \in \mathcal{N}_i} \mu_{j,t-1}^{a_{ij}}(\tilde{\theta})}. \quad (\text{I})$$

For any $\hat{\theta} \in \Theta$ we can define $\mathcal{L}(\hat{\theta}) = \sum_{i=1}^n \sum_{j=1}^{n_i} \log(\ell_i(\mathbf{s}_i^j|\hat{\theta}))$, then the global maximum likelihood estimate of θ given all the initial data points is any member of the set $\Theta^* := \arg \max_{\hat{\theta} \in \Theta} \mathcal{L}(\hat{\theta})$.

Theorem G.1 (Maximum Likelihood Estimation). *Under (I), $\lim_{t \rightarrow \infty} \mu_{i,t}(\tilde{\theta}) = 1/|\Theta^*|, \forall \tilde{\theta} \in \Theta^*$ and $\lim_{t \rightarrow \infty} \mu_{i,t}(\tilde{\theta}) = 0, \forall \tilde{\theta} \notin \Theta^*$ almost surely, for all $i \in [n]$. In particular, if $\Theta^* = \{\theta^*\}$ is a singleton, then $\lim_{t \rightarrow \infty} \mu_{i,t}(\theta^*) = 1$, almost surely for all i . Hence, after a large enough number of iterations any agent $i \in [n]$ can recover θ^* as $\theta^* = \arg \max_{\tilde{\theta} \in \Theta} \mu_{i,t}(\tilde{\theta})$.*

Proof. Consider the following log-belief ratio statistics:

$$\Phi_{i,t}(\hat{\theta}, \check{\theta}) := \log(\mu_{i,t}(\hat{\theta})/\mu_{i,t}(\check{\theta})), \lambda_i(\hat{\theta}, \check{\theta}) := \log(\gamma_i(\hat{\theta})/\gamma_i(\check{\theta})), \quad (\text{G.0.1})$$

and their vectorization

$$\bar{\Phi}_t(\hat{\theta}, \check{\theta}) := (\Phi_{1,t}(\hat{\theta}, \check{\theta}), \dots, \Phi_{n,t}(\hat{\theta}, \check{\theta})), \bar{\lambda}(\hat{\theta}, \check{\theta}) := (\lambda_1(\hat{\theta}, \check{\theta}), \dots, \lambda_n(\hat{\theta}, \check{\theta})). \quad (\text{G.0.2})$$

By forming the belief ratio $\mu_{i,t}(\hat{\theta})/\mu_{i,t}(\check{\theta})$, taking the logarithms of both sides, and using the vectorization in (G.0.2), we can rewrite the belief updates in (I) as a linear updated in terms of log ratios:

$$\bar{\Phi}_{t+1}(\hat{\theta}, \check{\theta}) = (I + A)\bar{\Phi}_t(\hat{\theta}, \check{\theta}) = (I + A)^{t+1}\bar{\Phi}_0(\hat{\theta}, \check{\theta}) = (I + A)^{t+1}\bar{\lambda}(\hat{\theta}, \check{\theta}). \quad (\text{G.0.3})$$

When the network graph \mathcal{G} is connected, the matrix $I + A$ is primitive. The Perron-Frobenius theory [115, Theorems 1.5 and 1.7] implies that $I + A$ has a simple positive real eigenvalue equal to its spectral radius $\rho(I + A) = 2$. Moreover, the left and right eigenspaces associated with this eigenvalue are both one-dimensional and the corresponding eigenvectors can be taken to be both equal to $(1/\sqrt{n})\mathbf{1}$. The magnitude of any other eigenvalue of $I + A$ is strictly less than 2. Hence, the eigenvalues of $I + A$ denoted by $\alpha_i := \lambda_i(I + A)$, $i \in [n]$, which are all real, can be ordered as follows: $-2 < \lambda_n(I + A) \leq \lambda_{n-1}(I + A) \leq \dots \leq \lambda_1(I + A) = 2$. Susequently, we can employ the eigendecomposition of $(I + A)$ to analyze the behavior of $(I + A)^{t+1}$ in (G.0.3). 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: $\|\bar{l}_i\|_2 = \|\bar{r}_i\|_2 = 1, \bar{l}_i^\top \bar{r}_i = 1$ for all i and $\bar{l}_i^\top \bar{r}_j = 0, i \neq j$; in particular, $\bar{l}_1 = \bar{r}_1 = (1/\sqrt{n})\mathbf{1}$. Moreover, we have that [111]:

$$(I + A)^t = 2^t \left(\frac{1}{n} \mathbf{1} \mathbf{1}^\top + \sum_{i=2}^n (\alpha_i/2)^t \bar{r}_i \bar{l}_i^\top \right). \quad (\text{G.0.4})$$

To proceed denote $\Lambda(\hat{\theta}, \check{\theta}) := \Lambda(\hat{\theta}) - \Lambda(\check{\theta})$ and note that $\Lambda(\hat{\theta}, \check{\theta}) = \mathbf{1}^\top \bar{\lambda}(\hat{\theta}, \check{\theta})$. We can use (G.0.4) and (G.0.3), together with the fact that $|\alpha_i| < 2$ for all $i > 1$, established above using the Perron-Frobenius theory, to conclude that $\bar{\Phi}_t(\hat{\theta}, \check{\theta}) \rightarrow (2^t/n)\mathbf{1}\Lambda(\hat{\theta}, \check{\theta})$ almost surely. Moreover, since Θ^* consists of the set of all maximizers of $\Lambda(\tilde{\theta})$, we have that $\Lambda(\hat{\theta}, \check{\theta}) < 0$ whenever $\check{\theta} \in \Theta^*$ and $\hat{\theta} \notin \Theta^*$. Hence, for all $\check{\theta} \in \Theta^*$ and any $\hat{\theta}$,

$\phi_{i,t}(\hat{\theta}, \tilde{\theta}) \rightarrow -\infty$ if $\hat{\theta} \notin \Theta^*$ and $\phi_{i,t}(\hat{\theta}, \tilde{\theta}) = 0$ whenever $\hat{\theta} \in \Theta^*$; or equivalently, $\mu_{i,t}(\hat{\theta})/\mu_{i,t}(\tilde{\theta}) \rightarrow 0$ for all $\hat{\theta} \notin \Theta^*$, while $\mu_{i,t}(\hat{\theta}) = \mu_{i,t}(\tilde{\theta})$ for any $\hat{\theta} \in \Theta^*$. The latter together with the fact that $\sum_{\theta \in \Theta} \mu_{i,t}(\theta) = 1$ for all t implies that with probability one: $\lim_{t \rightarrow \infty} \mu_{i,t}(\tilde{\theta}) = 1/|\Theta^*|, \forall \tilde{\theta} \in \Theta^*$ and $\lim_{t \rightarrow \infty} \mu_{i,t}(\tilde{\theta}) = 0, \forall \tilde{\theta} \notin \Theta^*$ as claimed.

Furthermore, we can use (G.0.4) and (G.0.3) to bound the distance between $\phi_{i,t}(\hat{\theta}, \check{\theta})$ and $(2^t/n)\mathcal{L}(\hat{\theta}, \check{\theta})$ for any i , as follows:

$$\begin{aligned} \left| \phi_{i,t}(\hat{\theta}, \check{\theta}) - \frac{2^t}{n} \mathcal{L}(\hat{\theta}, \check{\theta}) \right| &\leq \left\| \overline{\phi}_t(\hat{\theta}, \check{\theta}) - \frac{2^t}{n} \mathcal{L}(\hat{\theta}, \check{\theta}) \mathbf{1} \right\|_2 \\ &= \left\| \sum_{i=2}^n \left(\frac{\alpha_i}{2} \right)^t \bar{l}_i \bar{r}_i^T \bar{\lambda}(\hat{\theta}, \check{\theta}) \right\|_2 \leq \sum_{i=2}^n \left| \frac{\alpha_i}{2} \right|^t |\bar{r}_i^T \bar{\lambda}(\hat{\theta}, \check{\theta})| \|\bar{l}_i\|_2 \\ &\leq \sum_{i=2}^n \left| \frac{\alpha_i}{2} \right|^t \|\bar{\lambda}(\hat{\theta}, \check{\theta})\|_2 \|\bar{r}_i\|_2 \|\bar{l}_i\|_2. \end{aligned} \quad (\text{G.0.5})$$

Orthonormality of the eigenvectors yields that $\|\bar{r}_i\|_2 = \|\bar{l}_i\|_2 = 1$; also by monotonicity of the ℓ_p norm we get that

$$\|\bar{\lambda}(\hat{\theta}, \check{\theta})\|_2 \leq \|\bar{\lambda}(\hat{\theta}, \check{\theta})\|_1 = \sum_{i=1}^n |\lambda_i(\hat{\theta}, \check{\theta})| \leq 2n\mathbf{L}_n,$$

where $\mathbf{L}_n = \max_{i \in [n]} \max_{\theta \in \Theta} |\log(\gamma_i(\tilde{\theta}))|$ is the largest absolute log of product likelihoods that is achieved in the initial data sets, so that $|\lambda_i(\hat{\theta}, \check{\theta})| = |\log(\gamma_i(\hat{\theta})) - \log(\gamma_i(\check{\theta}))| < 2\mathbf{L}_n$ for all i . Subsequently, (G.0.5) becomes

$$\left| \phi_{i,t}(\hat{\theta}, \check{\theta}) - \frac{2^t}{n} \mathcal{L}(\hat{\theta}, \check{\theta}) \right| \leq 2\mathbf{L}_n n(n-1)(\alpha^*)^t,$$

where $\alpha^* = (1/2) \max\{\alpha_2, |\alpha_n|\}$. Hence,

$$\log(\mu_{i,t}(\hat{\theta})) \leq \log\left(\frac{\mu_{i,t}(\hat{\theta})}{\mu_{i,t}(\tilde{\theta})}\right) = \phi_{i,t}(\hat{\theta}, \tilde{\theta}) \leq \frac{2^t}{n} \mathcal{L}(\hat{\theta}, \tilde{\theta}) + 2\mathbf{L}_n n(n-1)(\alpha^*)^t,$$

for all i and any $\tilde{\theta} \in \Theta^*$. Next suppose that the maximum likelihood estimator is unique so that $\Theta^* = \{\theta^*\}$ and let $\mathbf{I}_\Theta = \min_{\hat{\theta} \neq \theta^*} |\mathcal{L}(\hat{\theta}, \theta^*)|$. Then for any $\hat{\theta} \neq \theta^*$ and all agents i we can bound the belief on $\hat{\theta}$ as follows:

$$\log(\mu_{i,t}(\hat{\theta})) \leq -\frac{2^t}{n} \mathbf{I}_\Theta + 2\mathbf{L}_n n(n-1)(\alpha^*)^t.$$

Therefore, if we take

$$\mathbf{T} = \max \left\{ 1 + \frac{\log \left(\frac{n \log(n-1)}{I_n} \right)}{\log 2}, \frac{\log \left(\frac{\log(n-1)}{2L_n n(n-1)} \right)}{\log(\alpha^*)} \right\}, \quad (\text{G.0.6})$$

then for all $t > \mathbf{T}$, $\log(\boldsymbol{\mu}_{i,t}(\hat{\theta})) \leq -\log(n-1)$ so that $\boldsymbol{\mu}_{i,t}(\hat{\theta}) < \frac{1}{n-1} < \boldsymbol{\mu}_{i,t}(\theta^*)$ for all $\hat{\theta} \neq \theta^*$ and any $i \in [n]$; whence, any agent $i \in [n]$ can recover θ^* as $\theta^* = \arg \max_{\tilde{\theta} \in \Theta} \boldsymbol{\mu}_{i,t}(\tilde{\theta})$ at all $t > \mathbf{T}$. \square

In the proof of Theorem G.1, we also give a more detailed description of the claimed belief convergence result; in particular, we characterize a finite time \mathbf{T} , such that any agent $i \in [n]$ can recover θ^* by $\theta^* = \arg \max_{\tilde{\theta} \in \Theta} \boldsymbol{\mu}_{i,t}(\tilde{\theta})$ at all $t > \mathbf{T}$, cf. (G.0.6) of the appendix and the explanations therein.

We next consider a network of agents that make streams of observations intermittently and communicate their beliefs at every time period. At any time t , agent i makes $\mathbf{n}_{i,t}$ i.i.d. observations $\mathbf{s}_{i,t}^1, \dots, \mathbf{s}_{i,t}^{\mathbf{n}_{i,t}}$ that are distributed according to $\ell(\cdot|\theta)$; and the numbers of observations at each time period: $\{\mathbf{n}_{i,t}, t \in \mathbb{N}\}$ constitute a sequence of i.i.d. signals with mean $\mathbb{E}\{\mathbf{n}_{i,t}\} = \nu_i$. The agents aim to determine the true state θ from their stream of observations.

Every time $t \in \mathbb{N}_0$, each agent forms the likelihood product of the signals that it has received at that time-period: $\boldsymbol{\gamma}_{i,t}(\tilde{\theta}) = \prod_{j=1}^{\mathbf{n}_{i,t}} \ell_i(\mathbf{s}_{i,t}^j|\tilde{\theta})$, if $\mathbf{n}_{i,t} \geq 1$, and $\boldsymbol{\gamma}_{i,t}(\tilde{\theta}) = 1$ if $\mathbf{n}_{i,t} = 0$. It then updates its belief according to:

$$\boldsymbol{\mu}_{i,t}(\hat{\theta}) = \frac{\boldsymbol{\gamma}_{i,t}(\hat{\theta}) \boldsymbol{\mu}_{i,t-1}^{\text{a}_{ii}}(\hat{\theta}) \prod_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}^{\text{a}_{ij}}(\hat{\theta})}{\sum_{\tilde{\theta} \in \Theta} \boldsymbol{\gamma}_{i,t}(\tilde{\theta}) \boldsymbol{\mu}_{i,t-1}^{\text{a}_{ii}}(\tilde{\theta}) \prod_{j \in \mathcal{N}_i} \boldsymbol{\mu}_{j,t-1}^{\text{a}_{ij}}(\tilde{\theta})}, \quad (\text{II})$$

initialized by: $\boldsymbol{\mu}_{i,0}(\hat{\theta}) = \boldsymbol{\gamma}_{i,0}(\hat{\theta}) / \sum_{\tilde{\theta} \in \Theta} \boldsymbol{\gamma}_{i,0}(\tilde{\theta})$.

Theorem G.2 (Learning from Intermittent Streams). *For all $i \in [n]$, and any pair of states $\hat{\theta}, \check{\theta} \in \Theta$, let $\Lambda_i(\hat{\theta}, \check{\theta}) = \mathbb{E}_{\theta} \{\log(\ell_i(\mathbf{s}_{i,0}|\hat{\theta})/\ell_i(\mathbf{s}_{i,0}|\check{\theta}))\}$. If $\sum_{i=1}^n \nu_i \Lambda_i(\hat{\theta}, \theta) < 0$ for all $\hat{\theta} \neq \theta$, then under (II), $\lim_{t \rightarrow \infty} \boldsymbol{\mu}_{i,t}(\theta) = 1$, for all i . Moreover, the learning is asymptotically exponentially fast with the rate equal to $\min_{\hat{\theta} \neq \theta} \{(-1/n) \sum_{i=1}^n \nu_i \Lambda_i(\hat{\theta}, \theta)\}$.*

Proof. The belief update rule proposed in (II) is the same as the time-invariant log-linear update with weighted self-beliefs considered in [164, Equation (13)]; except that here at every round each agent is receiving a random number of signals. Hence, the proof of convergence in [164, Equation (13)] can be applied here and with minor modifications.

Specifically, we let $\bar{\Phi}_t(\hat{\theta}, \check{\theta})$ be the vectorized log belief ratio statistics as defined in (G.0.1) and (G.0.2), and define the log ratio of the likelihood products of the received signals: $\lambda_{i,t}(\hat{\theta}, \check{\theta}) = \log(\gamma_{i,t}(\hat{\theta})/\gamma_{i,t}(\check{\theta}))$, and its vectorization $\bar{\lambda}_t(\hat{\theta}, \check{\theta}) = (\lambda_{1,t}(\hat{\theta}, \check{\theta}), \dots, \lambda_{n,t}(\hat{\theta}, \check{\theta}))$. Then after forming the log belief ratios, (II) in vectorized form yields that: $\bar{\Phi}_t(\hat{\theta}, \check{\theta}) = \Lambda \bar{\Phi}_{t-1}(\hat{\theta}, \check{\theta}) + \bar{\lambda}_t(\hat{\theta}, \check{\theta}) = \sum_{\tau=0}^t \Lambda^\tau \bar{\lambda}_{t-\tau}(\hat{\theta}, \check{\theta})$ and the latter converges almost surely to $((t/n) \mathbb{1}^\top \mathbb{E}\{\bar{\lambda}_0(\hat{\theta}, \check{\theta})\}) \mathbb{1}$, as $t \rightarrow \infty$; this is a simple consequence of the Cesàro mean together with the strong law of large numbers. The proof follows since $\mathbb{E}\{\bar{\lambda}_0(\hat{\theta}, \check{\theta})\} = (\Lambda_1(\hat{\theta}, \check{\theta}), \dots, \Lambda_n(\hat{\theta}, \check{\theta}))^\top$; in particular, $\lim_{t \rightarrow \infty} \frac{1}{t} \bar{\Phi}_t(\hat{\theta}, \check{\theta}) = ((1/n) \sum_{i=1}^n \Lambda_i(\hat{\theta}, \check{\theta})) \mathbb{1}$, with probability one and whenever $\sum_{i=1}^n \Lambda_i(\hat{\theta}, \check{\theta}) < 0$, the agents learn the truth asymptotically exponentially fast, at the rate $\min_{\hat{\theta} \neq \theta} \{(-1/n) \sum_{i=1}^n \Lambda_i(\hat{\theta}, \theta)\}$. \square

To understand the nature of the convergence result and learning rate in Theorem G.2, consider the special case where each agent at every time t may or may not have access to a sample point $\mathbf{s}_{i,t}$ and the accessibility of the new measurement $\mathbf{s}_{i,t}$ is determined by the outcome of an independent coin flip with success probability p_i , i.e. $\{\mathbf{n}_{i,t}, t \in \mathbb{N}\}$ are i.i.d Bernoulli(p_i) variables. Then the convergence rate in Theorem G.2 becomes $\min_{\hat{\theta} \neq \theta} \{(-1/n) \sum_{i=1}^n p_i \Lambda_i(\hat{\theta}, \theta)\}$, which decreases linearly with the decreasing probability of making new observations. Also note that $\Lambda_i(\hat{\theta}, \theta) = \mathbb{E}_\theta \{\log(\ell_i(\mathbf{s}_{i,0}|\hat{\theta})/\ell_i(\mathbf{s}_{i,0}|\theta))\} := -D_{\text{KL}}(\ell_i(\cdot|\hat{\theta})\|\ell_i(\cdot|\theta)) \leq 0$, where $D_{\text{KL}}(\cdot\|\cdot) \geq 0$ is the Kullback-Leibler divergence. It measures a pseudo-distance between the two distributions and it is strictly positive whenever $\ell_i(\cdot|\hat{\theta}) \neq \ell_i(\cdot|\theta)$, i.e. the two distributions disagree over a non-trivial (nonzero measure) set [193, Theorem 2.6.3]; hence, the closer the alternative distributions are to the true distributions the slower is the rate.

Bibliography

- [1] J. Surowiecki, *The Wisdom of Crowds*. Knopf Doubleday Publishing Group, 2005.
- [2] F. Galton, “Vox populi (the wisdom of crowds),” *Nature*, vol. 75, no. 7, pp. 450–451, 1907.
- [3] M. T. Maloney and J. H. Mulherin, “The complexity of price discovery in an efficient market: the stock market reaction to the challenger crash,” *Journal of corporate finance*, vol. 9, no. 4, pp. 453–479, 2003.
- [4] I. L. Janis, *Groupthink: Psychological studies of policy decisions and fiascoes*. Houghton Mifflin Boston, 1982.
- [5] K. Eliaz, D. Ray, and R. Razin, “Choice shifts in groups: A decision-theoretic basis,” *The American economic review*, pp. 1321–1332, 2006.
- [6] J. A. Stoner, “Risky and cautious shifts in group decisions: The influence of widely held values,” *Journal of Experimental Social Psychology*, vol. 4, no. 4, pp. 442–459, 1968.
- [7] D. J. Isenberg, “Group polarization: a critical review and meta-analysis,” *Journal of personality and social psychology*, vol. 50, no. 6, p. 1141, 1986.
- [8] N. Roux and J. Sobel, “Group polarization in a model of information aggregation,” *American Economic Journal: Microeconomics*, vol. 7, no. 4, pp. 202–32, 2015.
- [9] D. Kahneman, *Thinking, fast and slow*. Farrar Straus Giroux, 2011.
- [10] G. Klein, “Performing a project premortem,” *Harvard Business Review*, vol. 85, no. 9, pp. 18–19, 2007.
- [11] C. R. Sunstein, “Deliberative trouble? why groups go to extremes,” *The Yale Law Journal*, vol. 110, no. 1, pp. 71–119, 2000.
- [12] R. J. Aumann, “Agreeing to disagree,” *The annals of statistics*, pp. 1236–1239, 1976.
- [13] J. D. Geanakoplos and H. M. Polemarchakis, “We can’t disagree forever,” *Journal of Economic Theory*, vol. 28, no. 1, pp. 192–200, 1982.

- [14] D. Gale and S. Kariy, “Bayesian learning in social networks,” *Games and Economic Behavior*, vol. 45, pp. 329–346, 2003.
- [15] D. Rosenberg, E. Solan, and N. Vieille, “Informational externalities and emergence of consensus,” *Games and Economic Behavior*, vol. 66, no. 2, pp. 979–994, 2009.
- [16] J. Marschak and R. Radner, *Economic theory of teams*. New Haven: Yale University Press, 1972.
- [17] F. A. Csaszar and J. Eggers, “Organizational decision making: An information aggregation view,” *Management Science*, vol. 59, no. 10, pp. 2257–2277, 2013.
- [18] Y. M. Ioannides, “On the architecture of complex organizations,” *Economics Letters*, vol. 25, no. 3, pp. 201–206, 1987.
- [19] R. Gibbons, “Incentives in organizations,” *Journal of Economic Perspectives*, vol. 12, no. 4, pp. 115–132, 1998.
- [20] P. Bolton and M. Dewatripont, “The firm as a communication network,” *The Quarterly Journal of Economics*, vol. 109, no. 4, pp. 809–839, 1994.
- [21] B. Visser, “Organizational communication structure and performance,” *Journal of economic behavior & organization*, vol. 42, no. 2, pp. 231–252, 2000.
- [22] S. Vassilakis, “Accelerating new product development by overcoming complexity constraints,” *Journal of Mathematical Economics*, vol. 28, no. 3, pp. 341–373, 1997.
- [23] A. V. Banerjee, “A simple model of herd behavior,” *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, 1992.
- [24] S. Bikhchandani, D. Hirshleifer, and I. Welch, “Learning from the behavior of others: Conformity, fads, and informational cascades,” *The Journal of Economic Perspectives*, vol. 12, no. 3, pp. pp. 151–170, 1998.
- [25] M. Mueller-Frank, “A general framework for rational learning in social networks,” *Theoretical Economics*, vol. 8, no. 1, pp. 1–40, 2013.
- [26] E. Eyster and M. Rabin, “Naive herding in rich-information settings,” *American economic journal: microeconomics*, vol. 2, no. 4, pp. 221–243, 2010.
- [27] V. Krishnamurthy and W. Hoiles, “Online reputation and polling systems: Data incest, social learning, and revealed preferences,” *IEEE Transactions on Computational Social Systems*, vol. 1, no. 3, pp. 164–179, 2014.

- [28] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, “Bayesian learning in social networks,” *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.
- [29] A. Jadbabaie, E. Mossel, and M. A. Rahimian, “Bayesian group decisions: algorithms and complexity,” *arXiv preprint*, 2017.
- [30] G. Gigerenzer and W. Gaissmaier, “Heuristic decision making,” *Annual review of psychology*, vol. 62, pp. 451–482, 2011.
- [31] G. Gigerenzer and D. G. Goldstein, “Reasoning the fast and frugal way: models of bounded rationality,” *Psychological review*, vol. 103, no. 4, p. 650, 1996.
- [32] G. Gigerenzer and P. M. Todd, *Simple heuristics that make us smart*. Oxford University Press, USA, 1999.
- [33] J. S. B. Evans, “In two minds: dual-process accounts of reasoning,” *Trends in cognitive sciences*, vol. 7, no. 10, pp. 454–459, 2003.
- [34] J. S. B. Evans and K. E. Stanovich, “Dual-process theories of higher cognition advancing the debate,” *Perspectives on psychological science*, vol. 8, no. 3, pp. 223–241, 2013.
- [35] E. J. Masicampo and R. F. Baumeister, “Toward a physiology of dual-process reasoning and judgment: Lemonade, willpower, and expensive rule-based analysis,” *Psychological Science*, vol. 19, no. 3, pp. 255–260, 2008.
- [36] A. D. Redish, *The mind within the brain: How we make decisions and how those decisions go wrong*. Oxford University Press, 2013.
- [37] D. Kahneman, “Maps of bounded rationality: Psychology for behavioral economics,” *The American economic review*, vol. 93, no. 5, pp. 1449–1475, 2003.
- [38] R. H. Thaler, *Misbehaving: The making of behavioral economics*. WW Norton & Company, 2015.
- [39] R. Hegselmann and U. Krause, “Opinion dynamics driven by various ways of averaging,” *Computational Economics*, vol. 25, no. 4, pp. 381–405, 2005.
- [40] V. Grimm and F. Mengel, “An experiment on belief formation in networks,” *Available at SSRN 2361007*, 2014.
- [41] A. G. Chandrasekhar, H. Larreguy, and J. P. Xandri, “Testing models of social learning on networks: Evidence from a lab experiment in the field,” National Bureau of Economic Research, Tech. Rep., 2015.

- [42] M. H. DeGroot, “Reaching a consensus,” *Journal of American Statistical Association*, vol. 69, pp. 118 – 121, 1974.
- [43] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, “Non-bayesian social learning,” *Games and Economic Behavior*, vol. 76, no. 1, pp. 210 – 225, 2012.
- [44] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, “Information heterogeneity and the speed of learning in social networks,” *Revise and Resubmit, Review of Economic Studies*, 2013.
- [45] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, “Foundations of non-bayesian social learning,” *Columbia Business School Research Paper*, 2016.
- [46] M. Mueller-Frank and C. Neri, “A general model of boundedly rational observational learning: Theory and experiment,” *Available at SSRN 2566210*, 2015.
- [47] P. M. DeMarzo, D. Vayanos, and J. Zwiebel, “Persuasion bias, social influence, and unidimensional opinions,” *The Quarterly Journal of Economics*, vol. 118, pp. 909–968, 2003.
- [48] J. Sobel, “On the relationship between individual and group decisions,” *Theoretical Economics*, vol. 9, no. 1, pp. 163–185, 2014.
- [49] J. P. Lightle, J. H. Kagel, and H. R. Arkes, “Information exchange in group decision making: The hidden profile problem reconsidered,” *Management Science*, vol. 55, no. 4, pp. 568–581, 2009.
- [50] D. Fudenberg, *The Theory of Learning in Games*. MIT Press, 1998.
- [51] H. P. Young, *Strategic learning and its limits*. Oxford university press, 2004.
- [52] N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*. New York, NY, USA: Cambridge University Press, 2006.
- [53] E. Mossel, A. Sly, and O. Tamuz, “Strategic learning and the topology of social networks,” *Econometrica*, vol. 83, no. 5, pp. 1755–1794, 2015.
- [54] ———, “Asymptotic learning on bayesian social networks,” *Probability Theory and Related Fields*, vol. 158, no. 1-2, pp. 127–157, 2014.
- [55] E. Mossel and O. Tamuz, “Efficient bayesian learning in social networks with gaussian estimators,” *arXiv preprint arXiv:1002.0747*, 2010.
- [56] M. Harel, E. Mossel, P. Strack, and O. Tamuz, “When more information reduces the speed of learning,” *arXiv preprint arXiv:1412.7172*, 2014.

- [57] Y. Kanoria and O. Tamuz, “Tractable bayesian social learning on trees,” *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 4, pp. 756–765, April 2013.
- [58] J. H. Aguilar and A. A. Buckareff, *Causing Human Actions*. The MIT Press, 2010.
- [59] P. Shah and P. A. Parrilo, “ ℓ_2 -optimal decentralized control over posets: A state-space solution for state-feedback,” *IEEE Transactions on Automatic Control*, vol. 58, no. 12, pp. 3084–3096, 2013.
- [60] E. Mossel and O. Tamuz, “Making consensus tractable,” *ACM Transactions on Economics and Computation*, vol. 1, no. 4, p. 20, 2013.
- [61] J. N. Tsitsiklis and M. Athans, “On the complexity of decentralized decision making and detection problems,” *IEEE Transactions on Automatic Control*, vol. 30, no. 5, pp. 440–446, 1985.
- [62] R. Radner, “Team decision problems,” *The Annals of Mathematical Statistics*, vol. 33, no. 3, pp. 857–881, 1962.
- [63] C. H. Papadimitriou and J. N. Tsitsiklis, “The complexity of markov decision processes,” *Mathematics of operations research*, vol. 12, no. 3, pp. 441–450, 1987.
- [64] H. Simon, *Reason in Human Affairs*. Stanford University Press, 1990.
- [65] K. Velupillai, *Computable economics: the Arne Ryde memorial lectures*. Oxford University Press on Demand, 2000, vol. 5.
- [66] S. Aaronson, “The complexity of agreement,” in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. ACM, 2005, pp. 634–643.
- [67] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, “The complexity of computing a nash equilibrium,” *SIAM Journal on Computing*, vol. 39, no. 1, pp. 195–259, 2009.
- [68] K. Etessami and M. Yannakakis, “On the complexity of nash equilibria and other fixed points,” *SIAM Journal on Computing*, vol. 39, no. 6, pp. 2531–2597, 2010.
- [69] X. Chen and X. Deng, “3-nash is ppad-complete,” in *Electronic Colloquium on Computational Complexity*, vol. 134. Citeseer, 2005.
- [70] C. Daskalakis and C. H. Papadimitriou, “Three-player games are hard,” in *Electronic colloquium on computational complexity*, vol. 139, 2005, pp. 81–87.
- [71] X. Chen, X. Deng, and S.-H. Teng, “Settling the complexity of computing two-player nash equilibria,” *Journal of the ACM (JACM)*, vol. 56, no. 3, p. 14, 2009.

- [72] J. Nash, “Non-cooperative games,” *Annals of mathematics*, pp. 286–295, 1951.
- [73] S. Hart and Y. Mansour, “How long to equilibrium? the communication complexity of uncoupled equilibrium procedures,” *Games and Economic Behavior*, vol. 69, no. 1, pp. 107–126, 2010.
- [74] M. A. Rahimian, P. Molavi, and A. Jadbabaie, “(Non-) bayesian learning without recall,” *IEEE Conference on Decision and Control (CDC)*, pp. 5730–5735, 2014.
- [75] G. Levy and R. Razin, “Information diffusion in networks with the bayesian peer influence heuristic,” *preprint*, 2016.
- [76] G. Levy, I. M. de Barreda, and R. Razin, “Persuasion with correlation neglect: Media bias and media power through correlation of news content,” *preprint*, 2017.
- [77] M. Mueller-Frank and C. Neri, “Quasi-bayesian updating in social networks,” *Available at SSRN 2933411*, 2017.
- [78] H. A. Simon, “A behavioral model of rational choice,” *The Quarterly Journal of Economics*, pp. 99–118, 1955.
- [79] ———, “Rational choice and the structure of the environment,” *Psychological review*, vol. 63, no. 2, p. 129, 1956.
- [80] ———, “Rational decision making in business organizations,” *The American economic review*, vol. 69, no. 4, pp. 493–513, 1979.
- [81] V. Krishnamurthy, *Partially Observed Markov Decision Processes*. Cambridge University Press, 2016.
- [82] M. Luca, “Reviews, reputation, and revenue: The case of yelp. com,” *Com (September 16, 2011)*. *Harvard Business School NOM Unit Working Paper*, no. 12-016, 2011.
- [83] D. Blackwell and L. Dubins, “Merging of opinions with increasing information,” *The Annals of Mathematical Statistics*, vol. 33, pp. 882 – 886, 1962.
- [84] E. Lehrer and R. Smorodinsky, “Merging and learning,” *Lecture Notes-Monograph Series*, pp. 147–168, 1996.
- [85] S. N. Ali, “Social learning with endogenous information,” *Technical report*, 2014.
- [86] L. E. Blume, “The statistical mechanics of strategic interaction,” *Games and Economic Behavior*, vol. 5, no. 3, pp. 387–424, 1993.

- [87] J. R. Marden and J. S. Shamma, “Revisiting log-linear learning: Asynchrony, completeness and payoff-based implementation,” *Games and Economic Behavior*, vol. 75, no. 2, pp. 788–808, 2012.
- [88] M. Rahimian, A. Jadbabaie, and E. Mossel, “Complexity of bayesian belief exchange over a network,” *56th IEEE Conference on Decision and Control (CDC)*, 2017.
- [89] J. Cave, “Learning to agree,” *Economics Letters*, vol. 12, no. 2, pp. 147–152, 1983.
- [90] M. Bacharach, “Some extensions of a claim of aumann in an axiomatic model of knowledge,” *Journal of Economic Theory*, vol. 37, no. 1, pp. 167–190, 1985.
- [91] R. Parikh and P. Krasucki, “Communication, consensus, and knowledge,” *Journal of Economic Theory*, vol. 52, no. 1, pp. 178–189, 1990.
- [92] R. K. Sah and J. E. Stiglitz, “Human fallibility and economic organization,” *The American Economic Review*, vol. 75, no. 2, pp. 292–297, 1985.
- [93] ———, “The architecture of economic systems: Hierarchies and polyarchies,” *The American Economic Review*, pp. 716–727, 1986.
- [94] ———, “Committees, hierarchies and polyarchies,” *The Economic Journal*, vol. 98, no. 391, pp. 451–470, 1988.
- [95] M. Christensen and T. Knudsen, “Design of decision-making organizations,” *Management Science*, vol. 56, no. 1, pp. 71–89, 2010.
- [96] J. Bang-Jensen and G. Z. Gutin, *Digraphs: theory, algorithms and applications*. Springer Science & Business Media, 2008.
- [97] E. Eyster and M. Rabin, “Extensive imitation is irrational and harmful,” *The Quarterly Journal of Economics*, p. qju021, 2014.
- [98] T. Wang, Q. Berthet, R. J. Samworth *et al.*, “Statistical and computational tradeoffs in estimation of sparse principal components,” *The Annals of Statistics*, vol. 44, no. 5, pp. 1896–1930, 2016.
- [99] Y. Chen and J. Xu, “Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices,” *Journal of Machine Learning Research*, vol. 17, no. 27, pp. 1–57, 2016.
- [100] A. Bogdanov, L. Trevisan *et al.*, “Average-case complexity,” *Foundations and Trends® in Theoretical Computer Science*, vol. 2, no. 1, pp. 1–106, 2006.

- [101] R. Motwani and P. Raghavan, *Randomized algorithms*. Chapman & Hall/CRC, 2010.
- [102] S. Arora and B. Barak, *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- [103] M. A. Rahimian and A. Jadbabaie, “Bayesian heuristics for group decisions,” *preprint*, 2017.
- [104] ———, “Learning without recall from actions of neighbors,” in *American Control Conference (ACC), 2016*. IEEE, 2016, pp. 1060–1065.
- [105] ———, “Group decision making and social learning,” in *Decision and Control (CDC), 2016 IEEE 55th Conference on*. IEEE, 2016, pp. 6783–6794.
- [106] A. Tversky and D. Kahneman, “Judgment under uncertainty: Heuristics and biases,” in *Utility, probability, and human decision making*. Springer, 1975, pp. 141–162.
- [107] M. A. Rahimian and A. Jadbabaie, “Bayesian learning without recall,” *IEEE Transactions on Signal and Information Processing over Networks*, 2016.
- [108] V. Bala and S. Goyal, “Learning from neighbours,” *The Review of Economic Studies*, vol. 65, no. 3, pp. 595–621, 1998.
- [109] G. Levy and R. Razin, “Correlation neglect, voting behavior, and information aggregation,” *The American Economic Review*, vol. 105, no. 4, pp. 1634–1645, 2015.
- [110] L. Farina and S. Rinaldi, *Positive linear systems: theory and applications*. John Wiley & Sons, 2011, vol. 50.
- [111] G. James and V. Rumchev, “Stability of positive linear discrete-time systems,” *Bulletin of the Polish Academy of Sciences. Technical Sciences*, vol. 53, no. 1, pp. 1–8, 2005.
- [112] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian data analysis*, 3rd ed., ser. Texts in Statistical Science Series. CRC Press, Boca Raton, FL, 2014.
- [113] R. E. Kass and L. Wasserman, “The selection of prior distributions by formal rules,” *Journal of the American Statistical Association*, vol. 91, no. 435, pp. 1343–1370, 1996.
- [114] C. Robert, *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*, ser. Springer Texts in Statistics. Springer, 2001.

- [115] E. Seneta, *Non-negative matrices and Markov chains*. Springer, 2006.
- [116] N. E. Friedkin, “A formal theory of social power,” *Journal of Mathematical Sociology*, vol. 12, no. 2, pp. 103–126, 1986.
- [117] J. R. French Jr, “A formal theory of social power.” *Psychological review*, vol. 63, no. 3, pp. 181–194, 1956.
- [118] F. Harary, “A criterion for unanimity in french’s theory of social power,” in *Studies in Social Power*. University of Michigan, 1959, pp. 168–182, d. Cartwright, Ed.
- [119] J. Baikerikar, S. Surve, and S. Prabhu, “Consensus based dynamic load balancing for a network of heterogeneous workstations,” in *Advances in Computing, Communication and Control*. Springer, 2011, pp. 116–124.
- [120] V. Borkar and P. Varaiya, “Asymptotic agreement in distributed estimation,” *IEEE Transactions on Automatic Control*, vol. 27, no. 3, pp. 650–655, Jun 1982.
- [121] J. N. Tsitsiklis and M. Athans, “Convergence and asymptotic agreement in distributed decision problems,” *Automatic Control, IEEE Transactions on*, vol. 29, no. 1, pp. 42–50, 1984.
- [122] S. Kar and J. M. Moura, “Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise,” *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 355–369, 2009.
- [123] S. Katragadda, J. C. SanMiguel, and A. Cavallaro, “Consensus protocols for distributed tracking in wireless camera networks,” in *17th International Conference on Information Fusion (FUSION)*. IEEE, 2014, pp. 1–8.
- [124] L. Xiao, S. Boyd, and S. Lall, “A scheme for robust distributed sensor fusion based on average consensus,” in *Fourth International Symposium on Information Processing in Sensor Networks (IPSN)*, 2005, pp. 63–70.
- [125] ———, “A space-time diffusion scheme for peer-to-peer least-squares estimation,” in *Proceedings of the 5th international conference on Information Processing in Sensor Networks*, 2006, pp. 168–176.
- [126] A. Jadbabaie, J. Lin, and A. S. Morse, “Coordination of groups of mobile autonomous agents using nearest neighbor rules,” *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.
- [127] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press, 2010.

- [128] J. Tsitsiklis, D. Bertsekas, and M. Athans, “Distributed asynchronous deterministic and stochastic gradient optimization algorithms,” *IEEE Transactions on Automatic Control*, vol. 31, no. 9, pp. 803–812, 1986.
- [129] R. Lyons and Y. Peres, *Probability on Trees and Networks*. Cambridge Series in Statistical and Probabilistic Mathematics, 2016.
- [130] B. Golub and M. O. Jackson, “Naïve Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, vol. 2, no. 1, pp. 112–149, Feb. 2010.
- [131] G. L. Gilardoni and M. K. Clayton, “On reaching a consensus using DeGroot’s iterative pooling,” *The Annals of Statistics*, pp. 391–401, 1993.
- [132] M. J. Rufo, J. Martin, C. J. Pérez *et al.*, “Log-linear pool to combine prior distributions: A suggestion for a calibration-based approach,” *Bayesian Analysis*, vol. 7, no. 2, pp. 411–438, 2012.
- [133] S. Shahrampour and A. Jadbabaie, “Exponentially fast parameter estimation in networks using distributed dual averaging,” *52nd IEEE Conference on Decision and Control (CDC)*, pp. 6196–6201, 2013.
- [134] K. Rahnama Rad and A. Tahbaz-Salehi, “Distributed parameter estimation in networks,” *49th IEEE Conference on Decision and Control (CDC)*, pp. 5050–5055, 2010.
- [135] A. Nedić, A. Olshevsky, and C. A. Uribe, “Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs,” *arXiv preprint arXiv:1410.1977*, 2014.
- [136] A. Lalitha, A. Sarwate, and T. Javidi, “Social learning and distributed hypothesis testing,” *IEEE International Symposium on Information Theory*, pp. 551–555, 2014.
- [137] S. Bandyopadhyay and S.-J. Chung, “Distributed estimation using bayesian consensus filtering,” *American Control Conference*, pp. 634–641, 2014.
- [138] G. Battistelli, L. Chisci, S. Morrocchi, and F. Papi, “An information-theoretic approach to distributed state estimation,” *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 12 477–12 482, 2011.
- [139] G. Battistelli and L. Chisci, “Kullback–leibler average, consensus on probability densities, and distributed state estimation with guaranteed stability,” *Automatica*, vol. 50, no. 3, pp. 707–718, 2014.

- [140] A. Nedić, A. Olshevsky, and C. Uribe, “Distributed learning with infinitely many hypotheses,” *arXiv preprint arXiv:1605.02105*, 2016.
- [141] A. Nedić, A. Olshevsky, and C. A. Uribe, “Fast convergence rates for distributed non-bayesian learning,” *arXiv preprint arXiv:1508.05161*, 2015.
- [142] A. Zellner, “Optimal information processing and bayes’s theorem,” *The American Statistician*, vol. 42, no. 4, pp. 278–280, 1988.
- [143] M. A. Rahimian and A. Jadbabaie, “Naive social learning in ising networks,” *American Control Conference*, 2016.
- [144] N. H. Anderson, “Integration theory and attitude change.” *Psychological review*, vol. 78, no. 3, p. 171, 1971.
- [145] ———, *Foundations of information integration theory*. Academic Press, 1981.
- [146] J. Glazer and A. Rubinstein, “Motives and implementation: On the design of mechanisms to elicit opinions,” *Journal of Economic Theory*, vol. 79, no. 2, pp. 157–173, 1998.
- [147] J.-F. Laslier and J. Weibull, “Committee decisions: optimality and equilibrium,” *cahier de recherche 2008-24. v3 du halshs-00121741*, 2008.
- [148] G. Desanctis and R. B. Gallupe, “A foundation for the study of group decision support systems,” *Management Science*, vol. 33, no. 5, pp. 589–609, May 1987.
- [149] P. E. Slater, “Contrasting correlates of group size,” *Sociometry*, vol. 21, no. 2, pp. 129–139, 1958.
- [150] I. D. Steiner, *Group Processes and Productivity, Social psychology: a series of monographs, treatises, and texts*. Academic Press, 1972.
- [151] G. E. Littlepage, “Effects of group size and task characteristics on group performance: A test of steiner’s model,” *Personality and Social Psychology Bulletin*, vol. 17, no. 4, pp. 449–456, 1991.
- [152] N. L. K. Bray, Robert M. and R. S. Atkin, “Effects of group size, problem difficulty, and sex on group performance and member reactions,” *Journal of Personality and Social Psychology*, vol. 36, no. 11, p. 1224, 1978.
- [153] J. R. Hackman and N. Vidmar, “Effects of size and task type on group performance and member reactions,” *Sociometry*, pp. 37–54, 1970.

- [154] F. Frank and L. R. Anderson, “Effects of task and group size upon group productivity and member satisfaction,” *Sociometry*, vol. 34, no. 1, pp. 135–149, 1971.
- [155] H. JR, *Groups that work (and those that don’t)*. San Francisco: Jossey-Bass, 1990.
- [156] R. A. Guzzo and G. P. Shea, “Group performance and intergroup relations in organizations,” *Handbook of industrial and organizational. Palo Alto, CA, US: Consulting Psychologists Press*, pp. 269–313, 1992.
- [157] M. Gradstein, S. Nitzan, and J. Paroush, “Collective decision making and the limits on the organization’s size,” *Public Choice*, vol. 66, no. 3, pp. 279–291, 1990.
- [158] C. F. Thomas, Edwin J.; Fink, “Effects of group size,” *Psychological Bulletin*, vol. 60, no. 4, pp. 371–384, 1963.
- [159] R. Z. Gooding and J. A. Wagner, “A meta-analytic review of the relationship between size and performance: The productivity and efficiency of organizations and their subunits,” *Administrative Science Quarterly*, vol. 30, no. 4, pp. 462–481, 1985.
- [160] L. A. Currell, R. H. Forrester, J. F. Dawson, and M. A. West, “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*, vol. 10, no. 2, pp. 187–204, 2001.
- [161] M. Hoegl, “Smaller teams-better teamwork: How to keep project teams small,” *Business Horizons*, vol. 48, no. 3, pp. 209–214, 2005.
- [162] W. School, “Is your team too big? too small? whats the right number?” *Knowledge@Wharton*, Jul. 2006.
- [163] A. Mao, W. Mason, S. Suri, and D. J. Watts, “An experimental study of team size and performance on a complex task,” *PloS one*, vol. 11, no. 4, p. e0153048, 2016.
- [164] M. A. Rahimian and A. Jadbabaie, “Learning without recall: A case for log-linear learning,” *5th IFAC Workshop on Distributed Estimation and Control in Networked Systems*, 2015.
- [165] ———, “Learning without recall in directed circles and rooted trees,” *American Control Conference*, pp. 4222–4227, 2015.
- [166] M. A. Rahimian, S. Shahrapour, and A. Jadbabaie, “Learning without recall by random walks on directed graphs,” *IEEE Conference on Decision and Control (CDC)*, 2015.

- [167] M. A. Rahimian and A. Jadbabaie, “Distributed estimation and learning over heterogeneous networks,” in *Communication, Control, and Computing (Allerton), 2016 54th Annual Allerton Conference on.* IEEE, 2016, pp. 1314–1321.
- [168] C. P. Chamley, *Rational Herds: Economic Models of Social Learning.* Cambridge University Press, 2004.
- [169] A. Wilson, “Bounded memory and biases in information processing,” *Econometrica*, vol. 82, no. 6, pp. 2257–2294, 2014.
- [170] T. M. Cover, “A note on the two-armed bandit problem with finite memory,” *Information and Control*, vol. 12, no. 5, pp. 371–377, 1968.
- [171] ———, “Hypothesis testing with finite statistics,” *The Annals of Mathematical Statistics*, pp. 828–835, 1969.
- [172] L. Kontorovich, “Statistical estimation with bounded memory,” *Statistics and Computing*, vol. 22, no. 5, pp. 1155–1164, 2012.
- [173] M. E. Hellman and T. M. Cover, “Learning with finite memory,” *The Annals of Mathematical Statistics*, pp. 765–782, 1970.
- [174] T. M. Cover, M. A. Freedman, and M. E. Hellman, “Optimal finite memory learning algorithms for the finite sample problem,” *Information and Control*, vol. 30, no. 1, pp. 49 – 85, 1976.
- [175] T. M. Cover and M. E. Hellman, “The two-armed-bandit problem with time-invariant finite memory,” *IEEE Transactions on Information Theory*, vol. 16, no. 2, pp. 185–195, 1970.
- [176] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis, “On learning with finite memory,” *IEEE Transactions on Information Theory*, vol. 59, no. 10, pp. 6859–6872, 2013.
- [177] P. Billingsley, *Probability and Measure*, 3rd ed. Wiley-Interscience, 1995.
- [178] R. O’Donnell, *Analysis of boolean functions.* Cambridge University Press, 2014.
- [179] I. Benjamini, G. Kalai, and O. Schramm, “Noise sensitivity of boolean functions and applications to percolation,” *Publications Mathématiques de l’Institut des Hautes Études Scientifiques*, vol. 90, no. 1, pp. 5–43, 1999.
- [180] E. Mossel, R. O’Donnell, and K. Oleszkiewicz, “Noise stability of functions with low influences: invariance and optimality,” *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 21–30, 2005.

- [181] Y. Peres, “Noise stability of weighted majority,” *arXiv preprint math/0412377*, 2004.
- [182] W. S. McCulloch and W. Pitts, “A logical calculus of the ideas immanent in nervous activity,” *The bulletin of mathematical biophysics*, vol. 5, no. 4, pp. 115–133, 1943.
- [183] J. J. Hopfield, “Neural networks and physical systems with emergent collective computational abilities,” *Proceedings of the national academy of sciences*, vol. 79, no. 8, pp. 2554–2558, 1982.
- [184] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times*. American Mathematical Society, 2009.
- [185] F. Martinelli, *Lectures on Glauber Dynamics for Discrete Spin Models*. Springer, 1999.
- [186] M. Ostilli, E. Yoneki, I. X. Y. Leung, J. F. F. Mendes, P. Lió, and J. Crowcroft, “Statistical mechanics of rumour spreading in network communities,” *Procedia Computer Science*, vol. 1, no. 1, pp. 2331–2339, 2010.
- [187] J.-P. Nadal, D. Phan, M. B. Gordon, and J. Vannimenus, “Multiple equilibria in a monopoly market with heterogeneous agents and externalities,” *Quantitative Finance*, vol. 5, no. 6, pp. 557–568, 2005.
- [188] M. H. Afrasiabi, R. Guerin, and S. S. Venkatesh, “Opinion formation in ising networks,” *Information Theory and Applications Workshop*, pp. 1–10, 2013.
- [189] D. V. Lindley, *Introduction to probability and statistics from bayesian viewpoint. part I: Probability, Part II: inference*. CUP Archive, 1965.
- [190] M. H. DeGroot, *Optimal Statistical Decisions*, ser. Wiley Classics Library. New York:McGraw-Hill, 1969.
- [191] L. J. Savage, *The foundations of statistics*. Courier Dover Publications, 1972.
- [192] J. O. Berger, *Statistical decision theory and Bayesian analysis*. Springer, 1985.
- [193] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, ser. A Wiley-Interscience publication. Wiley, 2006.
- [194] R. A. Brualdi, *The mutually beneficial relationship of graphs and matrices*. American Mathematical Soc., 2011, no. 115.
- [195] T. Gagnon-Bartsch and M. Rabin, “Naive social learning, mislearning, and unlearning,” *Working Paper*, 2016.

- [196] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Randomized gossip algorithms,” *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2508–2530, 2006.
- [197] A. Tahbaz-Salehi and A. Jadbabaie, “Consensus over ergodic stationary graph processes,” *IEEE Transactions on Automatic Control*, vol. 55, no. 1, pp. 225–230, 2010.
- [198] J. R. Norris, *Markov Chains*. Cambridge University Press, 1999.
- [199] M. Newman, *Networks: An Introduction*. Oxford University Press, 2010.
- [200] J. G. Kemeny and J. L. Snell, *Finite Markov Chains*. New York: Van Nostrand, 1960.
- [201] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 2014.
- [202] K. P. Murphy, Y. Weiss, and M. I. Jordan, “Loopy belief propagation for approximate inference: An empirical study,” in *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 1999, pp. 467–475.
- [203] V. Chandrasekaran, N. Srebro, and P. Harsha, “Complexity of inference in graphical models,” *arXiv preprint arXiv:1206.3240*, 2012.
- [204] R. G. Cowell, P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter, *Probabilistic networks and expert systems: Exact computational methods for Bayesian networks*. Springer, 2006.
- [205] S. Kar and J. Moura, “Consensus+ innovations distributed inference over networks: cooperation and sensing in networked systems,” *Signal Processing Magazine, IEEE*, vol. 30, no. 3, pp. 99–109, 2013.
- [206] N. E. Leonard and A. Olshevsky, “Cooperative learning in multiagent systems from intermittent measurements,” *SIAM Journal on Control and Optimization*, vol. 53, no. 1, pp. 1–29, 2015.
- [207] J.-F. Mertens and S. Zamir, “Formulation of bayesian analysis for games with incomplete information,” *International Journal of Game Theory*, vol. 14, no. 1, pp. 1–29, 1985.
- [208] A. Brandenburger and E. Dekel, “Hierarchies of beliefs and common knowledge,” *Journal of Economic Theory*, vol. 59, no. 1, pp. 189–198, 1993.

- [209] M. Rabin and J. L. Schrag, “First impressions matter: A model of confirmatory bias*,” *The Quarterly Journal of Economics*, vol. 114, no. 1, pp. 37–82, 1999.
- [210] S. Mahdavi and M. A. Rahimian, “Hindsight bias impedes learning,” in *Proceedings of the NIPS 2016 Workshop on Imperfect Decision Makers*, ser. Proceedings of Machine Learning Research, vol. 58, 2017, pp. 111–127.
- [211] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York, NY, USA: W. H. Freeman & Co., 1979.
- [212] R. M. Karp, “Reducibility among combinatorial problems,” in *Complexity of computer computations*. Springer, 1972, pp. 85–103.
- [213] P. J. Bickel and K. A. Doksum, *Mathematical Statistics: Basic Ideas and Selected Topics, volume I*. CRC Press, 2015.
- [214] P. Diaconis and D. Ylvisaker, “Conjugate priors for exponential families,” *The Annals of statistics*, vol. 7, no. 2, pp. 269–281, 1979.
- [215] G. R. Brightwell and P. Winkler, “Graph homomorphisms and phase transitions,” *Journal of Combinatorial Theory, Series B*, vol. 77, no. 2, pp. 221–262, 1999.
- [216] D. Aldous and J. Fill, “Reversible markov chains and random walks on graphs,” *Book in progress*, 2002.
- [217] W. Feller, *An Introduction to Probability Theory and Its Applications*, 3rd ed. Wiley, 1968, vol. 1.
- [218] P. Molavi, C. Eksin, A. Ribeiro, and A. Jadbabaie, “Learning to coordinate in social networks,” *Operations Research*, vol. 64, no. 3, pp. 605–621, 2015.
- [219] J. N. Tsitsiklis, “Decentralized detection,” *Advances in Statistical Signal Processing*, vol. 2, no. 2, pp. 297–344, 1993.
- [220] J.-F. Chamberland and V. V. Veeravalli, “Decentralized detection in sensor networks,” *IEEE Transactions on Signal Processing*, vol. 51, no. 2, pp. 407–416, 2003.
- [221] R. Olfati-Saber and J. Shamma, “Consensus filters for sensor networks and distributed sensor fusion,” *IEEE Conference on Decision and Control*, pp. 6698 – 6703, 2005.
- [222] S. Kar, J. Moura, and K. Ramanan, “Distributed parameter estimation in sensor networks: Nonlinear observation models and imperfect communication,” *IEEE Transactions on Information Theory*, vol. 58, no. 6, pp. 3575–3605, 2012.

- [223] N. A. Atanasov, J. Le Ny, and G. J. Pappas, “Distributed algorithms for stochastic source seeking with mobile robot networks,” *Journal of Dynamic Systems, Measurement, and Control*, 2014.
- [224] N. Atanasov, R. Tron, V. M. Preciado, and G. J. Pappas, “Joint estimation and localization in sensor networks,” *IEEE Conference on Decision and Control (CDC)*, pp. 6875–6882, 2014.
- [225] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, “Distributed tsitsiklisection: Finite-time analysis and impact of network topology,” *arXiv preprint arXiv:1409.8606*, 2014.
- [226] V. Krishnamurthy and H. V. Poor, “Social learning and bayesian games in multiagent signal processing: How do local and global decision makers interact?” *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 43–57, 2013.
- [227] Y. Wang and P. M. Djuric, “Social learning with bayesian agents and random decision making,” *IEEE Transactions on Signal Processing*, vol. 63, no. 12, pp. 3241–3250, 2015.
- [228] A. D. Dimakis, A. D. Sarwate, and M. J. Wainwright, “Geographic gossip: Efficient averaging for sensor networks,” *Signal Processing, IEEE Transactions on*, vol. 56, no. 3, pp. 1205–1216, 2008.
- [229] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*. Princeton University Press, 2009.
- [230] S. Boyd, P. Diaconis, and L. Xiao, “Fastest mixing markov chain on a graph,” *SIAM review*, vol. 46, no. 4, pp. 667–689, 2004.
- [231] G. Casella and R. L. Berger, *Statistical inference*. Duxbury Pacific Grove, CA, 2002, vol. 2.
- [232] A. Olshevsky, “Linear time average consensus on fixed graphs and implications for decentralized optimization and multi-agent control,” *arXiv preprint arXiv:1411.4186*, 2014.
- [233] P. Diaconis and D. Stroock, “Geometric bounds for eigenvalues of markov chains,” *The Annals of Applied Probability*, pp. 36–61, 1991.
- [234] L. J. Billera and P. Diaconis, “A geometric interpretation of the metropolis-hastings algorithm,” *Statistical Science*, pp. 335–339, 2001.