Title
Social Learning over Weak Graphs

Author
Salami, Hawraa

Publication Date
2019

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Social Learning over Weak Graphs

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Electrical Engineering

by

Hawraa Salami

2019

ABSTRACT OF THE DISSERTATION

Social Learning over Weak Graphs

by

Hawraa Salami
Doctor of Philosophy in Electrical Engineering
University of California, Los Angeles, 2019
Professor Ali H. Sayed, Chair

In this dissertation, we study diffusion social learning over weakly-connected graphs and reveal several interesting properties characterizing the flow of information over such networks. We discover that the asymmetric flow of information hinders the learning ability of certain agents regardless of their local observations. Under some circumstances that we clarify in this work, a scenario of total influence (or "mind-control") arises where a set of influential agents ends up shaping the beliefs of non-influential agents. We derive useful closed-form expressions that characterize this influence, and then analyze this control mechanism more closely to highlight some critical properties. In particular, we use the theoretical analysis to address two main questions: (a) First, how much freedom do influential agents have in controlling the beliefs of the receiving agents? That is, can influential agents drive receiving agents to arbitrary beliefs or does the network structure limit the scope of control by the influential agents? and (b) second, even if there is a limit to what influential agents can accomplish, how can they ensure that receiving agents will end up with particular beliefs? These questions raise interesting possibilities about belief control. Once addressed, we end up with design procedures that allow influential agents to drive other agents to endorse particular beliefs regardless of their convictions. We illustrate the theoretical findings and results by means of several examples and numerical simulations.

The dissertation of Hawraa Salami is approved.

Adnan Youssef Darwiche

Vwani P. Roychowdhury

Christina Panagio Fragouli

Ali H. Sayed, Committee Chair

University of California, Los Angeles

2019

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

My Ph.D. journey was certainly an enriching experience for me that allowed me to grow on so many levels. I learned a great deal from all the people I have met and worked with, and to whom I am deeply thankful.

I am grateful first for my advisor, Prof. Ali Sayed, for giving me the opportunity to be a member in his lab. His dedication, perseverance and passion for his work will always keep inspiring me. His encouragement, patience and insistence on a high-quality work certainly made something out of me and challenged me to do better. I always recall him saying to me, in a moment when I doubted myself: "You need to be more ambitious than this". His valuable advice and great vision made all of this work possible. Moreover, his special way of explaining and conveying ideas definitely played a role in improving my communication skills.

I would like to thank Prof. Christina Fragouli, Prof. Vwani Roychowdhury and Prof. Adnan Darwiche for agreeing to serve on my doctoral committee, and for their valuable remarks. I cannot also thank enough my lab-mate, Bicheng Ying, for his helpful comments, great discussions, and for always having the time to help. His ready answers, vast knowledge and sharp mind provided me with a lot of insightful help. I cannot also forget to thank all my lab-mates: Kun Yuan, Stefan Vlaski, Sulaiman Alghunaim, Lucas Cassano and Chung-Kai Yu, for being great colleagues and good friends during my Ph.D. experience.

My time at UCLA would not have been easy without having great and supportive friends: Yasmine Badr, a real big sister, Salma Elmalaki, the inspiring mom, and Sara Azzam, the deep thinker, with whom I share a lot of memorable moments. Being away from home was not always easy, but I was blessed with the presence of the beautiful family: Areej Abdallah, Ghaith Hattab and Rayyon, who were always like a small family for me.

I am mostly thankful for what I am now to the endless love and support of my parents, brother and sisters. I am thankful for all the efforts and sacrifices of my parents and for them pushing me to learn and become stronger. I am thankful to my father, a self-made man, from whom I learned to rely on myself. I am thankful to my mother, who always keeps me in her prayers, listens to me and understands me more than anyone else can do. I am thankful for the encouragement of my brother Ali and my sisters Souad, Zahraa and Latifa, and for their laughters and warm hugs. I am especially grateful for having by my side my husband, Anas Alanqar, who did not miss any opportunity to support and encourage me, and with whom I share precious UCLA memories. Finally and most importantly, all praise be to Allah!

At the end, we would like to mention that the work in this dissertation was supported in part by NSF grants ECCS-1407712 and CCF-1524250. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

VITA

2009–2013	B.E. in Computer and Communications Engineering, American University of Beirut, Lebanon.
2013–2014	M.S. in Electrical Engineering, University of California, Los Angeles.
2014–2019	Graduate Student Researcher/ Teaching Assistant and Associate, Department of Electrical Engineering, University of California, Los Angeles.

PUBLICATIONS

1. H. Salami, B. Ying, and A. H. Sayed, "Diffusion social learning over weakly-connected graphs," in *Proc. IEEE ICASSP*, Shanghai, China, March 2016, pp. 4119–4123.

2. H. Salami and A. H. Sayed, "Belief control strategies for interactions over weak graphs," in *Proc. IEEE ICASSP*, New Orleans, LA, March 2017, pp. 4232–4236.

3. S. Basir-Kazeruni, S. Vlaski, H. Salami, A. H. Sayed, and D. Markovic, "A blind adaptive stimulation artifact rejection (ASAR) engine for closed-loop implantable neuromodulation systems," in *Proc. Inter. IEEE EMBS Neural Engineering Conference*, Shanghai, China, May 2017, pp. 186-189.

4. H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly-connected graphs," in *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 2, pp. 222–238, June 2017.

5. H. Salami, B. Ying, and A. H. Sayed, "Belief Control Strategies for Interactions over Weak Graphs," in *IEEE Trans. Information Theory*, submitted.

CHAPTER 1

Introduction

Private information and social interactions among networked agents influence the beliefs of agents about the state of nature. This latter represents an unknown social state that agents would like to discover; for instance, agents might be interested in knowing who is the best candidate to vote for, what movie to watch, which restaurant to choose, what diet to follow, or whether vaccines are safe or not. In deciding whether the state of nature, denoted by $\boldsymbol{\theta}$, is either $\boldsymbol{\theta} = 1$ or $\boldsymbol{\theta} = 0$, an agent k observes some private data, which can represent past experiences, personal knowledge, or own convictions. Additionally, agent k observes the actions of its neighboring agents (e.g., friends, family, colleagues, ...) or consults with them about their opinion on the most plausible value for $\boldsymbol{\theta}$. By combining local measurements with information from neighbors, agents update their belief about $\boldsymbol{\theta}$. In updating their beliefs, agents might also be affected by some external influence through media, advertisement, or celebrities, which can play a role in creating a form of disagreement. Many models in the literature of social learning have been proposed to analyze how agents socially interact and aggregate information to form their opinions, and to investigate if agents can reach an agreement or if they can successfully aggregate their information.

Two main categories of models have been proposed in the literature to examine social learning. The first one is based on a Bayesian approach [1–17], where an agent relies on some priors, some observations and its understanding of the world, to perform a Bayesian operation and update its opinion. The second category of models consists of non-Bayesian learning models, which traditionally do not involve any Bayesian update step, and describe how each agent interacts locally with its neighbors and aggregates their opinions to form its own [17–25]. Both approaches provide insights into the formation of some interesting phenomena over social networks, but they both have their own shortcomings. Some recent emerging models of social learning involve a Bayesian update step without being fully Bayesian [26–40]; these models take into consideration the incoming of new information for agents, as well as their interaction with their neighbors. Our work focuses on one type of these recent models, and examines through it the external influence on opinion formation. In this chapter, we provide a summary of some of the social learning models, and then we outline our work and main contributions.

Notation: We use lowercase letters to denote vectors, uppercase letters for matrices, plain letters for deterministic variables, and boldface for random variables. We also use $(.)^{\mathsf{T}}$ for transposition, $(.)^{-1}$ for matrix inversion, and $\rho(.)$ for the spectral radius of a matrix. We use $\preceq$ and $\succeq$ for vector element-wise comparisons.

1.1 Bayesian and non-Bayesian Learning

In this section, we present models from both Bayesian and non-Bayesian frameworks. We describe first the basic model of herding that arises with fully Bayesian agents and discuss some of its extensions. We then present a classical non-Bayesian model and some of its variants.

1.1.1 Bayesian Approach

Some of the earliest Bayesian appear in [1] and [2]. Based on observational learning, these models describe how each agent learns by observing the actions of others, and how it incorporates its observations with its own private signals to decide whether to adopt a specific action or not. In these models, agents sequentially make their decisions given their private signals and all past actions of previous agents. More specifically, by observing all past actions, each agent infers the private signals received previously by other agents, then performs a Bayesian update for its belief about the state of nature and finally chooses the decision that maximizes its payoff. The key result in [1] and [2] is the possible emergence of the interesting phenomenon of herding, where at some point agents stop relying on their private signals and start following others' actions. We illustrate this situation by the following example.

Consider an infinitely countable number of agents, indexed by $k \in \mathbb{N}$, which sequentially make decisions related to an unknown state $\boldsymbol{\theta}$. There are two possible values for $\boldsymbol{\theta}$: 1 (good state) and 0 (bad state). For instance, if there is a new fashion that agents would like to decide whether to follow it or not, $\boldsymbol{\theta} = 1$ means that the new fashion is good to follow and $\boldsymbol{\theta} = 0$ means otherwise. Each agent k receives a binary signal $\boldsymbol{\xi}_k \in \{0, 1\}$ generated according to the following likelihood function:

$$L(\boldsymbol{\xi}_k = 1|\boldsymbol{\theta} = 1) = L(\boldsymbol{\xi}_k = 0|\boldsymbol{\theta} = 0) = q > \frac{1}{2} \tag{1.1}$$

The observational signal that takes the value of one is labeled as a positive signal, and the signal that takes the value of zero is labeled as a negative signal. Moreover, the signals received by agents are independent. The decision of each agent k is denoted by $\boldsymbol{z}_k \in \{0, 1\}$, where $\boldsymbol{z}_k = 1$ means to adopt, and $\boldsymbol{z}_k = 0$ means to reject. When making decisions, agents choose the action that maximizes their expected payoff, denoted by the following function:

$$u(z_k, \theta) = \begin{cases} 1, & \text{if } z_k = 1 \text{ and } \theta = 1 \\ -1, & \text{if } z_k = 1 \text{ and } \theta = 0 \\ 0, & \text{otherwise} \end{cases} \tag{1.2}$$

All agents start by assuming that the two possible states $\boldsymbol{\theta} = 1$ and $\boldsymbol{\theta} = 0$ are equally likely to occur, i.e., all agents start with the following prior: $\Pr(\boldsymbol{\theta} = 0) = \Pr(\boldsymbol{\theta} = 1) = \frac{1}{2}$. We analyze next what happens given that the first agent has received a positive signal.

Assume that agent 1 receives a positive signal, i.e., $\boldsymbol{\xi}_1 = 1$, then agent 1 updates its belief about the state of nature by computing the following posterior probability:

$$\Pr(\boldsymbol{\theta} = 1|\boldsymbol{\xi}_1 = 1) = \frac{L(\boldsymbol{\xi}_1 = 1|\boldsymbol{\theta} = 1)\Pr(\boldsymbol{\theta} = 1)}{L(\boldsymbol{\xi}_1 = 1|\boldsymbol{\theta} = 1)\Pr(\boldsymbol{\theta} = 1) + L(\boldsymbol{\xi}_1 = 1|\boldsymbol{\theta} = 0)\Pr(\boldsymbol{\theta} = 0)} = q \tag{1.3}$$

Agent 1 has two options, either to accept or to reject. In order to decide, it computes its expected payoff for each possible action and chooses the action that maximizes its expected payoff. If agent 1 decides to accept, i.e., $\boldsymbol{z}_1 = 1$, then its expected payoff is:

$$u(\boldsymbol{z}_1 = 1, \boldsymbol{\theta} = 1)\Pr(\boldsymbol{\theta} = 1|\boldsymbol{\xi}_1 = 1) + u(\boldsymbol{z}_1 = 1, \boldsymbol{\theta} = 0)\Pr(\boldsymbol{\theta} = 0|\boldsymbol{\xi}_1 = 1) = q - (1-q) = 2q - 1. \tag{1.4}$$

Since $q > \frac{1}{2}$, if agent 1 decides $\boldsymbol{z}_1 = 1$, the expected reward is $2q - 1 > 0$. On the other hand, if agent 1 decides to reject, i.e., $\boldsymbol{z}_1 = 0$, its expected payoff is 0. Therefore, agent 1 decides to take the action: $\boldsymbol{z}_1 = 1$.

Assume now that agent 2 also receives a positive signal: $\boldsymbol{\xi}_2 = 1$. Additionally, agent 2 observes the action of agent 1. Since agent 1 decided to accept, agent 2 deduces that agent 1 has received a positive signal, because otherwise agent 1 would not choose to accept. Therefore, according to this assumption, agent 2 updates its belief as follows:

$$\Pr(\boldsymbol{\theta} = 1|\boldsymbol{\xi}_1 = 1, \boldsymbol{\xi}_2 = 1) = \frac{q^2}{q^2 + (1-q)^2} > \frac{1}{2} \text{ (since } q > \frac{1}{2}). \tag{1.5}$$

Similarly to agent 1, agent 2 will also decide to accept, i.e., $\boldsymbol{z}_2 = 1$, since it is the action that maximizes the expected payoff. On the other hand, if agent 2 receives a negative signal, then agent 2 will be indifferent between adopting and rejecting. This is because, the posterior probability of agent 2 will be in this case:

$$\Pr(\boldsymbol{\theta} = 1|\boldsymbol{\xi}_1 = 1, \boldsymbol{\xi}_2 = 0) = \frac{q(1-q)}{q(1-q) + (1-q)q} = \frac{1}{2} \tag{1.6}$$

Then in this case, agent 2 can choose to either adopt or reject.

Agent 3 observes the actions of agents 1 and 2. According to what we have just analyzed, there are two possible scenarios for what agent 3 can observe: agents 1 and 2 have both accepted or agent 1 has accepted while agent 2 has rejected.

First Scenario (Agents 1 and 2 have accepted): if agent 3 observes that both agents have accepted, then agent 3 concludes that agent 1 has definitely received a positive signal and that agent 2 has most probably received a positive signal. Now, we have two cases for the signal of agent 3; if agent 3 receives a positive signal, then after updating its belief, it will also choose $\boldsymbol{z}_3 = 1$, because this action will maximize its payoff. On the other hand, if agent 3 receives a negative signal, then it will also decide on $\boldsymbol{z}_3 = 1$ despite its negative signal. This is because, agent 3 first updates its belief:

$$\Pr(\boldsymbol{\theta} = 1 | \boldsymbol{\xi}_1 = 1, \boldsymbol{\xi}_2 = 1, \boldsymbol{\xi}_3 = 0) = \frac{q^2(1-q)}{q^2(1-q) + (1-q)^2 q} = \frac{q}{q+1-q} = q > \frac{1}{2} \tag{1.7}$$

Therefore by choosing $\boldsymbol{z}_3 = 1$, agent 3 maximizes its expected payoff. If we continue the same reasoning for any agent k where $k \geq 4$, we see that all remaining agents will end up choosing to accept, no matter what private signals they receive. We can similarly show that if the first two agents reject (which can happen when agents 1 and 2 both receive negative signals), all remaining agents will end up rejecting. Therefore, all agents end up deciding on the same action, even though it might not be the correct action that they all should be taking. In other words, if the underlying true state is $\boldsymbol{\theta} = 0$ (bad state), and if the first two agents receive two positive signals (which can happen with probability $(1-q)^2$ in this case), then all agents will end up on deciding to accept rather than to reject.

Second Scenario (Agent 1 *has accepted and agent* 2 *has rejected)*: in this case, agent 3 concludes that agent 1 has definitely received a positive signal and that agent 2 has most probably received a negative signal. In this case, the history of past actions provides no information for agent 3, which will now follow its signal. To see this, if for instance agent 3 receives a positive signal, it updates its belief as follows:

$$\Pr(\boldsymbol{\theta} = 1 | \boldsymbol{\xi}_1 = 1, \boldsymbol{\xi}_2 = 0, \boldsymbol{\xi}_3 = 1) = \frac{q(1-q)q}{q(1-q)q + (1-q)q(1-q)} = q > \frac{1}{2} \tag{1.8}$$

Therefore, agent 3 will choose to accept, i.e., $\boldsymbol{z}_3 = 1$. Otherwise it will choose to decline. For agent 4, similarly, the actions of the first two agents provide no information, by observing the action of agent 3 and its private signals, agent 4 will choose either to accept or to decline. We are then back to the same reasoning we started with. If now agents 3 and 4 accept, all remaining agents will start to accept despite the value of their private signals.

Through this example, we see how the herding phenomenon is possible to arise with fully Bayesian agents, where agents start to imitate each other and ignore their private signal. It is also possible for a mistaken herding to happen, where agents fail to aggregate the dispersed information. In [3], the authors extended these results to the case where private signals are unbounded (i.e., the support of the log likelihood ratio is unbounded). In this case, agents receive a richer set of informational signals, and the authors showed that with unbounded signals, wrong herd does not occur almost surely and agents succeed in asymptotically learning the underlying true state.

Various extensions were considered in [4–8]. In [6], the authors assumed the agents do not observe all past actions, but they randomly sample from past actions. In [7], the authors considered the case where agents observe the actions of a subset of agents rather than the whole network. And in [8], the authors considered a Bayesian model based on communication learning, where Bayesian agents learn by communicating with their neighbors. One of the problems of the fully-Bayesian approach is that it is computationally intractable on the part of agents. This is because it requires from each agent to infer the private information in the network by observing the actions of others, which involves some sophisticated and demanding reasoning on the side of the agents. We next present the second category of social learning models, which require less complex reasoning on the part of agents.

1.1.2 Non-Bayesian Approach

We present in this section some of the non-Bayesian models. We start by the classical DeGroot model and discuss some of its variants that take into consideration the presence of stubborn agents and the spread of misinformation. In these models, agents are assumed to be interconnected through a network topology and each agent communicates with its neighbors. Note that the models of this section assume that the underlying true state θ is not binary as in the example of the previous section, it is instead a real value. Note also that while the beliefs in the previous section are represented as a probability for the possible values of the state, the beliefs in this section are scalar estimates of the true state. The beliefs here can be deterministic or random variable depending on the assumed model. We use boldface for random variables and plain letters for deterministic variables.

DeGroot Model: The DeGroot model [18] is a classical non-Bayesian model based on a simple learning technique where each agent linearly combines the opinions of its neighbors. Consider a network of N agents connected by some graph, indexed by $\mathcal{N} = \{1, 2, \dots, N\}$. Let the scalar $a_{\ell k}$ represent the weight with which agent k scales the data arriving from agent ℓ and, similarly, for $a_{k\ell}$. The weight $a_{\ell k}$ can represent the level of trust that agent k has for the data coming from agent ℓ. Let $\mathcal{N}_k$ denote the neighborhood of agent k, which consists of all agents connected to k. Each agent k scales data arriving from its neighbors in a convex manner, i.e.,

$$a_{\ell k} \geq 0, \quad \sum_{\ell \in \mathcal{N}_k} a_{\ell k} = 1, \quad a_{\ell k} = 0 \text{ if } \ell \notin \mathcal{N}_k \tag{1.9}$$

We collect the weights $\{a_{\ell k}\}$ into an $N \times N$ matrix A. Note that, according to (1.9), the matrix A is left-stochastic so that its spectral radius is equal to one. Assume that there is an underlying unknown true state $\theta \in \mathbb{R}$. Each agent k starts with an estimate of the true state denoted by $x_{k,0} \in \mathbb{R}$, and then at each instant $i > 0$, agent k communicates with its neighbors to update its belief as follows:

$$x_{k,i} = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \, x_{\ell,i-1} \tag{1.10}$$

Note that in the case where the underlying state is discrete, the belief of agent k will be defined as a probability distribution over the possible values of the state and the same update rule (1.10) is used. Let x_i represent the vector of beliefs at time i of all agents, i.e., $x_i = [x_{1,i}\ , \dots ,\ x_{N,i}]^\mathsf{T}$, then from (1.10), we have:

$$x_i = A^\mathsf{T} x_{i-1} \implies x_i = (A^\mathsf{T})^i x_0 \tag{1.11}$$

The objective is to examine when agents reach consensus, i.e., $\lim_{i\to\infty} x_i$ exists and is identical for any agent k. It can be shown that when the network is strongly-connected, i.e., there exists a path with non-zero weights connecting any two agents and, moreover, there is at least one self-loop, i.e., $a_{kk} > 0$ for some agent k, consensus is reached [18,19]. When the network is strongly-connected, matrix A is primitive [41,42], and it then follows from the Perron-Frobenius Theorem [43], [44] that A has a single eigenvalue at one while all other eigenvalues are strictly inside the unit disc. We denote the right-eigenvector of A that corresponds to the eigenvalue at one by y, and all entries of this vector will be strictly positive. We normalize the entries of y to add up to one, so that y satisfies the following conditions:

$$Ay = y, \quad \mathbb{1}^{\mathsf{T}} y = \mathbb{1}, \quad y \succ 0 \tag{1.12}$$

We refer to y as the Perron eigenvector of A. Then, in this case, we have the following result [19]:

$$\lim_{i \to \infty} x_{k,i} = y^{\mathsf{T}} x_0, \text{ for any agent } k \tag{1.13}$$

In other words, all agents will end up with the same belief, which is a weighted average of their initial beliefs. More analysis to the model was provided in [19,20]. Also, time-varying weights were considered in [22,23].

Variants of DeGroot Model: In [24], the authors proposed a variation to DeGroot model, which takes into consideration the presence of forceful agents that can sometimes impose their beliefs without changing their own. Consider again N agents, where each agent k starts with an initial belief or estimate $x_{k,0} \in \mathbb{R}$ about the unknown underlying state $\theta \in \mathbb{R}$. The authors assumed that the information available through the initial beliefs is sufficient to know the underlying state, by assuming the following:

$$\theta = \frac{1}{N} \sum_{k=1}^{N} x_{k,0} \tag{1.14}$$

The agents are assumed to meet according to an asynchronous time model (randomized gossip algorithm studied earlier in [45]), where agents meet at times defined by a rate one Poisson process. At each time slot i, one agent, denoted by k, will be active to meet another agent, denoted by ℓ, where the probability of meeting is $p_{\ell k}$. Agents k and ℓ update their beliefs according to three possibilities:

1. Agents k and ℓ agree on reaching a pairwise consensus, i.e.,

$$\boldsymbol{x}_{k,i} = \boldsymbol{x}_{\ell,i} = \frac{\boldsymbol{x}_{k,i-1} + \boldsymbol{x}_{\ell,i-1}}{2} \tag{1.15}$$

   This possibility happens with probability $\beta_{\ell k}$, labeled as the averaging probability.

2. Agent ℓ influences agent k to change its belief as follows:

$$\boldsymbol{x}_{k,i} = \epsilon \boldsymbol{x}_{k,i-1} + (1-\epsilon)\boldsymbol{x}_{\ell,i-1}, \quad \boldsymbol{x}_{\ell,i} = \boldsymbol{x}_{\ell,i-1} \tag{1.16}$$

   where $\epsilon \in (0, 1/2]$. This possibility happens with probability $\alpha_{\ell k}$, labeled as the influence probability.

3. Agents prefer to stick to their own beliefs,

$$\boldsymbol{x}_{k,i} = \boldsymbol{x}_{k,i-1}, \quad \boldsymbol{x}_{\ell,i} = \boldsymbol{x}_{\ell,i-1} \tag{1.17}$$

   The probability of this possibility is $1 - \beta_{\ell k} - \alpha_{\ell k}$

The advantage of this model over the DeGroot model is that it takes into consideration the presence of influential agents that force others to change their beliefs. Before stating the main result of this work, we list the assumptions considered. First, the probabilities of meeting satisfy the following for any agent k:

$$p_{kk} = 0, \;\; p_{\ell k} \geq 0, \;\; \sum_{\ell=1}^{N} p_{\ell k} = 1 \tag{1.18}$$

Second, consider the directed graph $(\mathcal{N}, \mathcal{E})$, where the edges $\mathcal{E}$ are induced from the probabilities of meeting, i.e., a link from agent ℓ to agent k exists if $p_{\ell k} > 0$. The graph $(\mathcal{N}, \mathcal{E})$ is assumed to be connected, i.e., there exists a path between any two agents in both directions.

Third, we have that:

$$\alpha_{\ell k} + \beta_{\ell k} > 0, \quad \text{for all } (k, \ell) \in \mathcal{E} \tag{1.19}$$

The third assumption is to ensure that forceful agents change their beliefs at some point. Under these assumptions, it was shown in [24] that agents are able to asymptotically reach consensus despite the presence of forceful agents. More specifically, there exists a random variable $\overline{\boldsymbol{x}}$ such that for any agent k,

$$\lim_{i\to\infty} \boldsymbol{x}_{k,i} = \overline{\boldsymbol{x}}, \quad \text{with propbability 1} \tag{1.20}$$

where $\overline{\boldsymbol{x}}$ is a convex combination of the initial beliefs, i.e.,

$$\overline{\boldsymbol{x}} = \sum_{k=1}^{N} \boldsymbol{\pi}_k x_{k,0} \tag{1.21}$$

where the variables $\boldsymbol{\pi}_k$ are random for any k and satisfy: $\boldsymbol{\pi}_k \geq 0$ and $\sum_{k=1}^{N} \boldsymbol{\pi}_k = 1$. It is true that all agents are able to reach an agreement, however they do not effectively aggregate the information from their initial beliefs. However, in the absence of forceful agents ($\alpha_{\ell k} = 0$), it was shown in [24] the following:

$$\lim_{i\to\infty} \boldsymbol{x}_{k,i} = \frac{1}{N} \sum_{k=1}^{N} x_{k,0} = \theta \tag{1.22}$$

In other words, agents are able to find the true state in the absence of forceful agents. More characterization of the results can be found in [24].

Another variant to DeGroot model that also considers the presence of agents that stick to their initial beliefs is analyzed in [25]. Consider a set of N agents, partitioned into K anchors and M sensors ($N = K + M$). Anchor agents are the agents whose beliefs are fixed and the sensor agents are the agents that update their beliefs. Let κ denote the set of anchors and Ω the set of sensors. A sensor k starts with an initial belief denoted by $x_{k,0} \in \mathbb{R}$ and an anchor k starts with an initial belief denoted by $u_{k,0}$. Then sensor agents update synchronously their beliefs as follows:

$$x_{k,i} = \sum_{\ell \in \Omega} d_{\ell k}\, x_{\ell,i-1} + \sum_{\ell \in \kappa} b_{\ell k}\, u_{\ell,0} \tag{1.23}$$

where $d_{\ell k} \in \mathbb{R}$ represents the weight with which agent k scales the data from sensor agent ℓ and $b_{\ell k} \in \mathbb{R}$ is the weight for the data from anchor agent ℓ. If we collect the weights $d_{\ell k}$ into the matrix D and the weights $b_{\ell k}$ into the matrix B, then from (1.23) we have:

$$x_i = D^{\mathsf{T}} x_{i-1} + B^{\mathsf{T}} u_0 \tag{1.24}$$

Then, according to [25], if the spectral radius of D satisfies:

$$\rho(D) < 1 \tag{1.25}$$

then the limiting belief of the sensors is given by:

$$\lim_{i \to \infty} x_i = (I - D^{\mathsf{T}})^{-1} B^{\mathsf{T}} u_0 \tag{1.26}$$

In other words, the beliefs of the sensors converge to a linear combination of the beliefs of the anchors, which captures a leader-follower relationship between sensors and anchors. Note here that the agents do not reach an agreement, as each agent ends up with a different estimate about the underlying state. Involving the presence of anchors in [25] and forceful agents in [24] is used to model the effect of influential agents on the opinion formation. The presence of influential agents was also investigated in [46, 47], wherein influential agents are described as malicious agents that affect the opinion formation by modifying the information they fuse.

Some other models of social learning use tools from statistical physics [48–61] and rely on classical results of particles' interaction, to study the conditions for consensus or opinion fragmentation. For instance, the *voter model* [49, 50] focuses on studying the transition of a collection of agents from a disordered condition to an ordered state (i.e., consensus). In this model, agents are assumed to be in one of two possible states ± 1 and they interact in a pairwise way continuously with time. During each interaction, an agent k is randomly selected along with one of its neighbors ℓ, and then agent k changes its state to the same state of its neighbor ℓ. The questions are whether agents can reach consensus (all agents end up having the same state) and how long it takes for agents to reach consensus. The analysis of this type of models is done by setting up some differential equations that describe the probabilistic evolution of the states of agents, and reveals that consensus depends on the graph structure and the interaction rules between agents. In [51], it was shown that when there is an infinite number of agents placed in a lattice of dimension d, consensus is reached only when $d \leq 2$. On the other hand, if there is a finite number of agents placed in a grid of any dimension [50] or interconnected through a strongly-connected graph [52], consensus is reached. In [53], the voter model is analyzed over different types of networks. A summary of the variants of the voter model can be found in [48].

Another model that also studies the possible emergence of an ordered condition for the agents is the *majority rule model* [54–56]. Similarly to the voter model, agents are assumed to be in one of the two possible states ± 1. The states of agents evolve as follows: at each time i, a group of n_i agents is randomly selected and then agents of the selected group adopt the state of the majority (n_i might differ over time). Again, the focus of the model's analysis is to figure out if agents reach consensus and how fast they reach it. Let p_0 denote the proportion of agents with initial state $+1$. Then according to [55], there exists p_f such that if $p_0 > p_f$, all agents will end up in the state $+1$. On the other hand, if $p_0 < p_f$, then all agents will end up in the state -1. In [54], the majority rule model was studied where the size of the selected group was assumed to be fixed over time, i.e. n_i is constant, and it was shown that in this case consensus time is proportional to $\ln(N)$ where N is the total number of agents.

While the voter and majority rule models focus on modeling simple interactions between agents, another class of models focuses on incorporating the abilities of agents to persuade others. These models correspond to the class of social impact theory. We describe here the model of [57]. Consider a set of N agents, let $\boldsymbol{\sigma}_{k,i} \in \{+1, -1\}$ denote the state of agent k at time i. Each agent k is characterized by two random parameters: $\boldsymbol{p}_k$ and $\boldsymbol{s}_k$. The parameter $\boldsymbol{p}_k$ represents the persuasiveness of agent k, i.e., the ability of agent k to persuade others about a different opinion. The parameter $\boldsymbol{s}_k$ represents the supportiveness of agent k, i.e., the degree at which agent k can support a specific topic. At each time i, the total influence exerted on agent k and denoted by $\boldsymbol{I}_{k,i}$ is defined as follows:

$$\boldsymbol{I}_{k,i} = \sum_{\ell=1}^{N} \frac{\boldsymbol{p}_\ell}{g(d_{\ell k})} (1 - \boldsymbol{\sigma}_{k,i}\boldsymbol{\sigma}_{\ell,i}) - \sum_{\ell=1}^{N} \frac{\boldsymbol{s}_\ell}{g(d_{\ell k})} (1 + \boldsymbol{\sigma}_{k,i}\boldsymbol{\sigma}_{\ell,i}) \tag{1.27}$$

where $d_{\ell k}$ denotes the distance between agents k and ℓ whose value depends on the assumed geometry, and the function $g(.)$ is some increasing function. The first term in (1.27) represents the total influence on agent k from agents of opposite state, while the second term in (1.27) represents the total influence on agent k from agents of same state. Given the total influence $\boldsymbol{I}_{k,i}$, agent k updates its state as follows:

$$\boldsymbol{\sigma}_{k,i+1} = -\text{sgn}\left(\boldsymbol{\sigma}_{k,i}\boldsymbol{I}_{k,i}\right) \tag{1.28}$$

In other words, if the total influence exerted by agents of opposite state is greater than that exerted by agents of same state, agent k switches its state. Otherwise, agent k remains in its state. In [57], the model was analyzed when the graph is complete. In this case, $g(d_{\ell k})$ was set to N and it was shown that the system has infinitely many stationary states and in general those stationary states might not represent a complete consensus. The state's update in (1.28) can be also updated to take into account external influence as follows:

$$\boldsymbol{\sigma}_{k,i+1} = -\text{sgn}\left(\boldsymbol{\sigma}_{k,i}\boldsymbol{I}_{k,i} + \boldsymbol{h}_{k,i}\right) \tag{1.29}$$

where the variables $\{\boldsymbol{h}_{k,i}\}$ are random variables independent across time and agents, and represent any source of influence other than the social impact, e.g., media and celebrities. It was shown in [57] that with the presence of noise, the only stationary state is the system's state that is nearly uniform in opinion (almost all agents are in the same state).

Not all models considered a binary state for agents. Other models also assumed the case where the opinion of agents can be continuous. For instance, in the Deffuant model [58, 59], each agent k starts with an initial opinion $x_{k,0} \in [0, 1]$. At each instant i, two agents k and ℓ are randomly selected to interact as follows. If the difference in their opinion is less than a threshold ϵ, i.e., if $|x_{k,i} - x_{\ell,i}| < \epsilon$, they do not update their opinions. Otherwise, they update their opinions as follows:

$$\begin{aligned} x_{k,i+1} &= x_{k,i} + \mu\left(x_{\ell,i} - x_{k,i}\right) \\ x_{\ell,i+1} &= x_{\ell,i} + \mu\left(x_{k,i} - x_{\ell,i}\right) \end{aligned} \tag{1.30}$$

where $\mu \in [0, 0.5]$. The results on Deffuant dynamics were mostly derived through numerical simulations [59], where it was observed that, for large values of threshold ϵ, consensus is reached where the opinions of all agents end up to be the average of their initial opinions. On the other hand, for lower values of threshold several clusters of different opinions are observed. A model similar to Deffuant model was proposed in [60], where each agent k communicates with its neighbors and updates its opinion as follows:

$$x_{k,i+1} = \frac{\sum_{\ell:|x_{k,i}-x_{\ell,i}|<\epsilon} a_{\ell k} x_{\ell,i}}{\sum_{\ell:|x_{k,i}-x_{\ell,i}|<\epsilon} a_{\ell k}} \tag{1.31}$$

Numerical simulations were also carried out to analyze the model and similar observations was also found with this model. In particular, for large value of threshold ϵ one cluster of agents with same opinion is observed. For lower values of threshold, more clusters are observed. A summary of more models for opinion formation based on statistical physics can be found in [48].

To sum up, non-Bayesian models provide a less complex approach than that of Bayesian models and are widely used for studying information manipulation. However, in these models, the only source of information is assumed to be in the starting beliefs of the agents, or the focus is on examining different forms of interactions between agents and what might result from them. We present next a set of models that do not belong to a fully-Bayesian approach, but do involve a Bayesian update step that takes into consideration the continuous flow of new information.

1.2 Non-Bayesian Learning with Continuous Information Flow

Classical non-Bayesian models focus only on modeling the interactions between networked agents. In the work [26], the authors considered incorporating the private information of non-Bayesian agents into their process of learning. In this new class of non-Bayesian models, agents do not only consult with their neighbors about the state of nature, but also continuously receive private observations related to the true state.

1.2.1 Consensus Non-Bayesian Learning

Consider a set of N agents indexed by $\mathcal{N} = \{1, 2 \cdots, N\}$. Similarly to the notation we used in presenting the DeGroot model, we denote by $a_{\ell k}$ the weight with which agent k scales data from agent ℓ, and that these weights satisfy (1.9). We denote by Θ the finite set of all possible values for the state of nature. Let $\theta^\circ \in \Theta$ represent the underlying true state that is unknown for agents. At each instant i, each agent k receives an observational signal $\boldsymbol{\xi}_{k,i}$ generated according to a likelihood function denoted by $L_k(.|\theta^\circ)$. Each agent starts with a prior belief about the state of nature, modeled as a probability distribution over Θ. We denote the prior belief of agent k at any $\theta \in \Theta$ by: $\boldsymbol{\mu}_{k,0}(\theta)$.

In [26], the authors proposed a consensus-type construction to update the agents' beliefs. In this construction, instead of combining the opinions of the neighbors in a fully Bayesian manner, each agent follows the Bayes' rule to obtain an intermediate belief (using it private signal) and subsequently combines it with the old beliefs of its neighbors (through convex combination). More specifically, at instant i agent k updates its belief according to following rule:

$$\begin{cases} \boldsymbol{\psi}_{k,i}(\theta) = \dfrac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \boldsymbol{\mu}_{k,i-1}(\theta') L_k(\boldsymbol{\xi}_{k,i}|\theta')} \quad \text{(intermediate beleif)} \\ \\ \boldsymbol{\mu}_{k,i}(\theta) = a_{kk} \boldsymbol{\psi}_{k,i}(\theta) + \sum\limits_{\ell \in \mathcal{N}_k, \ell \neq k} a_{\ell k}\, \boldsymbol{\mu}_{\ell,i-1}(\theta) \end{cases} \tag{1.32}$$

The first step consists of the Bayesian update and the second step is the aggregation step. Since this approach does not belong to a fully Bayesian approach, the authors classified it as non-Bayesian social learning. Under some technical assumptions related to the structure of the network and the information provided by the observational signals, it was shown in [26] that agents following this model can asymptotically learn the true state, i.e., for any agent k, we have:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta^\circ) \overset{a.s.}{=} 1 \tag{1.33}$$

The assumptions taken are to ensure the strong connectivity of the network and that the observational signals are informative for the whole network. In [27,28], the authors provided further analysis to the consensus-based model, by showing that agents (asymptotically) learn the true state exponentially fast and that the rate of learning depends on the relative entropy of agents' signal structures and their eigenvector centralities. More specifically, at instant i, the level of uncertainty across the agents can be measured by:

$$\boldsymbol{e}_i \triangleq \frac{1}{2}\sum_{k=1}^{N} ||\boldsymbol{\mu}_{k,i}(.) - \mathbb{1}_{\theta^\circ}(.)||_1 = \sum_{k=1}^{N}\sum_{\theta\neq\theta^\circ} \boldsymbol{\mu}_{k,i}(\theta) \tag{1.34}$$

which represents the total variation between the belief $\boldsymbol{\mu}_{k,i}(.)$ at instant i and the ultimate belief $\mathbb{1}_{\theta^\circ}(.)$ (where $\mathbb{1}_{\theta^\circ}(.)$ denote a vector whose entries are zero at any $\theta \neq \theta^\circ$ and one at $\theta = \theta^\circ$). Then, according to [27], the following results were shown for the consensus-based model:

- Let

$$\gamma \triangleq \lim_{i\to\infty} \frac{1}{i} \sup \log \boldsymbol{e}_i \tag{1.35}$$

then γ is finite and $\gamma < 0$;

- The rate of learning $|\gamma|$ is upper-bounded as follows:

$$|\gamma| \leq \alpha \min_{\theta\neq\theta^\circ} \sum_{k=1}^{N} y_k D_{KL}(L_k(.|\theta^\circ)||L_k(.|\theta)) \tag{1.36}$$

where y_k is the $k-$th element of the Perron eigenvector y of matrix A, α is the weight to all self-loops in the network, i.e., $a_{kk} = \alpha$ for all agents k, and $D_{KL}(L_k(.|\theta^\circ)||L_k(.|\theta))$ is the Kullback-Leibler divergence between $L_k(.|\theta^\circ)$ and $L_k(.|\theta)$.

In [29], the authors studied the consensus-based model of [26] over a specific type of time-varying undirected graphs, where instead of assuming the fixed weights in A, they considered the following time-varying weight matrix:

$$A(i) = (1 - \eta(i))I + \eta(i)A \tag{1.37}$$

where I is the identity matrix and the parameter $\eta(i) \in (0, 1]$. As in [26], under some technical assumptions, it was shown in [29] that agents learn the true state asymptotically almost surely.

1.2.2 Diffusion Non-Bayesian Learning

An alternative to the consensus mechanism was proposed in [30] by relying on diffusion strategies, due to their enhanced performance and stability range [41, 62]. In the diffusion-based model, each agent combines its intermediate belief with the updated (rather than old) beliefs of its neighbors in a convex manner. More specifically, each agent k updates its belief according to the following update rule:

$$\begin{cases} \boldsymbol{\psi}_{k,i}(\theta) = \dfrac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \boldsymbol{\mu}_{k,i-1}(\theta') L_k(\boldsymbol{\xi}_{k,i}|\theta')} \\ \boldsymbol{\mu}_{k,i}(\theta) = \sum\limits_{\ell \in \mathcal{N}_k} a_{\ell k}\, \boldsymbol{\psi}_{\ell,i-1}(\theta) \end{cases} \tag{1.38}$$

Results in [30] established that agents are also able to asymptotically learn the underlying state under the diffusion strategy. We are going to discuss this model and its results in greater detail in Chapter 2. In [31], the authors proposed a variation to the diffusion-based learning model, where each agent averages the log beliefs of its neighbors, instead of using convex combination as in (1.38) of [30]. In other words, the second step of the model of [31] is defined as:

$$\boldsymbol{\mu}_{k,i}(\theta) = \frac{\exp\left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k}\, \log \boldsymbol{\psi}_{\ell,i-1}(\theta)\right)}{\sum_{\theta' \in \Theta} \exp\left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k}\, \log \boldsymbol{\psi}_{\ell,i-1}(\theta')\right)} \tag{1.39}$$

The authors of [31] showed the exponentially fast convergence of agents' beliefs to the true state with probability one.

1.2.3 Various Non-Bayesian Models

Other models for non-Bayesian social learning were proposed [32–40]. For instance, in [33], the authors proposed a non-Bayesian model where, at each time, each agent selects randomly one of its neighbors to communicate with. Let $\boldsymbol{\sigma}_{k,i} \in \mathcal{N}_k$ represent the index of the neighbor that agent k selects at time i. Then agent k uses the past belief of the selected neighbor as a prior, and updates its own belief as follows:

$$\boldsymbol{\mu}_{k,i}(\theta) = \frac{\boldsymbol{\mu}_{\boldsymbol{\sigma}_{k,i},i-1}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \boldsymbol{\mu}_{\boldsymbol{\sigma}_{k,i},i-1}(\theta') L_k(\boldsymbol{\xi}_{k,i}|\theta')} \tag{1.40}$$

Results in [33] established that agents learn the true state asymptotically almost surely, given that the observational signals are globally informative and that the network is connected.

Optimization Characterization of Bayesian update: In [34], the authors used an optimization characterization for the Bayesian update to motivate their proposed model. More specifically, the Bayes update given by:

$$\boldsymbol{\psi}_{k,i}(\theta) = \frac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \boldsymbol{\mu}_{k,i-1}(\theta') L_k(\boldsymbol{\xi}_{k,i}|\theta')} \tag{1.41}$$

can be shown yo be the closed-form solution to the following optimization problem [63]:

$$\arg \min_{\psi(.) \in \Omega} \left\{ D_{KL}(\psi(.)||\boldsymbol{\mu}_{k,i-1}(.)) - \sum_{\theta \in \Theta} \log \big(L_k(\boldsymbol{\xi}_{k,i}|\theta) \big) \, \psi(\theta) \right\} \tag{1.42}$$

where Ω is the probability simplex given by:

$$\Omega = \{\psi(.) \mid \sum_{\theta \in \Theta} \psi(\theta) = 1, \psi(\theta) \geq 0\} \tag{1.43}$$

and $D_{KL}(\psi(.)||\boldsymbol{\mu}_{k,i-1}(.))$ is the Kullback-Leibler divergence between $\psi(.)$ and $\boldsymbol{\mu}_{k,i-1}(.)$. In other words, the posterior distribution obtained by performing the Bayesian update is the solution to the problem of maximizing the log-likelihood function of the observed signal, regularized by its Kullback-Leibler divergence from the prior. Using this characterization of the Bayesian update, the authors in [34] suggested a non-Bayesian learning model inspired by Nesterov's dual averaging method, which we describe here. In this work, the agents are assumed to communicate through a gossip scheme [45] similarly to [24], and the underlying network is assumed to be undirected. During each instant i defined by a rate one Poisson process, two agents meet to aggregate their accumulated observations by sharing the log-likelihood of their observational signals. Let $\boldsymbol{z}_{k,i}(\theta) \in \mathbb{R}$ denote the log-likelihood of the accumulated observational signals for agent k up to time $i-1$ given any $\theta \in \Theta$, and $\boldsymbol{z}_{k,0}(\theta) = 0$ for any agent k and any θ. Suppose that at i, agent k is active and meets agent ℓ with probability $p_{\ell k}$, and together they update $\boldsymbol{z}_{k,i-1}(\theta)$ and $\boldsymbol{z}_{\ell,i}(\theta)$ as follows:

$$\begin{aligned} \boldsymbol{z}_{k,i}(\theta) &= \frac{\boldsymbol{z}_{k,i-1}(\theta) + \boldsymbol{z}_{\ell,i-1}(\theta)}{2} + \log \big(L_k(\boldsymbol{\xi}_{k,i-1}|\theta) \big) \\ \boldsymbol{z}_{\ell,i}(\theta) &= \frac{\boldsymbol{z}_{k,i-1}(\theta) + \boldsymbol{z}_{\ell,i-1}(\theta)}{2} + \log \big(L_\ell(\boldsymbol{\xi}_{\ell,i-1}|\theta) \big) \end{aligned} \tag{1.44}$$

At the same time instant, any other agent $m \notin \{\ell, k\}$ updates $\boldsymbol{z}_{m,i-1}(\theta)$ as follows:

$$\boldsymbol{z}_{m,i}(\theta) = \boldsymbol{z}_{m,i-1}(\theta) + \log\left(L_m(\boldsymbol{\xi}_{m,i-1}|\theta)\right) \tag{1.45}$$

Then, all agent k update their belief by performing the following Bayes' like step:

$$\boldsymbol{\mu}_{k,i}(.) = \arg\min_{\psi(.)\in\Omega}\left\{D_{KL}(\psi(.)||\mu_{k,0}(.)) - \sum_{\theta\in\Theta}\boldsymbol{z}_{k,i}(\theta)\psi(\theta)\right\} \tag{1.46}$$

It was shown in [34] that the beliefs of agents converge in the probability sense to an impulse of size one at the location $\theta = \theta^\circ$. Additionally, the learning happens exponentially fast with high probability. More specifically, it was shown that for large enough i and for any $\epsilon > 0$, we have the following for any agent k:

$$|\boldsymbol{\mu}_{k,i}(\theta^\circ) - 1| \leq \mathcal{K}\exp\left[\left(-\min_{\theta\neq\theta^\circ} D(\theta) + \epsilon\right) i\right] \tag{1.47}$$

with probability at least $1 - \dfrac{C}{\epsilon^2 i}$, for some constants $C > 0$ and $\mathcal{K} > 0$, and

$$D(\theta) = \frac{1}{N}\sum_{k=1}^{N} D_{KL}(L_k(.|\theta^\circ)||L_k(.|\theta)). \tag{1.48}$$

The authors in [35] extended the work of [34] to directed networks, where the interactions between agents are captured through a weight matrix A. In this work, the agents synchronously share the log-likelihood of the observed signals with their neighbors, and update their beliefs as follows:

$$\boldsymbol{\phi}_{k,i}(\theta) = \sum_{\ell\in\mathcal{N}_k} a_{\ell k}\boldsymbol{\phi}_{\ell,i-1}(\theta) + \log\left(L_k(\boldsymbol{\xi}_{k,i}|\theta)\right) \tag{1.49}$$

$$\boldsymbol{\mu}_{k,i}(.) = \arg\min_{\psi(.)\in\Omega}\left\{\frac{1}{\eta}D_{KL}(\psi(.)||\mu_{k,0}(.)) - \sum_{\theta\in\Theta}\boldsymbol{\phi}_{k,i}(\theta)\psi(\theta)\right\} \tag{1.50}$$

where $\boldsymbol{\phi}_{k,i}(\theta)$ represents, for agent k at time i, the combined log-likelihoods of observational signals given a state θ, and $\boldsymbol{\phi}_{k,0}(\theta) = 0$ for any agent k. Moreover, the parameter η is non-negative and all agents are assumed to start with a uniform prior belief. In [35], the authors provide a finite-time analysis for the model, and show that agents asymptotically learn the true state almost surely [35], under some technical conditions.

In [36–38], the authors assumed the problem of finding the best hypothesis that explains the observations received by the network's agents. More specifically, the observational signals are generated according to an unknown function denoted by f_k for each agent k, and it is not required that there exists $\theta \in \Theta$ such that $L(k(.|\theta)) = f_k$. The agents are assumed to interact over a time-varying network, where at each time i, the interaction between agents is captured by a weight matrix $A(i)$ that is doubly-stochastic and all its diagonal entries are strictly positive. The graphs need not to be strongly-connected at each time but it is assumed that there exists $B \geq 1$ such that the graph union of any sequence of B graphs is strongly-connected. The agents update their beliefs according to the following update rule:

$$\boldsymbol{\mu}_{k,i}(\theta) = \frac{1}{\boldsymbol{Z}_{k,i}} \prod_{\ell=1}^{N} L_k(\boldsymbol{\xi}_{k,i}|\theta)^{\boldsymbol{\beta}_{k,i}} \, \boldsymbol{\mu}_{\ell,i-1}(\theta)^{a_{\ell k}(i-1)} \tag{1.51}$$

where $Z_{k,i}$ is a normalization factor, and $\{\boldsymbol{\beta}_{k,i}\}$ are IID Bernoulli random variables where $\beta_{k,i} = 1$ means that agent k received an observational signal at time i and $\beta_{k,i} = 0$ means that agent k failed to receive an observational signal at time i. The proposed update rule is shown to be the solution to the following Bayes like step:

$$\boldsymbol{\mu}_{k,i}(.) = \arg\min_{\psi(.)\in\Omega} \left\{ \sum_{\ell=1}^{N} a_{\ell k}(i-1) \, D_{KL}\left(\psi(.)||\boldsymbol{\mu}_{\ell,i-1}(.)\right) - \boldsymbol{\beta}_{k,i} \sum_{\theta\in\Theta} \psi(\theta) \log\left(L_k(\boldsymbol{\xi}_{k,i}|\theta)\right) \right\} \tag{1.52}$$

The objective is for agents to learn the set of states that best explain the observational signals. More specifically, let Θ^* denote the optimal hypothesis set defined as follows:

$$\Theta^* = \arg\min_{\theta\in\Theta} \sum_{k=1}^{N} b_k \, D_{KL}(f_k||L_k(.|\theta)) \tag{1.53}$$

where b_k is the mean of the IID Bernoulli random variables $\{\beta_{k,i}\}$ for agent k. It was then shown in [36] that under some technical assumptions agents can asymptotically learn the optimal hypothesis set, i.e., for any agent k ,

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) \overset{a.s.}{=} 0, \quad \text{for any } \theta \notin \Theta^* \tag{1.54}$$

A geometric non-asymptotic characterization of the learning rate was also provided in [36].

Non-Bayesian Learning with Faulty Agents: In [39], the authors analyzed a non-Bayesian model in the presence of some faulty agents, i,e., agents that send incorrect, inconsistent or empty messages. The non-faulty agents do not know the identity of faulty agents but know that at most f agents are faulty. At time i, a non-faulty agent k sends its log-belief $\log(\boldsymbol{\mu}_{k,i-1}(\theta))$ to its neighbors and receive through its incoming links the log-beliefs of other agents. Since some of the log-beliefs received by agent k might be faulty, to mitigate this effect, agent k follows an algorithm that allows agent k to trim away any extreme messages and then combine a subset of the received log-beliefs into an updated log-belief denoted by $\boldsymbol{\eta}_{k,i}(\theta)$. The details of the algorithm can be found in [39]. Agent k then computes its final belief $\boldsymbol{\mu}_{k,i}(\theta)$ as follows:

$$\begin{aligned} \boldsymbol{L}^{\theta}_{k,1:i} &= L_k(\boldsymbol{\xi}_{k,i}|\theta)\ \boldsymbol{L}^{\theta}_{k,1:i-1} \\ \boldsymbol{\mu}_{k,i}(\theta) &= \frac{\boldsymbol{L}^{\theta}_{k,1:i}\ \exp(\boldsymbol{\eta}_{k,i}(\theta))}{\sum_{\theta'\in\Theta} \boldsymbol{L}^{\theta'}_{k,1:i}\ \exp(\boldsymbol{\eta}_{k,i}(\theta'))} \end{aligned} \tag{1.55}$$

where $L^{\theta}_{k,1:i}$ denotes for agent k the likelihood of the cumulative observations up to time i given θ (instead of the likelihood of the current observation only). Under some technical assumptions, agents are shown to asymptotically learn the true state almost surely.

A generic framework that further treats non-Bayesian social learning can be found in [40]. We have seen different models suggested for studying non-Bayesian learning in the presence of continuous flow of information. The focus was to show the ability of agents to learn the true underlying state. However, this might not be always the case as agents might get exposed to some external influence and end up disagreeing on the state of nature. We have seen models from classical non-Bayesian learning [24,25] that studied the presence of forceful or anchor agents and their effect on consensus reach. In this work, we focus on studying similar phenomena of external influence when non-Bayesian agents learn with continuous flow of information.

1.3 Diffusion Learning over Weak Graphs

In this dissertation[1], we are going to focus on the mechanism proposed in [30], which relies on diffusion strategies due to their enhanced performance and stability ranges, especially in scenarios that involve continuous learning [41, 62]. The models of social interaction studied in [26,30] assume *strongly-connected* graphs whereby a path with positive weights connecting any two agents is always possible and at least one agent has a self-loop. Over such graphs, social influences diffuse over time and all agents are able to learn asymptotically the true state of the environment. This is possible even when the local observations at the agents may be of varying quality with some agents being more informed than others.

1.3.1 Weakly-Connected Networks

In this work, we examine social learning over *weakly-connected* graphs, as opposed to strongly-connected graphs. Over a weak topology, there exist some select edges over which information flows in one direction only, with information never flowing back from the receiving agents to the originating agents. This scenario is common in practice, especially over social networks. For example, in Twitter networks, it is not unusual for some influential agents (e.g., celebrities) to have a large number of followers, while the influential agent itself may not consult information from most of these followers. A similar effect arises when social networks operate in the presence of stubborn agents [24, 25, 67]; these agents insist on their opinion regardless of the evidence provided by local observations or by neighboring agents. It turns out that weak graphs influence the evolution of the agents' beliefs in a critical manner. The objective of this work is to clarify this effect, its origin, and to quantify its implications by means of

[1]The material in this dissertation is based on work published in [64–66].

closed-form expressions.

1.3.2 Social Disagreement

In the previous works [67,68], the authors examined the influence of weak graphs on the solution of distributed *inference* problems, where agents are interested in learning a parameter of interest that minimizes an aggregate cost function. It was shown there that a leader-follower relationship develops among the agents with the performance of some agents being fully controlled by the performance of other agents. In the different context of social learning, this type of weak connectivity was briefly discussed in [69] where consensus social learning was analyzed over non-strongly connected networks. This work considered only the special case in which all agents in the network are interested in the *same* state of nature. A richer and more revealing dynamics arises when different clusters within the network monitor different state variables.

For example, consider a situation in which a weak graph consists of four sub-graphs (see future Fig. 3.1): the two top graphs are strongly-connected while the other two are weakly-connected to them. In this case, each of the first two sub-graphs is able to learn its truth asymptotically. However, the agents in the lower sub-graphs will be shown to reach a state of disarray in relation to their true state, with different agents reaching in general different conclusions and, moreover, with each of these conclusions being directly determined by the separate states of the two top sub-graphs. In this work we carry out a detailed analysis to show how influential agents dictate the performance of weak components in the network, and arrive at closed-form expressions that describe this influence in analytical form (suitable for subsequent design purposes). We will find that, under some conditions, non-influential agents will be forced to adopt beliefs centered around the true states of the influential agents. This situation is similar to the leader-follower relationship discussed in [67,68] in the context of decentralized inference and continuous adaptation. We will also find that these beliefs differ from one agent to another, which results in a disturbing form of social disagreement.

In some applications, the influential agents may be malicious as in [46, 47]. In contrast to these works, in our development, influential agents do not alter the information they are fusing, but the nature of what they are sending need not be consistent with the true state of the receiving agents. Moreover, our assumed model takes into consideration not only the interaction between agents, but also the information continuously received by each agent. We are going to find out how the quality of this information as well as the network's structure affect the learning ability of agents.

1.3.3 Enhancing Self-Awareness

Motivated by the results in the next sections, we will also incorporate an element of self-awareness into the social learning process of the network through the introduction of a scaling factor — see Eq. (3.58). This factor will enable agents in the network to assign more or less weight to their local information in comparison to the information received from their neighbors. This variation helps infuse into the network some elements of human behavior. For example, in an interactive social setting, a human agent may not be satisfied or convinced by an observation and prefers to give more weight to their prior belief based on accumulated experiences. This mode of operation was studied for *single* stand-alone agents in [70, 71] and was studied there as a mechanism for self-control. We will instead examine the influence of self-awareness in the challenging network setting, where the behavior of the various agents are coupled together. In particular, we will show that self-awareness helps agents converge towards a fixed belief distribution, rather than have their beliefs exhibit an undesired oscillatory behavior, which reflects their inability to settle on a decision — see Fig. 3.3.

1.3.4 Belief Control Strategies

Using the expressions that describes the effect of influential agents, we will analyze the control mechanism more closely. We have three main contributions. First, we show that the internal graph structure connecting the receiving agents imposes a form of resistance to manipulation, but only to a certain degree. Second, we characterize the set of states that can be imposed on receiving networks; while this set is large, it turns out that it is not unlimited. And, third, for any attainable state, we develop a control mechanism that allows sending agents to force the receiving agents to reach that state and behave in that manner.

1.4 Dissertation Outline and Contributions

As mentioned earlier, the dissertation is focused on analyzing the model of diffusion learning over weak graphs. In particular, we study how the asymmetric flow of information over this type of network affects the learning abilities of some agents. We next outline the dissertation and summarize the main contributions.

In Chapter 2, we review the model of diffusion social learning over strong graphs. We explain in details the components of the model, clarify the assumptions for the agents to learn the true state and summarize the results obtained.

In Chapter 3, we present weak graphs which consist of two types of sub-networks: sending and receiving sub-networks. We show that when the observational signals of receiving agents are not informative, receiving agents will not be able to find their own true state and their limiting beliefs will be concentrated around the true states of the sending agents. We also provide closed-form expressions for the limiting beliefs. We then consider a variation to the diffusion learning model to enable agents to give less or more weights to their observational signals. We show in this case that the total influence scheme can still occur even when the observational signals of receiving agents are informative.

In Chapter 4, we analyze the control mechanism more closely by exploring the expression of the limiting beliefs. We find that there are some limitations to what sending agents can control and provide design procedures to achieve a specific achievable control scheme.

CHAPTER 2

Diffusion Learning over Strong Networks

We first review strongly-connected networks and summarize the results derived earlier in [30] for this graph topology. Then, we explain in subsequent chapters how the results are affected when the underlying topology happens to be weak and show how a leader-follower relationship develops. We characterize in some detail the limiting behavior of this relation and identify the factors that influence the ability of the social agents to learn the truth or to follow other influential agents.

2.1 Network Model

Thus, consider a network of N agents connected by some graph. Let $\mathcal{N} = \{1, 2, \ldots, N\}$ denote the indexes of the agents in the network. We assign a pair of non-negative weights, $\{a_{k\ell}, a_{\ell k}\}$, to the edge connecting any two agents k and ℓ. The scalar $a_{\ell k}$ represents the weight with which agent k scales the data arriving from agent ℓ and, similarly, for $a_{k\ell}$ – see Fig. 2.1. The network is said to be strongly-connected if there exists a path with non-zero weights connecting any two agents and, moreover, there is at least one self-loop, i.e., $a_{kk} > 0$ for some agent k. Let $\mathcal{N}_k$ denote the neighborhood of agent k, which consists of all agents connected to k. Each agent k scales data arriving from its neighbors in a convex manner, i.e.,

$$a_{\ell k} \geq 0, \quad \sum_{\ell \in \mathcal{N}_k} a_{\ell k} = 1, \quad a_{\ell k} = 0 \text{ if } \ell \notin \mathcal{N}_k \tag{2.1}$$

We collect the weights $\{a_{\ell k}\}$ into an $N \times N$ matrix A. From condition (2.1), A is a left-stochastic matrix so that its spectral radius is equal to one, $\rho(A) = 1$. Since the network is strongly-connected, A is also a primitive matrix [41]. It then follows from the Perron-

Frobenius Theorem [43], [44] that A has a single eigenvalue at one while all other eigenvalues are strictly inside the unit disc. We denote the right-eigenvector of A that corresponds to the eigenvalue at one by y, and all entries of this vector will be strictly positive. We normalize the entries of y to add up to one, so that y satisfies the following conditions:

$$Ay = y, \quad \mathbb{1}^\mathsf{T} y = \mathbb{1}, \quad y \succ 0 \tag{2.2}$$

We refer to y as the Perron eigenvector of A. This network structure plays an important role in diffusing information across the network and helps agents in learning the true state. We describe next the mechanism of this learning.

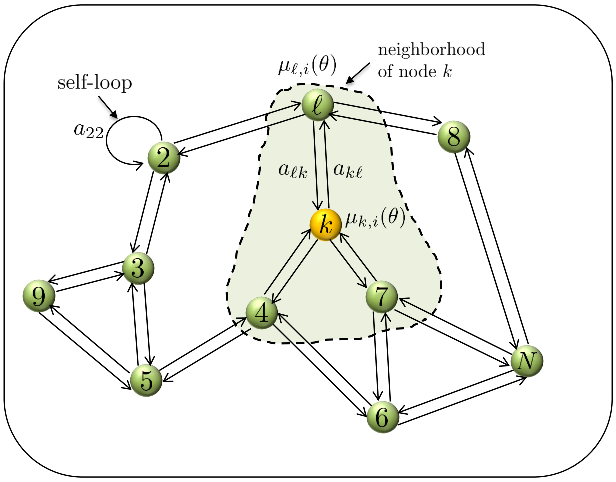

Figure 2.1: An example of a strongly-connected network where $\mu_{k,i}(\theta)$ denotes the belief (pdf) of agent k at time i.

2.2 Diffusion Social Learning

Let Θ denote a finite set of all possible events that can be detected by the network. Let $\theta^\circ \in \Theta$ denote the *unknown* true event that has happened, while the other elements in Θ represent possible variations of that event. The objective of the network is to learn the true state, θ°. For this purpose, agents will be continually updating their beliefs about the true state through a localized cooperative process. Initially, at time $i = 0$, each agent k starts from some prior belief, denoted by the function $\mu_{k,0}(\theta) \in [0, 1]$. This function represents the probability distribution over the events $\theta \in \Theta$. For instance, if $\theta_1 \in \Theta$ then

$$\mu_{k,0}(\theta_1) = \text{Prob}(\boldsymbol{\theta} = \theta_1), \qquad \text{at time } i = 0 \tag{2.3}$$

For subsequent time instants $i \geq 1$, the private belief of agent k is denoted by $\mu_{k,i}(\theta) \in [0, 1]$. All beliefs across all agents must be valid probability measures over Θ. That is, they must obey the normalization:

$$\sum_{\theta \in \Theta} \mu_{k,i}(\theta) = 1, \qquad \text{for any } i \geq 0 \text{ and } k \in \mathcal{N} \tag{2.4}$$

Figure 2.2 presents an example of a belief distribution $\mu_{k,i}(\theta)$ defined over $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$. The agents will update their private beliefs $\{\mu_{k,i}(\theta)\}$ over time based on the private signals they observe from the environment and the information shared by their social neighbors. We assume that, at each time $i \geq 1$, every agent k observes a realization of some signal, $\boldsymbol{\xi}_{k,i}$, whose probability distribution is dependent on the true event θ^o, namely, the process $\{\boldsymbol{\xi}_{k,i}\}$ is generated according to some known likelihood function $L_k(\cdot|\theta^\circ)$ – see Fig.2.3. We further assume that for each agent k, the signals $\{\boldsymbol{\xi}_{k,i}\}$ belong to a *finite* signal space denoted by Z_k and that these signals are independent over time.

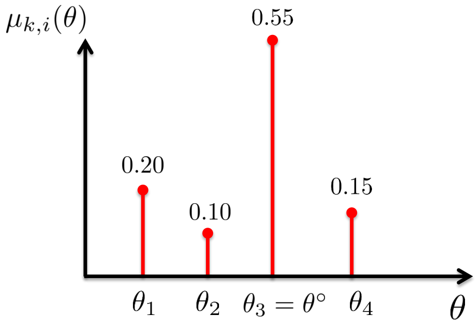

Figure 2.2: An example of a belief distribution $\mu_{k,i}(\theta)$.

Diffusion social learning, described in [30], provides a mechanism by which agents can process the information they receive from their private signals and from their neighbors. A consensus-based strategy can also be employed, as was done in [26]. We focus on the diffusion strategy due to its enhanced performance, as observed in [30] and as further explained in the treatments [41,62]. In diffusion learning, at every time $i \geq 1$, each agent k first updates its belief, $\mu_{k,i-1}(\theta)$, based on its observed private signal $\xi_{k,i}$ by means of the Bayesian rule:

$$\psi_{k,i}(\theta) = \frac{\mu_{k,i-1}(\theta) L_k(\xi_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \mu_{k,i-1}(\theta') L_k(\xi_{k,i}|\theta')} \tag{2.5}$$

This step leads to an intermediate belief $\psi_{k,i}(\theta)$. After learning from their observed signals, agents can then learn from their social neighbors through cooperation to compute:

$$\mu_{k,i}(\theta) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \, \psi_{\ell,i}(\theta) \tag{2.6}$$

Subsequently, agent k can use its updated belief, $\mu_{k,i}(\theta)$, to predict the probability of a certain signal $\zeta_k \in Z_k$ occurring in the next time instant $i+1$. This prediction or forecast is based on the following calculation:

$$m_{k,i}(\zeta_k) \triangleq \sum_{\theta \in \Theta} \mu_{k,i}(\theta) L_k(\zeta_k|\theta) = \text{Prob}(\boldsymbol{\xi}_{k,i+1} = \zeta_k) \tag{2.7}$$

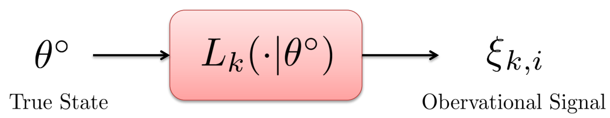

Figure 2.3: Generation of observational signals.

In the sequel, we will be interpreting the diffusion learning model as a *stochastic* system of interacting agents, especially since the operation of this mechanism is driven by the random observational signals. Thus, we rewrite (2.5) and (2.6) as follows by using boldface letters to refer to random variables.

$$\begin{cases} \boldsymbol{\psi}_{k,i}(\theta) = \dfrac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \boldsymbol{\mu}_{k,i-1}(\theta') L_k(\boldsymbol{\xi}_{k,i}|\theta')} \\ \\ \boldsymbol{\mu}_{k,i}(\theta) = \displaystyle\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \, \boldsymbol{\psi}_{\ell,i}(\theta) \end{cases} \tag{2.8}$$

2.3 Correct Forecasting

When agents in strongly-connected networks follow model (2.8) to update their beliefs, the agents will eventually learn the truth according to the results established in [30]. The argument there is based on an identifiability condition similar to the one used in [26], and which is motivated as follows. We assume first that the agents' private signals $\{\boldsymbol{\xi}_{k,i}\}$ do not hold enough information about the true state, so that individual agents cannot rely solely on their observations to identify θ° and are motivated to cooperate. More specifically, this requirement amounts to assuming that each agent k has a subset of states $\Theta_k \subseteq \Theta$ for which:

$$L_k(\zeta_k|\theta) = L_k(\zeta_k|\theta^\circ), \quad \theta \in \Theta_k \tag{2.9}$$

for any $\zeta_k \in Z_k$. We refer to Θ_k as the set of indistinguishable states for agent k. We subsequently assume that through cooperation with their neighbors, agents are able to identify the true state by imposing the identifiability condition:

$$\bigcap_{k\in\mathcal{N}} \Theta_k = \{\theta^\circ\} \tag{2.10}$$

We refer to this case as θ° being globally identifiable. To prove that agents are able to learn the true state, the analysis in [30] is based on first showing that agents are able to learn the correct distribution of incoming signals.

Lemma 1 (Correct Forecasting [30])**.** *Assume that there exists at least one agent with a positive prior belief about the true state* θ°*, i.e.,* $\boldsymbol{\mu}_{k,0}(\theta^\circ) > 0$ *for some* $k \in \mathcal{N}$*. Then, agents are able to correctly predict the distribution of the incoming signals, namely, for any* $\zeta_k \in Z_k$ *and* $k \in \mathcal{N}$*:*

$$\lim_{i\to\infty} \boldsymbol{m}_{k,i}(\zeta_k) \overset{a.s.}{=} L_k(\zeta_k|\theta^\circ) \tag{2.11}$$

where $\overset{a.s.}{=}$ *denotes almost-sure convergence.* ■

This lemma does not require the identifiability condition (2.10). It explores forms of learning that were studied in [26, 70] and also in [72, 73], which dealt with either learning the true parameter θ° (similar to the setting we are considering) or learning the distribution of the incoming signal itself.

Correct forecasting does not always imply the ability of agents to learn the true parameter, θ°. However, in the case of strongly-connected networks, this conclusion is true under some conditions mentioned next (the same implication will not hold for weakly-connected networks; there, we will show that correct forecasting does not imply the ability of agents to learn the truth).

Theorem 1 (Truth Learning [30])**.** *Under the same conditions of Lemma 1, assume that there exists at least one prevailing signal ζ_k° for each agent k, namely, that*

$$L_k(\zeta_k^\circ|\theta^\circ) - L_k(\zeta_k^\circ|\theta) > 0, \quad \forall \theta \in \Theta \setminus \Theta_k \tag{2.12}$$

and assume as well that the true state θ° is globally identifiable as in (2.10). Then, all agents asymptotically learn the truth, i.e., for any $k \in \mathcal{N}$:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta^\circ) \overset{a.s.}{=} 1 \tag{2.13}$$

■

Figure 2.4 illustrates what it means for a prevailing signal to exist for an agent k. In this example, the true state θ° is assumed to be θ_1. Assume also that for agent k, the set of distinguishable states is $\bar{\Theta}_k = \Theta \setminus \Theta_k = \{\theta_2, \theta_3\}$ and the space of observational signals is $Z_k = \{\zeta_1, \zeta_2, \zeta_3, \zeta_4\}$. We see in the example that the signal ζ_1 plays the role of a prevailing signal. This is because when the true state is θ_1, the likelihood of ζ_1 is greater than its likelihood when the true state is θ_2 or θ_3, i.e.,

$$L_k(\zeta_1|\theta_1) > L_k(\zeta_1|\theta_2),\; L_k(\zeta_1|\theta_1) > L_k(\zeta_1|\theta_3) \tag{2.14}$$

These two conditions are not jointly satisfied for the other observational signals. The presence of a prevailing signal provides agent k with sufficient information to identify the distinguishable set $\bar{\Theta}_k = \Theta \setminus \Theta_k$. This means that agent k will be able to assign a zero probability to any θ in this set. Then, with the help of neighboring agents, and in the presence of the identifiability condition (2.10), agent k will be able to discover the true sate θ° in Θ_k.

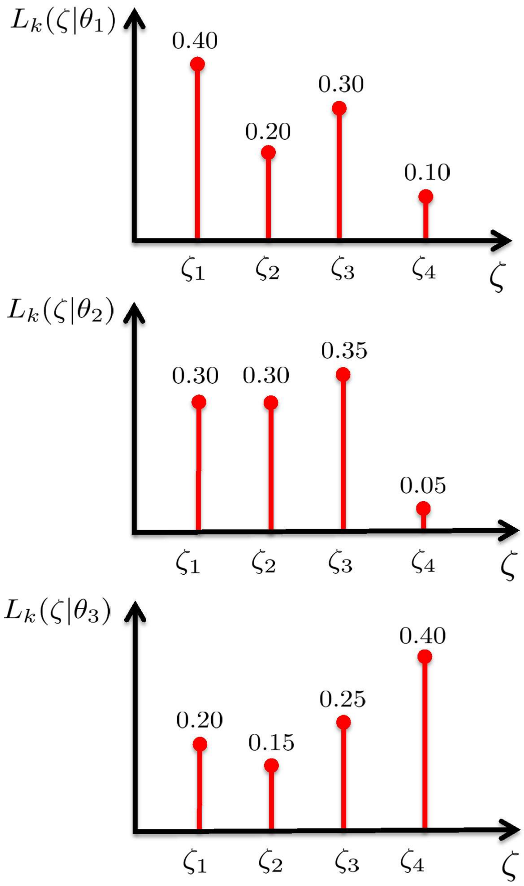

Figure 2.4: An example showing the existence of a prevailing signal ζ_1 for agent k.

2.4 Simulation Example

We illustrate the results with the following simulation example. Consider the social network shown in Fig. 2.5 which consists of $N = 8$ agents. We assume that there are 3 possible events $\Theta = \{\theta_1^\circ, \theta_2^\circ, \theta_3^\circ\}$, where θ_1° is the true event. We further assume that the observational signals of each agent k are binary and belong to $Z_k = \{H, T\}$ where H denotes head and T denotes tail.

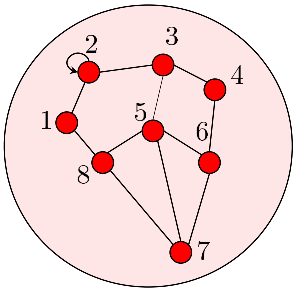

Figure 2.5: A strongly-connected network consisting of eight agents.

Agents are connected through the following combination matrix:

$$A = \begin{bmatrix} 0 & 0.3 & 0 & 0 & 0 & 0 & 0 & 0.3 \\ 0.4 & 0 & 0.3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.7 & 0 & 0.5 & 0.25 & 0 & 0 & 0 \\ 0 & 0 & 0.4 & 0 & 0 & 0.3 & 0 & 0 \\ 0 & 0 & 0.3 & 0 & 0 & 0.1 & 0.2 & 0.45 \\ 0 & 0 & 0 & 0.5 & 0.25 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0 & 0.3 & 0.6 & 0 & 0.25 \\ 0.6 & 0 & 0 & 0 & 0.2 & 0 & 0.7 & 0 \end{bmatrix} \tag{2.15}$$

The likelihood of the head signals for each agent k is selected as the following matrix:

$$L(H) = \begin{bmatrix} 5/8 & 3/4 & 1/6 & 1/2 & 1/3 & 1/5 & 4/5 & 1/2 \\ 5/8 & 3/4 & 1/6 & 2/3 & 1/2 & 1/5 & 2/3 & 1/2 \\ 1/4 & 3/4 & 1/3 & 1/2 & 1/4 & 1/5 & 4/5 & 1/3 \end{bmatrix} \tag{2.16}$$

where each (j, k)-th element of this matrix corresponds to $L_k(H/\theta_j)$, i.e., each column corresponds to one agent and each row to one network state. The likelihood of the tail signal is $L(T) = \mathbb{1}_{3\times 7} - L(H)$. We further assume that each agent starts at time $i = 0$ with an initial belief that is uniform over Θ and then updates it over time according to the model described in (2.8). Figures 2.6 and 2.6 show the evolution of $\boldsymbol{\mu}_{k,i}(\theta_1^\circ)$ of agents.

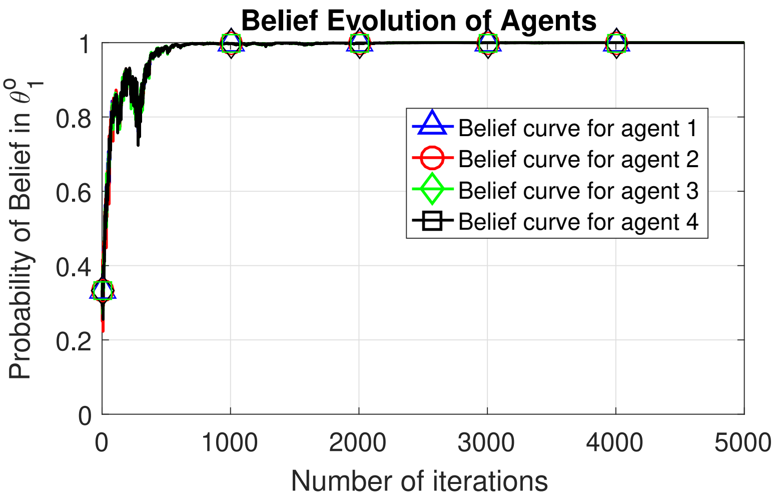

Figure 2.6: Evolution of agent k belief at θ_1° over time $(1 \leq k \leq 4)$.

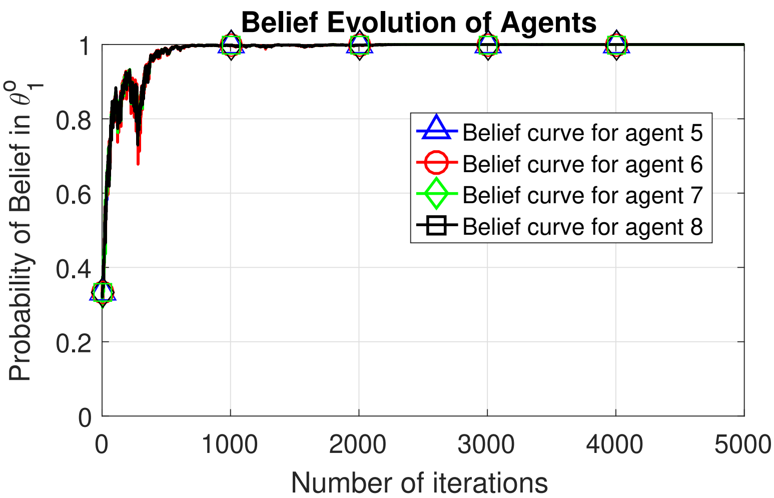

Figure 2.7: Evolution of agent k belief at θ_1° over time $(5 \leq k \leq 8)$.

2.5 Conclusions

In this chapter, we reviewed the model of diffusion social learning over strongly-connected graphs. Given that the observational signals are globally identifiable, agents are able to efficiently aggregate the information over the network and learn the true state. The exposition in this chapter provide a summary of the results in [30].

CHAPTER 3

Diffusion Learning over Weak Networks

We first review the main features of the weakly-connected network model from [67, 68]. Consider a network that consists of two types of sub-networks: S sub-networks and R sub-networks. Each sub-network in the S family has a strongly-connected topology. In contrast, each sub-network in the R family is only required to be connected. This means that any receiving sub-network has a path connecting any two agents without requiring any agent to have a self-loop. Moreover, the interaction between S and R sub-networks is not symmetric: information can flow from S ("sending") sub-networks to R ("receiving") sub-networks but not the other way around.

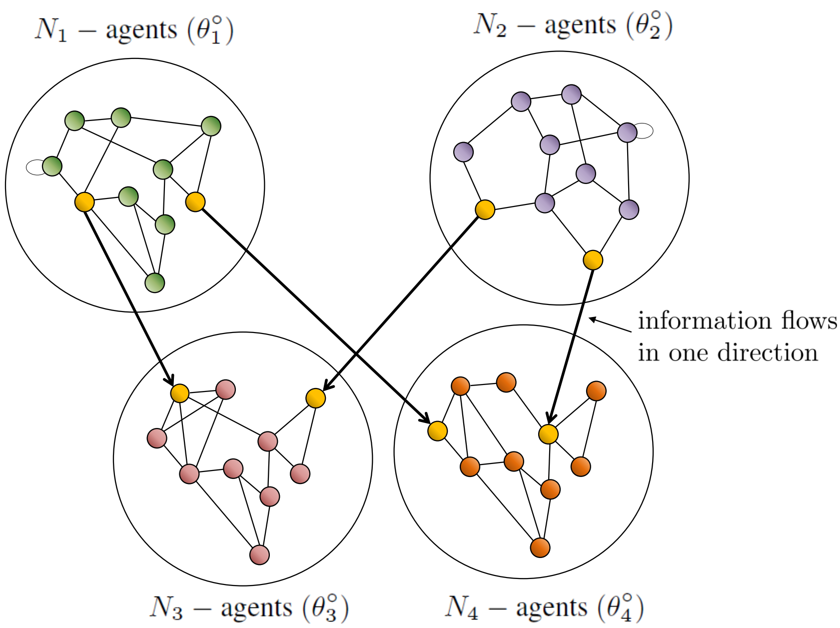

Figure 3.1: An example of a weakly connected network.

We index each strongly-connected sub-network by s where $s = \{1, 2, \cdots, S\}$. Similarly, we index each receiving sub-network by r where $r = \{S+1, \cdots, S+R\}$. Each sub-network s

has N_s agents, and the total number of agents in the S sub-networks is:

$$N_{gS} \triangleq N_1 + N_2 + \cdots + N_S \tag{3.1}$$

Similarly, each sub-network r has N_r agents, and the total number of agents in the R sub-networks is:

$$N_{gR} \triangleq N_{S+1} + N_{S+2} + \cdots + N_{S+R} \tag{3.2}$$

We still denote by N the total number of agents across all sub-networks, i.e., $N = N_{gS} + N_{gR}$. We continue to denote by $\mathcal{N} = \{1, 2, \cdots, N\}$ the indexes of the agents. We assume that the agents are numbered such that the indexes of $\mathcal{N}$ represent first the agents from the S sub-networks, followed by those from the R sub-networks. In this way, the structure of the network can be represented by a large $N \times N$ combination matrix A, which will have an upper block-triangular structure of the following form [67, 68]:

$$\overbrace{\qquad\qquad\qquad\qquad}^{\text{Subnetworks:}1,2,\ldots,S}\quad\overbrace{\qquad\qquad\qquad\qquad\qquad\qquad}^{\text{Subnetworks:}S+1,S+2,\ldots,S+R}$$

$$\left[\begin{array}{cccc|cccc}
A_1 & 0 & \ldots & 0 & A_{1,S+1} & A_{1,S+2} & \ldots & A_{1,S+R} \\
0 & A_2 & \ldots & 0 & A_{2,S+1} & A_{2,S+2} & \ldots & A_{2,S+R} \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\
0 & 0 & \ldots & A_S & A_{S,S+1} & A_{S,S+2} & \ldots & A_{S,S+R} \\
\hline
0 & 0 & \ldots & 0 & A_{S+1} & A_{S+1,S+2} & \ldots & A_{S+1,S+R} \\
0 & 0 & \ldots & 0 & 0 & A_{S+2} & \ldots & A_{S+2,S+R} \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\
0 & 0 & \ldots & 0 & 0 & 0 & \ldots & A_{S+R}
\end{array}\right] \tag{3.3}$$

The matrices $\{A_1, \cdots, A_S\}$ on the upper left corner are left-stochastic primitive matrices corresponding to the S strongly-connected sub-networks. Each of these matrices has spectral radius equal to one, $\rho(A_s) = 1$. Moreover, each A_s has a single eigenvalue at one and the corresponding right eigenvector has positive entries. We denote it by y_s and normalize its entries to add up to one, i.e., $\mathbb{1}^\mathsf{T} y_s = 1$.

Likewise, the matrices $\{A_{S+1}, \cdots, A_{S+R}\}$ in the lower right-most block correspond to the internal weights of the R sub-networks. These matrices are not necessarily left-stochastic because they do not include the coefficients over the links that connect the R sub-networks to the S sub-networks. Nevertheless, based on results from [74], it was shown in [67] that for any receiving subnetwork r, it holds that $\rho(A_r) < 1$. Moreover, since A_r has non-negative entries and sub-network r is connected, it follows from the Perron-Frobenius theorem [43,44] that A_r has a unique positive real eigenvalue λ_r, that is equal to its spectral radius $\rho(A_r)$, and the corresponding right eigenvector has positive entries. We denote this eigenvector by y_r. We again normalize the entries of y_r to add up to one, $\mathbb{1}^\mathsf{T} y_r = 1$:

$$A_r y_r = \lambda_r y_r, \quad \mathbb{1}^\mathsf{T} y_r = \mathbb{1}, \quad y_r \succ 0 \tag{3.4}$$

We denote the block structure of A in (3.3) by:

$$A \triangleq \left[\begin{array}{c|c} T_{SS} & T_{SR} \\ \hline 0 & T_{RR} \end{array}\right] \tag{3.5}$$

This specific structure has one useful property that we will exploit in the analysis.

Lemma 2 (Limiting Power of A [67]). *It holds that:*

$$A_\infty \triangleq \lim_{n\to\infty} A^n = \left[\begin{array}{c|c} E & EW \\ \hline 0 & 0 \end{array}\right] \tag{3.6}$$

where the $N_{gS} \times N_{gS}$ matrix E and the $N_{gS} \times N_{gR}$ matrix W are given by:

$$W \triangleq T_{SR}(I - T_{RR})^{-1} \tag{3.7}$$

$$E \triangleq \textit{blockdiag}\left\{y_1 \mathbb{1}_{N_1}^\mathsf{T}, \ldots, y_S \mathbb{1}_{N_S}^\mathsf{T}\right\} \tag{3.8}$$

The matrix W has non-negative entries and the sum of the entries in each column is equal to one. ■

We now examine the belief evolution of agents in weakly-connected networks. We still denote by Θ the set of all possible states, and we assume that Θ is uniform across all sub-networks. However, we allow each sub-network to have its own true state, which may differ from one sub-network to another. We denote by θ_s° the true state of sending sub-network s and by θ_r° the true state of receiving sub-network r, where both θ_s° and θ_r° are in Θ. Therefore, if agent k belongs to a sub-network s, its observational signals $\boldsymbol{\xi}_{k,i}$ will be generated according to the likelihood function $L_k(\cdot|\theta_s^\circ)$. On the other hand, if agent k belongs to a sub-network r, its observational signals $\boldsymbol{\xi}_{k,i}$ will be generated according to $L_k(\cdot|\theta_r^\circ)$.

We already know that the $S-$type sub-networks are strongly-connected, so that their agents can cooperate together to learn the truth. More specifically, according to Theorem 1, if agent k belongs to sub-network s, then it holds that:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_s^\circ) \overset{a.s.}{=} 1 \tag{3.9}$$

The question that we want to examine is how the beliefs of the agents in the receiving sub-networks are affected. These agents are now influenced by the beliefs of the $S-$type groups. Since this external influence carries information not related to the true state of each receiving sub-network, the receiving agents may not be able to learn their own true states. We will show that a leader-follower relationship develops.

3.1 Weak Graphs

We consider that all agents are following the diffusion strategy (2.8) for social learning. In a manner similar to (2.9), if agent k belongs to sub-network r, then we assume that there exists a subset of states $\Theta_k \subseteq \Theta$ such that:

$$L_k(\zeta_k|\theta) = L_k(\zeta_k|\theta_r^\circ) \tag{3.10}$$

for any $\zeta_k \in Z_k$ and $\theta \in \Theta_k$, i.e., Θ_k is the set of indistinguishable states for agent k. Moreover, we assume a scenario in which the private signals of agents in the receiving sub-networks are not informative enough to let their agents discover that the true states of the sending sub-networks do not represent their own truth. That is, we are assuming *for now* the following condition.

Assumption. *The true state θ_s°, of each sub-network $s \in \{1, 2, \dots, S\}$, belongs to the indistinguishable set Θ_k:*

$$\theta_s^\circ \in \Theta_k, \quad \textit{for any } k > N_{gS} \tag{3.11}$$

■

Under (3.11), we will now verify that the interaction with the S sub-networks ends up forcing the receiving agents to focus their beliefs on the true states of the $S-$type. Later, we will show that a similar conclusion continues to hold even when (3.11) is relaxed.

Thus, let $\Theta^\bullet = \{\theta_1^\circ, \cdots, \theta_S^\circ\}$ denote the set of all true states of the $S-$type sub-networks. We are assuming, for notational simplicity, that the true states $\{\theta_s^\circ\}$ are distinct from each other. Otherwise, we only include in $\Theta^\bullet$ the set of truly distinct states, which will be smaller than S in number. We denote the complement of $\Theta^\bullet$ by $\bar{\Theta}^\bullet$, such that $\Theta^\bullet \cap \bar{\Theta}^\bullet = \emptyset$ and $\Theta^\bullet \cup \bar{\Theta}^\bullet = \Theta$. We first show that as $i \to \infty$, each receiving agent k will assign zero belief to any event $\theta \in \bar{\Theta}^\bullet$. This means that receiving agents will end up searching for the truth within the set $\Theta^\bullet$.

Lemma 3 (Focus on True States of S Sub-Networks)**.** *Under (3.11), each agent k of any receiving sub-network r eventually identifies the set $\bar{\Theta}^\bullet$, namely, for any $\theta \in \bar{\Theta}^\bullet$:*

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) \overset{a.s.}{=} 0 \tag{3.12}$$

Proof: See Appendix 3.A. ■

This lemma implies that the receiving agents are still able to perform correct forecasting.

Lemma 4 (Correct Forecasting)**.** *Under (3.11), every agent k in sub-network r develops correct forecasting, namely,*

$$\lim_{i\to\infty} \boldsymbol{m}_{k,i}(\zeta_k) \overset{a.s.}{=} L_k(\zeta_k|\theta_r^\circ), \quad \textit{for any } \zeta_k \in Z_k \tag{3.13}$$

Proof: See Appendix 3.B. ■

Even with the external influence, agent k is still able to attain correct forecasting because any true state θ_s° of any sending sub-network, s, belongs to the indistinguishable set of agent k, i.e., $L_k(\zeta_k|\theta_r^\circ) = L_k(\zeta_k|\theta_s^\circ)$ from (3.11) and (3.10). Since agents zoom onto the set $\Theta^\bullet$, this fact enables correct forecasting but does not necessarily imply truth learning for weak graphs, as discussed in the sequel.

The previous two lemmas establish that the belief of each agent k in sub-network r will converge to a distribution whose support is limited to $\theta \in \Theta^\bullet$. The next question is to evaluate this distribution, which is the subject of the following main result. First let

$$\boldsymbol{\mu}_i^s(\theta) \triangleq col\left\{\boldsymbol{\mu}_{k_s(1),i}(\theta), \boldsymbol{\mu}_{k_s(2),i}(\theta), \ldots, \boldsymbol{\mu}_{k_s(N_s),i}(\theta)\right\} \tag{3.14}$$

$$\boldsymbol{\mu}_i^r(\theta) \triangleq col\left\{\boldsymbol{\mu}_{k_r(1),i}(\theta), \boldsymbol{\mu}_{k_r(2),i}(\theta), \ldots, \boldsymbol{\mu}_{k_r(N_r),i}(\theta)\right\} \tag{3.15}$$

collect all beliefs from agents that belong respectively to sub-network s and sub-network r, where the notation $k_s(n)$ denotes the index of the n-th agent within sub-network s, i.e.,

$$k_s(n) = \sum_{v=1}^{s-1} N_v + n \tag{3.16}$$

and $n \in \{1, 2, \cdots, N_s\}$ and the notation $k_r(n)$ denotes the index of the n-th agent within sub-network r, i.e.,

$$k_r(n) = N_{gS} + \sum_{v=S+1}^{r-1} N_v + n \tag{3.17}$$

and $n \in \{1, 2, \cdots, N_r\}$. Furthermore, let

$$\boldsymbol{\mu}_{\mathcal{S},i}(\theta) \triangleq col\left\{\boldsymbol{\mu}_i^1(\theta), \boldsymbol{\mu}_i^2(\theta), \ldots, \boldsymbol{\mu}_i^S(\theta)\right\} \tag{3.18}$$

$$\boldsymbol{\mu}_{\mathcal{R},i}(\theta) \triangleq col\left\{\boldsymbol{\mu}_i^{S+1}(\theta), \boldsymbol{\mu}_i^{S+2}(\theta), \ldots, \boldsymbol{\mu}_i^{S+R}(\theta)\right\} \tag{3.19}$$

collect all belief vectors respectively from all $S-$type sub-networks and from all $R-$type sub-networks.

Theorem 2 (Limiting Beliefs for Receiving Agents). *Under (3.11), it holds that*

$$\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) = W^{\mathsf{T}} \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta) \right) \tag{3.20}$$

Proof: See Appendix 3.C. ■

We expand (3.20) to clarify its meaning and to show how the beliefs are distributed among the elements of $\Theta^{\bullet}$. We already know from the result in Theorem 1 that, for each agent k of sending sub-network s, $\boldsymbol{\mu}_{k,i}(\theta)$ converges asymptotically to an impulse of size one at the location $\theta = \theta_s^{\circ}$. Thus, we write:

$$\lim_{i\to\infty} \boldsymbol{\mu}_i^s(\theta) = e_{\theta,\theta_s^{\circ}} \triangleq \begin{cases} \mathbb{1}_{N_s}, & if \quad \theta = \theta_s^{\circ} \\ \mathbf{0}_{N_s}, & otherwise \end{cases} \tag{3.21}$$

where $\mathbb{1}_{N_s}$ denotes a column vector of length N_s whose elements are all one. Similarly, $\mathbf{0}_{N_s}$ denotes a column vector of length N_s whose elements are all zero. Hence,

$$\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta) = col\left\{e_{\theta,\theta_1^{\circ}}, e_{\theta,\theta_2^{\circ}}, \dots, e_{\theta,\theta_S^{\circ}}\right\} \tag{3.22}$$

Now, let w_k^{T} denote the row of W^{T} that corresponds to agent k in sub-network[1] r. We partition it into

$$w_k^{\mathsf{T}} = \left[\begin{array}{ccc} w_{k,N_1}^{\mathsf{T}} & w_{k,N_2}^{\mathsf{T}} & \dots, w_{k,N_S}^{\mathsf{T}} \end{array} \right] \tag{3.23}$$

where the $\{N_1, N_2, \dots, N_S\}$ are the number of agents in each sub-network $s \in \{1, 2, \dots, S\}$. By examining (3.20), we conclude that the distribution for each agent k in an $R-$type sub-network converges to a combination of the various vectors $\{e_{\theta,\theta_s^{\circ}}\}$, namely,

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) = q_k(\theta) \triangleq \sum_{s=1}^{S} w_{k,N_s}^{\mathsf{T}} e_{\theta,\theta_s^{\circ}} \tag{3.24}$$

Observe that, from this equation, to get $q_k(\theta_s^{\circ})$, the elements of the corresponding block in w_k, i.e., w_{k,N_s}, should be summed. Now, if we consider that multiple sending sub-networks

[1]The real index of the row of W^{T} that corresponds to agent k is $k - N_{gS}$.

have the same true state, then to get $q_k(.)$ at this true state, the elements of all corresponding blocks in w_k will need to be summed. Note that this is a valid probability measure in view of Lemma 2, i.e.,

$$\sum_{\theta \in \Theta^\bullet} q_k(\theta) = 1 \tag{3.25}$$

Note also that if it happens that $\theta_s^\circ = \theta^\circ$ for all s, then $q_k(\theta^\circ) = 1$ and $q_k(\theta) = 0$ for all $\theta \neq \theta^\circ$, and in this case, sending agents can be seen as helping receiving agents to find the true state. We also observe that the beliefs of agents in the receiving sub-networks differ from one agent to another, since for each agent k, $q_k(\theta)$ depends on w_k. This means that the external influence has created social disagreement in the receiving sub-networks.

We therefore established that the beliefs of receiving agents converge to a distribution whose support is limited to the true states of the sending sub-networks. We will refer to this situation as a *total influence* or "mind-control" scenario where the learning of the $R-$subnetworks is fully dictated by the $S-$subnetworks. When all agents follow model (2.8) and when assumption (3.11) is satisfied, this total influence scenario arises. Although the private signals of the receiving agents are supposed to hold information regarding their own true state, however, under assumption (3.11), these signals are not informative enough, so that agents are naturally driven to be under the influence of the sending sub-networks.

We are interested now in knowing whether this total influence situation can still occur when assumption (3.11) is not satisfied anymore. When this is the case, sending agents may not be able to totally control the beliefs of receiving agents anymore. Before establishing the analytical results, and before showing how self-awareness can alter this dynamics, we provide an illustrative example.

3.2 Implications of Violating Condition (3.11)

We consider a network consisting of three agents, with the first two playing the role of influential agents and the third one acting as a receiving agent. The combination matrix is chosen as follows:

$$A = \left[\begin{array}{cc|c} 1 & 0 & 0.1 \\ 0 & 1 & 0.2 \\ \hline 0 & 0 & 0.7 \end{array}\right] \tag{3.26}$$

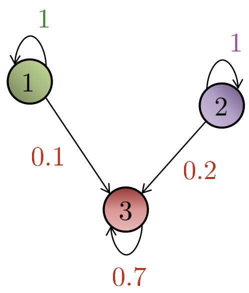

Figure 3.2: A weakly connected network and the corresponding combination policy (3.26).

We denote by θ_1° the true state for agent 1, by θ_2° the true state for agent 2, and by θ_3° the true state for agent 3 so that $\Theta = \{\theta_1^\circ, \theta_2^\circ, \theta_3^\circ\}$. The observational signal for all three agents is either a head "H" or a tail "T". In order for agents 1 and 2 to learn their true states asymptotically, we need to ensure that the conditions of Theorem 1 are satisfied. One of these requirements is the identifiability condition (2.10), which requires that the intersection of the indistinguishable sets (2.9) of all agents in a given sending sub-network s must be the singleton $\{\theta_s^\circ\}$. In this example, each sending sub-network consists of only one agent, so that condition (2.10) reduces to $\Theta_1 = \{\theta_1^\circ\}$ for the first agent and $\Theta_2 = \{\theta_2^\circ\}$ for the second agent. In other words, since agents 1 and 2 do not have neighbors to communicate with, they must rely solely on their observational signals to learn the truth. This is feasible when for agents 1 and 2 no state is observationally equivalent to their true state (or indistinguishable). Using the definition of the indistinguishable set (2.9), $\Theta_1 = \{\theta_1^\circ\}$ translates into the following requirement for agent 1:

$$L_1(\zeta_1|\theta_1^\circ) \neq L_1(\zeta_1|\theta_2^\circ) \;\; \text{and} \;\; L_1(\zeta_1|\theta_1^\circ) \neq L_1(\zeta_1|\theta_3^\circ) \tag{3.27}$$

for any $\zeta_1 \in \{H, T\}$. Similarly, $\Theta_2 = \{\theta_2^\circ\}$ translates into the following requirement for agent

2:

$$L_1(\zeta_2|\theta_2^\circ) \neq L_1(\zeta_2|\theta_1^\circ) \;\text{ and }\; L_1(\zeta_2|\theta_2^\circ) \neq L_1(\zeta_2|\theta_3^\circ) \tag{3.28}$$

for any $\zeta_2 \in \{H, T\}$. For this example, we are choosing the likelihood functions arbitrarily but satisfying (3.27) for agent 1 and (3.28) for agent 2. For instance, we select for agent 1,

$$L_1(H|\theta_1^\circ) = 0.10, \; L_1(H|\theta_2^\circ) = 0.35, \; L_1(H|\theta_3^\circ) = 0.45 \tag{3.29}$$

and set $L_1(T|\theta) = 1 - L_1(H|\theta)$ for any $\theta \in \Theta$. Likewise for agent 2, we select

$$L_2(H|\theta_1^\circ) = 0.10, \; L_2(H|\theta_2^\circ) = 0.20, \; L_2(H|\theta_3^\circ) = 0.30 \tag{3.30}$$

and set $L_2(T|\theta) = 1 - L_2(H|\theta)$ for any $\theta \in \Theta$. Before analyzing the beliefs of agent 3 when (3.11) is not satisfied, we consider first the case in which this assumption is satisfied. In this way, we will be able to compare what is happening in both cases. More specifically, following (3.11), we consider first that θ_1^o and θ_2^o belong to the indistinguishable set of agent 3 denoted by Θ_3, i.e., $\{\theta_1^\circ, \theta_2^\circ\} \in \Theta_3$. This means, according to the definition of the indistinguishable set (2.9), that

$$L_3(\zeta_3|\theta_1^\circ) = L_3(\zeta_3|\theta_3^\circ) \;\text{ and }\; L_3(\zeta_3|\theta_2^\circ) = L_3(\zeta_3|\theta_3^\circ) \tag{3.31}$$

for any $\zeta_3 \in \{H, T\}$. According to model (2.8), the intermediate belief of agent 3 is given by:

$$\boldsymbol{\psi}_{3,i}(\theta) = \frac{\boldsymbol{\mu}_{3,i-1}(\theta) L_3(\boldsymbol{\xi}_{3,i}|\theta)}{\left(\sum_{\theta' \in \Theta} \boldsymbol{\mu}_{3,i-1}(\theta')\right) L_3(\boldsymbol{\xi}_{3,i}|\theta)} = \boldsymbol{\mu}_{3,i-1}(\theta) \tag{3.32}$$

We observe in this example that the private signals of agent 3 end up not contributing to its intermediate belief. As a result, it is only the beliefs of agents 1 and 2 that affect the belief of agent 3, so that:

$$\boldsymbol{\mu}_{3,i}(\theta) = a_{13}\boldsymbol{\psi}_{1,i}(\theta) + a_{23}\boldsymbol{\psi}_{2,i}(\theta) + a_{33}\boldsymbol{\psi}_{3,i}(\theta) = a_{13}\boldsymbol{\mu}_{1,i}(\theta) + a_{23}\boldsymbol{\mu}_{2,i}(\theta) + a_{33}\boldsymbol{\mu}_{3,i-1}(\theta) \tag{3.33}$$

In writing (3.33), we used the fact that the intermediate beliefs for agents 1 and 2 coincide with their updated beliefs since, in this example, agents 1 and 2 have no neighbors. Thus, since $a_{33} < 1$,

$$\lim_{i\to\infty} \boldsymbol{\mu}_{3,i}(\theta) = \left(\frac{a_{13}}{1-a_{33}}\right) \lim_{i\to\infty} \boldsymbol{\mu}_{1,i}(\theta) + \left(\frac{a_{23}}{1-a_{33}}\right) \lim_{i\to\infty} \boldsymbol{\mu}_{2,i}(\theta) \tag{3.34}$$

from which we conclude that

$$\lim_{i\to\infty} \boldsymbol{\mu}_{3,i}(\theta_1^\circ) = \frac{a_{13}}{1-a_{33}} \tag{3.35}$$

$$\lim_{i\to\infty} \boldsymbol{\mu}_{3,i}(\theta_2^\circ) = \frac{a_{12}}{1-a_{33}} \tag{3.36}$$

$$\lim_{i\to\infty} \boldsymbol{\mu}_{3,i}(\theta_3^\circ) = 0 \tag{3.37}$$

This total influence result is expected to occur according to Theorem 2, when assumption (3.11) is satisfied.

Let us consider now the case in which assumption (3.11) is not satisfied. This means that θ_1° and θ_2° do not need to both belong to the indistinguishable set Θ_3° of agent 3, i.e.,

$$L_3(\zeta_k|\theta_1^\circ) \neq L_3(\zeta_k|\theta_3^\circ) \quad \text{or} \quad L_3(\zeta_k|\theta_2^\circ) \neq L_3(\zeta_k|\theta_3^\circ) \tag{3.38}$$

for any $\zeta_k \in \{H, T\}$. In this example, we study the worst case scenario in which both conditions in (3.38) are met (even if we consider other situations in which only one of these conditions is met, we still arrive at a similar conclusion, namely, the belief of agent 3 will not reach a fixed distribution). We select arbitrarily the values for the likelihood function of agent 3, but in a way that these values satisfy both conditions in (3.38). For instance, we select

$$L_3(H|\theta_1^\circ) = 0.4, \quad L_3(H|\theta_2^\circ) = 0.3, \quad L_3(H|\theta_3^\circ) = 0.8 \tag{3.39}$$

In this case, the belief for agent 3 will be updated as:

$$\begin{aligned}\boldsymbol{\mu}_{3,i}(\theta) &= a_{13}\boldsymbol{\mu}_{1,i}(\theta) + a_{23}\boldsymbol{\mu}_{2,i}(\theta) + a_{33}\boldsymbol{\psi}_{3,i-1}(\theta) \\ &= a_{13}\boldsymbol{\mu}_{1,i}(\theta) + a_{23}\boldsymbol{\mu}_{2,i}(\theta) + a_{33}\frac{L_3(\boldsymbol{\xi}_{3,i}|\theta)\boldsymbol{\mu}_{3,i-1}(\theta)}{\sum_{\theta'\in\Theta}\boldsymbol{\mu}_{3,i-1}(\theta')L_3(\boldsymbol{\xi}_{3,i}|\theta')}\end{aligned} \tag{3.40}$$

We see here how this equality is different from (3.33), where the last term $\boldsymbol{\psi}_{3,i-1}(\theta)$ holds information about θ_3° that contradicts with the information held in the other terms. We now show by contradiction that in this case, agent 3 will not converge to a fixed distribution. Assume, to the contrary, that the beliefs of agent 3 reach the following distribution:

$$\lim_{i\to\infty}\boldsymbol{\mu}_{3,i}(\theta_1^\circ) = b,\ \lim_{i\to\infty}\boldsymbol{\mu}_{3,i}(\theta_2^\circ) = c,\ \lim_{i\to\infty}\boldsymbol{\mu}_{3,i}(\theta_3^\circ) = d \tag{3.41}$$

for some fixed non-negative constants b, c and d satisfying

$$b + c + d = 1 \tag{3.42}$$

We know that, as $i \to \infty$, agents 1 and 2 approach their true states so that by evaluating (3.40) at θ_1° when $i \to \infty$, we get:

$$b = a_{13} + \frac{a_{33}L_3(\boldsymbol{\xi}_{3,i}|\theta_1^\circ)b}{bL_3(\boldsymbol{\xi}_{3,i}|\theta_1^\circ) + cL_3(\boldsymbol{\xi}_{3,i}|\theta_2^\circ) + dL_3(\boldsymbol{\xi}_{3,i}|\theta_3^\circ)} \tag{3.43}$$

Evaluating (3.40) at θ_2° when $i \to \infty$:

$$c = a_{23} + \frac{a_{33}L_3(\boldsymbol{\xi}_{3,i}|\theta_2^\circ)c}{bL_3(\boldsymbol{\xi}_{3,i}|\theta_1^\circ) + cL_3(\boldsymbol{\xi}_{3,i}|\theta_2^\circ) + dL_3(\boldsymbol{\xi}_{3,i}|\theta_3^\circ)} \tag{3.44}$$

Evaluating (3.40) at θ_3° when $i \to \infty$:

$$d = \frac{a_{33}L_3(\boldsymbol{\xi}_{3,i}|\theta_3^\circ)d}{bL_3(\boldsymbol{\xi}_{3,i}|\theta_1^\circ) + cL_3(\boldsymbol{\xi}_{3,i}|\theta_2^\circ) + dL_3(\boldsymbol{\xi}_{3,i}|\theta_3^\circ)} \tag{3.45}$$

Then, from (3.45), we have:

$$d = \frac{0.56d}{0.4b + 0.3c + 0.8d}, \quad \text{if observation is H} \tag{3.46}$$

$$d = \frac{0.14d}{0.6b + 0.7c + 0.2d}, \quad \text{if observation is T} \tag{3.47}$$

Then, either $d = 0$ or

$$0.4b + 0.3c + 0.8d = 0.56 \text{ and } 0.6b + 0.7c + 0.2d = 0.14 \tag{3.48}$$

However, conditions (3.48) contradict the fact that we must have

$$(0.4b + 0.3c + 0.8d) + (0.6b + 0.7c + 0.2d) = b + c + d \overset{(3.42)}{=} 1 \tag{3.49}$$

We conclude that $d = 0$. Thus, condition (3.42) reduces to:

$$b + c = 1 \tag{3.50}$$

With regards to the values of b and c, we know from (3.43) that

$$b = 0.1 + \frac{0.28b}{0.4b + 0.3c}, \quad \text{if observation is H} \tag{3.51}$$

$$b = 0.1 + \frac{0.42b}{0.6b + 0.7c}, \quad \text{if observation is T} \tag{3.52}$$

That is, the scalars b and c must satisfy

$$\frac{0.28}{0.4b + 0.3c} = \frac{0.42}{0.6b + 0.7c} \tag{3.53}$$

The denominators are related as follows:

$$(0.4b + 0.3c) + (0.6b + 0.7c) = b + c \overset{(3.50)}{=} 1 \tag{3.54}$$

Thus,

$$\frac{0.28}{0.4b + 0.3c} = \frac{0.42}{1 - (0.4b + 0.3c)} \tag{3.55}$$

This leads to

$$0.4b + 0.3c = \frac{0.28}{0.28 + 0.42} = 0.4 \tag{3.56}$$

so that from (3.51), we have

$$b = 0.1 + \frac{0.28}{0.4}b \tag{3.57}$$

Thus, $b = \frac{1}{3}$, and since $0.4b + 0.3c = 0.4$, then $c = \frac{8}{9}$. However, $b + c = \frac{11}{9}$, which contradicts (3.50). We conclude that the beliefs of agent 3 cannot reach a fixed distribution. This conclusion is illustrated in Fig. 3.3, which plots the evolution of beliefs of agent 3 for all $\theta \in \Theta$. It is clear from the figure how the contradictory information conveyed by the influential agents and the private signals do not lead agent 3 to approach a fixed belief. This also means that agents 1 and 2 cannot fully control agent 3. However, if agent 3 decides

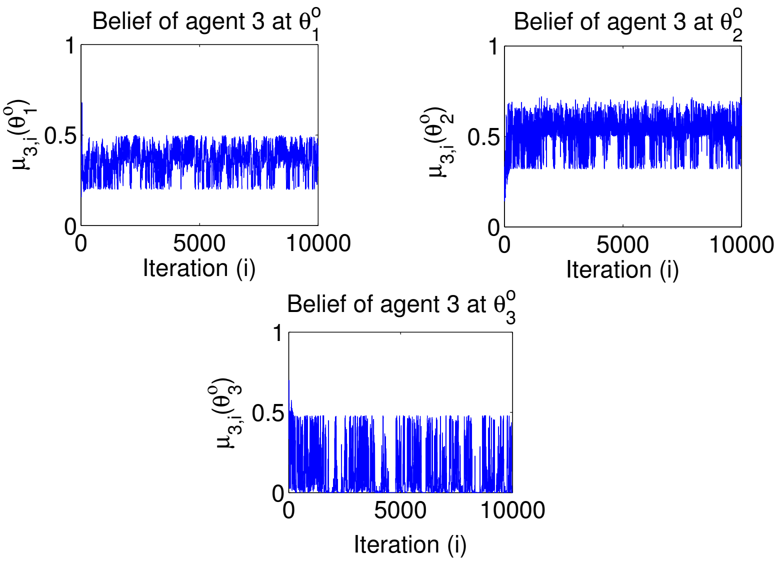

Figure 3.3: Evolution of the beliefs of agent 3 over time for the case in which condition (3.11) is not satisfied.

to limit the contribution of its private signal on the update of its intermediate belief, will agents 1 and 2 be able to totally influence agent 3? In other words, will the total influence scenario arise again even if assumption (3.11) is not satisfied? We show next that this is possible by incorporating an element of self-awareness into the learning process.

3.3 Diffusion Learning with Self-Awareness

We are therefore now motivated to modify the diffusion strategy (2.8) by incorporating a non-negative convex combination $\gamma_{k,i}$. This factor enables agents to assign more or less weight to their local information in comparison to the information received from their neighbors. Specifically, we modify (2.8) as follows:

$$\begin{cases} \boldsymbol{\psi}_{k,i}(\theta) = (1-\gamma_{k,i})\,\boldsymbol{\mu}_{k,i-1}(\theta) + \gamma_{k,i}\dfrac{\boldsymbol{\mu}_{k,i-1}(\theta)L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta'\in\Theta}\boldsymbol{\mu}_{k,i-1}(\theta')L_k(\boldsymbol{\xi}_{k,i}|\theta')} \\ \boldsymbol{\mu}_{k,i}(\theta) = \displaystyle\sum_{\ell\in\mathcal{N}_k} a_{\ell k}\boldsymbol{\psi}_{\ell,i}(\theta) \end{cases} \tag{3.58}$$

where $\gamma_{k,i}\in[0,1]$ is a scalar variable. Observe that the intermediate belief $\boldsymbol{\psi}_{k,i}(\theta)$ of agent k is now a combination of its prior belief, $\boldsymbol{\mu}_{k,i-1}(\theta)$, and the Bayesian update. The scalar $\gamma_{k,i}$ represents the amount of trust that agent k gives to its private signal and how it is balancing this trust between the new observation and its own past belief. This weight can also model the lack of an observational signal at time i.

Model (3.58) helps capture some elements of human behavior. For example, in an interactive social setting, a human agent may not be satisfied or convinced by an observation and prefers to give more weight to their prior belief based on accumulated experiences. This model was studied for single agents in [70, 71] and was motivated as a mechanism for self-control and temptation. The agent might observe a private signal at some time that can move this agent away from its current conviction. The agent can control this temptation by increasing the weight given to its prior belief or it can change its opinion by giving more weight to its Bayesian update, which is based on the private signal.

We next analyze model (3.58) over weakly-connected graphs and establish two results. The first result is related to the sending agents and the second result is related to the receiving agents.

Lemma 5 (Correct Forecasting with Self-Awareness)**.** *Assume that* $\lim_{i\to\infty} \gamma_{k,i} \neq 0$ *and the same conditions of Lemma 1. Then, self-aware sending agents develop correct forecasts of the incoming signals, namely, result (2.11) continues to hold.*

Proof: See Appendix 3.D. ∎

Theorem 3 (Truth Learning by Self-Aware Sending Agents)**.** *Under the same assumptions of Theorem 1, self-aware sending agents learn the truth asymptotically and condition (2.13) continues to hold.*

Proof: See Appendix 3.E. ∎

We therefore find that sending agents, whether self-aware or not, are always able to learn the truth. With regards to receiving agents, we now have the following conclusion. For each agent k in a receiving sub-network r, we write $\gamma_{k,i} = \tau_{k,i}\gamma_{\max}$, where $\gamma_{\max}$ are both positive scalars less than 1, and $\gamma_{\max} = \sup_{k,i} \gamma_{k,i}$.

Theorem 4 (Learning by Self-Aware Receiving Agents)**.** *The beliefs of self-aware receiving agents are confined as follows:*

$$\limsup_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) \preceq W^{\mathsf{T}} \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta)\right) + \gamma_{\max} C \mathbb{1}_{N_{gR}} \tag{3.59}$$

$$\liminf_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) \succeq W^{\mathsf{T}} \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta)\right) - \gamma_{\max} C \mathbb{1}_{N_{gR}} \tag{3.60}$$

where $C \triangleq (I - T_{RR}^{\mathsf{T}})^{-1}$ *is an* $N_{gR} \times N_{gR}$ *matrix.*

Proof: See Appendix 3.F. ∎

This final result coincides with that of Theorem 2, but with an additional $O(\gamma_{\max})$ term. This means that if each receiving agent chooses the $\gamma-$coefficient to be small enough, then its belief converges to the same distribution (3.24) of Theorem 2. When agent k gives a small weight to its Bayesian update, it means that it is giving its current signal $\boldsymbol{\xi}_{k,i}$ a reduced role to play in affecting its belief formation at time i, and it is instead relying more heavily on its prior belief $\boldsymbol{\mu}_{k,i-1}(\theta)$ and on its communication with its neighbors. When agent k continues to give less importance to any current signal it is receiving, its belief update will be mainly affected by its interaction with influential agents and its neighbors that are also under the influence of sending agents. Therefore, over time, these circumstances will help establish a leader-follower relationship in the network. In other words, the receiving sub-networks will be driven away from the truth and be under total indoctrination by the influential agents.

3.4 Simulation Results

We illustrate the previous results for weakly-connected networks. We assume that the social network has $N = 8$ agents interconnected as shown in Fig. 3.4, which corresponds to the following combination matrix:

$$A = \left[\begin{array}{ccccc|ccc} 0.2 & 0.2 & 0.8 & 0 & 0 & 0 & 0 & 0 \\ 0.5 & 0.4 & 0.1 & 0 & 0 & 0.2 & 0 & 0.4 \\ 0.3 & 0.4 & 0.1 & 0 & 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0 & 0.4 & 0.3 & 0.3 & 0 & 0 \\ 0 & 0 & 0 & 0.6 & 0.7 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0.2 & 0.3 & 0.2 \\ 0 & 0 & 0 & 0 & 0 & 0.1 & 0.5 & 0.3 \\ 0 & 0 & 0 & 0 & 0 & 0.1 & 0.2 & 0.1 \end{array}\right] \tag{3.61}$$

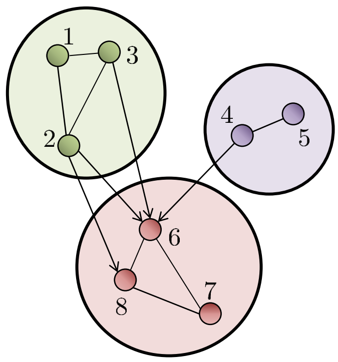

Figure 3.4: A weakly connected network consisting of three sub-networks and the corresponding combination policy (3.61).

We assume that there are 3 possible events $\Theta = \{\theta_1^\circ, \theta_2^\circ, \theta_3^\circ\}$, where θ_1° is the true event for the first sending sub-network, θ_2° is the true event for the second sending sub-network, and θ_3° is the true event for the receiving sub-network. We further assume that the observational signals of each agent k are binary and belong to $Z_k = \{H, T\}$ where H denotes head and T denotes tail. We consider two cases. In the first case, we assume that agents update their belief according to the model described in (2.8) and that assumption (3.11) is met. In the second case, we assume that agents follow the second model described in (3.58) where assumption (3.11) is not met.

3.4.1 First Case

In this first case, the likelihood of the head signals for each agent k is selected as the following 3×8 matrix:

$$L(H) = \begin{bmatrix} 5/8 & 3/4 & 1/3 & 7/8 & 5/8 & 1/3 & 1/4 & 5/8 \\ 5/8 & 1/4 & 1/6 & 7/8 & 2/3 & 1/3 & 1/4 & 5/8 \\ 1/4 & 3/4 & 1/6 & 1/3 & 2/3 & 1/3 & 1/4 & 5/8 \end{bmatrix}$$

where each (j,k)-th element of this matrix corresponds to $L_k(H/\theta_j)$, i.e., each column corresponds to one agent and each row to one network state. The likelihood of the tail signal is $L(T) = \mathbb{1}_{3\times 8} - L(H)$. We observe from $L(H)$ that assumption (3.11) is met here where for agent k in the receiving sub-network ($k > 5$) we have $L_k(\zeta_k|\theta_1^\circ) = L_k(\zeta_k|\theta_2^\circ) = L_k(\zeta_k|\theta_3^\circ)$ for both cases in which ζ_k is either head or tail. Assumption (3.11) is met here because the true state of the first sending sub-network θ_1° belongs to the indistinguishable set of any receiving agent k in the receiving sub-network 3, i.e., $L_k(\zeta_k|\theta_1^\circ) = L_k(\zeta_k|\theta_3^\circ)$, and the true state of the second sending sub-network θ_2° belongs to the indistinguishable set of any receiving agent k, i.e., $L_k(\zeta_k|\theta_2^\circ) = L_k(\zeta_k|\theta_3^\circ)$, where $k = 6, 7, 8$. We further assume that each agent starts at time $i = 0$ with an initial belief that is uniform over Θ and then updates it over time according to the model described in (2.8). Then, we know from [30] that $\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_1^\circ) = 1$ for $k = 1, 2, 3$ and $\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_2^\circ) = 1$ for $k = 4, 5$. Now for the agents of the receiving sub-network, we need first to compute:

$$\begin{aligned} W^{\mathsf{T}} &= (I - T_{RR}^{\mathsf{T}})^{-1} T_{SR}^{\mathsf{T}} \\ &= \begin{bmatrix} 0 & 0.4045 & 0.1489 & 0.4466 & 0 \\ 0 & 0.5267 & 0.1183 & 0.3550 & 0 \\ 0 & 0.7099 & 0.0725 & 0.2176 & 0 \end{bmatrix} \end{aligned} \tag{3.62}$$

The first row of W^{T} corresponds to agent 6, the second row to agent 7 and the third row to agent 8. Now each row is partitioned into two blocks: the first block is of length $N_1 = 3$ that corresponds to sub-network 1 of true state θ_1° and the second block is of length $N_2 = 2$ that corresponds to sub-network 2 of true state θ_2°. Then, according to Theorem 2, we can compute the belief at θ_1° for each receiving agent at steady state, by taking the first block in the agent's corresponding row and summing its elements:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_1^\circ) = \begin{cases} 0 + 0.4045 + 0.1489 = 0.5534, & k = 6 \\ 0 + 0.5267 + 0.1183 = 0.6450, & k = 7 \\ 0 + 0.7099 + 0.0725 = 0.7824, & k = 8 \end{cases}$$

Likewise, we can compute the belief at θ_2° for each receiving agent at steady state, by taking the second block in the agent's corresponding row and summing its elements:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_2^\circ) = \begin{cases} 0.4466 + 0 = 0.4466, & k = 6 \\ 0.3550 + 0 = 0.3550, & k = 7 \\ 0.2176 + 0 = 0.2176, & k = 8 \end{cases}$$

We run this example for 7000 time iterations. We assigned to each agent an initial belief that is uniform over $\{\theta_1^\circ, \theta_2^\circ, \theta_3^\circ\}$. Figures 3.5 shows the evolution of $\boldsymbol{\mu}_{k,i}(\theta_1^\circ)$ and $\boldsymbol{\mu}_{k,i}(\theta_2^\circ)$ of agents in the receiving sub-network ($k = 6, 7, 8$). These figures show the convergence of the beliefs of the agents in the receiving sub-networks to the same probability distribution already computed according to the results of Theorem 2. Figure 3.6 shows this limiting distribution over Θ for all receiving agents.

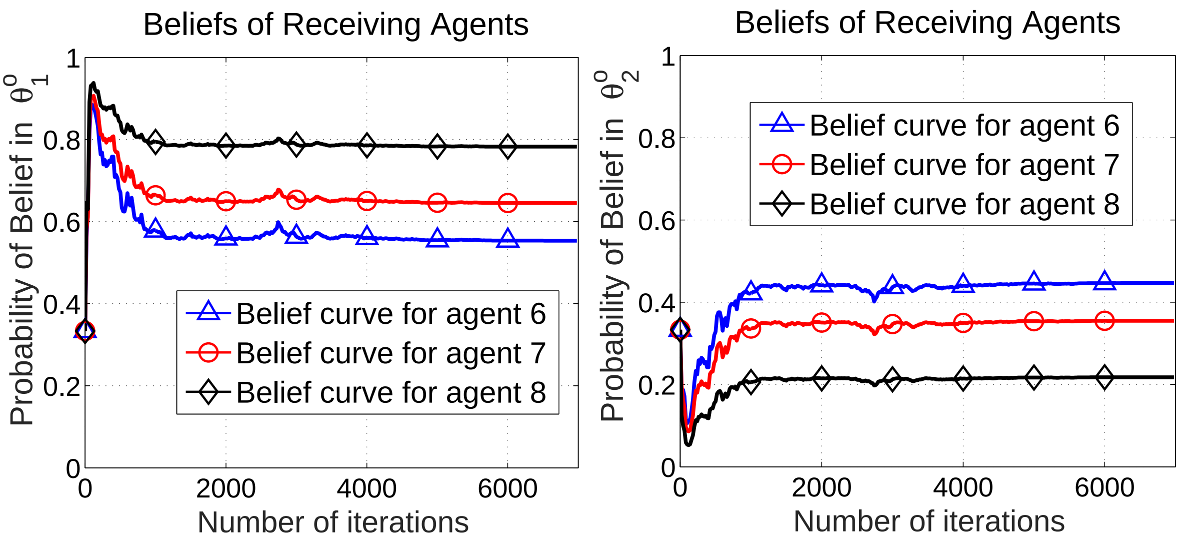

Figure 3.5: Evolution of agent k belief over time for $k = 6, 7, 8$ in the first case.

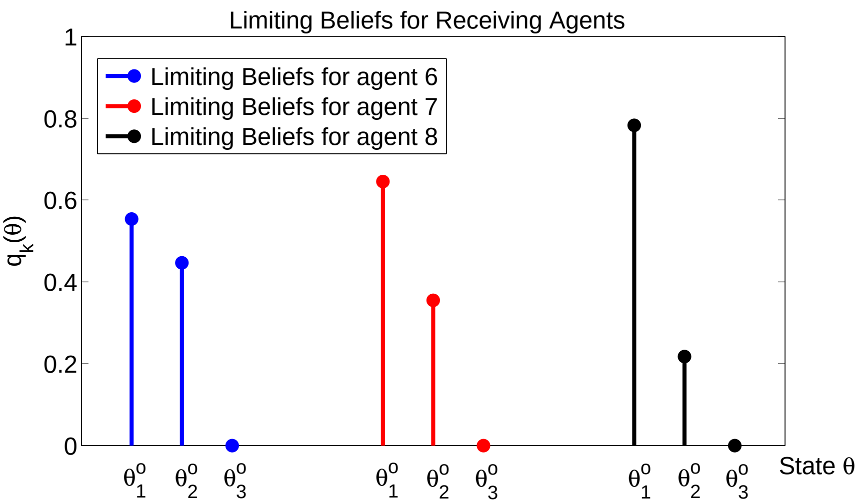

Figure 3.6: Limiting distribution of agent k, $q_k(\theta)$, over Θ for $k = 6, 7, 8$

3.4.2 Second Case

We now assume that the likelihood of the head signals for each agent k is selected as the following 3×8 matrix:

$$L(H) = \begin{bmatrix} 5/8 & 3/4 & 1/3 & 7/8 & 5/8 & 1/2 & 2/3 & 3/8 \\ 5/8 & 1/4 & 1/6 & 7/8 & 2/3 & 1/3 & 3/5 & 5/7 \\ 1/4 & 3/4 & 1/6 & 1/3 & 2/3 & 2/5 & 1/4 & 1/3 \end{bmatrix}$$

We observe now from $L(H)$ that assumption (3.11) is not met here where for agent k in the receiving sub-network ($k > 5$) we have $L_k(\zeta_k|\theta_1^\circ) \neq L_k(\zeta_k|\theta_2^\circ) \neq L_k(\zeta_k|\theta_3^\circ)$ for both cases in which ζ_k is either head or tail. Assumption (3.11) is not met here because θ_1° does not belong to the indistinguishable set of any receiving agent k in the receiving sub-network 3, i.e., $L_k(\zeta_k|\theta_1^\circ) \neq L_k(\zeta_k|\theta_3^\circ)$, and θ_2° does not belong to the indistinguishable set of any receiving agent k, i.e., $L_k(\zeta_k|\theta_2^\circ) \neq L_k(\zeta_k|\theta_3^\circ)$, where $k = 6, 7, 8$. We further assume that agents now update their beliefs according to the model described in (3.58). We choose $\gamma_{k,i} = 0.4$ for $k = 1, 2, 3$ (agents of the first sending sub-network) at any i, $\gamma_{k,i} = 0.5$ for $k = 4, 5$ (agents of the second sending sub-network) at any i and $\gamma_{k,i} = 0.1$ for $k = 6, 7, 8$ (agents of the receiving sub-network) at any i. We also assume that each agent starts at time $i = 0$ with an initial belief that is uniform over Θ. Then, we know from Theorem 3 that $\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_1^\circ) = 1$ for $k = 1, 2, 3$ and $\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_2^\circ) = 1$ for $k = 4, 5$. Figure 3.7 shows the evolution of $\boldsymbol{\mu}_{k,i}(\theta_1^\circ)$ and $\boldsymbol{\mu}_{k,i}(\theta_2^\circ)$ of agents in the receiving sub-network ($k = 6, 7, 8$). These figures show how the beliefs of the receiving agents are confined around the probability distribution already computed in the previous case.

3.5 Conclusions

In this chapter, we studied diffusion social learning over weakly-connected networks. We examined the circumstances under which receiving agents come under the total influence of sending agents. This total influence is reflected by forcing the receiving agents to focus

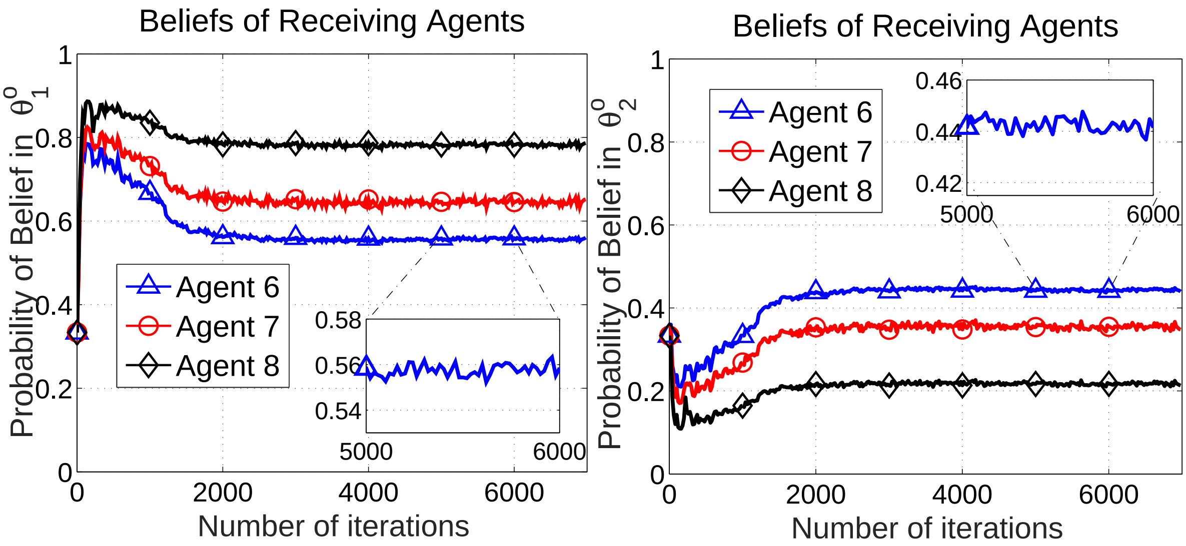

Figure 3.7: Evolution of agent k belief over time for $k = 6, 7, 8$ in the second case.

their beliefs on the set of true states for the sending sub-networks. We determined for each receiving agent what the exact probability distribution is in steady-state. We also illustrated the results with examples. Future work will focus on how the network can be designed so that receiving agents adopt specific limiting beliefs, and how receiving agents can detect the external influence and limit it. Results of this chapter are based on [64].

3.A Proof of Lemma 3

The proof is based on showing first that for any receiving agent k, it holds that

$$\lim_{i\to\infty} \sum_{\theta\in\Theta^{\bullet}} \boldsymbol{\mu}_{k,i}(\theta) = 1 \tag{3.63}$$

From this result, we will conclude that $\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) = 0$ for all $\theta \in \bar{\Theta}^{\bullet}$. To examine the evolution of agents' beliefs toward $\Theta^{\bullet}$, we associate with each agent k the following regret function:

$$Q^{W}(\boldsymbol{\mu}_{k,i}) \triangleq -\log\left(\sum_{\theta\in\Theta^{\bullet}} \boldsymbol{\mu}_{k,i}(\theta)\right) \tag{3.64}$$

We view $\boldsymbol{\mu}_{k,i}(\theta)$ as a stochastic process that depends on the sequence of random observations $\{\boldsymbol{\xi}_{k,j}\}$ over all k and for all $j \leq i$. Therefore, we shall examine agent k's individual performance by taking the expectation of $Q^{W}(\boldsymbol{\mu}_{k,i})$ over these observations. More specifically, we define agent k's risk at time i as

$$J^{W}(\boldsymbol{\mu}_{k,i}) \triangleq \mathbb{E}_{\mathcal{F}_i} Q^{W}(\boldsymbol{\mu}_{k,i}) \tag{3.65}$$

where $\mathcal{F}_i$ denotes the of sequence $\{\boldsymbol{\xi}_{k,j}\}$ over all k and for all $j \leq i$.

Proof of Lemma 3. We start with agent k's risk at time i defined in (3.65), where $k > N_{gS}$.

Recall that $N = N_{gS} + N_{gR}$ represents the total number of agents in the whole network:

$$
\begin{aligned}
J^W(\boldsymbol{\mu}_{k,i}) &= -\mathbb{E}_{\mathcal{F}_i} \log\left(\sum_{\theta\in\Theta^\bullet} \boldsymbol{\mu}_{k,i}(\theta)\right) \\
&\overset{(2.8)}{=} -\mathbb{E}_{\mathcal{F}_i} \log\left[\sum_{\ell=1}^{N}\sum_{\theta\in\Theta^\bullet} a_{\ell k}\boldsymbol{\psi}_{\ell,i}(\theta)\right] \\
&\overset{(2.5)}{=} -\mathbb{E}_{\mathcal{F}_i} \log\left[\sum_{\ell=1}^{N_{gS}} a_{\ell k}\sum_{\theta\in\Theta^\bullet}\boldsymbol{\psi}_{\ell,i}(\theta) + \sum_{\ell=N_{gS}+1}^{N} a_{\ell k}\frac{\sum_{\theta\in\Theta^\bullet}\boldsymbol{\mu}_{\ell,i-1}(\theta)L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right] \\
&\overset{(a)}{=} -\mathbb{E}_{\mathcal{F}_i} \log\left[\sum_{\ell=1}^{N_{gS}} a_{\ell k}\sum_{\theta\in\Theta^\bullet}\boldsymbol{\psi}_{\ell,i}(\theta) + \sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}\frac{\sum_{\theta\in\Theta^\bullet}\boldsymbol{\mu}_{\ell,i-1}(\theta)L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right] \\
&\overset{(b)}{=} -\mathbb{E}_{\mathcal{F}_i} \log\left[\sum_{\ell=1}^{N_{gS}} a_{\ell k}\sum_{\theta\in\Theta^\bullet}\boldsymbol{\psi}_{\ell,i}(\theta) + \sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}\frac{\sum_{\theta\in\Theta^\bullet}\boldsymbol{\mu}_{\ell,i-1}(\theta)L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_r^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right] \\
&\overset{(c)}{\le} -\mathbb{E}_{\mathcal{F}_i}\left[\sum_{\ell=1}^{N_{gS}} a_{\ell k}\log\left(\sum_{\theta\in\Theta^\bullet}\boldsymbol{\psi}_{\ell,i}(\theta)\right)\right. \\
&\quad\left. + \sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}\log\frac{\left(\sum_{\theta\in\Theta^\bullet}\boldsymbol{\mu}_{\ell,i-1}(\theta)\right)L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_r^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right] \\
&= -\sum_{\ell=1}^{N_{gS}} a_{\ell k}\mathbb{E}_{\mathcal{F}_i}\log\left(\sum_{\theta\in\Theta^\bullet}\boldsymbol{\psi}_{\ell,i}(\theta)\right) - \sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}\mathbb{E}_{\mathcal{F}_i}\log\left(\sum_{\theta\in\Theta^\bullet}\boldsymbol{\mu}_{\ell,i-1}(\theta)\right) \\
&\quad - \sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}\mathbb{E}_{\mathcal{F}_i}\log\left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_r^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right) \\
&\overset{(3.65)}{=} \sum_{\ell=1}^{N_{gS}} a_{\ell k}J^W(\boldsymbol{\psi}_{\ell,i}) + \sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}J^W(\boldsymbol{\mu}_{\ell,i-1}) \\
&\quad - \sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}\mathbb{E}_{\mathcal{F}_i}\log\left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_r^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right) \\
&\overset{(d)}{=} \sum_{\ell=1}^{N_{gS}} a_{\ell k}J^W(\boldsymbol{\psi}_{\ell,i}) + \sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}J^W(\boldsymbol{\mu}_{\ell,i-1}) - \mathbb{E}_{\mathcal{F}_{i-1}}\left(\sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}\right. \\
&\qquad\qquad \left.\mathbb{E}_{\boldsymbol{\xi}_{\ell,i}}\left(\log\left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_r^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right)|\mathcal{F}_{i-1}\right)\right) \\
&\overset{(e)}{\le} \sum_{r=S+1}^{S+R}\sum_{\ell\in\mathcal{I}_r} a_{\ell k}J^W(\boldsymbol{\mu}_{\ell,i-1}) + \sum_{\ell=1}^{N_{gS}} a_{\ell k}J^W(\boldsymbol{\psi}_{\ell,i}) \qquad (3.66)
\end{aligned}
$$

where

- in the third equality, we only expanded the second term that corresponds to receiving agents in order to study its behavior. We did not do the same thing with the first term because it corresponds to sending agents and we already know how that $\psi_{\ell,i}(\theta)$ will converge with time for any sending agent ℓ, as later shown in (3.74).

- in step (a), we split the second summation corresponding to receiving agents into R groups, with each group corresponding to one receiving sub-network. Moreover, the symbol $\mathcal{I}_r$ denotes the set of indexes of agents that belong to receiving sub-network r;

- in step (b), we replaced $L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta)$ by $L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta_r^\circ)$. This follows from assumption (3.11): for any θ that is in $\Theta^\bullet$, $L_\ell(\zeta_\ell|\theta) = L_\ell(\zeta_\ell|\theta_r^\circ)$, for any $\zeta_\ell \in Z_\ell$;

- in step (c), we applied the convexity property of $-\log(.)$ since the elements $\{a_{\ell k}\}$ form a convex combination for each agent k;

- in step (d), we applied the conditional expectation property

$$\left(\mathbb{E}_X[g(X)] = \mathbb{E}_Y[\mathbb{E}_{X|Y}[g(X)|Y]]\right)$$

  as follows:

$$\begin{aligned}
\mathbb{E}_{\mathcal{F}_i} \log\left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right) &= \mathbb{E}_{\mathcal{F}_{i-1}}\left(\mathbb{E}_{\mathcal{F}_i|\mathcal{F}_{i-1}}\left(\log\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}|\mathcal{F}_{i-1}\right)\right) \\
&= \mathbb{E}_{\mathcal{F}_{i-1}}\left(\mathbb{E}_{\boldsymbol{\xi}_{\ell,i}}\left(\log\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}|\mathcal{F}_{i-1}\right)\right) \qquad (3.67)
\end{aligned}$$

- in step (e), we replaced the previous expression in (d) by an upper bound using the non-negativity of the KL-divergence from $L_\ell(.|\theta_r^\circ)$ to $\boldsymbol{m}_{\ell,i-1}(.)$ [75].

To continue with the argument we collect the risk values of $S-$agents and $R-$agents into two vectors as follows:

$$J^W(\boldsymbol{\psi}_{\mathcal{S},i}) \triangleq col\left\{J^W(\boldsymbol{\psi}_{1,i}),\ldots,J^W(\boldsymbol{\psi}_{N_{gS},i})\right\} \tag{3.68}$$

$$J^W(\boldsymbol{\mu}_{\mathcal{R},i}) \triangleq col\left\{J^W(\boldsymbol{\mu}_{N_{gS}+1,i}),\ldots,J^W(\boldsymbol{\mu}_{N,i})\right\} \tag{3.69}$$

Then, from (3.66), we write the vector inequality:

$$J^W(\boldsymbol{\mu}_{\mathcal{R},i}) \preceq T_{RR}^{\mathsf{T}} J^W(\boldsymbol{\mu}_{\mathcal{R},i-1}) + T_{SR}^{\mathsf{T}} J^W(\boldsymbol{\psi}_{\mathcal{S},i}) \tag{3.70}$$

We now establish the convergence of this inequality. We first consider the term $J^W(\boldsymbol{\psi}_{\mathcal{S},i})$. We know that agents in the sending sub-networks can learn the truth if the assumptions mentioned in Lemma 1 and Theorem 1 are met. One of the assumptions is that at least one agent in each strongly-connected sub-network s starts with a non-zero prior belief at θ_s°. Let us denote this agent by ℓ_o. As shown in [30], this condition guarantees that for large enough i, $\boldsymbol{\mu}_{k,i}(\theta_s^\circ) > 0$ for all k in this sub-network. Accordingly, it also holds that for large enough i agents in this sub-network will have nonzero intermediate beliefs at θ_s°, i.e., $\boldsymbol{\psi}_{k,i}(\theta_s^\circ) > 0$. This implies that $J^W(\boldsymbol{\psi}_{\mathcal{S},i}) \succeq 0$ and $T_{SR}^{\mathsf{T}} J^W(\boldsymbol{\psi}_{\mathcal{S},i}) \succeq 0$ for large enough i since the elements of T_{SR} are all non-negative. Let us now consider agent k' of a receiving sub-network r, which has agent ℓ' from sending sub-network s in its neighborhood. After large enough i,

$$\boldsymbol{\mu}_{k',i}(\theta_s^\circ) = \sum_{\ell\in\mathcal{N}_{k'}} a_{\ell k'}\boldsymbol{\psi}_\ell(\theta_s^\circ) \geq a_{\ell' k'}\boldsymbol{\psi}_{\ell'}(\theta_s^\circ) > 0 \tag{3.71}$$

Then, in the next time step, all agents of sub-network r that have agent k' in their neighborhood will have non-zero belief at θ_s°. Since the received sub-network r is connected, it follows that after large enough i,

$$\boldsymbol{\mu}_{k,i}(\theta_s^\circ) > 0 \implies \sum_{\theta\in\Theta^\bullet} \boldsymbol{\mu}_{k,i}(\theta) > 0 \tag{3.72}$$

for all agents k that belong to sub-network r. We employ the same argument for all other receiving sub-networks. Therefore, $\sum_{\theta\in\Theta^\bullet} \boldsymbol{\mu}_{k,i}(\theta) > 0$ for any $k > N_{gR}$ so that $J^W(\boldsymbol{\mu}_{\mathcal{R},i}) \succeq 0$

for large enough i. Thus,

$$0 \preceq J^W(\boldsymbol{\mu}_{\mathcal{R},i}) \preceq T_{RR}^{\mathsf{T}} J^W(\boldsymbol{\mu}_{\mathcal{R},i-1}) + T_{SR}^{\mathsf{T}} J^W(\boldsymbol{\psi}_{\mathcal{S},i}) \tag{3.73}$$

Furthermore, any agent k in any sending sub-network s can learn asymptotically its own true state, so that $\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_s^\circ) \overset{a.s.}{=} 1$ implies

$$\lim_{i\to\infty} \boldsymbol{\psi}_{k,i}(\theta_s^\circ) = \lim_{i\to\infty} \frac{\boldsymbol{\mu}_{k,i}(\theta_s^\circ) L_k(\boldsymbol{\xi}_{k,i}|\theta_s^\circ)}{\sum_{\theta\in\Theta} \boldsymbol{\mu}_{k,i}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)} = \lim_{i\to\infty} \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta_s^\circ)}{L_k(\boldsymbol{\xi}_{k,i}|\theta_s^\circ)} \overset{a.s.}{=} 1 \tag{3.74}$$

The denominator in the second equality follows from the fact that

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_s^\circ) \overset{a.s.}{=} 1 \tag{3.75}$$

for any agent k of sending sub-network s. It follows that

$$\lim_{i\to\infty} \sum_{\theta\in\Theta^\bullet} \boldsymbol{\psi}_{k,i}(\theta) \overset{a.s.}{=} 1 \tag{3.76}$$

for any $k \le N_{gS}$. Therefore, $\lim\limits_{i\to\infty} J^W(\boldsymbol{\psi}_{S,i}) = 0$. Moreover, since $\rho(T_{RR}) < 1$ [67], we conclude that

$$\lim_{i\to\infty} J^W(\boldsymbol{\mu}_{\mathcal{R},i}) = 0 \implies \lim_{i\to\infty} J^W(\boldsymbol{\mu}_{k,i}) = 0, \quad \forall\, k > N_{gS} \tag{3.77}$$

As previously discussed after large enough i,

$$\sum_{\theta\in\Theta^\bullet} \boldsymbol{\mu}_{k,i}(\theta) > 0 \implies -\log\left(\sum_{\theta\in\Theta^\bullet} \boldsymbol{\mu}_{k,i}(\theta)\right) \ge 0 \tag{3.78}$$

Using the definition of $J^W(\boldsymbol{\mu}_{k,i})$ in (3.65), it holds that $J^W(\boldsymbol{\mu}_{k,i})$ represents the expectation over $\mathcal{F}_i$ of non-negative quantities. Hence, result (3.77) implies

$$\lim_{i\to\infty} -\log\left(\sum_{\theta\in\Theta^\bullet} \boldsymbol{\mu}_{k,i}(\theta)\right) = 0 \,;\, \lim_{i\to\infty} \sum_{\theta\in\Theta^\bullet} \boldsymbol{\mu}_{k,i}(\theta) \overset{a.s.}{=} 1 \tag{3.79}$$

■

3.B Proof of Lemma 4

Assume agent k belongs to sub-network r and $\zeta_k \in \boldsymbol{Z}_k$:

$$\lim_{i\to\infty} \boldsymbol{m}_{k,i}(\zeta_k) = \lim_{i\to\infty} \sum_{\theta\in\Theta} \boldsymbol{\mu}_{k,i}(\theta) L_k(\zeta_k|\theta) \overset{(a)}{=} \lim_{i\to\infty} \sum_{\theta\in\Theta^\bullet} \boldsymbol{\mu}_{k,i}(\theta) L_k(\zeta_k|\theta)$$
$$\overset{(b)}{=} \left(\lim_{i\to\infty} \sum_{\theta\in\Theta^\bullet} \boldsymbol{\mu}_{k,i}(\theta)\right) L_k(\zeta_k|\theta_r^\circ) \overset{a.s.}{=} L_k(\zeta_k|\theta_r^\circ) \tag{3.80}$$

where step (a) follows from the result of Lemma 3 and step (b) follows from assumption (3.11). ∎

3.C Proof of Theorem 2

The intermediate belief of any agent k is given by:

$$\boldsymbol{\psi}_{k,i}(\theta) = \frac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\boldsymbol{m}_{k,i-1}(\boldsymbol{\xi}_{k,i})} \tag{3.81}$$

Let us assume that agent k belongs to receiving sub-network r. Using Lemma 4, we have for any $\theta \in \Theta^\bullet$:

$$\lim_{i\to\infty} \boldsymbol{\psi}_{k,i}(\theta) = \lim_{i\to\infty} \frac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\boldsymbol{m}_{k,i-1}(\boldsymbol{\xi}_{k,i})} = \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) \tag{3.82}$$

We can establish the same property for any agent in a sending sub-network because (3.80) was already proven for sending agents in [30]. It follows that, for any agent k,

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) = \lim_{i\to\infty} \sum_{\ell\in\mathcal{N}_k} a_{\ell k} \boldsymbol{\mu}_{k,i-1}(\theta) \tag{3.83}$$

for any $\theta \in \Theta^{\bullet}$. We defined the vectors $\boldsymbol{\mu}_{\mathcal{S},i}(\theta)$ in (3.19) and $\boldsymbol{\mu}_{\mathcal{R},i}(\theta)$ in (3.15). Then,

$$\lim_{i\to\infty} \begin{bmatrix} \boldsymbol{\mu}_{\mathcal{S},i}(\theta) \\ \boldsymbol{\mu}_{\mathcal{R},i}(\theta) \end{bmatrix} = A^{\mathsf{T}} \left(\lim_{i\to\infty} \begin{bmatrix} \boldsymbol{\mu}_{\mathcal{S},i-1}(\theta) \\ \boldsymbol{\mu}_{\mathcal{R},i-1}(\theta) \end{bmatrix} \right) \tag{3.84}$$

from which we obtain using the structure of A in (3.5):

$$\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) = T_{SR}^{\mathsf{T}} \lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta) + T_{RR}^{\mathsf{T}} \lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) \tag{3.85}$$

We then conclude that

$$\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) = (I - T_{RR}^{\mathsf{T}})^{-1} T_{SR}^{\mathsf{T}} \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta) \right) = W^{\mathsf{T}} \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta) \right) \tag{3.86}$$

■

3.D Proof of Lemma 5

We start by introducing some notation and definitions. Since we are now interested in examining the evolution of the agents' beliefs toward the true state, let us introduce the true probability mass function $p(\theta)$ defined over Θ, namely:

$$p(\theta) = \delta_{\theta,\theta^\circ} \triangleq \begin{cases} 1, & \text{if} \quad \theta = \theta^\circ \\ 0, & \text{otherwise} \end{cases} \tag{3.87}$$

The evolution of the belief of agent k toward the true state can be analyzed by computing the KL divergence of $\boldsymbol{\mu}_{k,i}(\theta)$ from $p(\theta)$ at each time instant i. We therefore introduce the new regret function for agent k at time i as:

$$Q(\boldsymbol{\mu}_{k,i}) \triangleq D_{KL}(p||\boldsymbol{\mu}_{k,i}) = \sum_{\theta\in\Theta} p(\theta) \log \left(\frac{p(\theta)}{\boldsymbol{\mu}_{k,i}(\theta)} \right) = -\log \boldsymbol{\mu}_{k,i}(\theta^\circ) \tag{3.88}$$

where we used the convention that $0 \log 0 = 0$. We shall again define agent k's individual risk at time i as

$$J(\boldsymbol{\mu}_{k,i}) \triangleq \mathbb{E}_{\mathcal{F}_i} Q(\boldsymbol{\mu}_{k,i}) = -\mathbb{E}_{\mathcal{F}_i} \log \boldsymbol{\mu}_{k,i}(\theta^\circ) \tag{3.89}$$

where $\mathcal{F}_i$ denotes the history of $\{\boldsymbol{\xi}_{k,j}\}$ over all k and for all $j \le i$. We then assess the overall network performance by considering the weighted aggregate risk:

$$J(\boldsymbol{\mu}_i) \triangleq \sum_{k=1}^{N} y(k) J(\boldsymbol{\mu}_{k,i}) \tag{3.90}$$

where the $\{y(k)\}$ denote the entries of the Perron vector, y, of the primitive left-stochastic matrix A, as defined by (2.2). To prove Lemma 5, namely, the ability of agents to arrive at correct forecasts, we prove first the convergence of the sequence $\{J(\boldsymbol{\mu}_i)\}$ as $i \to \infty$. This convergence will then imply the correct forecasting by agents.

Proof of Lemma 5: We assumed in the statement of the lemma that at least one agent ℓ_o starts with a non-zero prior belief at θ°, i.e., $\boldsymbol{\mu}_{\ell_o,0}(\theta^\circ) > 0$. As shown in [30], this condition guarantees that for large enough i, $\boldsymbol{\mu}_{k,i}(\theta^\circ) > 0$ for all $k \in \mathcal{N}$, which implies that the terms of the time sequence $\{Q(\boldsymbol{\mu}_{k,i})\}$ assume nonnegative values for large i and for any agent k. Thus, the time sequences $\{J(\boldsymbol{\mu}_{k,i})\}$ and $\{J(\boldsymbol{\mu}_i)\}$ are non-negative for large enough i. Let us now expand agent k's risk for large time i:

$$\begin{aligned}
J(\boldsymbol{\mu}_{k,i}) &= -\mathbb{E}_{\mathcal{F}_i} \log \boldsymbol{\mu}_{k,i}(\theta^\circ) = -\mathbb{E}_{\mathcal{F}_i} \log \left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \boldsymbol{\psi}_{\ell,i}(\theta^\circ) \right) \\
&\overset{(a)}{\leq} -\mathbb{E}_{\mathcal{F}_i} \left[\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \left(\boldsymbol{\psi}_{\ell,i}(\theta^\circ) \right) \right] \\
&\overset{(3.58)}{=} -\mathbb{E}_{\mathcal{F}_i} \left[\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \left((1-\gamma_{\ell,i}) \left(\boldsymbol{\mu}_{\ell,i-1}(\theta^\circ) \right) + \gamma_{\ell,i} \left(\frac{\boldsymbol{\mu}_{\ell,i-1}(\theta^\circ) L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) \right) \right] \\
&\overset{(b)}{\leq} -\mathbb{E}_{\mathcal{F}_i} \left[\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \left((1-\gamma_{\ell,i}) \log \left(\boldsymbol{\mu}_{\ell,i-1}(\theta^\circ) \right) + \gamma_{\ell,i} \log \left(\frac{\boldsymbol{\mu}_{\ell,i-1}(\theta^\circ) L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) \right) \right] \\
&= -\mathbb{E}_{\mathcal{F}_i} \left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \left(\boldsymbol{\mu}_{\ell,i-1}(\theta^\circ) \right) \right) - \mathbb{E}_{\mathcal{F}_i} \left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \gamma_{\ell,i} \log \left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) \right) \\
&\overset{(c)}{=} -\mathbb{E}_{\mathcal{F}_i} \left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \left(\boldsymbol{\mu}_{\ell,i-1}(\theta^\circ) \right) \right) \\
&\quad - \mathbb{E}_{\mathcal{F}_{i-1}} \left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \gamma_{\ell,i} \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log \left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) |\mathcal{F}_{i-1} \right] \right) \\
&\overset{(d)}{\leq} -\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \mathbb{E}_{\mathcal{F}_i} \log \left(\boldsymbol{\mu}_{\ell,i-1}(\theta^\circ) \right) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} J(\boldsymbol{\mu}_{\ell,i-1})
\end{aligned} \tag{3.91}$$

where

- steps (a) and (b) follow from the convexity of $-\log(.)$;
- step (c) follows from the conditional expectation property

$$\left(\mathbb{E}_X[g(X)] = \mathbb{E}_Y[\mathbb{E}_{X|Y}[g(X)|Y]] \right)$$

  as in (3.67);
- step (d) follows by replacing the expression in (c) by an upper bound using the non-negativity of the KL divergence from $L_\ell(.|\theta^\circ)$ to $\boldsymbol{m}_{\ell,i-1}(.)$ according to Gibb's inequality [75].

Accordingly, the overall performance at time i, satisfies:

$$J(\boldsymbol{\mu}_i) \overset{(a)}{\leq} \sum_{k=1}^{N} y(k) \sum_{\ell \in \mathcal{N}_k} a_{\ell k} J(\boldsymbol{\mu}_{\ell,i-1}) \overset{(b)}{=} \sum_{\ell=1}^{N} y(\ell) J(\boldsymbol{\mu}_{\ell,i-1}) = J(\boldsymbol{\mu}_{i-1}) \tag{3.92}$$

where step (a) follows from (3.91), and step (b) follows from (2.2). Therefore, the sequence $\{J(\boldsymbol{\mu}_i)\}$ is a decreasing sequence. But, since this sequence is non-negative, we conclude that $\{J(\boldsymbol{\mu}_i)\}$ converges to a real number according to the monotone convergence theorem of real numbers [76].

We now establish the ability of agents to attain correct predictions. From step (c) in (3.91), we get

$$J(\boldsymbol{\mu}_{k,i}) \leq \sum_{\ell \in \mathcal{N}_k} a_{\ell k} J(\boldsymbol{\mu}_{\ell,i-1}) - \mathbb{E}_{\mathcal{F}_{i-1}} \left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \gamma_{\ell,i} \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log \left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) |\mathcal{F}_{i-1} \right] \right) \tag{3.93}$$

Then, rearranging terms,

$$\sum_{\ell \in \mathcal{N}_k} a_{\ell k} J(\boldsymbol{\mu}_{\ell,i-1}) - J(\boldsymbol{\mu}_{k,i}) \geq \mathbb{E}_{\mathcal{F}_{i-1}} \left(\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \gamma_{\ell,i} \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log \left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) |\mathcal{F}_{i-1} \right] \right) \tag{3.94}$$

Scaling by $y(k)$, summing over k, and using (2.2) we get:

$$\sum_{\ell=1}^{N} y(\ell) J(\boldsymbol{\mu}_{\ell,i-1}) - \sum_{k=1}^{N} y(k) J(\boldsymbol{\mu}_{k,i}) \geq \mathbb{E}_{\mathcal{F}_{i-1}} \left(\sum_{\ell \in \mathcal{N}_k} y(\ell) \gamma_{\ell,i} \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log \left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) |\mathcal{F}_{i-1} \right] \right) \tag{3.95}$$

Then,

$$J(\boldsymbol{\mu}_{i-1}) - J(\boldsymbol{\mu}_i) \geq \mathbb{E}_{\mathcal{F}_{i-1}} \left(\sum_{\ell \in \mathcal{N}_k} y(\ell) \gamma_{\ell,i} \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log \left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) |\mathcal{F}_{i-1} \right] \right) \tag{3.96}$$

Since $\{J(\boldsymbol{\mu}_i)\}$ is a convergent sequence, it is also a Cauchy sequence [76] and, therefore,

$$0 = \lim_{i\to\infty} \left[J(\boldsymbol{\mu}_{i-1}) - J(\boldsymbol{\mu}_i)\right] \geq \lim_{i\to\infty} \mathbb{E}_{\mathcal{F}_{i-1}} \left(\sum_{\ell\in\mathcal{N}_k} y(\ell)\gamma_{\ell,i} \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log\left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) |\mathcal{F}_{i-1} \right] \right) \geq 0 \tag{3.97}$$

where the rightmost inequality follows from the non-negativity of the KL-divergence. We conclude that:

$$\lim_{i\to\infty} \mathbb{E}_{\mathcal{F}_{i-1}} \left(\sum_{\ell\in\mathcal{N}_k} y(\ell)\gamma_{\ell,i} \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log\left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) |\mathcal{F}_{i-1} \right] \right) = 0 \tag{3.98}$$

Since we assumed that $\lim\limits_{i\to\infty} \gamma_{k,i} \neq 0$ for any k, $y(\ell) > 0$ from (2.2), and

$$\mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log\left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) |\mathcal{F}_{i-1} \right] \geq 0$$

from the non-negativity of the KL-divergence, then

$$\lim_{i\to\infty} \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log\left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^\circ)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \right) |\mathcal{F}_{i-1} \right] = 0 \tag{3.99}$$

Thus,

$$\lim_{i\to\infty} \sum_{\zeta_\ell\in Z_\ell} L_\ell(\zeta_\ell|\theta^\circ) \log\left(\frac{L_\ell(\zeta_\ell|\theta^\circ)}{m_{\ell,i-1}(\zeta_\ell)} \right) = 0 \tag{3.100}$$

Let

$$f_{\ell,i-1} \triangleq \sum_{\zeta_\ell\in Z_\ell} L_\ell(\zeta_\ell|\theta^\circ) \log\left(\frac{L_\ell(\zeta_\ell|\theta^\circ)}{m_{\ell,i-1}(\zeta_\ell)} \right) \tag{3.101}$$

where $f_{\ell,i}$ represents the KL-divergence of $m_{\ell,i}(.)$ from $L_\ell(.|\theta^\circ)$. We know from Gibb's inequality [75] that the KL-divergence of a probability distribution from another distribution achieves the value zero only when the two distributions are equal. Since the KL-divergence $f_{\ell,i}$ converges to zero as $i \to \infty$ and $L_\ell(.|\theta^\circ)$ is a fixed distribution, this implies that $m_{\ell,i}(.)$

should converge, i.e., its limit exists and it takes the following value:

$$\lim_{i\to\infty} m_{\ell,i}(\zeta_\ell) = L_\ell(\zeta_\ell|\theta^\circ) \tag{3.102}$$

for any $\zeta_\ell \in Z_\ell$. Since this result is achieved for any realization of observational signals $\mathcal{F}_{i-1}$, we conclude that:

$$\lim_{i\to\infty} \boldsymbol{m}_{\ell,i}(\zeta_\ell) \overset{a.s.}{=} L_\ell(\zeta_\ell|\theta^\circ) \tag{3.103}$$

for any $\ell \in \mathcal{N}$ and any $\zeta_\ell \in \boldsymbol{Z}_\ell$. ■

3.E Proof of Theorem 3

We have established that self-aware agents are able to make correct forecast, i.e.,

$$\lim_{i\to\infty} \boldsymbol{m}_{k,i}(\zeta_k) \overset{a.s.}{=} L_k(\zeta_k|\theta^\circ) \tag{3.104}$$

for any $k \in \mathcal{N}$ and any $\zeta_k \in \boldsymbol{Z}_k$. From (3.104), we have:

$$\sum_{\theta\in\Theta} L_k(\zeta_k|\theta) \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) \overset{a.s.}{=} L_k(\zeta_k|\theta^\circ) \tag{3.105}$$

which implies the following for any observational signal $\zeta_k \in Z_k$:

$$\begin{aligned}
\sum_{\theta\in\Theta_k} L_k(\zeta_k|\theta) \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) + \sum_{\theta\in\bar{\Theta}_k} L_k(\zeta_k|\theta) \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) &\overset{a.s.}{=} L_k(\zeta_k|\theta^\circ) \\
L_k(\zeta_k|\theta^\circ) \sum_{\theta\in\Theta_k} \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) + \sum_{\theta\in\bar{\Theta}_k} L_k(\zeta_k|\theta) \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) &\overset{a.s.}{=} L_k(\zeta_k|\theta^\circ) \\
L_k(\zeta_k|\theta^\circ) \left(\sum_{\theta\in\Theta_k} \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) - 1\right) + \sum_{\theta\in\bar{\Theta}_k} L_k(\zeta_k|\theta) \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) &\overset{a.s.}{=} 0 \\
-L_k(\zeta_k|\theta^\circ) \sum_{\theta\in\bar{\Theta}_k} \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) + \sum_{\theta\in\bar{\Theta}_k} L_k(\zeta_k|\theta) \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) &\overset{a.s.}{=} 0 \\
\sum_{\theta\in\bar{\Theta}_k} \left[L_k(\zeta_k|\theta) - L_k(\zeta_k|\theta^\circ)\right] \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) &\overset{a.s.}{=} 0 \qquad (3.106)
\end{aligned}$$

We know according to (2.12), that there exists a prevailing signal ζ_k° where

$$L_k(\zeta_k^\circ|\theta^\circ) - L_k(\zeta_k^\circ|\theta) > 0, \quad \forall \theta \in \bar{\Theta}_k \tag{3.107}$$

Applying (3.106) to the prevailing signal, we have:

$$\sum_{\theta \in \bar{\Theta}_k} [L_k(\zeta_k^\circ|\theta) - L_k(\zeta_k^\circ|\theta^\circ)] \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) \overset{a.s.}{=} 0 \tag{3.108}$$

From (3.108) and (3.107), we conclude that:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) \overset{a.s.}{=} 0, \ \text{ for any } \theta \in \bar{\Theta}_k \tag{3.109}$$

Now since,

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \lim_{i\to\infty} \boldsymbol{\psi}_{\ell,i}(\theta) \tag{3.110}$$

we have for any neighbor ℓ of agent k,

$$\lim_{i\to\infty} \boldsymbol{\psi}_{\ell,i}(\theta) \overset{a.s.}{=} 0 \implies \lim_{i\to\infty} \boldsymbol{\mu}_{\ell,i}(\theta) \overset{a.s.}{=} 0, \ \text{ for any } \theta \in \bar{\Theta}_k \text{ and any } \ell \in \mathcal{N}_k \tag{3.111}$$

Sine the network is strongly-connected, by propagating the same argument, we end up having:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{\ell,i}(\theta) \overset{a.s.}{=} 0, \ \text{ for any } \theta \in \bar{\Theta}_k \text{ and any } \ell \in \mathcal{N} \tag{3.112}$$

By repeating the same steps for all agent k, all agents will assign a zero probability to the following set:

$$\bigcup_{k \in \mathcal{N}} \bar{\Theta}_k \implies \overline{\bigcap_{k \in \mathcal{N}} \Theta_k} \implies \Theta \backslash \{\theta^\circ\} \tag{3.113}$$

We therefore have for any agent k:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) \overset{a.s.}{=} 0, \quad \text{for any } \theta \neq \theta^{\circ} \implies \lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta^{\circ}) \overset{a.s.}{=} 1 \tag{3.114}$$

■

3.F Proof of Theorem 4

According to model (3.58), the intermediate belief of any agent k in a receiving group can be written as follows:

$$\boldsymbol{\psi}_{k,i}(\theta) = \boldsymbol{\mu}_{k,i-1}(\theta) + \gamma_{k,i}\left[\boldsymbol{\mu}_{k,i-1}(\theta)\left(\frac{L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta'} \boldsymbol{\mu}_{k,i-1}(\theta')L_k(\boldsymbol{\xi}_{k,i}|\theta')} - 1\right)\right] \tag{3.115}$$

We assume that $\gamma_{k,i} = \tau_{k,i}\gamma_{\max}$, where $\tau_{k,i}$ and $\gamma_{\max}$ are both nonnegative scalars less than one. Then,

$$\boldsymbol{\psi}_{k,i}(\theta) = \boldsymbol{\mu}_{k,i-1}(\theta) + \gamma_{\max}\left[\tau_{k,i}\boldsymbol{\mu}_{k,i-1}(\theta)\left(\frac{L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta'} \boldsymbol{\mu}_{k,i-1}(\theta')L_k(\boldsymbol{\xi}_{k,i}|\theta')} - 1\right)\right] \tag{3.116}$$

We define the auxiliary function:

$$\boldsymbol{h}_{k,i}(\theta, \zeta_k) \;\triangleq\; \tau_{k,i}\boldsymbol{\mu}_{k,i-1}(\theta)\left(\frac{L_k(\zeta_k|\theta)}{\sum_{\theta'} \boldsymbol{\mu}_{k,i-1}(\theta')L_k(\zeta_k|\theta')} - 1\right) \tag{3.117}$$

where $\theta \in \Theta$ and $\zeta_k \in Z_k$, so that

$$\boldsymbol{\psi}_{k,i}(\theta) = \boldsymbol{\mu}_{k,i-1}(\theta) + \gamma_{\max}\boldsymbol{h}_{k,i}(\theta, \boldsymbol{\xi}_{k,i}) \tag{3.118}$$

Therefore,

$$\boldsymbol{\mu}_{k,i}(\theta) = \sum_{\ell\in\mathcal{N}_k} a_{\ell k}\boldsymbol{\psi}_{\ell,i}(\theta) = \sum_{\ell\in\mathcal{N}_k} a_{\ell k}\boldsymbol{\mu}_{\ell,i-1}(\theta) + \gamma_{\max}\sum_{\ell\in\mathcal{N}_k} a_{\ell k}\boldsymbol{h}_{\ell,i}(\theta, \boldsymbol{\xi}_{\ell,i}) \tag{3.119}$$

Let us introduce the vectors:

$$\boldsymbol{h}_{\mathcal{S},i}(\theta, \boldsymbol{\xi}_{S,i}) \triangleq \text{col}\left\{\boldsymbol{h}_{1,i}(\theta, \boldsymbol{\xi}_{1,i}), \ldots, \boldsymbol{h}_{N_{gS},i}(\theta, \boldsymbol{\xi}_{N_{gS},i})\right\} \tag{3.120}$$

$$\boldsymbol{h}_{\mathcal{R},i}(\theta, \boldsymbol{\xi}_{R,i}) \triangleq \text{col}\left\{\boldsymbol{h}_{N_{gS}+1,i}(\theta, \boldsymbol{\xi}_{N_{gS}+1,i}), \ldots, \boldsymbol{h}_{N,i}(\theta, \boldsymbol{\xi}_{N,i})\right\} \tag{3.121}$$

and,

$$\boldsymbol{\xi}_{\mathcal{S},i} \triangleq \text{col}\left\{\boldsymbol{\xi}_{1,i}, \ldots, \boldsymbol{\xi}_{N_{gS},i}\right\} \tag{3.122}$$

$$\boldsymbol{\xi}_{\mathcal{R},i} \triangleq \text{col}\left\{\boldsymbol{\xi}_{N_{gS}+1,i}, \ldots, \boldsymbol{\xi}_{N,i}\right\} \tag{3.123}$$

Recall that we defined the vectors $\boldsymbol{\mu}_{\mathcal{S},i}(\theta)$ in (3.19) and $\boldsymbol{\mu}_{\mathcal{R},i}(\theta)$ in (3.15). Then, we have

$$\begin{bmatrix} \boldsymbol{\mu}_{\mathcal{S},i}(\theta) \\ \boldsymbol{\mu}_{\mathcal{R},i}(\theta) \end{bmatrix} = A^{\mathsf{T}} \left(\begin{bmatrix} \boldsymbol{\mu}_{\mathcal{S},i-1}(\theta) \\ \boldsymbol{\mu}_{\mathcal{R},i-1}(\theta) \end{bmatrix} + \gamma_{\max} \begin{bmatrix} \boldsymbol{h}_{\mathcal{S},i}(\theta, \boldsymbol{\xi}_{\mathcal{S},i}) \\ \boldsymbol{h}_{\mathcal{R},i}(\theta, \boldsymbol{\xi}_{\mathcal{R},i}) \end{bmatrix} \right) \tag{3.124}$$

Using the structure of A in (3.5), it follows that

$$\boldsymbol{\mu}_{\mathcal{R},i}(\theta) = T_{RR}^{\mathsf{T}} \boldsymbol{\mu}_{\mathcal{R},i-1}(\theta) + T_{SR}^{\mathsf{T}} \boldsymbol{\mu}_{\mathcal{S},i-1}(\theta) + \gamma_{\max}\left(T_{SR}^{\mathsf{T}} \boldsymbol{h}_{\mathcal{S},i}(\theta, \boldsymbol{\xi}_{\mathcal{S},i}) + T_{RR}^{\mathsf{T}} \boldsymbol{h}_{\mathcal{R},i}(\theta, \boldsymbol{\xi}_{\mathcal{R},i})\right) \tag{3.125}$$

We study the convergence of this recursion. Let

$$\boldsymbol{\zeta}_{\mathcal{S}} \triangleq \text{col}\left\{\boldsymbol{\zeta}_1, \ldots, \boldsymbol{\zeta}_{N_{gS}}\right\}, \ \boldsymbol{\zeta}_{\mathcal{R}} \triangleq \text{col}\left\{\boldsymbol{\zeta}_{N_{gS}+1}, \ldots, \boldsymbol{\zeta}_N\right\} \tag{3.126}$$

We will first establish that

$$\gamma_{\max}\left(T_{SR}^{\mathsf{T}} \boldsymbol{h}_{\mathcal{S},i}(\theta, \zeta_{\mathcal{S}}) + T_{RR}^{\mathsf{T}} \boldsymbol{h}_{\mathcal{R},i}(\theta, \zeta_{\mathcal{R}})\right) = O(\gamma_{\max}) \tag{3.127}$$

for any θ, $\zeta_{\mathcal{S}}$ and $\zeta_{\mathcal{R}}$.

Lemma 6. *For any $k \in \mathcal{N}$, $i \geq 0$, $\theta \in \Theta$ and $\zeta_k \in Z_k$, it holds that*

$$|\boldsymbol{h}_{k,i}(\theta, \zeta_k)| \leq 1 \tag{3.128}$$

Proof. From (3.117),

$$\boldsymbol{h}_{k,i}(\theta, \zeta_k) = \tau_{k,i} \left(\frac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\zeta_k|\theta)}{\sum_{\theta'} \boldsymbol{\mu}_{k,i-1}(\theta') L_k(\zeta_k|\theta')} - \boldsymbol{\mu}_{k,i-1}(\theta) \right) \tag{3.129}$$

Since $\tau_{k,i}$ is a nonnegative scalar that is less than one, and since

$$0 \leq \frac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\zeta_k|\theta)}{\sum_{\theta'} \boldsymbol{\mu}_{k,i-1}(\theta') L_k(\zeta_k|\theta')} \leq 1 \tag{3.130}$$

for any $k \in \mathcal{N}$, $i \geq 0$, $\theta \in \Theta$ and $\zeta_k \in Z_k$, we conclude that

$$\boldsymbol{h}_{k,i}(\theta, \zeta_k) \geq -\tau_{k,i} \boldsymbol{\mu}_{k,i-1}(\theta) \geq -\boldsymbol{\mu}_{k,i-1}(\theta) \tag{3.131}$$

and

$$\boldsymbol{h}_{k,i}(\theta, \zeta_k) \leq \tau_{k,i} \left(1 - \boldsymbol{\mu}_{k,i-1}(\theta)\right) \leq 1 - \boldsymbol{\mu}_{k,i-1}(\theta) \tag{3.132}$$

Moreover, we know that $0 \leq \boldsymbol{\mu}_{k,i}(\theta) \leq 1$ for all k, i and θ. We then conclude that

$$-1 \leq \boldsymbol{h}_{k,i}(\theta, \zeta_k) \leq 1 \tag{3.133}$$

□

From (3.128), we get for any $i \geq 0$ and $\theta \in \Theta$,

$$\begin{aligned} \left|T_{SR}^{\mathsf{T}} \boldsymbol{h}_{\mathcal{S},i}(\theta, \zeta_{\mathcal{S}}) + T_{RR}^{\mathsf{T}} \boldsymbol{h}_{\mathcal{R},i}(\theta, \zeta_{\mathcal{R}})\right| &\preceq T_{SR}^{\mathsf{T}} |\boldsymbol{h}_{\mathcal{S},i}(\theta, \zeta_{\mathcal{S}})| + T_{RR}^{\mathsf{T}} |\boldsymbol{h}_{\mathcal{R},i}(\theta, \zeta_{\mathcal{R}})| \\ &\overset{(a)}{\preceq} T_{SR}^{\mathsf{T}} \mathbb{1}_{N_{gS}} + T_{RR}^{\mathsf{T}} \mathbb{1}_{N_{gR}} \overset{(b)}{=} \mathbb{1}_{N_{gR}} \end{aligned} \tag{3.134}$$

where (a) follows from (3.128) and (b) follows from the left-stochasticity of the combination matrix A. Note that the above inequality, as well as the absolute value operator, are element-wise. Moreover, $\mathbb{1}_{N_{gR}}$ is a vector of all ones of size N_{gR} and $\mathbb{1}_{N_{gS}}$ is a vector of all ones of size N_{gS}. Thus,

$$\gamma_{\max} \left| T_{SR}^{\mathsf{T}} \boldsymbol{h}_{\mathcal{S},i}(\theta, \zeta_{\mathcal{S}}) + T_{RR}^{\mathsf{T}} \boldsymbol{h}_{\mathcal{R},i}(\theta, \zeta_{\mathcal{R}}) \right| \preceq \gamma_{\max} \mathbb{1}_{N_{gR}} \tag{3.135}$$

for all $i \geq 0$. This fact leads to the desired conclusion (3.127). In this way, equality (3.125) implies:

$$\begin{aligned} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) &\preceq T_{RR}^{\mathsf{T}} \boldsymbol{\mu}_{\mathcal{R},i-1}(\theta) + T_{SR}^{\mathsf{T}} \boldsymbol{\mu}_{\mathcal{S},i-1}(\theta) + \gamma_{\max} \mathbb{1}_{N_{gR}} \\ \boldsymbol{\mu}_{\mathcal{R},i}(\theta) &\succeq T_{RR}^{\mathsf{T}} \boldsymbol{\mu}_{\mathcal{R},i-1}(\theta) + T_{SR}^{\mathsf{T}} \boldsymbol{\mu}_{\mathcal{S},i-1}(\theta) - \gamma_{\max} \mathbb{1}_{N_{gR}} \end{aligned} \tag{3.136}$$

We have $\rho(T_{RR}^{\mathsf{T}}) < 1$ and $\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta)$ exists since agents of sending sub-networks can learn asymptotically the truth. Then,

$$\begin{aligned} \limsup_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) &\preceq T_{RR}^{\mathsf{T}} \left(\limsup_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i-1}(\theta) \right) + T_{SR}^{\mathsf{T}} \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i-1}(\theta) \right) + \gamma_{\max} \mathbb{1}_{N_{gR}} \\ \liminf_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) &\succeq T_{RR}^{\mathsf{T}} \left(\liminf_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i-1}(\theta) \right) + T_{SR}^{\mathsf{T}} \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i-1}(\theta) \right) - \gamma_{\max} \mathbb{1}_{N_{gR}} \end{aligned} \tag{3.137}$$

It follows that

$$\begin{aligned} \limsup_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) &\preceq W^{\mathsf{T}} \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta) \right) + \gamma_{\max} C \mathbb{1}_{N_{gR}} \\ \liminf_{i\to\infty} \boldsymbol{\mu}_{\mathcal{R},i}(\theta) &\succeq W^{\mathsf{T}} \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta) \right) - \gamma_{\max} C \mathbb{1}_{N_{gR}} \end{aligned} \tag{3.138}$$

where $C = (I - T_{RR}^{\mathsf{T}})^{-1}$. ■

CHAPTER 4

Belief Control Mechanism

Expression (3.20) shows how the limiting distributions of the sending sub-networks determine the limiting distributions of the receiving sub-networks through the matrix W^{T}. In other words, it indicates how influential agents (from within the sending sub-networks) can control the steady-state beliefs of receiving agents. Two critical questions arise at this stage: (a) first, how much freedom do influential agents have in controlling the beliefs of the receiving agents? That is, can receiving agents be driven to arbitrary beliefs or does the network structure limit the scope of control by the influential agents? and (b) second, even if there is a limit to what influential agents can accomplish, how can they ensure that receiving agents will end up with particular beliefs? Questions (a) and (b) raise interesting possibilities about belief (or what we will sometimes refer to as "mind") control. In the next sections, we will address these questions and we will end up with the conditions that allow influential agents to drive other agents to endorse particular beliefs regardless of their local observations (or "convictions").

Observe from expression (3.24) that the limiting beliefs of receiving agents depend on the columns of $W = T_{SR}\,(I - T_{RR})^{-1}$. Note also that the entries of W are determined by the internal combination weights within the receiving networks (i.e., T_{RR}), and the combination weights from the S to the R sub-networks (i.e., T_{SR}). The question we would like to examine now is that given a set of desired beliefs for the receiving agents, is this set always attainable? Or does the internal structure of the receiving sub-networks impose limitations on where their beliefs can be driven to? To answer this useful question, we consider the following problem setting. Let $q_k(\theta)$ denote some desired limiting distribution for receiving agent k (i.e., $q_k(\theta)$ denotes what we desire the limiting distribution $\boldsymbol{\mu}_{k,i}(\theta)$ in (3.24) to become as $i \to \infty$). We would like to examine whether it is possible to force agent k to converge to *any* $q_k(\theta)$, i.e., whether it is possible to find a matrix T_{SR} so that the belief of receiving agent k converges to this specific $q_k(\theta)$.

4.1 Motivation

In this first approach, we are interested in designing T_{SR} while T_{RR} is assumed fixed and known. This scenario allows us to understand in what ways the internal structure of the receiving networks limits the effect of external influence by the sending sub-networks. This approach also allows us to examine the range of belief control over the receiving sub-networks (i.e., how much freedom the sending sub-networks have in selecting these beliefs). Note that the entries of T_{SR} correspond to weights by which the receiving agents scale information from the sending sub-networks. These weights are set by the receiving agents and, therefore, are not under the direct control of the sending sub-networks. As such, it is fair to question whether it is useful to pursue a design procedure for selecting T_{SR} since its entries are not under the direct control of the designer or the sending sub-networks. The useful point to note here, however, is that the entries of T_{SR}, although set by the receiving agents, can still be interpreted as a measure of the level of trust that receiving agents have in the sending agents they are connected to. The higher this level of confidence is between two agents, the larger the value of the scaling weight on the link connecting them. In many applications, these levels of confidence (and, therefore, the resulting scaling weights) can be influenced by external campaigns (e.g., through advertisement or by way of reputation). In this way, we can interpret the problem of designing T_{SR} as a way to guide the campaign that influences receiving agents to set their scaling weights to desirable values. The argument will show that by influencing and knowing T_{SR}, sending agents end up controlling the beliefs of receiving agents in desirable ways. For the analysis in the sequel, note that by fixing T_{RR} and designing T_{SR}, we are in effect fixing the sum of each column of T_{SR} and, accordingly, fixing the overall external influence on each receiving agent. In this way, the problem of designing T_{SR} amounts to deciding on how much influence each individual sub-network should have in driving the beliefs of the receiving sub-networks.

4.2 Conditions for Attainable Beliefs

Given these considerations, let us now show how to design T_{SR} to attain certain beliefs. As is already evident from (3.24), the desired belief $q_k(\theta)$ at any agent k needs to be a probability distribution defined over the true states of all *sending* sub-networks, $\Theta^\bullet = \{\theta_1^\circ, \theta_2^\circ, \dots, \theta_S^\circ\}$. We assume, without loss of generality, that the true states of the sending sub-networks are distinct, so that $|\Theta^\bullet| = S$. If two or more sending sub-networks have the same true state, we can merge them together and treat them as corresponding to one sending sub-network; although this enlarged component is not necessarily connected, it nevertheless consists of strongly-connected elements and the same arguments and conclusions will apply.

We collect the desired limiting beliefs for all receiving agents into the vector:

$$q_{\mathcal{R}}(\theta) \triangleq \begin{bmatrix} q_{N_{gS}+1}(\theta) \\ q_{N_{gS}+2}(\theta) \\ \vdots \\ q_N(\theta) \end{bmatrix} \tag{4.1}$$

which has length N_{gR}. Then, from (3.20), we must have:

$$q_{\mathcal{R}}^{\mathsf{T}}(\theta) = \left(\lim_{i\to\infty} \boldsymbol{\mu}_{\mathcal{S},i}(\theta)\right)^{\mathsf{T}} W \tag{4.2}$$

Evaluating this expression at the successive states $\{\theta_1^\circ, \theta_2^\circ, \dots, \theta_S^\circ\}$, we get

$$\underbrace{\begin{bmatrix} q_{\mathcal{R}}^{\mathsf{T}}(\theta_1^\circ) \\ q_{\mathcal{R}}^{\mathsf{T}}(\theta_2^\circ) \\ \vdots \\ q_{\mathcal{R}}^{\mathsf{T}}(\theta_S^\circ) \end{bmatrix}}_{\triangleq Q_{S\times N_{gR}}} = \underbrace{\begin{bmatrix} \mathbb{1}_{N_1}^{\mathsf{T}} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbb{1}_{N_2}^{\mathsf{T}} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbb{1}_{N_S}^{\mathsf{T}} \end{bmatrix}}_{\triangleq E_{S\times N_{gS}}} W \tag{4.3}$$

where Q is the $S \times N_{gR}$ matrix that collects the desired beliefs for all receiving agents. Using (3.7), we rewrite (4.3) more compactly in matrix form as:

$$E\, T_{SR} = Q\,(I - T_{RR}) \tag{4.4}$$

Therefore, given Q and T_{RR}, the design problem becomes one of finding a matrix T_{SR} that satisfies (4.4) subject to the following constraints:

$$\mathbb{1}^{\mathsf{T}} T_{SR} + \mathbb{1}^{\mathsf{T}} T_{RR} = \mathbb{1}^{\mathsf{T}} \tag{4.5}$$

$$T_{SR} \succcurlyeq 0 \tag{4.6}$$

$$t_{SR,k}(j) = 0, \text{ if receiving agent } k \text{ is not connected to sending agent } j \tag{4.7}$$

The first condition (4.5) is because the entries on each column of A defined in (3.5) add up to one. The second condition (4.6) ensures that each element of T_{SR} is a non-negative combination weight. The third condition (4.7) takes into account the network structure, where $t_{SR,k}$ represents the column of T_{SR} that corresponds to receiving agent k, and $t_{SR,k}(j)$ represents the j^{th} entry of this column (which corresponds to sending agent j–see Fig. 4.2.1). In other words, if receiving agent k is not connected to sending agent j, the corresponding entry in T_{SR} should be zero.

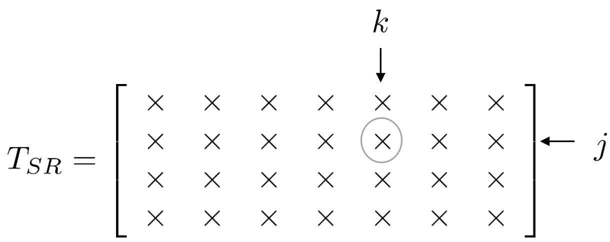

Figure 4.2.1: An illustration of the $k-$th column of T_{SR} and the $j-$th entry on that column.

It is useful to note that condition (4.5) is actually unnecessary and can be removed. This is because if we can find T_{SR} that satisfies (4.4), then condition (4.5) will be automatically satisfied. To see this, we first sum the elements of the columns on the left-hand side of (4.4) and observe that

$$\mathbb{1}_S^\mathsf{T} E T_{SR} = \mathbb{1}^\mathsf{T} T_{SR} \tag{4.8}$$

We then sum the elements of the columns on the right-hand side of (4.4) to get

$$\mathbb{1}_S^\mathsf{T} (Q - QT_{RR}) = \mathbb{1}^\mathsf{T} - \mathbb{1}^\mathsf{T} T_{RR} \tag{4.9}$$

This is because $\mathbb{1}_S^\mathsf{T} Q = \mathbb{1}_{N_{gR}}^\mathsf{T}$ since the entries on each column of Q add up to one. Thus, equating (4.8) and (4.9), we find that (4.5) must hold. The problem we are attempting to solve is then equivalent to finding T_{SR} that satisfies (4.4) subject to

$$T_{SR} \succcurlyeq 0 \tag{4.10}$$

$$t_{SR,k}(j) = 0, \text{ if receiving agent } k \text{ is not connected to sending agent } j \tag{4.11}$$

To find T_{SR} that satisfies (4.4) under the constraints (4.10)-(4.11), we can solve separately for each column of T_{SR}. Let $t_{RR,k}$ and q_k, respectively, denote the columns of T_{RR} and Q that correspond to receiving agent k. Then, relations (4.4) and (4.10)–(4.11) imply that column $t_{SR,k}$ must satisfy:

$$E t_{SR,k} = q_k - Q t_{RR,k} \tag{4.12}$$

subject to

$$t_{SR,k} \succcurlyeq 0 \tag{4.13}$$

$$t_{SR,k}(j) = 0, \text{ if receiving agent } k \text{ is not connected to sending agent } j \tag{4.14}$$

The problem is then equivalent to finding $t_{SR,k}$ for each receiving agent k such that $t_{SR,k}$

satisfies (4.12)-(4.14). For Q to be attainable (i.e., for the beliefs of all receiving agents to converge to the desired beliefs), finding such $t_{SR,k}$ should be possible for each receiving agent k. However, finding $t_{SR,k}$ that satisfies (4.12) under the constraints (4.13)-(4.14) may not be always possible. The desired belief matrix Q will need to satisfy certain conditions so that it is not possible to drive the receiving agents to any belief matrix Q. Before stating these conditions, we introduce two auxiliary matrices. We define first the following difference matrix, which appears on the right-hand side of (4.4) — this matrix is known:

$$V \triangleq Q(I - T_{RR}) \tag{4.15}$$

Note that V has dimensions $S \times N_{gR}$. The $k-$th column of V, which we denote by v_k appears on the right-hand side of (4.12), i.e.,

$$v_k = q_k - Qt_{RR,k} \tag{4.16}$$

The $(s, k)-$th entry of V is then:

$$v_k(s) = q_k(\theta_s^\circ) - \sum_{\ell=1}^{N_{gR}} t_{RR,k}(\ell) q_{N_{gS}+\ell}(\theta_s^\circ) \tag{4.17}$$

Each $(s, k)-$th entry of V represents the difference between the desired limiting belief at θ_s° of receiving agent k and a weighted combination of the desired limiting beliefs of its neighboring receiving agents. We remark that this sum includes agent k if $t_{RR,k}(k)$ is not zero. Similarly, it includes any receiving agent ℓ if $t_{RR,k}(\ell)$ is not zero. In this way, the sum runs only over the neighbors of agent k, because any agent ℓ that is not a neighbor of agent k has its corresponding entry in $t_{RR,k}$ as zero.

Let C denote an $S \times N_{gR}$ binary matrix, with as many rows as the number of sending sub-networks and as many columns as the number of receiving agents. The matrix C is an indicator matrix that specifies whether a receiving agent is connected or not to a sending sub-network. The $(s, k)-$th entry of C is one if receiving agent k is connected to sending sub-network s; otherwise, it is zero. We are now ready to state when a given set of desired beliefs is attainable.

Theorem 5. *(**Attainable Beliefs**) A given belief matrix Q is attainable if, and only if, the entries of V will be zero wherever the entries of C are zero, and the entries of V will be positive wherever the entries of C are one.* ■

Before proving theorem 1, we first clarify its statement. For Q to be achievable, the matrices V and C must have the same structure with the unit entries of C translated into positive entries in V. This theorem reveals two possible cases for each receiving agent k and gives, for each case, the condition required for the desired beliefs to be attainable.

In the first case, receiving agent k is not connected to any agent of sending sub-network s (the $(s,k)-$th entry of C is zero). Then, according to Theorem 1, receiving agent k achieves its desired limiting belief $q_k(\theta_s^\circ)$ if, and only if,

$$v_k(s) = q_k(\theta_s^\circ) - \sum_{\ell=1}^{N_{gR}} t_{RR,k}(\ell) q_{N_{gS}+\ell}(\theta_s^\circ) = 0 \tag{4.18}$$

That is, the cumulative influence from the agent's neighbors must match the desired limiting belief.

In the second case, receiving agent k is connected to at least one agent of sending sub-network s (the $(s,k)-$th entry of C is one). Now, according to Theorem 1 again, receiving agent k achieves its desired limiting belief $q_k(\theta_s^\circ)$ if, and only if,

$$v_k(s) = q_k(\theta_s^\circ) - \sum_{\ell=1}^{N_{gR}} t_{RR,k}(\ell) q_{N_{gS}+\ell}(\theta_s^\circ) > 0 \tag{4.19}$$

Proof of Theorem 1. We start by first proving that if Q is attainable, then V and C have the same structure. If Q is attainable, then there exists $t_{SR,k}$ for each receiving agent k that satisfies (4.12)-(4.14). Using the definition of E in (4.3), the $s-$th row on the left-hand side of (4.12) is:

$$\sum_{j \in \mathcal{I}_s} t_{SR,k}(j) \tag{4.20}$$

where $\mathcal{I}_s$ represents the set of indexes of sending agents that belong to sending sub-network s. Expression (4.20) represents the sum of the elements of the block of $t_{SR,k}$ that correspond to sending sub-network s. Therefore, if Q is attainable, then the $s-$th row of (4.12) satisfies the following relation:

$$\sum_{j \in \mathcal{I}_s} t_{SR,k}(j) = v_k(s) \tag{4.21}$$

From this relation, we see that if agent k is not connected to any agent in sub-network s, then $\sum_{j \in \mathcal{I}_s} t_{SR,k}(j) = 0$ which implies that $v_k(s)$ is zero. On the other hand, if agent k is connected to sub-network s, then $\sum_{j \in \mathcal{I}_s} t_{SR,k}(j) > 0$ which implies that $v_k(s) > 0$. In other words, C and V have the same structure.

Conversely, if C and V have the same structure, then it is possible to find $t_{SR,k}$ for each receiving agent k that satisfies (4.12)-(4.14). In particular, if agent k is not connected to sub-network s, then the (s, k)-th entry of C is zero. Since C and V have the same structure, then $v_k(s) = 0$. By setting to zero the entries of $t_{SR,k}$ that correspond to sending sub-network s, relation (4.21) is satisfied. On the other hand, if agent k is connected to sub-network s (connected to at least one agent in sub-network s), then the (s, k)-th entry of C is one. Since C and V have the same structure, we get $v_k(s) > 0$. Therefore, since the entries of $t_{SR,k}$ must be non-negative, we first set to zero the entries of $t_{SR,k}$ that correspond to agents of sub-network s that are not connected to agent k and the remaining entries can be set to non-negative values such that relation (4.21) is satisfied. That is, if C and V have the same structure, then Q is attainable. □

We next move to characterize the set of solutions, i.e., how we can design $t_{SR,k}$ assuming the conditions on V are met.

4.3 Characterizing the Set of Possible Solutions

In the sequel, we assume that the conditions on V from Theorem 1 are satisfied. That is, if receiving agent k is not connected to sub-network s, then $v_k(s) = 0$. Otherwise, $v_k(s) > 0$. The desired beliefs are then attainable. This means that for each receiving agent k, we can find $t_{SR,k}$ that satisfies (4.12)-(4.14). Many solutions may exist. In this section, we characterize the set of possible solutions.

First of all, to meet (4.11), we set the required entries of $t_{SR,k}$ to zero. We then remove the corresponding columns of E, and label the reduced E by E_k. Similarly, we remove the zero elements of $t_{SR,k}$ and label the reduced $t_{SR,k}$ by $t'_{SR,k}$. On the other hand, if agent k is not connected to some sub-network s, then the corresponding row in E will be removed and E_k will have fewer number of rows, denoted by S'. Without loss of generality, we assume agent k is connected to the first S' sending sub-networks. We denote by N_s^k the number of agents of sending sub-network s that are connected to receiving agent k and by N_{gS}^k the total number of all sending agents connected to agent k. The matrix E_k will then have the form (this matrix is obtained from E by removing rows and columns with zero entries; the resulting dimensions are now denoted by S' and N_{gS}^k):

$$E_k = \begin{bmatrix} \mathbb{1}^{\mathsf{T}}_{N_1^k} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbb{1}^{\mathsf{T}}_{N_2^k} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbb{1}^{\mathsf{T}}_{N_{S'}^k} \end{bmatrix}_{S' \times N_{gS}^k} \tag{4.22}$$

Note that if receiving agent k is connected to all sending sub-networks, then E and E_k will have the same number of rows, $S' = S$. In the case where agent k is not connected to some sub-network s, condition (4.18) should be satisfied, and the corresponding row in $q_k - Qt_{RR,k}$ should be removed to obtain the reduced vector $q'_k - Q't_{RR,k}$. We are therefore reduced to determining $t'_{SR,k}$ by solving a system of equations of the form:

$$E_k t'_{SR,k} = q'_k - Q' t_{RR,k} \tag{4.23}$$

subject to

$$t'_{SR,k} \succcurlyeq 0 \tag{4.24}$$

We can still have some of the entries of the solution $t'_{SR,k}$ turn out to be zero. Now note that the number of rows of E_k is S' (number of sending sub-networks connected to k), which is always smaller than or equal to N^k_{gS}. Moreover, the rows of E_k are linearly independent and thus E_k is a right-invertible matrix. Its right-inverse is given by [77]:

$$E^R_k = E^\intercal_k \left(E_k E^\intercal_k\right)^{-1} \tag{4.25}$$

Therefore, if we ignore condition (4.24) for now, then equation (4.23) has an infinite number of solutions parametrized by the expression [77]:

$$t'_{SR,k} = E^R_k \left(q'_k - Q' t_{RR,k}\right) + (I - E^R_k E_k) y \tag{4.26}$$

where y is an arbitrary vector of length N^k_{gS}. We still need to satisfy condition (4.24). Let

$$v'_k \triangleq q'_k - Q' t_{RR,k} \tag{4.27}$$

and note that

$$E^R_k v'_k = \begin{bmatrix} \frac{v'_k(1)}{N^k_1} \mathbb{1}_{N^k_1} \\ \frac{v'_k(2)}{N^k_2} \mathbb{1}_{N^k_2} \\ \vdots \\ \frac{v'_k(S')}{N^k_{S'}} \mathbb{1}_{N^k_{S'}} \end{bmatrix} \tag{4.28}$$

where $v'_k(i)$ represents the i^{th} entry of vector v'_k. Likewise,

$$I - E_k^{\mathsf{T}} \left(E_k E_k^{\mathsf{T}}\right)^{-1} E_k = \text{diag}\left\{I_{N_1^k} - \frac{1}{N_1^k}\mathbb{1}_{N_1^k}\mathbb{1}_{N_1^k}^{\mathsf{T}}, I_{N_2^k} - \frac{1}{N_2^k}\mathbb{1}_{N_2^k}\mathbb{1}_{N_2^k}^{\mathsf{T}}, \cdots I_{N_{S'}^k} - \frac{1}{N_{S'}^k}\mathbb{1}_{N_{S'}^k}\mathbb{1}_{N_{S'}^k}^{\mathsf{T}}\right\} \tag{4.29}$$

and if we partition y into sub-vectors as

$$y = \begin{bmatrix} y_{N_1^k} \\ y_{N_2^k} \\ \vdots \\ y_{N_{S'}^k} \end{bmatrix} \tag{4.30}$$

then expression (4.26) becomes:

$$t'_{SR,k} = \begin{bmatrix} \frac{v'_k(1)}{N_1^k}\mathbb{1}_{N_1^k} \\ \frac{v'_k(2)}{N_2^k}\mathbb{1}_{N_2^k} \\ \vdots \\ \frac{v'_k(S')}{N_{S'}^k}\mathbb{1}_{N_{S'}^k} \end{bmatrix} + \begin{bmatrix} \left(I_{N_1^k} - \frac{1}{N_1^k}\mathbb{1}_{N_1^k}\mathbb{1}_{N_1^k}^{\mathsf{T}}\right) y_{N_1^k} \\ \left(I_{N_2^k} - \frac{1}{N_2^k}\mathbb{1}_{N_2^k}\mathbb{1}_{N_2^k}^{\mathsf{T}}\right) y_{N_2^k} \\ \vdots \\ \left(I_{N_{S'}^k} - \frac{1}{N_{S'}^k}\mathbb{1}_{N_{S'}^k}\mathbb{1}_{N_{S'}^k}^{\mathsf{T}}\right) y_{N_{S'}^k} \end{bmatrix} \tag{4.31}$$

This represents the general form of all possible solutions, but from these solutions we want only those which are nonnegative in order to satisfy condition (4.24). From (4.31), the vector $t'_{SR,k}$ is partitioned into multiple blocks, where each block has the form:

$$\frac{v'_k(s)}{N_s^k}\mathbb{1}_{N_s^k} + \left(I_{N_s^k} - \frac{1}{N_s^k}\mathbb{1}_{N_s^k}\mathbb{1}_{N_s^k}^{\mathsf{T}}\right) y_{N_s^k} \tag{4.32}$$

We already have from the conditions of attainable beliefs (4.19) that $v'_k(s) > 0$. Therefore, we can choose $y_{N_s^k}$ as zero or set it to arbitrary values as long as (4.32) stays non-negative. We also know that for the beliefs to be attainable, we cannot have $v'_k(s) < 0$. Otherwise, no solution can be found. Indeed, if $v'_k(s) < 0$, then to make (4.32) non-negative, we would need to select $y_{N_s^k}$ such that:

$$\left(I_{N_s^k} - \frac{1}{N_s^k}\mathbb{1}_{N_s^k}\mathbb{1}_{N_s^k}^\mathsf{T}\right) y_{N_s^k} \succeq -\frac{v_k'(s)}{N_s^k}\mathbb{1}_{N_s^k} \tag{4.33}$$

However, there is no $y_{N_s^k}$ that satisfies this relation because if we sum the elements of the vector on the left-hand side of (4.33), we obtain:

$$\mathbb{1}_{N_s^k}^\mathsf{T}\left(I_{N_s^k} - \frac{1}{N_s^k}\mathbb{1}_{N_s^k}\mathbb{1}_{N_s^k}^\mathsf{T}\right) y_{N_s^k} = 0 \tag{4.34}$$

While if we sum the elements of the vector on the right-hand side of (4.33), we obtain:

$$-\frac{v_k'(s)}{N_s^k}\mathbb{1}_{N_s^k}^\mathsf{T}\mathbb{1}_{N_s^k} = -v_k'(s) > 0 \tag{4.35}$$

This means that we cannot find $t'_{SR,k}$ such that $t'_{SR,k} \succeq 0$ when any of the entries of v_k' or $q_k' - Q' t_{RR,k}$ is negative.

In summary, we have established the validity of the following statement.

Theorem 6. *Assume receiving agent k is connected to N_s^k agents in sending sub-network s. If $v_k(s) > 0$, then all possible choices for the weights from sending agents in network s to receiving agent k are parameterized as:*

$$\frac{v_k'(s)}{N_s^k}\mathbb{1}_{N_s^k} + \left(I_{N_s^k} - \frac{1}{N_s^k}\mathbb{1}_{N_s^k}\mathbb{1}_{N_s^k}^\mathsf{T}\right) y_{N_s^k} \tag{4.36}$$

where $y_{N_s^k}$ is an arbitrary vector of length N_s^k chosen so that (4.36) stays non-negative. ■

4.4 Enforcing Uniform Beliefs

In this section, we explore one special case of attainable beliefs, which is driving all receiving agents towards the *s*ame belief. In this case, Q is of the following form:

$$Q = q\mathbb{1}_{N_{gR}}^\mathsf{T} \tag{4.37}$$

for some column q that represents the desired limiting belief (the entries of q are non-negative and add up to one). We verify that the conditions that ensure that uniform beliefs are attainable by all receiving agents. In this case, v_k is of the following form:

$$v_k = q_k - Qt_{RR,k} = (1 - \mathbb{1}_{N_{gR}}^{\top} t_{RR,k})q \tag{4.38}$$

and the (s, k)-th entry of V is:

$$v_k(s) = (1 - \mathbb{1}_{N_{gR}}^{\top} t_{RR,k})q(\theta_s^{\circ}) \tag{4.39}$$

Now we know that $1 - \mathbb{1}_{N_{gR}}^{\top} t_{RR,k} > 0$ when agent k is connected to at least one agent from any sending sub-network, and that $1 - \mathbb{1}_{N_{gR}}^{\top} t_{RR,k} = 0$ when it is not connected to any sending sub-network. In the second case where $1 - \mathbb{1}_{N_{gR}}^{\top} t_{RR,k} = 0$, expression (4.39) implies that $v_k(s) = 0$ for any s. Therefore, in this case, we have agent k not connected to any sending sub-network s and $v_k(s) = 0$ for any s, and condition (4.18) is satisfied. In the first case where $1 - \mathbb{1}_{N_{gR}}^{\top} t_{RR,k} > 0$ (i.e., agent k is connected to some sending sub-networks but not necessarily to all of them), expression (4.39) implies that $v_k(s) > 0$ no matter whether agent k is connected or not to sending sub-network s. However, when agent k is not connected to sending sub-network s, condition (4.18) requires that $v_k(s) = 0$ for agent k to achieve its desired belief at θ_s°. In summary, we arrive at the following conclusion.

Lemma 7. *For the scenario of uniform beliefs to be attainable, agent k should be connected either to all sending sub-networks or to none of them.*

We next illustrate the results with two examples.

4.4.1 Example 1

Consider the network shown in Fig. 4.4.1. It consists of $N = 8$ agents, two sending sub-networks and one receiving sub-network, with the following combination matrix:

$$A = \left[\begin{array}{ccccc|ccc} 0.2 & 0.2 & 0.8 & 0 & 0 & 0 & 0 & \times \\ 0.5 & 0.4 & 0.1 & 0 & 0 & \times & 0 & 0 \\ 0.3 & 0.4 & 0.1 & 0 & 0 & \times & \times & 0 \\ 0 & 0 & 0 & 0.4 & 0.3 & \times & 0 & \times \\ 0 & 0 & 0 & 0.6 & 0.7 & 0 & \times & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0.2 & 0.3 & 0.2 \\ 0 & 0 & 0 & 0 & 0 & 0.1 & 0.2 & 0.3 \\ 0 & 0 & 0 & 0 & 0 & 0.1 & 0.2 & 0.1 \end{array}\right] \tag{4.40}$$

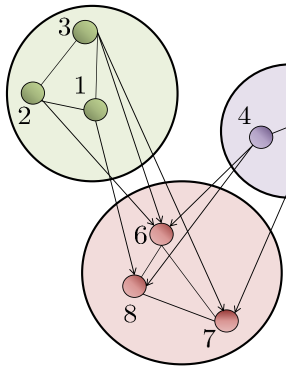

Figure 4.4.1: A weakly connected network consisting of three sub-networks in a broadband influence scenario.

We assume that there are 3 possible states $\Theta = \{\theta_1^\circ, \theta_2^\circ, \theta_3^\circ\}$, where θ_1° is the true event for the first sending sub-network, θ_2° is the true event for the second sending sub-network, and θ_3° is the true event for the receiving sub-network.

Let us first design T_{SR} so that all receiving agents' beliefs converge to the same belief over $\{\theta_1^o, \theta_2^o\}$, say:

$$q = \begin{bmatrix} 0.2 \\ 0.8 \end{bmatrix} \tag{4.41}$$

We determine the columns of T_{SR} one at a time. Starting with agent 6, we focus on the first column of T_{SR}. The vector v_6 defined in (4.27) is given by (4.38) for the case of uniform beliefs. Therefore,

$$v_6 = 0.6q = \begin{bmatrix} 0.12 \\ 0.48 \end{bmatrix} \tag{4.42}$$

Thus, according to (4.31),

$$t'_{SR,6} = \begin{bmatrix} \frac{v_6(1)}{2}\mathbb{1}_2 \\ v_6(2) \end{bmatrix} + \begin{bmatrix} I_2 - \frac{1}{2}\mathbb{1}_2\mathbb{1}_2^\mathsf{T} \\ 1 - \frac{1}{1}\mathbb{1}_1\mathbb{1}_1^\mathsf{T} \end{bmatrix} y \tag{4.43}$$

where y is an arbitrary vector of length 3. Note that $t'_{SR,6}$ represents respectively the coefficients of agents 2, 3 and 4 that are linked to agent 6. It follows that

$$t'_{SR,6} = \begin{bmatrix} 0.06 \\ 0.06 \\ 0.48 \end{bmatrix} + \begin{bmatrix} \frac{1}{2}y(1) - \frac{1}{2}y(2) \\ -\frac{1}{2}y(1) + \frac{1}{2}y(2) \\ 0 \end{bmatrix} \tag{4.44}$$

Let $\alpha_6 \triangleq \frac{1}{2}y(1) - \frac{1}{2}y(2)$ so that

$$t'_{SR,6} = \begin{bmatrix} 0.06 + \alpha_6 \\ 0.06 - \alpha_6 \\ 0.48 \end{bmatrix} \tag{4.45}$$

In order to have positive entries for $t'_{SR,6}$, we can choose $|\alpha_6| \leq 0.06$.

Now for agent 7, the vector v_7 is given by

$$v_7 = 0.3q = \begin{bmatrix} 0.06 \\ 0.24 \end{bmatrix} \tag{4.46}$$

so that

$$t'_{SR,7} = \begin{bmatrix} v_7(1) \\ v_7(2) \end{bmatrix} = \begin{bmatrix} 0.06 \\ 0.24 \end{bmatrix} \tag{4.47}$$

The entries $t'_{SR,7}$ represent respectively the coefficients of agents 3 and 5 that are linked to agent 7.

Next for agent 8, we have

$$v_8 = 0.4q = \begin{bmatrix} 0.08 \\ 0.32 \end{bmatrix} \tag{4.48}$$

so that

$$t'_{SR,8} = \begin{bmatrix} v_8(1) \\ v_8(2) \end{bmatrix} = \begin{bmatrix} 0.08 \\ 0.32 \end{bmatrix} \tag{4.49}$$

Therefore, one possible solution is

$$T_{SR} = \begin{bmatrix} 0 & 0 & 0.08 \\ 0.06 & 0 & 0 \\ 0.06 & 0.06 & 0 \\ 0.48 & 0 & 0.32 \\ 0 & 0.24 & 0 \end{bmatrix} \tag{4.50}$$

To verify that the beliefs of the receiving agents converge in this case to the desired belief, we compute the matrix $W^\top$ from (3.7):

$$\begin{aligned} W^\top &= (I - T_{RR}^\top)^{-1} T_{SR}^\top \\ &= \begin{bmatrix} 0.0169 & 0.0839 & 0.0992 & 0.7390 & 0.0610 \\ 0.0322 & 0.0394 & 0.1284 & 0.4441 & 0.3559 \\ 0.1034 & 0.0318 & 0.0648 & 0.6678 & 0.1322 \end{bmatrix} \end{aligned} \tag{4.51}$$

Then, according to (3.24), we can compute the belief at θ_1° for each receiving agent at steady state, by taking the first block in the agent's corresponding row and summing its elements:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_1^\circ) = \begin{cases} 0.0169 + 0.0839 + 0.0992 = 0.2, & k = 6 \\ 0.0322 + 0.0394 + 0.1284 = 0.2, & k = 7 \\ 0.1034 + 0.0318 + 0.0648 = 0.2, & k = 8 \end{cases}$$

Likewise, we can compute the belief at θ_2° for each receiving agent at steady state, by taking the second block in the agent's corresponding row and summing its elements:

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_2^\circ) = \begin{cases} 0.7390 + 0.0610 = 0.8, & k = 6 \\ 0.4441 + 0.3559 = 0.8, & k = 7 \\ 0.6678 + 0.1322 = 0.8, & k = 8 \end{cases}$$

Let us now consider the case where we want to design T_{SR} so that the desired limiting beliefs are not necessarily uniform but rather

$$Q = \begin{bmatrix} 0.8 & 0.7 & 0.75 \\ 0.2 & 0.3 & 0.25 \end{bmatrix} \tag{4.52}$$

Note that now the beliefs are different from an agent to another, but they are still close. Computing,

$$v_k = q_k - Qt_{RR,k} \tag{4.53}$$

for each receiving agent k, we obtain:

$$v_6 = q_6 - Qt_{RR,6} = \begin{bmatrix} 0.495 \\ 0.105 \end{bmatrix} \tag{4.54}$$

$$v_7 = q_7 - Qt_{RR,7} = \begin{bmatrix} 0.17 \\ 0.13 \end{bmatrix} \tag{4.55}$$

$$v_8 = q_8 - Qt_{RR,8} = \begin{bmatrix} 0.305 \\ 0.195 \end{bmatrix} \tag{4.56}$$

Therefore, one possible T_{SR} is

$$T_{SR} = \begin{bmatrix} 0 & 0 & 0.305 \\ 0.495/2 & 0 & 0 \\ 0.495/2 & 0.17 & 0 \\ 0.105 & 0 & 0.195 \\ 0 & 0.13 & 0 \end{bmatrix} \tag{4.57}$$

Let us now consider the case where the desired limiting beliefs are more dispersed, such as

$$Q = \begin{bmatrix} 0.8 & 0.2 & 0.3 \\ 0.2 & 0.8 & 0.7 \end{bmatrix} \tag{4.58}$$

In this case for agent 7, we have

$$v_7 = q_7 - Qt_{RR,7} = \begin{bmatrix} -0.14 \\ 0.44 \end{bmatrix} \tag{4.59}$$

with a negative first entry. Therefore, the desired belief for agent 7 cannot be attained.

4.4.2 Example 2

Consider now the network shown in Fig. 4.4.2 with the following combination matrix

$$A = \left[\begin{array}{ccccc|ccc} 0.2 & 0.2 & 0.8 & 0 & 0 & 0 & 0 & \times \\ 0.5 & 0.4 & 0.1 & 0 & 0 & \times & 0 & \times \\ 0.3 & 0.4 & 0.1 & 0 & 0 & \times & \times & 0 \\ 0 & 0 & 0 & 0.4 & 0.3 & \times & 0 & 0 \\ 0 & 0 & 0 & 0.6 & 0.7 & 0 & \times & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0.2 & 0.3 & 0.1 \\ 0 & 0 & 0 & 0 & 0 & 0.1 & 0.2 & 0.6 \\ 0 & 0 & 0 & 0 & 0 & 0.1 & 0.2 & 0 \end{array}\right] \tag{4.60}$$

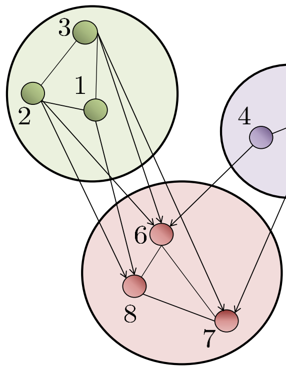

Figure 4.4.2: A weakly connected network consisting of three sub-networks.

Let us consider the case where we want to design T_{SR} so that the desired limiting beliefs are as follows:

$$Q = \begin{bmatrix} 0.8 & 0.7 & 0.8 \\ 0.2 & 0.3 & 0.2 \end{bmatrix} \tag{4.61}$$

Computing,

$$v_k = q_k - Qt_{RR,k} \tag{4.62}$$

for each receiving agent k, we obtain:

$$v_6 = q_6 - Qt_{RR,6} = \begin{bmatrix} 0.49 \\ 0.11 \end{bmatrix} \tag{4.63}$$

$$v_7 = q_7 - Qt_{RR,7} = \begin{bmatrix} 0.16 \\ 0.14 \end{bmatrix} \tag{4.64}$$

$$v_8 = q_8 - Qt_{RR,8} = \begin{bmatrix} 0.3 \\ 0 \end{bmatrix} \tag{4.65}$$

Note that in this example, agent 8 is not connected to the second sending sub-network, but the controlling scheme can still work because condition (4.18) is satisfied. Therefore, one possible choice for T_{SR} is the following:

$$T_{SR} = \begin{bmatrix} 0 & 0 & 0.3/2 \\ 0.49/2 & 0 & 0.3/2 \\ 0.49/2 & 0.16 & 0 \\ 0.11 & 0 & 0 \\ 0 & 0.14 & 0 \end{bmatrix} \tag{4.66}$$

To verify that the beliefs of the agents converge in this case to the desired belief, we compute W^{T} from (3.7) and use (3.24) to determine the limiting beliefs at θ_1^o and θ_2^o at the receiving agents. This calculation gives

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_1^\circ) = \begin{cases} 0.0309 + 0.3737 + 0.3954 = 0.8, & k = 6 \\ 0.0586 + 0.2200 + 0.4214 = 0.7, & k = 7 \\ 0.1883 + 0.3193 + 0.2924 = 0.8, & k = 8 \end{cases}$$

and

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_2^\circ) = \begin{cases} 0.1539 + 0.046 = 0.2, & k = 6 \\ 0.0724 + 0.2276 = 0.3, & k = 7 \\ 0.0588 + 0.1412 = 0.2, & k = 8 \end{cases}$$

4.5 Joint Design of T_{RR} and T_{SR}

In the previous sections, we analyzed the conditions that drive receiving agents to desired beliefs. The approach relies on determining the entries of the weighting matrix T_{SR} from knowledge of Q (the desired beliefs) and T_{RR} (the internal weighting structure within the receiving sub-networks). We saw how there is limitation to where the beliefs of receiving agents can converge. In particular, the internal combination of receiving sub-networks contribute to this limitation. We now examine the problem of designing T_{SR} and T_{RR} jointly, to see whether by having more freedom in choosing the coefficients of T_{RR}, we still encounter limitations on how to influence the receiving agents. We assume that we know the number of receiving sub-networks and the number of agents in each of these sub-networks. Using (4.4), we have

$$\underbrace{\begin{bmatrix} E & Q \end{bmatrix}}_{\triangleq B} \begin{bmatrix} T_{SR} \\ T_{RR} \end{bmatrix} = Q \tag{4.67}$$

Therefore, given Q (the desired limiting beliefs of the receiving agents), the design problem becomes one of finding matrices T_{SR} and T_{RR} that satisfy (4.67) subject to the following constraints:

$$\mathbb{1}^{\mathsf{T}} T_{SR} + \mathbb{1}^{\mathsf{T}} T_{RR} = \mathbb{1}^{\mathsf{T}} \tag{4.68}$$

$$T_{SR,k}(j) = 0, \text{ if sending agent } j \text{ does not feed into } k$$

$$T_{SR,k}(j) \geq 0, \text{ otherwise} \tag{4.69}$$

$$T_{RR,k}(j) = 0, \text{ if receiving agent } j \text{ does not feed into } k$$

$$T_{RR,k}(j) > 0, \text{ otherwise} \tag{4.70}$$

In the last condition (4.70), we are requiring $T_{RR,k}(j)$ to be strictly positive if receiving agent j feeds into k. This is in order to avoid solutions where the receiving sub-networks become unconnected. For instance, consider the example shown in Fig. 4.5.1. This figure shows a case where agent k is connected to all sending sub-networks, and it depicts only the incoming links into agent k. Let us assume that the desired limiting belief for agent k is

$$\begin{bmatrix} q_k(\theta_1^\circ) \\ q_k(\theta_2^\circ) \end{bmatrix} = \begin{bmatrix} 0.1 \\ 0.9 \end{bmatrix} \tag{4.71}$$

Then a possible solution to (4.67) is to assign zero as weights for the data originating from its receiving neighbors, 0.1 for the data received from sending agent 1, and 0.9 for the data received from sending agent 2. Then, for this example,

$$E = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \; Q = \begin{bmatrix} 0.1 & q_4(\theta_1^\circ) & q_5(\theta_1^\circ) \\ 0.9 & q_4(\theta_2^\circ) & q_5(\theta_2^\circ) \end{bmatrix}, \; T_{SR} = \begin{bmatrix} 0.1 \\ 0.9 \end{bmatrix}, \; T_{RR} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \tag{4.72}$$

so that (4.67) is satisfied. However, this solution affects the connectedness of the receiving sub-network of agent k, because there will be no path that leads to this agent.

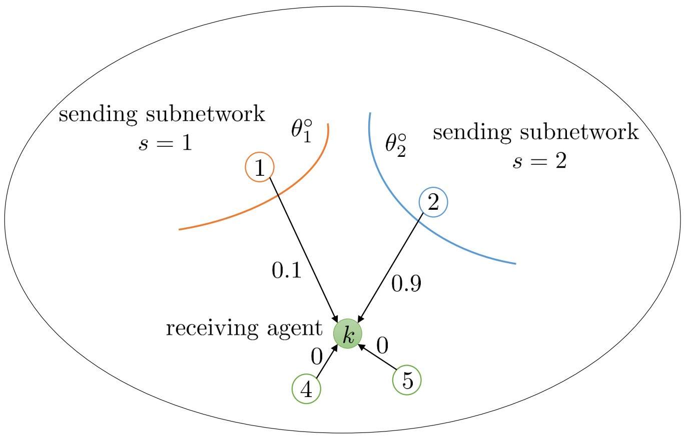

Figure 4.5.1: An example where the receiving network of agent k ends up being disconnected.

To find T_{SR} and T_{RR} satisfying (4.67)-(4.70), we can solve separately for each of their columns. If it is possible to find a solution for each column, then Q is attainable. We explore next the possibility of finding solutions for each column. Similarly to the previous section, $t_{SR,k}$ and $t_{RR,k}$ respectively represent the columns of T_{SR} and T_{RR} that correspond to receiving agent k, and $t_{SR,k}(j)$ and $t_{RR,k}(j)$ respectively represent the $j-$th entries of this

$t_{SR,k}$ and $t_{RR,k}$. Also q_k denote the column of Q that corresponds to receiving agent k. Then, relations (4.67) and (4.69)–(4.70) imply that the columns $t_{SR,k}$ and $t_{RR,k}$ must satisfy:

$$\underbrace{\begin{bmatrix} E & Q \end{bmatrix}}_{=B} \begin{bmatrix} t_{SR,k} \\ t_{RR,k} \end{bmatrix} = q_k \tag{4.73}$$

subject to

$$\mathbb{1}^{\mathsf{T}} t_{SR,k} + \mathbb{1}^{\mathsf{T}} t_{RR,k} = 1 \tag{4.74}$$

$$\begin{aligned} t_{SR,k}(j) &= 0, \text{ if } j \text{ does not feed } k \\ t_{SR,k}(j) &\geq 0, \text{ otherwise} \end{aligned} \tag{4.75}$$

$$\begin{aligned} t_{RR,k}(j) &= 0, \text{ if } j \text{ does not feed } k \\ t_{RR,k}(j) &> 0, \text{ otherwise} \end{aligned} \tag{4.76}$$

Since the connections within the sending and receiving networks are known, but not the combination weights T_{SR} and T_{RR} whose values we are seeking, we can then set to zero the entries of $t_{SR,k}$ and $t_{RR,k}$ that correspond to unlinked agents. We remove these zero entries and relabel the vectors as $t'_{SR,k}$ and $t'_{RR,k}$. We also remove the corresponding columns of E and Q, and label the modified E and Q by E_k and Q_k. We are therefore reduced to determining $t'_{SR,k}$ and $t'_{RR,k}$ by solving a system of equations of the form:

$$\underbrace{\begin{bmatrix} E_k & Q_k \end{bmatrix}}_{\triangleq B_k} \begin{bmatrix} t'_{SR,k} \\ t'_{RR,k} \end{bmatrix} = q_k \tag{4.77}$$

subject to

$$\mathbb{1}^\mathsf{T} t'_{SR,k} + \mathbb{1}^\mathsf{T} t'_{RR,k} = 1 \tag{4.78}$$

$$t'_{SR,k} \succcurlyeq \mathbf{0} \tag{4.79}$$

$$t'_{RR,k} \succ \mathbf{0} \tag{4.80}$$

Formulation (4.77)-(4.80) has the following interpretation. After some sufficient time $i \geq I$, we know that the beliefs of all agents will approach some limiting beliefs, and based on the results of the previous work [64], the belief update (2.8) approaches for $i \geq I$,

$$\begin{cases} \boldsymbol{\psi}_{k,i+1}(\theta) = \boldsymbol{\mu}_{k,i}(\theta) \\ \boldsymbol{\mu}_{k,i+1}(\theta) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k}\, \boldsymbol{\psi}_{\ell,i+1}(\theta) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k}\, \boldsymbol{\mu}_{\ell,i}(\theta) \end{cases} \tag{4.81}$$

This means that:

$$\lim_{i \to \infty} \boldsymbol{\mu}_{k,i}(\theta) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \left(\lim_{i \to \infty} \boldsymbol{\mu}_{\ell,i}(\theta) \right) \tag{4.82}$$

In other words, if we want the beliefs of the receiving agents to converge to some belief vector q, then we need to make sure that these desired beliefs satisfy the relationship:

$$q_k(\theta) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} q_\ell(\theta) \tag{4.83}$$

for any $\theta \in \{\theta_1^\circ, \cdots, \theta_S^\circ\}$ and for all receiving agents k. In other words, given the set of desirable beliefs, we would like to know if it is possible to express the desired limiting belief for each receiving agent k as a convex combination of the limiting beliefs of its receiving neighbors and the limiting beliefs of the sending agents to which agent k is connected. If this is possible for each agent k, then Q is attainable, i.e., all receiving agents can reach their desired limiting beliefs. This is precisely what the formulation (4.77)–(4.80) is attempting to enforce, by finding suitable coefficients such that (4.83) is satisfied. Finding $t'_{SR,k}$ and $t'_{RR,k}$ that satisfy (4.77) and constraints (4.78)-(4.80) might not be always possible. Since each agent k can be connected to all sending sub-networks, or to some of them or to none of them, the matrix E_k that appears in (4.77) will have a different form for each of these cases, which will affect the possibility of finding a solution. Before analyzing how the three possible cases affect the possibility of finding a solution, we summarize first the results:

1. Agent k is connected to all sending sub-networks: the problem reduces to finding $t'_{RR,k}$ that satisfies (4.92a) and (4.92b, which always has a solution;

2. Agent k is connected to some sending sub-networks: the problem reduces to finding $t'_{RR,k}$ that satisfies conditions (4.100a)-(4.100c), which may not always have a solution;

3. Agent k is not connected to any sending sub-network: the problem reduces to finding $t'_{RR,k}$ that satisfies conditions (4.107a)-(4.107c), which may not always have a solution.

Note that relations (4.92a) and (4.100a) are what condition (4.19) required when we wanted to design T_{SR}, for the case where agent k is connected to sending sub-network s, when T_{RR} was given. Similarly, relations (4.100b) and (4.107a) are what condition (4.18) required when we wanted to design $t_{SR,k}$, for the case where agent k is not connected to sending sub-network s, when T_{RR} was given. In the earlier section, we had to make sure that the given T_{RR} satisfies (4.18) and (4.19) for Q to be attainable. Here, we are designing for T_{RR} as well, and we need to make sure that the entries we choose satisfy these conditions. We now analyze each case in detail.

Case 1: Agent k is connected to all sending sub-networks

We discuss first the case where agent k is connected to at least one agent from each sending sub-network. In this case, E_k will have the following form:

$$E_k = \begin{bmatrix} \mathbb{1}^{\mathsf{T}}_{N_1^k} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbb{1}^{\mathsf{T}}_{N_2^k} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbb{1}^{\mathsf{T}}_{N_S^k} \end{bmatrix} \tag{4.84}$$

and relation (4.77) is then:

$$\begin{bmatrix} \mathbb{1}^{\mathsf{T}}_{N_1^k} & \mathbf{0} & \dots & \mathbf{0} & q_{k(1)}(\theta_1^\circ) & \dots & q_{k(N_{gR}^k)}(\theta_1^\circ) \\ \mathbf{0} & \mathbb{1}^{\mathsf{T}}_{N_2^k} & \dots & \mathbf{0} & q_{k(1)}(\theta_2^\circ) & \dots & q_{k(N_{gR}^k)}(\theta_2^\circ) \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \dots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbb{1}^{\mathsf{T}}_{N_S^k} & q_{k(1)}(\theta_s^\circ) & \dots & q_{k(N_{gR}^k)}(\theta_s^\circ) \end{bmatrix} \begin{bmatrix} t'_{SR,k} \\ t'_{RR,k} \end{bmatrix} = q_k \tag{4.85}$$

where $q_{k(j)}(\theta_s^\circ)$ represents the desired limiting belief at θ_s° for the j^{th} receiving neighbor of agent k, and N_{gR}^k is the total number of receiving agents that are neighbors of agent k. The problem here is to find $t'_{SR,k}$ and $t'_{RR,k}$ that satisfy (4.85) subject to the constraints (4.78)-(4.79). It is useful to note that if we can find $t'_{SR,k}$ and $t'_{RR,k}$ that satisfy (4.85), then condition (4.78) will be automatically satisfied. To see this, we first sum the elements of the vector on the left-hand side of (4.85) and observe that

$$\mathbb{1}_S^{\mathsf{T}} B_k \begin{bmatrix} t'_{SR,k} \\ t'_{RR,k} \end{bmatrix} = \mathbb{1}^{\mathsf{T}} \begin{bmatrix} t'_{SR,k} \\ t'_{RR,k} \end{bmatrix} \tag{4.86}$$

This is because $\mathbb{1}_S^{\mathsf{T}} B_k = \mathbb{1}^{\mathsf{T}}$ since the entries on each column of B_k add up to one. We then sum the elements of the vector on the right-hand side of (4.85) to get

$$\mathbb{1}_S^{\mathsf{T}} q_k = 1 \tag{4.87}$$

Thus, equating (4.86) and (4.87), we obtain (4.78). The problem we are attempting to solve is then equivalent to finding $t'_{SR,k}$ and $t'_{RR,k}$ that satisfy (4.85) subject to

$$t'_{SR,k} \succcurlyeq \mathbf{0} \tag{4.88}$$

$$t'_{RR,k} \succ \mathbf{0} \tag{4.89}$$

Now, note that (4.85) consists of S equations and note that the number of variables (i.e., the total number of entries of $t'_{SR,k}$ and $t'_{RR,k}$) is greater than the number of equations. Each equation relates the entries of $t'_{SR,k}$ that correspond to agents of one of the sending sub-networks to all entries of $t'_{RR,k}$. In particular, the equation that corresponds to sending sub-network s has the following form:

$$\sum_{\ell \in \mathcal{I}_s} t'_{SR,k}(\ell) = q_k(\theta_s^\circ) - \sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ) t'_{RR,k}(j) \tag{4.90}$$

Equation (4.90) shows how the entries of $t'_{SR,k}$ that correspond to agents of sending sub-network s, are related to the entries of $t'_{RR,k}$ through the values of the desired beliefs at θ_s°. Therefore, the set of all possible solutions to (4.85) consist of vectors whose entries satisfy (4.90) for each s. In other words, by arbitrarily fixing the entries of $t'_{RR,k}$, we compute the entries of $t'_{SR,k}$ using (4.90) for each s to obtain a solution to (4.85). This is because (4.85) is made of S equations that only indicate how the entries of $t'_{SR,k}$ that correspond to each sending sub-network s are related to $t'_{RR,k}$ without having any additional equation for the entries of $t'_{RR,k}$. Note that it does not matter how the individual entries of $t'_{SR,k}$ that correspond to sub-network s are chosen as long as their sum satisfies (4.90). However, in the problem we are trying to solve, we are not interested in the entire set of solutions to (4.85). This is because we have two additional constraints (4.88) and (4.89). Therefore, in our problem we cannot arbitrarily fix the entries of $t'_{RR,k}$ to any values as we need to also satisfy (4.88) and (4.89). Constraint (4.88) implies that (4.90) should be non-negative for each sending sub-network s, i.e.,

$$q_k(\theta_s^\circ) \geq \sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ) t'_{RR,k}(j) \tag{4.91}$$

Therefore, the problem reduces to finding $t'_{RR,k}$ that satisfies:

$$q_k(\theta_s^\circ) \geq \sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ) t'_{RR,k}(j), \quad \forall s \tag{4.92a}$$

$$t'_{RR,k} \succ 0 \tag{4.92b}$$

If it possible to find $t'_{RR,k}$ that satisfies (4.92a) and (4.92b), then $t'_{SR,k}$ can de determined using (4.90) and therefore a solution for agent k is found. Finding $t'_{RR,k}$ that satisfies (4.92a) and (4.92b) is always possible. By appropriately attenuating the entries of $t'_{RR,k}$, we can have the right-hand side of (4.92a) smaller than $q_k(\theta_s^\circ)$. For instance, one solution is to assign the same value $\epsilon_k > 0$ to all entries of $t'_{RR,k}$. Then from (4.92a), we have for each s:

$$q_k(\theta_s^\circ) \geq \epsilon_k \sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ) \tag{4.93}$$

which means that ϵ_k should be chosen so that:

$$0 < \epsilon_k \leq \min_s \left\{ \frac{q_k(\theta_s^\circ)}{\sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ)} \right\} \tag{4.94}$$

We mentioned that, after finding $t'_{RR,k}$ that satisfies (4.92a) and (4.92b), $t'_{SR,k}$ can be determined using (4.90). We can alternatively express the solutions of $t'_{SR,k}$ using the same approach of the previous section. This is because after choosing the entries of $t'_{RR,k}$, the problem is now similar to the previous problem of finding $t_{SR,k}$ while $t_{RR,k}$ is given. Therefore, the solutions for $t'_{SR,k}$ can be also given by (4.31). Note that (4.31) is expressed in terms of v'_k to take into account that agent k may not be connected to some sending sub-networks, in the earlier section. Since in this case we are focusing on agent k connected to all sending sub-networks, the solution for $t'_{SR,k}$ is given by (4.31 where v_k is used instead of v'_k.

In summary, when agent k is connected to all sending sub-networks, the problem can have an infinite number of solutions. We first find $t'_{RR,k}$ that satisfies (4.92a) and (4.92b). Then, the entries of $t'_{SR,k}$ are nonnegative values chosen to satisfy (4.90). In other words, when a receiving agent k is under the direct influence of all sending sub-networks, it is relatively straightforward to affect its beliefs, especially since the influence from its receiving neighbors can be attenuated as much as needed through the choice ϵ_k.

Case 2: Agent k is connected to some sending sub-networks

We now consider the case where agent k is influenced by only a subset of the sending networks. Without loss of generality, we assume it is connected to the first s' sending sub-networks. In this case, E_k will have the following form:

$$E_k = \begin{bmatrix} \mathbb{1}^{\mathsf{T}}_{N_1^k} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbb{1}^{\mathsf{T}}_{N_2^k} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbb{1}^{\mathsf{T}}_{N_{s'}^k} \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \end{bmatrix} \tag{4.95}$$

and relation (4.77) becomes:

$$\begin{bmatrix} \mathbb{1}^{\mathsf{T}}_{N_1^k} & \mathbf{0} & \dots & \mathbf{0} & q_{k(1)}(\theta_1^\circ) & \dots & q_{k(N_{gR}^k)}(\theta_1^\circ) \\ \mathbf{0} & \mathbb{1}^{\mathsf{T}}_{N_2^k} & \dots & \mathbf{0} & q_{k(1)}(\theta_2^\circ) & \dots & q_{k(N_{gR}^k)}(\theta_2^\circ) \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbb{1}^{\mathsf{T}}_{N_{s'}^k} & q_{k(1)}(\theta_{s'}^\circ) & \dots & q_{k(N_{gR}^k)}(\theta_{s'}^\circ) \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & q_{k(1)}(\theta_{s'+1}^\circ) & \dots & q_{k(N_{gR}^k)}(\theta_{s'+1}^\circ) \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & q_{k(1)}(\theta_S^\circ) & \dots & q_{k(N_{gR}^k)}(\theta_S^\circ) \end{bmatrix} \begin{bmatrix} t'_{SR,k} \\ t'_{RR,k} \end{bmatrix} = q_k \tag{4.96}$$

The problem now is to find $t'_{SR,k}$ and $t'_{RR,k}$ that satisfy (4.96) subject to constraints (4.78)-(4.79). As before, if we can find $t'_{SR,k}$ and $t'_{RR,k}$ that satisfy (4.96), then condition (4.78) will be automatically satisfied. Note now that (4.96) consists of s' equations that relate the entries of $t'_{SR,k}$ to the entries of $t'_{RR,k}$, and $S - s'$ equations that involve the entries of $t'_{RR,k}$.

Therefore, any vector that satisfies (4.96) will have the following property:

$$\sum_{\ell=1}^{N_1^k} t'_{SR,k}(\ell) = q_k(\theta_s^\circ) - \sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ) t'_{RR,k}(j) \tag{4.97}$$

but only for $s \leq s'$. In other words, the entries of $t'_{SR,k}$ that correspond to sub-network $s \leq s'$ are expressed in terms of $t'_{RR,k}$ through (4.97). In addition, and differently from case 1, any solution to (4.96) should also satisfy:

$$\sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ) t'_{RR,k}(j) = q_k(\theta_s^\circ) \tag{4.98}$$

for any $s > s'$. Likewise, constraint (4.88) implies that (4.97) should be non-negative for each sending sub-network s where $s \leq s'$, i.e.,

$$q_k(\theta_s^\circ) \geq \sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ) t'_{RR,k}(j) \tag{4.99}$$

for any $s \leq s'$. Therefore, the problem reduces to finding $t'_{RR,k}$ that satisfies:

$$q_k(\theta_s^\circ) \geq \sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ) t'_{RR,k}(j), \quad s \leq s' \tag{4.100a}$$

$$q_k(\theta_s^\circ) = \sum_{j=1}^{N_{gR}^k} q_{k(j)}(\theta_s^\circ) t'_{RR,k}(j), \quad s > s' \tag{4.100b}$$

$$t'_{RR,k} \succ 0 \tag{4.100c}$$

If it possible to find $t'_{RR,k}$ that satisfies (4.100a)-(4.100c), then $t'_{SR,k}$ can de determined using (4.97) or alternatively using (4.31). However, in contrast to the case studied in the previous case, finding $t'_{RR,k}$ that satisfies conditions (4.100a)-(4.100c) may not be always possible. For instance, consider agent k shown in Fig. 4.5.2, which is connected to only the first sending sub-network but not to the other two sending sub-networks.

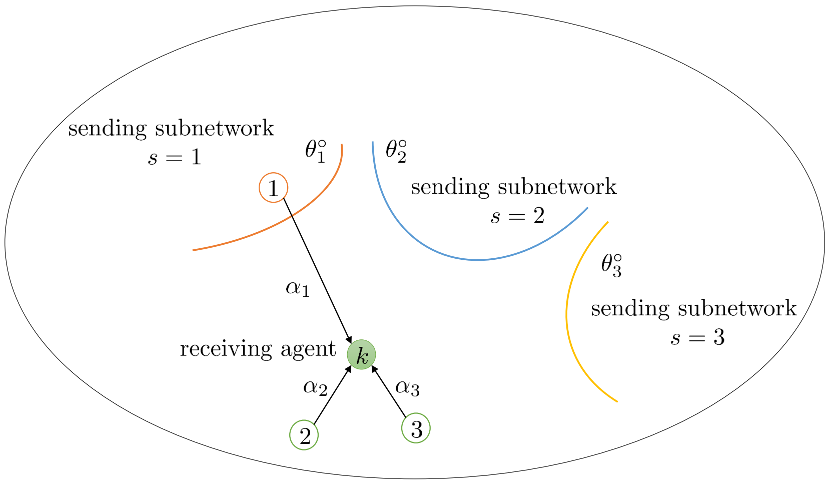

Figure 4.5.2: An example where receiving agent k is only connected to one sending sub-network.

Let us consider its desired limiting belief as

$$\begin{bmatrix} q_k(\theta_1^\circ) \\ q_k(\theta_2^\circ) \\ q_k(\theta_3^\circ) \end{bmatrix} = \begin{bmatrix} 0.1 \\ 0.45 \\ 0.45 \end{bmatrix} \tag{4.101}$$

while the desired limiting beliefs for its neighbors are:

$$\begin{bmatrix} q_2(\theta_1^\circ) \\ q_2(\theta_2^\circ) \\ q_2(\theta_3^\circ) \end{bmatrix} = \begin{bmatrix} 0.2 \\ 0.5 \\ 0.3 \end{bmatrix}, \quad \begin{bmatrix} q_3(\theta_1^\circ) \\ q_3(\theta_2^\circ) \\ q_3(\theta_3^\circ) \end{bmatrix} = \begin{bmatrix} 0.1 \\ 0.4 \\ 0.5 \end{bmatrix} \tag{4.102}$$

Then, from (4.100a), we should have:

$$q_k(\theta_1^\circ) \geq \alpha_2 q_2(\theta_1^\circ) + \alpha_3 q_3(\theta_1^\circ) \implies 0.1 \geq 0.2\alpha_2 + 0.1\alpha_3 \tag{4.103}$$

and from (4.100b),

$$q_k(\theta_2^\circ) = \alpha_2 q_2(\theta_2^\circ) + \alpha_3 q_3(\theta_2^\circ) \implies 0.45 = 0.5\alpha_2 + 0.4\alpha_3 \tag{4.104}$$

$$q_k(\theta_3^\circ) = \alpha_2 q_2(\theta_3^\circ) + \alpha_3 q_3(\theta_3^\circ) \implies 0.45 = 0.3\alpha_2 + 0.5\alpha_3 \tag{4.105}$$

Solving (4.104) and (4.105) gives the following solution: $\alpha_2 = 0.3462$ and $\alpha_3 = 0.6923$. However, $0.2\alpha_2 + 0.1\alpha_3 = 0.1385$, which violates (4.103). Still, we can have cases where all conditions (4.100a)-(4.100c) can be met (we are going to provide one example in a later section), then in these cases, we choose $t'_{SR,k}$ according to (4.97).

We observe from this case that the fewer the sending networks that influence agent k, the harder it is to affect its limiting belief. This emphasizes again the idea that the structure of the receiving sub-networks helps in limiting external manipulation.

Case 3: Agent k is not connected to any sending sub-networks

When agent k is not connected to any sending sub-network, relation (4.77) reduces to:

$$Q_k t'_{RR,k} = q_k \tag{4.106}$$

The problem is then to find $t'_{RR,k}$ that satisfies:

$$\boxed{\begin{aligned} Q_k t'_{RR,k} &= q_k && \text{(4.107a)} \\ \mathbb{1}^\top t'_{RR,k} &= 1 && \text{(4.107b)} \\ t'_{RR,k} &\succ 0 && \text{(4.107c)} \end{aligned}}$$

This problem might not have an exact solution. For instance, we discuss two examples in Appendix 4.A, where in the second example, we have an agent that is not connected to any sending sub-network and its desired belief cannot be expressed as a convex combination of the desired beliefs of its neighbors.

Comment and analysis

Since the problem of finding T_{SR} and T_{RR} satisfying (4.67)-(4.70) is separable, we studied the possibility of finding a solution for each column of T_{SR} and T_{RR}. We analyzed the problem for 3 cases and discovered that for the first case (when agent k is connected to at least one agent from each sending sub-network), problem (4.77)–(4.80) always has a solution. That is, if an agent k is connected to all sending sub-networks and *given knowledge of the limiting beliefs of its neighbors*, we can always find the weight combination for agent k such that (4.83) is satisfied. For the second case (when agent k is connected to some sending sub-networks) and the third case (when agent k is not connected to any sending sub-network), we found out that problem (4.77)–(4.80) might not always have a solution, i.e., it is not always possible to satisfy (4.83). These scenarios reinforce again the idea that the internal structure of receiving agents can resist some of the external influence.

However, for Q to be achievable (i.e., for the beliefs of all receiving agents converge to the desired beliefs), a solution must exist for each agent k. If the desired limiting belief of any receiving agent cannot be written as a convex combination of the limiting beliefs of its neighbors (i.e., a solution cannot be found for problem (4.77)–(4.80)), the whole scenario is not achievable. Even if it is possible for agent k to find its appropriate weights $t'_{SR,k}$ and $t'_{RR,k}$, finding this solution is based on the knowledge of the desired limiting beliefs of its neighbors. However, if one of the receiving neighbors cannot reach its desired belief, agent k will not be able anymore to reach its desired belief. Therefore, for Q to be attainable, a solution for problem (4.77)–(4.80) must exist for each receiving agent k. If Q is not attainable, then the desired scenario should be modified to an attainable scenario, by taking into consideration the limitation provided by the internal connection of the receiving sub-networks. Or an approximate least-squares solution for the weights can be found. That is, we can instead seek to solve

$$\min_{t'_{SR,k},t'_{RR,k}} \left\| B_k \begin{bmatrix} t'_{SR,k} \\ t'_{RR,k} \end{bmatrix} - q_k \right\|^2 \tag{4.108}$$

subject to

$$\mathbb{1}^\mathsf{T} t'_{SR,k} + \mathbb{1}^\mathsf{T} t'_{RR,k} = 1 \tag{4.109}$$

$$t'_{SR,k} \succeq \mathbf{0} \tag{4.110}$$

$$t'_{RR,k} \succ \mathbf{0} \tag{4.111}$$

The last condition can be relaxed to the following:

$$t'_{RR,k} \succeq \epsilon_k \mathbb{1} \tag{4.112}$$

where $0 < \epsilon_k < 1$. Clearly, when we solve problem (4.108)–(4.112), this does not mean that the objective function (4.108) will be zero at this solution. Note further that the optimization problem (4.108)–(4.112) is a quadratic convex problem: its objective function is quadratic, and it has a convex equality constraint (4.109) and inequality constraints (4.110) and (4.112). The inequality constraints are element-wise, i.e., $t'_{RR,k}(j) \geq \epsilon_k$ for all j, which can be equivalently written as $e_j^\top t'_{RR,k} \geq \epsilon_k$ for all j where e_j is a vector where all its elements are zero expect for the j^{th} element that is one. In this way, the problem becomes a classic constrained convex optimization problem, which can be solved numerically (using for instance interior point methods).

4.6 Simulation Results

We illustrate the previous results with the following simulation example. Consider the social network shown in Fig. 4.6.1 which consists of $N = 23$ agents.

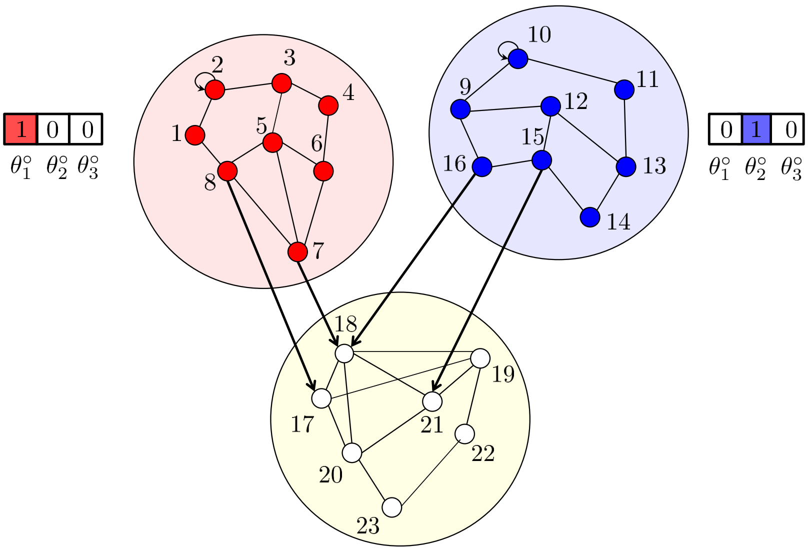

Figure 4.6.1: A weakly-connected network consisting of three sub-networks.

We assume that there are 3 possible events $\Theta = \{\theta_1^\circ, \theta_2^\circ, \theta_3^\circ\}$, where θ_1° is the true event for the first sending sub-network, θ_2° is the true event for the second sending sub-network, and θ_3° is the true event for the receiving sub-network. We further assume that the observational signals of each agent k are binary and belong to $Z_k = \{H, T\}$ where H denotes head and T denotes tail.

Agents of the first sending sub-network are connected through the following combination matrix:

$$A_1 = \begin{bmatrix} 0 & 0.3 & 0 & 0 & 0 & 0 & 0 & 0.3 \\ 0.4 & 0 & 0.3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.7 & 0 & 0.5 & 0.25 & 0 & 0 & 0 \\ 0 & 0 & 0.4 & 0 & 0 & 0.3 & 0 & 0 \\ 0 & 0 & 0.3 & 0 & 0 & 0.1 & 0.2 & 0.45 \\ 0 & 0 & 0 & 0.5 & 0.25 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0 & 0.3 & 0.6 & 0 & 0.25 \\ 0.6 & 0 & 0 & 0 & 0.2 & 0 & 0.7 & 0 \end{bmatrix} \tag{4.113}$$

Agents of the second sending sub-network are connected through the following combination matrix:

$$A_2 = \begin{bmatrix} 0 & 0.35 & 0 & 0.3 & 0 & 0 & 0 & 0.25 \\ 0.1 & 0.25 & 0.5 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.4 & 0 & 0 & 0.8 & 0 & 0 & 0 \\ 0.1 & 0 & 0 & 0 & 0.1 & 0 & 0.6 & 0 \\ 0 & 0 & 0.5 & 0.3 & 0 & 0.45 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.1 & 0 & 0.3 & 0 \\ 0 & 0 & 0 & 0.4 & 0 & 0.55 & 0 & 0.75 \\ 0.8 & 0 & 0 & 0 & 0 & 0 & 0.1 & 0 \end{bmatrix} \tag{4.114}$$

The matrices T_{SR} and T_{RR} are going to be designed so that the desired limiting beliefs for receiving agents are as follows:

$$Q_1 = \begin{bmatrix} 0.55 & 0.5 & 0.5 & 0.5 & 0.45 & 0.5 & 0.5 \\ 0.45 & 0.5 & 0.5 & 0.5 & 0.55 & 0.5 & 0.5 \end{bmatrix} \tag{4.115}$$

In other words, the weights are going to be designed so that θ_1° and θ_2° are almost equally probable for the receiving agents. This illustrates the case when the receiving agents listen to two different perspectives from two media sources that are trustworthy for them, which leaves them undecided regarding which true state to choose.

The likelihood of the head signals for each receiving agent k is selected as the following matrix:

$$L_R(H) = \begin{bmatrix} 5/8 & 3/4 & 1/6 & 7/8 & 2/3 & 1/3 & 1/4 \\ 5/8 & 3/4 & 1/6 & 7/8 & 2/3 & 1/3 & 1/4 \\ 5/8 & 3/4 & 1/6 & 7/8 & 2/3 & 1/3 & 1/4 \end{bmatrix}$$

where each (j,k)-th element of this matrix corresponds to $L_k(H/\theta_j)$, i.e., each column corresponds to one agent and each row to one network state. The likelihood of the tail signal is $L(T) = \mathbb{1}_{3\times 7} - L(H)$. The likelihood of the head signals for each sending agent k of the first sending sub-network is selected as the following matrix:

$$L_1(H) = \begin{bmatrix} 5/8 & 3/4 & 1/6 & 1/2 & 1/3 & 1/5 & 4/5 & 1/2 \\ 5/8 & 3/4 & 1/6 & 2/3 & 1/2 & 1/5 & 2/3 & 1/2 \\ 1/4 & 3/4 & 1/3 & 1/2 & 1/4 & 1/5 & 4/5 & 1/3 \end{bmatrix} \tag{4.116}$$

and the likelihood of the head signals of agents of the second sending sub-network is:

$$L_2(H) = \begin{bmatrix} 7/8 & 5/8 & 1/4 & 1/2 & 1/2 & 1/2 & 6/7 & 1/4 \\ 7/8 & 2/3 & 5/8 & 1/3 & 1/2 & 1/2 & 8/9 & 1/4 \\ 1/3 & 2/3 & 5/8 & 1/4 & 1/2 & 1/5 & 8/9 & 1/4 \end{bmatrix} \tag{4.117}$$

Design and Result Simulation

To achieve Q_1, we design T_{SR} and T_{RR} using the results in the previous section. The details of the numerical derivation are omitted for brevity. The non-zero weights in T_{SR} are shown in Fig. 4.6.2, and T_{RR} is given as follows:

$$T_{RR} = \begin{bmatrix} 0 & 0.1 & 0.25 & 0.25 & 0 & 0 & 0 \\ 0.3 & 0 & 0.25 & 0.25 & 0.3 & 0 & 0 \\ 0.3 & 0.1 & 0 & 0 & 0.3 & 0.5 & 0 \\ 0.3 & 0.1 & 0 & 0 & 0.3 & 0 & 0.5 \\ 0 & 0.1 & 0.25 & 0.25 & 0 & 0 & 0 \\ 0 & 0 & 0.25 & 0 & 0 & 0 & 0.5 \\ 0 & 0 & 0 & 0.25 & 0 & 0.5 & 0 \end{bmatrix} \tag{4.118}$$

We run this example for 7000 time iterations. We assigned to each agent an initial belief that is uniform over $\{\theta_1^\circ, \theta_2^\circ, \theta_3^\circ\}$. Figures 4.6.3 and 4.6.4 show the evolution of $\boldsymbol{\mu}_{k,i}(\theta_1^\circ)$ and $\boldsymbol{\mu}_{k,i}(\theta_2^\circ)$ of agents in the receiving sub-network. These figures show the convergence of the beliefs of the agents in the receiving sub-networks to the desired beliefs in Q_1. Figure 4.6.2 illustrates with color the limiting beliefs of receiving agents.

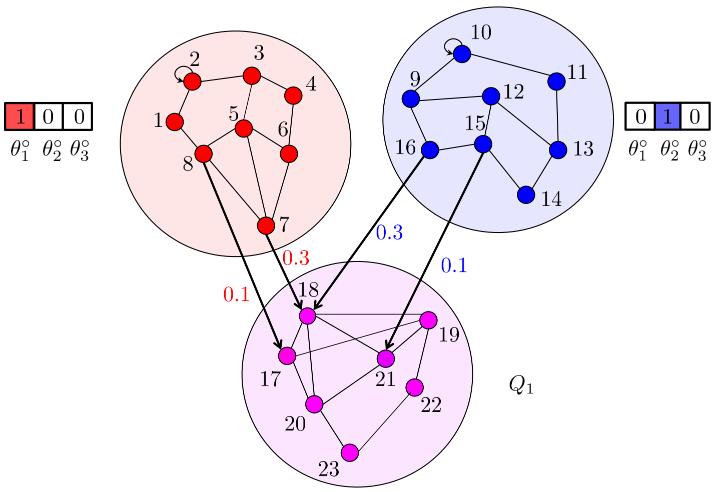

Figure 4.6.2: Illustration of the limiting beliefs of receiving agents

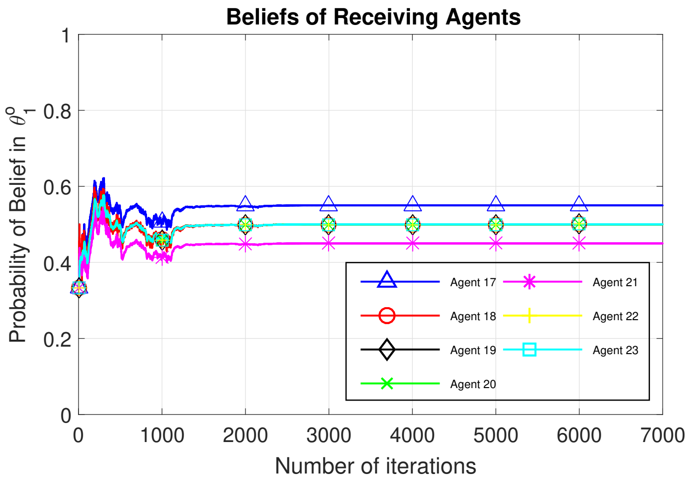

Figure 4.6.3: Evolution of the beliefs of the receiving agents at θ_1° over time

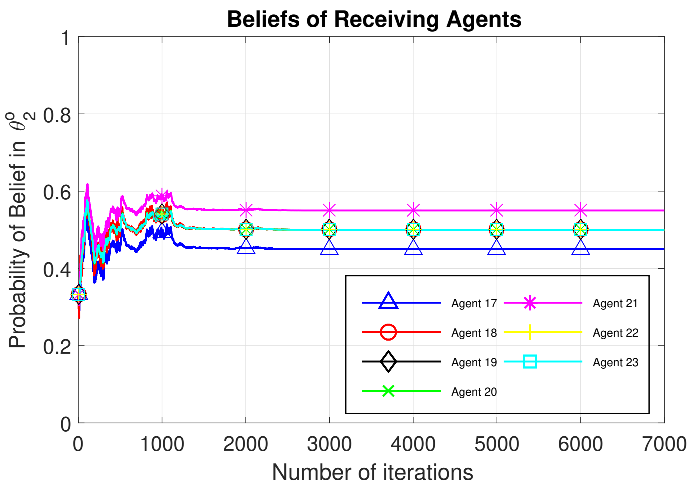

Figure 4.6.4: Evolution of the beliefs of the receiving agents at θ_2° over time

4.7 Conclusions

In this chapter, we characterized the set of beliefs that can be imposed on non-influential agents and clarified how the graph topology of these latter agents helps resist manipulation but only to a certain degree. We also derived a design procedures that allow influential agents to drive the beliefs of non-influential agents to desirable attainable states. The results of this chapter are based on [66].

4.A Two Revealing Examples for the Design Procedure (4.77)-(4.80)

Example I: Cases 1 and 2 (k is influenced by sending networks)

Consider the network shown in Fig. 4.A.1. It consists of $N = 8$ agents, two sending sub-networks and one receiving sub-network, with the following combination matrix:

$$A = \left[\begin{array}{ccccc|ccc} 0.2 & 0.2 & 0.8 & 0 & 0 & \times & 0 & 0 \\ 0.5 & 0.4 & 0.1 & 0 & 0 & 0 & \times & 0 \\ 0.3 & 0.4 & 0.1 & 0 & 0 & \times & 0 & \times \\ 0 & 0 & 0 & 0.4 & 0.3 & \times & \times & 0 \\ 0 & 0 & 0 & 0.6 & 0.7 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & \times & \times \\ 0 & 0 & 0 & 0 & 0 & \times & 0 & \times \\ 0 & 0 & 0 & 0 & 0 & \times & \times & \times \end{array}\right] \tag{4.119}$$

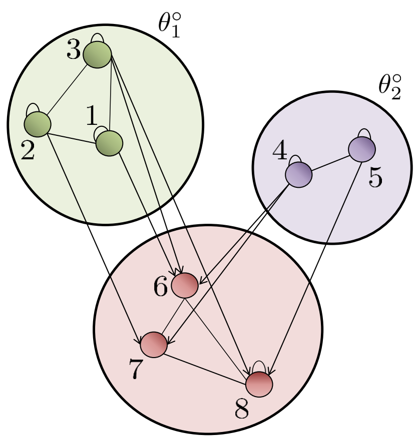

Figure 4.A.1: A weakly connected network consisting of three sub-networks. In this example, receiving agents 6 and 7 are influenced by both sending networks, while agent 8 is only influenced by the first sending network.

We assume that there are 3 possible states $\Theta = \{\theta_1^\circ, \theta_2^\circ, \theta_3^\circ\}$, where θ_1° is the true event for the first sending sub-network, θ_2° is the true event for the second sending sub-network, and θ_3° is the true event for the receiving sub-network. Let us consider the case where we want to design T_{SR} and T_{RR} to attain the desired limiting beliefs

$$Q = \begin{bmatrix} 0.2 & 0.3 & 0.5 \\ 0.8 & 0.7 & 0.5 \end{bmatrix} \tag{4.120}$$

The matrix B is therefore of the following form:

$$B = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0.2 & 0.3 & 0.5 \\ 0 & 0 & 0 & 1 & 1 & 0.8 & 0.7 & 0.5 \end{bmatrix} \tag{4.121}$$

We start with agent 6. After eliminating entries to satisfy the sparsity in the connections, we are reduced to finding $t'_{SR,6}$ and $t'_{RR,6}$ that satisfy

$$\underbrace{\begin{bmatrix} 1 & 1 & 0 & 0.3 & 0.5 \\ 0 & 0 & 1 & 0.7 & 0.5 \end{bmatrix}}_{\triangleq B_6} \begin{bmatrix} t'_{SR,6} \\ t'_{RR,6} \end{bmatrix} = \underbrace{\begin{bmatrix} 0.2 \\ 0.8 \end{bmatrix}}_{\triangleq q_6} \tag{4.122}$$

Let

$$\begin{bmatrix} t'_{SR,6} \\ t'_{RR,6} \end{bmatrix} \triangleq \begin{bmatrix} \alpha_1 & \alpha_2 & \alpha_3 & \alpha_4 & \alpha_5 \end{bmatrix}^\mathsf{T} \tag{4.123}$$

Agent 6 is connected to the two sending sub-networks (case 1). Therefore, the problem has a solution, where $t'_{SR,6}$ (α_1, α_2 and α_3) can be expressed in terms of $t'_{RR,6}$ (α_4 and α_5). More precisely, from (4.122) and (4.90), we have:

$$\alpha_1 + \alpha_2 = 0.2 - 0.3\alpha_4 - 0.5\alpha_5 \tag{4.124}$$

$$\alpha_3 = 0.8 - 0.7\alpha_4 - 0.5\alpha_5 \tag{4.125}$$

According to (4.91), to ensure that α_1, α_2 and α_3 can be chosen as nonnegative numbers, the scalars α_4 and α_5 should be chosen to satisfy

$$0.3\alpha_4 + 0.5\alpha_5 \leq 0.2 \tag{4.126}$$

$$0.7\alpha_4 + 0.5\alpha_5 \leq 0.8 \tag{4.127}$$

Note that what matters for scalars α_1 and α_2 (the weights with which the data received from sending sub-network 1 is scaled) is that their sum should be equal to $0.2 - 0.3\alpha_4 - 0.5\alpha_5$ according to (4.124). In other words, when a receiving agent is connected to many agents from the same sending sub-network, it does not matter how much weight is given to each of these agents as long as the sum of these weights takes the required value. This is because the beliefs of agents of the same sending sub-networks will converge to the same final distribution. An alternative way to express (4.124) is to set α_1 and α_2 to the following:

$$\alpha_1 = \frac{1}{2}(0.2 - 0.3\alpha_4 - 0.5\alpha_5) + \beta \tag{4.128}$$

$$\alpha_2 = \frac{1}{2}(0.2 - 0.3\alpha_4 - 0.5\alpha_5) - \beta \tag{4.129}$$

where

$$|\beta| \leq \frac{1}{2}(0.2 - 0.3\alpha_4 - 0.5\alpha_5) \tag{4.130}$$

This choice of β ensures that α_1 and α_2 are non-negative and less than $0.2 - 0.3\alpha_4 - 0.5\alpha_5$. Moreover, we can check from (4.128) and (4.129) that their sum satisfies (4.124). Therefore, the solution has the following form:

$$\begin{bmatrix} t'_{SR,6} \\ t'_{RR,6} \end{bmatrix} = \begin{bmatrix} \frac{1}{2}(0.2 - 0.3\alpha_4 - 0.5\alpha_5) + \beta \\ \frac{1}{2}(0.2 - 0.3\alpha_4 - 0.5\alpha_5) - \beta \\ 0.8 - 0.7\alpha_4 - 0.5\alpha_5 \\ \alpha_4 \\ \alpha_5 \end{bmatrix} \tag{4.131}$$

where

$$0.3\alpha_4 + 0.5\alpha_5 \leq 0.2 \tag{4.132}$$

$$0.7\alpha_4 + 0.5\alpha_5 \leq 0.8 \tag{4.133}$$

$$\alpha_4 > 0,\ \alpha_5 > 0 \tag{4.134}$$

$$|\beta| \leq \frac{1}{2}(0.2 - 0.3\alpha_4 - 0.5\alpha_5) \tag{4.135}$$

For example, one solution is to assign the same value ϵ_6 for α_4 and α_5. Then, from (4.132), (4.133) and (4.94), we have:

$$0 < \epsilon_6 \leq \min\left\{\frac{0.2}{0.5+0.3}, \frac{0.8}{0.7+0.5}\right\} = 0.25 \tag{4.136}$$

Let $\epsilon_6 = 0.1 = \alpha_4 = \alpha_5$, then

$$\alpha_1 + \alpha_2 = 0.2 - 0.3\alpha_4 - 0.5\alpha_5 = 0.12 \tag{4.137}$$

$$\alpha_3 = 0.8 - 0.7\alpha_4 - 0.5\alpha_5 = 0.68 \tag{4.138}$$

We can choose $\alpha_1 = 0.1$ and $\alpha_2 = 0.02$. Therefore, a possible solution for $t_{SR,6}$ is:

$$t_{SR,6} = \begin{bmatrix} 0.1 & 0 & 0.02 & 0.68 & 0 & 0 & 0.1 & 0.1 \end{bmatrix}^{\mathsf{T}} \tag{4.139}$$

We follow a similar procedure for agent 7 and obtain:

$$\begin{bmatrix} t'_{SR,7} \\ t'_{RR,7} \end{bmatrix} \triangleq \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \end{bmatrix} = \begin{bmatrix} 0.3 - 0.2\beta_3 - 0.5\beta_4 \\ 0.7 - 0.8\beta_3 - 0.5\beta_4 \\ \beta_3 \\ \beta_4 \end{bmatrix} \tag{4.140}$$

where

$$0.2\beta_3 + 0.5\beta_4 \leq 0.3 \tag{4.141}$$

$$0.8\beta_3 + 0.5\beta_4 \leq 0.7 \tag{4.142}$$

$$\beta_3 > 0,\ \beta_4 > 0 \tag{4.143}$$

For this agent, we can choose for instance as a solution $\beta_3 = 0.2$ and $\beta_4 = 0.1$ (as they both satisfy (4.141) and (4.142)). In this case,

$$\beta_1 = 0.3 - 0.2\beta_3 - 0.5\beta_4 = 0.21 \tag{4.144}$$

$$\beta_2 = 0.7 - 0.8\beta_3 - 0.5\beta_4 = 0.49 \tag{4.145}$$

Therefore, a possible solution for $t_{SR,7}$ is:

$$t_{SR,7} = \begin{bmatrix} 0 & 0.21 & 0 & 0.49 & 0 & 0.2 & 0 & 0.1 \end{bmatrix}^{\mathsf{T}} \tag{4.146}$$

Agent 8 is connected to the first sending sub-network only (case 2). For this agent, we have:

$$\underbrace{\begin{bmatrix} 1 & 0.2 & 0.3 & 0.5 \\ 0 & 0.8 & 0.7 & 0.5 \end{bmatrix}}_{\triangleq B_8} \begin{bmatrix} t'_{SR,8} \\ t'_{RR,8} \end{bmatrix} = \underbrace{\begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}}_{\triangleq q_8} \tag{4.147}$$

Let

$$\begin{bmatrix} t'_{SR,8} \\ t'_{RR,8} \end{bmatrix} \triangleq \begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \gamma_3 \\ \gamma_4 \end{bmatrix} \tag{4.148}$$

Therefore, from (4.147), (4.97) and (4.98), we have:

$$\gamma_1 = 0.5 - 0.2\gamma_2 - 0.3\gamma_3 - 0.5\gamma_4 \tag{4.149}$$

$$0.8\gamma_2 + 0.7\gamma_3 + 0.5\gamma_4 = 0.5 \tag{4.150}$$

and any vector that satisfies (4.147) has the following form:

$$\begin{bmatrix} t'_{SR,8} \\ t'_{RR,8} \end{bmatrix} = \begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \gamma_3 \\ \gamma_4 \end{bmatrix} = \begin{bmatrix} 0.5 - 0.2\gamma_2 - 0.3\gamma_3 - 0.5\gamma_4 \\ \gamma_2 \\ \gamma_3 \\ \gamma_4 \end{bmatrix} \tag{4.151}$$

where

$$0.8\gamma_2 + 0.7\gamma_3 + 0.5\gamma_4 = 0.5 \tag{4.152}$$

Now to ensure that γ_1 is non-negative, γ_2, γ_3 and γ_4 should be chosen as follows (as in (4.99)):

$$0.2\gamma_2 + 0.3\gamma_3 + 0.5\gamma_4 \leq 0.5 \tag{4.153}$$

Therefore, a solution in this case should satisfy (4.151) subject to

$$0.8\gamma_2 + 0.7\gamma_3 + 0.5\gamma_4 = 0.5 \tag{4.154}$$

$$0.2\gamma_2 + 0.3\gamma_3 + 0.5\gamma_4 \leq 0.5 \tag{4.155}$$

$$\gamma_2 > 0,\ \gamma_3 > 0,\ \gamma_4 > 0 \tag{4.156}$$

For this example, finding γ_2, γ_3 and γ_4 that satisfy (4.154)-(4.155) is always possible. To see this, for any choice of γ_2, γ_3 and γ_4 that satisfy (4.154), condition (4.155) is automatically satisfied. Indeed, if (4.154) is satisfied then

$$0.5\gamma_2 + 0.5\gamma_3 + 0.5\gamma_4 \leq 0.8\gamma_2 + 0.7\gamma_3 + 0.5\gamma_4 = 0.5 \tag{4.157}$$

$$\implies 0.5(\gamma_2 + \gamma_3 + \gamma_4) \leq 0.5 \implies \gamma_2 + \gamma_3 + \gamma_4 \leq 1 \tag{4.158}$$

Therefore,

$$\gamma_2 + \gamma_3 + \gamma_4 - 0.8\gamma_2 - 0.7\gamma_3 - 0.5\gamma_4 \leq 1 - 0.5 \tag{4.159}$$

$$\implies 0.2\gamma_2 + 0.3\gamma_3 + 0.5\gamma_4 \leq 0.5 \tag{4.160}$$

For instance, one possible choice for γ_2, γ_3 and γ_4 that satisfies (4.154) is

$$\gamma_2 = \gamma_3 = \gamma_4 = \frac{0.5}{0.8 + 0.7 + 0.5} = 0.25 \tag{4.161}$$

Then,

$$\gamma_1 = 0.5 - 0.2\gamma_2 - 0.3\gamma_3 - 0.5\gamma_4 = 0.25 \tag{4.162}$$

Therefore, a possible solution for $t_{SR,8}$ is:

$$t_{SR,8} = \begin{bmatrix} 0 & 0 & 0.25 & 0 & 0 & 0.25 & 0.25 & 0.25 \end{bmatrix}^{\mathsf{T}} \tag{4.163}$$

Thus, the overall solution is:

$$\begin{bmatrix} T_{SR} \\ \hline T_{RR} \end{bmatrix} = \left[\begin{array}{ccc} 0.1 & 0 & 0 \\ 0 & 0.21 & 0 \\ 0.02 & 0 & 0.25 \\ 0.68 & 0.49 & 0 \\ 0 & 0 & 0 \\ \hline 0 & 0.2 & 0.25 \\ 0.1 & 0 & 0.25 \\ 0.1 & 0.1 & 0.25 \end{array}\right] \tag{4.164}$$

To verify that the beliefs of the receiving agents converge to the desired beliefs, we compute W^T from (3.7) and use (3.24) to determine the limiting beliefs at θ_1^o and θ_2^o at the receiving agents. This calculation gives

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_1^\circ) = \begin{cases} 0.1070 + 0.0310 + 0.0620 = 0.2, & k = 6 \\ 0.0258 + 0.2247 + 0.0494 = 0.3, & k = 7 \\ 0.0443 + 0.0852 + 0.3705 = 0.5, & k = 8 \end{cases}$$

and

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_2^\circ) = \begin{cases} 0.8, & k = 6 \\ 0.7, & k = 7 \\ 0.5, & k = 8 \end{cases}$$

Example II: Case 3 (agent k not influenced by sending networks)

Consider the network shown in Fig. 4.A.2, with the following combination matrix:

$$A = \left[\begin{array}{ccccc|ccc} 0.2 & 0.2 & 0.8 & 0 & 0 & \times & 0 & 0 \\ 0.5 & 0.4 & 0.1 & 0 & 0 & 0 & \times & 0 \\ 0.3 & 0.4 & 0.1 & 0 & 0 & \times & 0 & 0 \\ 0 & 0 & 0 & 0.4 & 0.3 & \times & \times & 0 \\ 0 & 0 & 0 & 0.6 & 0.7 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & \times & \times \\ 0 & 0 & 0 & 0 & 0 & \times & 0 & \times \\ 0 & 0 & 0 & 0 & 0 & \times & \times & 0 \end{array}\right] \tag{4.165}$$

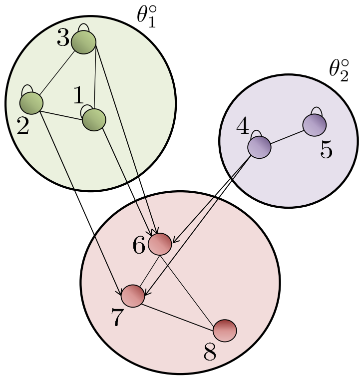

Figure 4.A.2: A weakly connected network consisting of three sub-networks. In this case, agent 8 is not influenced by any sending network.

What is different now is that agent 8 does not have is not connected to agent 3 (that is, agent 8 is not connected to any sending network). We are still assuming in this example that we have the same desired limiting beliefs:

$$Q = \begin{bmatrix} 0.2 & 0.3 & 0.5 \\ 0.8 & 0.7 & 0.5 \end{bmatrix} \tag{4.166}$$

For agents 6 and 7, the solutions found for their corresponding columns are still valid here. However, in this example, $t_{SR,8}$ should have all its elements equal to zero and $t_{RR,8}$ should have its third element equal to zero. Therefore, for agent 8, the problem reduces to finding

$t'_{RR,8}$ that satisfies the following relationship:

$$\underbrace{\begin{bmatrix} 0.2 & 0.3 \\ 0.8 & 0.7 \end{bmatrix}}_{\triangleq B_8} t'_{RR,8} = \underbrace{\begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}}_{\triangleq q_8} \tag{4.167}$$

where the elements of $t'_{RR,8}$ should be positive and add up to 1. Any convex combination of 0.2 and 0.3 can only produce a number between 0.2 and 0.3, but not 0.5. This is why in this case, the problem does not have a solution. However, we can seek instead a least-squares solution for agent 8:

$$\min_{t'_{RR,8}} \left\| B_8 t'_{RR,8} - q_8 \right\|^2 \tag{4.168}$$

subject to

$$t'_{RR,8} \succeq \epsilon_8 \mathbb{1} \tag{4.169}$$

$$\mathbb{1}^\mathsf{T} t'_{RR,8} = 1 \tag{4.170}$$

By choosing $\epsilon_8 = 0.01$ and solving it numerically, we obtain:

$$t'_{RR,8} = \begin{bmatrix} 0.01 \\ 0.99 \end{bmatrix} \tag{4.171}$$

This solution can be also deduced directly because $[0.3; 0.7]$ is the closer distribution to $[0.5; 0.5]$ than any other distribution formed by a convex combination of $[0.2; 0.8]$ and $[0.3; 0.7]$. Because the entries should be strictly greater than 0, the lowest possible value is given to the first entry of $t'_{RR,8}$. Therefore, with this choice:

$$\left[\begin{array}{c} T_{SR} \\ \hline T_{RR} \end{array}\right] = \left[\begin{array}{ccc} 0.1 & 0 & 0 \\ 0 & 0.21 & 0 \\ 0.02 & 0 & 0 \\ 0.68 & 0.49 & 0 \\ 0 & 0 & 0 \\ \hline 0 & 0.2 & 0.01 \\ 0.1 & 0 & 0.99 \\ 0.1 & 0.1 & 0 \end{array}\right] \tag{4.172}$$

we verify the limiting beliefs of the agents as follows. We compute W^{T} from (3.7) and use (3.24) to determine the limiting beliefs at θ_1^o and θ_2^o at the receiving agents . This calculation gives

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_1^\circ) = \begin{cases} 0.174, & k = 6 \\ 0.272, & k = 7 \\ 0.271, & k = 8 \end{cases}$$

and

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta_2^\circ) = \begin{cases} 0.826, & k = 6 \\ 0.728, & k = 7 \\ 0.729, & k = 8 \end{cases}$$

It is expected that the beliefs of agents 6 and 7 would not converge to the desired beliefs, because the belief of agent 8 cannot converge to its desired belief, which will definitely affect the limiting beliefs of agents 6 and 7. We know that agent 8 will not converge to its desired limiting belief because [0.5;0.5] cannot be obtained by any convex combination of [0.2;0.8] and [0.3;0.7] (its neighbors' limiting beliefs, (4.83)).

CHAPTER 5

Future Works

In this dissertation, we studied diffusion social learning over weakly-connected networks. We showed how the asymmetric flow of information prevents the receiving agents from learning their true states. We also examined the possibility for a leader-follower relationship to develop in the network, where receiving agents end up having beliefs focused on the set of the true states of sending agents. Moreover, we showed that not every control scheme is possible and that the internal structure of receiving agents plays role in limiting some forms of manipulation. We clarified the set of attainable beliefs and derived network's design procedures to drive receiving agents toward some desired attainable states.

We next list some possible future directions related to this work and some other open questions related to non-Bayesian learning:

- We studied the circumstances under which the receiving agents come under the total influence of sending agents. The next step is to study how receiving agents can mitigate this external influence and what conditions help them in learning their true states. One possible solution is that receiving agents stop receiving data from sending agents. However, receiving agents might not be able to tell that the data received from sending agents is what is drifting them from their own truth, or there might be also cases where one sending sub-network knows the truth of receiving agents and tries to guide them to their truth. In other words, how can receiving agents be equipped with better learning abilities to reach their truth despite the presence of sending agents?

- By having considered the weak graph connectivity, we were able to show that, when agents follow the diffusion model, a form of social disagreement arises in the receiving sub-networks. Another way to study the conditions for social disagreement is not through the connectivity of the network but by considering the presence of forceful or malicious agents, while assuming a strong network. As in [24, 46, 47], these agents might alter the information they are sending or force other agents to change their beliefs. An interesting direction would be to study similar settings in the case where agents update their beliefs according to the model of diffusion social learning.

- Another extension to diffusion social learning would be to study the model in a more dynamical environment. For instance, instead of analyzing the model over fixed topology, we can also consider time-varying graphs as in [23, 38]. Also, instead of assuming that the underlying true state is static, we can consider that the underlying true state changes with time. For instance, what might be true for a society over a period of time might become wrong at some point. In this case, how can agents track the underlying change?

- An important component in the learning process of agents is the continuous observation of time-independent private signals. Now if we assume that agents have a finite number of observations, how would the limited number of observations affect the process of learning of agents? Moreover, what if the observational signals are dependent over time?

- In our work, we assumed that agents perfectly share their beliefs with their neighbors. But we can also take into account communication issues such as communication delays [23]. Moreover, an agent might not be able to share its complete belief, especially when there are many possible values for the underlying state. Therefore, we can also consider the scenario where agents share samples of their beliefs.

REFERENCES

[1] A. Banerjee, "A simple model of herd behavior," *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, 1992.

[2] 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. 151–170, 1998.

[3] L. Smith and P. Sorensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.

[4] X. Vives, "Learning from others: A welfare analysis," *Games and Economic Behavior*, vol. 20, no. 2, pp. 177 – 200, 1997.

[5] C. Chamley and D. Gale, "Information revelation and strategic delay in a model of investment," *Econometrica*, vol. 62, no. 5, pp. 1065–85, 1994.

[6] A. Banerjee and D. Fudenberg, "Word-of-mouth learning," *Games and Economic Behavior*, vol. 46, no. 1, pp. 1–22, 2004.

[7] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[8] D. Acemoglu, K. Bimpikis, and A. Ozdaglar, "Dynamics of information exchange in endogenous social networks," *Theoretical Economics*, vol. 9, no. 1, pp. 41–97, 2014.

[9] C. Chamley, A. Scaglione, and L. Li, "Models for the diffusion of beliefs in social networks: An overview," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 16–29, May 2013.

[10] P. Molavi, C. Eksin, A. Ribeiro, and A. Jadbabaie, "Learning to coordinate in social networks," *Operations Research*, vol. 64, no. 3, pp. 605–621, 2016.

[11] V. Krishnamurthy and H.V. Poor, "A tutorial on interactive sensing in social networks," *IEEE Transactions on Computational Social Systems*, vol. 1, no. 1, pp. 3–21, March 2014.

[12] V. Krishnamurthy and M. Hamdi, "Mis-information removal in social networks: Constrained estimation on dynamic directed acyclic graphs," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 333–346, April 2013.

[13] V. Krishnamurthy and W. Hoiles, "Information diffusion in social sensing," *Numerical Algebra, Control and Optimization*, vol. 6, no. 3, pp. 365–411, Sept. 2016.

[14] V. Krishnamurthy and H. V. Poor, "Social learning and Bayesian games in multiagent signal processing," *IEEE Signal Processing Magazine*, vol. 33, no. 3, pp. 43–757, 2013.

[15] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[16] C. Chamley, *Rational Herds: Economic Models of Social Learning*, Cambridge University Press, 2004.

[17] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, vol. 1, no. 1, pp. 3–49, 2011.

[18] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. pp. 118–121, 1974.

[19] B. Golub and M. Jackson, "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, vol. 2, no. 1, pp. 112–49, 2010.

[20] P. DeMarzo, D. Vayanos, and J. Zwiebel, "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly Journal of Economics*, vol. 118, no. 3, pp. 909–968, 2003.

[21] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione, "Binary opinion dynamics with stubborn agents," *ACM Trans. Econ. Comput.*, vol. 1, no. 4, pp. 19:1–19:30, Dec. 2013.

[22] U. Krause, "A discrete nonlinear and non-autonomous model of consensus formation," *Communications in Difference Equations*, July 2000.

[23] V. Blondel, J.M. Hendrickx, A. Olshevsky, and J. Tsitsiklis, "Convergence in multiagent coordination, consensus, and flocking," in *Proc. IEEE Conf. on Decision and Control*, Jan. 2006, vol. 2005, pp. 2996 – 3000.

[24] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, "Spread of (mis)information in social networks," *Games and Economic Behavior*, vol. 70, no. 2, pp. 194–227, 2010.

[25] U. A. Khan, S. Kar, and J. M. F. Moura, "Higher dimensional consensus: Learning in large-scale networks," *IEEE Transactions on Signal Processing*, vol. 58, no. 5, pp. 2836–2849, May 2010.

[26] 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.

[27] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Business School Research Paper*, pp. 13–28, May 2013.

[28] P. Molavi, K. R. Rad, A. Tahbaz-Salehi, and A. Jadbabaie, "On consensus and exponentially fast social learning," in *Proc. IEEE ACC*, Montréal, Canada, June 2012, pp. 2165–2170.

[29] Q. Liu, A. Fang, L. Wang, and X. Wang, "Social learning with time-varying weights," *Journal of Systems Science and Complexity*, vol. 27, no. 3, pp. 581–593, June 2014.

[30] X. Zhao and A. H. Sayed, "Learning over social networks via diffusion adaptation," in *Proc. Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, Nov 2012, pp. 709–713.

[31] A. Lalitha, A. Sarwate, and T. Javidi, "Social learning and distributed hypothesis testing," in *Proc. IEEE International Symposium on Information Theory*, Honolulu, HI, June 2014, pp. 551–555.

[32] A. Nedić, A. Olshevsky, and C. A. Uribe, "A tutorial on distributed (non-Bayesian) learning: Problem, algorithms and results," in *Proc. IEEE Conf. on Decision and Control*, Dec 2016, pp. 6795–6801.

[33] M. A. Rahimian, S. Shahrampour, and A. Jadbabaie, "Learning without recall by random walks on directed graphs," in *Proc. IEEE Conf. on Decision and Control*, Dec 2015, pp. 5538–5543.

[34] S. Shahrampour and A. Jadbabaie, "Exponentially fast parameter estimation in networks using distributed dual averaging," in *Proc. IEEE Conf. on Decision and Control*, Firenze, Italy, Dec 2013, pp. 6196–6201.

[35] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, "Distributed detection: Finite-time analysis and impact of network topology," *IEEE Transactions on Automatic Control*, vol. 61, no. 11, pp. 3256–3268, Nov 2016.

[36] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-Bayesian learning," *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538–5553, Nov 2017.

[37] A. Nedić, A. Olshevsky, and C. A. Uribe, "Distributed learning with infinitely many hypotheses," in *Proc. IEEE Conf. on Decision and Control*, Dec 2016, pp. 6321–6326.

[38] A. Nedić, A. Olshevsky, and C. A. Uribe, "Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs," in *Proc. IEEE ACC*, July 2015, pp. 5884–5889.

[39] L. Su and N. H. Vaidya, "Non-Bayesian learning in the presence of Byzantine agents," in *Proc. International Symposium on Distributed Computing*, Paris, France, September 2016, pp. 414–427, Springer.

[40] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "A theory of non-Bayesian social learning," *Econometrica*, vol. 86, no. 2, pp. 445–490, March 2018.

[41] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, 2014.

[42] A. H. Sayed, "Diffusion adaptation over networks," *Academic Press Library in Signal Processing*, vol. 3, R. Chellapa and S. Theodoridis, editors, pp. 323–454, Academic Press, Elsevier, 2014. Also available as arXiv:1205.4220, May 2012.

[43] R. Horn and C. Johnson, *Matrix Analysis*, Cambridge University Press, 2003.

[44] S.U. Pillai, T. Suel, and C. Seunghun, "The Perron-Frobenius theorem: Some of its applications," *IEEE Signal Processing Magazine*, vol. 22, no. 2, pp. 62–75, March 2005.

[45] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2508–2530, June 2006.

[46] S. Sundaram and C. N. Hadjicostis, "Distributed function calculation via linear iterations in the presence of malicious agents – part I: Attacking the network," in *Proc. IEEE ACC*, Washington, USA, June 2008, pp. 1350–1355.

[47] B. Kailkhura, S. Brahma, and P. K. Varshney, "Data falsification attacks on consensus-based detection systems," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 1, pp. 145–158, March 2017.

[48] C. Castellano, S. Fortunato, and V. Loreto, "Statistical physics of social dynamics," *Rev. Mod. Phys.*, vol. 81, pp. 591–646, May 2009.

[49] R. Holley and T. Liggett, "Ergodic theorems for weakly interacting infinite systems and the voter model," *The Annals of Probability*, vol. 3, no. 4, pp. 643–663, 1975.

[50] J. Cox, "Coalescing random walks and voter model consensus times on the torus in $\mathbb{Z}^d$," *The Annals of Probability*, vol. 17, no. 4, pp. 1333–1366, 1989.

[51] P. L. Krapivsky, "Kinetics of monomer-monomer surface catalytic reactions," *Phys. Rev. A*, vol. 45, pp. 1067–1072, Jan 1992.

[52] M. Serrano, K. Klemm, F. Vazquez, V. Eguiluz, and M. San Migue, "Conservation laws for voter-like models on random directed networks," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2009, no. 10, pp. P10024.1–10, Oct 2009.

[53] V. Sood, Tibor Antal, and S. Redner, "Voter models on heterogeneous networks," *Phys. Rev. E*, vol. 77, pp. 041121.1–13, Apr 2008.

[54] P. L. Krapivsky and S. Redner, "Dynamics of majority rule in two-state interacting spin systems," *Phys. Rev. Lett.*, vol. 90, pp. 238701.1–4, Jun 2003.

[55] S. Galam, "Minority opinion spreading in random geometry," *Eur. Phys. J. B*, vol. 25, pp. 403–406, Feb 2002.

[56] S. Galam, "Sociophysics: A review of galam models," *International Journal of Modern Physics C*, vol. 19, no. 03, pp. 409–440, 2008.

[57] M. Lewenstein, A. Nowak, and B. Latané, "Statistical mechanics of social impact," *Phys. Rev. A*, vol. 45, pp. 763–776, Jan 1992.

[58] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing beliefs among interacting agents," *Advances in Complex Systems*, vol. 3, no. 01n04, pp. 87–98, 2000.

[59] G. Weisbuch, G. Deffuant, F. Amblard, and J. Nadal, "Meet, discuss, and segregate!," *Complexity*, vol. 7, no. 3, pp. 55–63, 2002.

[60] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence: Models, analysis and simulation," *Journal of Artificial Societies and Social Simulation*, vol. 5, pp. 1–24, 2002.

[61] J. Holyst, K. Kacperski, and F. Schweitzer, "Phase transitions in social impact models of opinion formation," *Physica A: Statistical Mechanics and its Applications*, vol. 285, no. 1, pp. 199 – 210, 2000.

[62] A. H. Sayed, "Adaptive networks," *Proceedings of the IEEE*, vol. 102, no. 4, pp. 460–497, 2014.

[63] A. Zellner, "Optimal information processing and Bayes's theorem," *The American Statistician*, vol. 42, no. 4, pp. 278–280, 1988.

[64] H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly-connected graphs," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 2, pp. 222–238, June 2017.

[65] H. Salami, B. Ying, and A. H. Sayed, "Diffusion social learning over weakly-connected graphs," in *Proc. IEEE ICASSP*, Shanghai, China, March 2016, pp. 4119–4123.

[66] H. Salami and A. H. Sayed, "Belief control strategies for interactions over weak graphs," in *Proc. IEEE ICASSP*, New Orleans, LA, March 2017, pp. 4232–4236.

[67] B. Ying and A. H. Sayed, "Information exchange and learning dynamics over weakly connected adaptive networks," *IEEE Transactions on Information Theory*, vol. 62, no. 3, pp. 1396–1414, March 2016.

[68] B. Ying and A. H. Sayed, "Learning by weakly-connected adaptive agents," in *Proc. IEEE ICASSP*, Brisbane, Australia, April 2015, pp. 5788–5792.

[69] P. Molavi, A. Jadbabaie, K. R. Rad, and A. Tahbaz-Salehi, "Reaching consensus with increasing information," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 358–369, 2013.

[70] L. G. Epstein, J. Noor, and A. Sandroni, "Non-Bayesian learning," *The B.E. Journal of Theoretical Economics*, vol. 10, no. 1, pp. 1–20, 2010.

[71] L. G. Epstein, "An axiomatic model of non-Bayesian updating," *Rev. Econ. Stud.*, vol. 73, no. 2, pp. 413–436, 2006.

[72] Y. Nyarko, "Bayesian learning leads to correlated equilibria in normal form games," *Economic Theory*, vol. 4, no. 6, pp. 821–841, 1994.

[73] A. Sandroni, "Necessary and sufficient conditions for convergence to nash equilibrium: The almost absolute continuity hypothesis," *Games and Economic Behavior*, vol. 22, no. 1, pp. 121–147, 1998.

[74] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*, SIAM PA, 2000.

[75] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley, NJ, 2006.

[76] W. Rudin, *Principles of Mathematical Analysis*, McGraw-Hill, New York, 1976.

[77] A. J. Laub, *Matrix Analysis For Scientists And Engineers*, SIAM, Philadelphia, PA, USA, 2004.