Multi-agent Topological Models for Evidence Diffusion

MSc Thesis *(Afstudeerscriptie)*

written by

Teodor-Ștefan Zotescu
(born December 18th, 2000 in Bucharest, Romania)

under the supervision of **Prof. Dr. Sonja Smets** and **Dr. Aybüke Özgün**,
and submitted to the Board of Examiners in partial fulfillment of the
requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:	Members of the Thesis Committee:
September 23rd, 2024	Dr. Malvin Gattinger (chair)
	Prof. Dr. Sonja Smets
	Dr. Aybüke Özgün
	Dr. Maria Aloni
	Dr. Luca Incurvati

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis explores the effect of social structure on the beliefs and knowledge of agents who reason in an evidence-based manner. This is done by formally analysing the mechanisms of information flow in the networks formed by the agents. This thesis is a contribution to the formal foundation of the idea that a distinct Social Epistemology is needed to account for real-world doxastic and epistemic phenomena. The first part of the original contribution of this work is bridging the gap between multi-agent Topological Evidence Models (Partitional Models) and Threshold models for diffusion. The resulting framework, which we call Evidence Diffusion Models, allows for the formal analysis of multi-agent evidence-based reasoning in the context of a social network where evidence pieces are being communicated between agents via threshold-limited diffusion. The second part of the original contribution of this work is this formal analysis *per se*. We show that, in our setting, network structures are expressible and known both defeasibly and infallibly by the agents. We prove that defeasible knowledge and (defeasible) distributed knowledge are easily lost under the diffusion of pieces of evidence in a network, whereas distributed evidence is strongly robust. Further, we obtain so-called Cluster Theorems characterising the evidential and network conditions for evidence cascades to form, and for individuals and groups to obtain knowledge in the diffusion process. Finally, we prove results about four special networks — the Total network, the Star network, the Cycle, and the Wheel network — characterising their speed of evidence diffusion and epistemic reliability, and probe the generalisability of the results that were obtained by Zollman about the same networks, but whilst working in a different paradigm.

Acknowledgements

First and foremost I would like to thank my supervisors, Sonja Smets and Aybüke Özgün, for their continued support throughout the conception and writing of this thesis. Without your quick replies, insightful comments and generous guidance, I could not have written this thesis. You are both inspiring academics that I hope to emulate one day.

A huge thank you to all my friends from in and outside of Amsterdam, who have kept my brain nourished with love and encouragement. So as not to exclude anyone, I am thankful first and foremost for all the memories made on Hogevecht street, both near the Nests of the Eagles, and at the Base of that Magnificent Tower. I owe a debt of gratitude to Bijlmer at large. I'm thankful for orange liquor, very very thankful. I am thankful to some barbecue on a balcony, to how the Bijlmer ArenA bus station is a good starting point for runs, to Cineville cards and small cappuccinos. Bucharest cannot be forgotten. I am thankful for the feeling of sitting on a bench after a long walk in IOR, to shawarmas, to tea houses and gifts I feel the need to gift back after. I am thankful for Utrecht, thank you specifically to some house and some wooden harbor there. The following objects I am thankful for are not stand-ins for people. I am thankful for how the water in Nelson Mandela Park reflects the sunset, I am thankful for running shoes and bike chains and 1.25kg plates, I am thankful for how well behaved the on/off switch of my night lamp is, for the birds on my balcony, for all the trees. Know that after I finished my thesis I did Lift my Skinny Fists Like Antennas to Heaven.

For as long as I draw breath, thank you mom, dad, brother, and aunt Dana. Without you, there is nothing.

Contents

1 Introduction **2**

2 Epistemology in a Group Setting **4**
2.1 A Historical Reason for Social Epistemology . . . 4
2.2 A Philosophical Reason for Social Epistemology . . . 5
2.3 A Formal Reason for Social Epistemology . . . 6

3 A Model for Evidential Networks **10**
3.1 Preliminaries . . . 10
3.1.1 Topological Preliminaries . . . 10
3.1.2 Epistemic Interpretation of Topology . . . 12
3.2 Topological Evidence Models . . . 13
3.3 Partitional Models as Multi-Agent Topo-e Models . . . 19
3.4 Threshold Models for Diffusion . . . 24
3.5 Evidence Diffusion Models . . . 28
3.5.1 Semantic Notions . . . 29
3.5.2 Dynamics . . . 31
3.5.3 Group Knowledge Modalities . . . 33

4 Epistemic and Evidential Phenomena on Evidence Diffusion Models **37**
4.1 Networks . . . 37
4.2 Diffusion . . . 42
4.3 Diffusion on Networks in General . . . 48
4.4 Diffusion on the Total, Star, Cycle, and Wheel Networks . . . 53
4.4.1 The Total Network . . . 55
4.4.2 The Star Network . . . 56
4.4.3 The Cycle Network . . . 58
4.4.4 The Wheel Network . . . 61
4.5 Zollman's Findings in This Model . . . 64

5 Limitations, Conclusion, and Further Research **67**

1 Introduction

Real-world reasoners are prone to failures of all kinds — from poor memory access or memory loss, to computational limitations, to the use of fallacious reasoning steps. Even in scientific practice, our highest and most stringent means to discover facts about the world, such deficiencies abound. A kind of failure that has come under more focused media and scientific attention in recent years has to do with the social aspect of science. On the one hand, the history of science is replete with examples where groundbreaking findings were delayed or dismissed due to the influence of established norms and hierarchies, and on the other there are many cases where results were established on shaky foundations that propagated through popularity rather than correctness. Let us briefly provide an illustration of the latter here.

The replication crisis in psychology [20] is a prime example of scientists acting as if scientific rumors, i.e. barely or poorly replicated experiments that become popular, are well-established results. The by now infamous social priming experiments purportedly showed that our behavior is strongly influenced by subtle environmental cues. Before it was shown that many of these experiments failed to replicate, they generated a great amount of subsequent research, and were even sold to the general public as undeniable parts of human psychology. The Nobel-prizewinning economist Daniel Kahneman wrote in his international best-seller *Thinking, Fast and Slow*:

> "When I describe priming studies to audiences, the reaction is often disbelief. [...] Disbelief is not an option. The results are not made up, nor are they statistical flukes. You have no choice but to accept that the major conclusions of these studies are true. More important, you must accept that they are true about *you*". [49]

Clearly, this is a case where the scientific community failed on social grounds, too much trust in one's fellow scientist and not enough cross-verification of results being two of the main factors of the crisis. Ideally, we would like to understand the underpinnings of such failures in order to avoid hurdles in the advancement of knowledge, misinforming the general public and tarnishing the authority of science.

In recent years, the intersection of epistemic logic and social epistemology has emerged as a significant area of inquiry [21, 22, 24, 23, 60, 72, 75, 91]. Also recently, advances in epistemic logic have permitted the formalisation of evidence based reasoning using topological tools [9, 70, 69, 36, 13, 15]. We therefore find ourselves at a pivotal moment in research, where the effects of network structure on the dynamics of evidence, belief, and knowledge can be rigorously modelled and understood. With this thesis, we aim to contribute to this understanding by leveraging the aforementioned topological methods in epistemic logic, as well as the logical modelling of social networks to model the propagation, or diffusion, of evidence in social networks and its effects on individual and group knowledge and belief. By doing this, we aim to shed light on the mechanisms that lead to both the successes and failures of collective reasoning in science, with the ultimate goal of enhancing the reliability and speed of scientific inquiry.

The thesis is structured as follows. In Chapter 2, we lay the Philosophical groundwork for our thesis by presenting three arguments for the need of Social Epistemology. In section 2.1, we analyse the historical case of Ignaz Semmelweis, where the failure of individual rationality leading to group rationality lead to many losses in human lives. In section 2.2 we present the philosophical concept of epistemic dependence as a reason for the need for Social Epistemology. In section 2.3 we present the so-called Independence Thesis — which states that rational individuals might form irrational groups, and rational groups might be formed of irrational individuals — as well as Zollman's formal results [92, 93, 94] pertaining to the trade-offs between the speed and reliability of scientific networks. We also set-up our thesis as a continuation of this line of work.

In Chapter 3, we develop the formal model that represents a great part of the novel contribution of this thesis. We begin by presenting the preliminary knowledge of topology required for the development of this model in Section 3.1, and by detailing the epistemic significance of topology. In sections 3.2, 3.3, and 3.4 we present the parent-frameworks of our model. Respectively, these are Topological evidence models for single-agent evidence-based knowledge and belief, Partitional models, the multi-agent extension of topological evidence models, and Threshold models for diffusion. In Section 3.5 we present the main contribution of this thesis, which we call Evidence Diffusion Models. These models allow for multi-agent evidence-based reasoning in the context of a social network where evidence pieces are being communicated between agents via threshold-limited diffusion.

In Chapter 4, we analyse the effect of social structure on the beliefs and knowledge of agents who reason in an evidence-based manner by focusing on the mechanisms of information flow in these networks. In Section 4.1, we provide graph-theoretic network measures that will be important in our analysis of connectivity's effect on knowledge, as well as highlight four networks which we pay special attention to. We also investigate our language's expressive power with regard to these networks. In Section 4.2, we zoom in on the process of evidence diffusion, and analyse what effects it has on individual or group knowledge gain, or loss. In Section 4.3, we obtained sweeping results regarding the necessary and sufficient conditions for evidence cascades (i.e. when pieces of evidence reach all agents) and individual and group learning. In Section 4.4, we focus on the four special networks and analyse their speed and reliability, whereafter in Section 4.5, we compare these results with Zollman's results on the speed and reliability of the same networks. Finally, we conclude, state some of the limitations of our model and outline further lines of research.

2 Epistemology in a Group Setting

Though it is individuals that hold knowledge, the use and acquisition of knowledge are processes of an unquestionably social nature. First, we rarely ever come to know things on our own. Most often, we learn about the world from our close community, the educational system, books and other media, and so on. Second, in big-scale projects like discovering the cure for a disease, it is often the case that almost every person is mission-critical, as no individual scientist would know how to advance the process of discovery on their own, but the group collectively knows.

Traditionally, epistemology asks questions that put the individual agent in focus: When does an agent possess knowledge, as opposed to belief? [48] What counts as good enough justification for beliefs to accede to knowledge? [78, §3] What are the sources of an agent's knowledge? [78, §5] etc. This epistemic individualism is said to be traceable to Descartes' *Meditations* [67, §1, 38, Preface]. A recent offshoot of this branch of philosophy is *social epistemology*, which adds a social dimension to the more traditional individual epistemology. Questions pertaining to social influence on knowledge have been asked before (for instance see [26, 54]), but it was Alvin Goldmann's book, *Knowledge in a Social World*, that laid the foundations for and kick-started contemporary social epistemology as the discipline that evaluates social practices by their causal contributions to knowledge or error [38, p. 69].

The links between the social and the epistemic have not been probed only by a few isolated questions, but by entire philosohpical frameworks [38, §1.2]. The approach pioneered by Goldman stands in contrast with these by its *veritisic* nature. That is to say that its main preoccupation is with the selection of the social practices which best advance knowledge, where we evaluate the goodness of a practice in terms of truth-related dimensions [38, p. 79], as opposed to, e.g. the assent or dissent produced by the practice [38, §3.1.A] or the utility produced by the practice [38, §3.1.B]. And since truth is central to individual epistemology [48, §1.1], veritistic social epistemology can be said to be the true social development of individual epistemology. We proceed to give three reasons for the study of social epistemology: one historical, one philosophical, and another formal.

2.1 A Historical Reason for Social Epistemology

It might not be immediately apparent why the social aspect of knowledge creation needs to be given explicit attention — can we not easily adapt the well-established accounts of individual rationality [9, 47, 59] by just allowing the agents to communicate? Incidents from the history of science — such as the cases of Ignaz Semmelweis [18] or the Einstein-de Haas effect [34, 85] — point to the fact that the reality of knowledge accrual and proliferation strays quite far from the classical picture of the lone reasoner deductively deriving more and more from a few first principles. It turns out there are several interfering mechanisms in the interaction between individual rationality and group rationality, resulting in differences between what it is best for individuals to do to acquire knowledge, compared to what is best for groups to do to acquire knowledge.

As an illustration, let us see how individual inquiry did not lead to group knowledge in the case of Semmelweis. The 19th Century physician discovered that chlorine handwashing quashed the mortality rate due to childbed fever at any institution where the protocols were implemented [18, p. 53]. The way he arrived at this conclusion was by critical analysis and rigorous hypothesis-testing — a perfectly rational process [18, pp. 45–

53]. However, his ideas did not spread and gain wide acceptance, but were rather shunned, and he was professionally berated [18, pp. 53–61]. The number of lives which could have been saved had chlorine hand-washing been implemented earlier is so great as to be difficult to estimate.

So why did the perfectly good scientific theory of Semmelweis take so long to be accepted, and why did a community of scientists block a perfectly rational answer? The answer lies in the social characteristics of the scientific community. Its authority structures, its resistance to new ideas, the general attitude towards evidence — all factors in the social realm — came in the way of rational inquiry [18, pp. 55–61]. If we are to prevent such situations from happening in current scientific practice, we ought to have a good grasp of exactly what goes wrong at the border between individuals and groups when it comes to knowledge.

2.2 A Philosophical Reason for Social Epistemology

The phenomenon of *epistemic dependence* offers another clear motivation for the study of the dynamics of knowledge and belief in social groups. In traditional epistemology, a reason R to believe a proposition P has as a necessary condition some notion of support between the content of R and P. But beliefs can be justified by more than content-related reasons. For instance, in order to justify one's belief in P one might appeal to an expert's competence and point to the fact that they believe P. As an illustration, consider Sam, an amateur bodybuilder, who is discussing muscle hypertrophy with Dr. Mike Israetel, a professor of exercise science: clearly it would be rational for Sam to allow his beliefs about building muscle to be steered by Dr. Mike's judgement, as his experience and knowledge of this domain assures a better chance at approximating the truth. For a discussion of such principles, often called deference to experts principles, in the setting of Bayesian epistemology, see [79, §5.2, 30]. If he did that, Sam would be manifesting epistemic dependence with respect to Dr. Mike.

This picture suggests the following intuitive definition for epistemic dependence.

> A agent i manifests social epistemic dependence just in case i's epistemic standing causally depends on other persons, or on other features of the broader social environment [42, p. 114].

Stated in terms of reasons, the previous discussion boils down to the following. It is completely tenable epistemically that an agent has the following reason R' to believe P: 'agent i is an expert in P's subject matter and i believes P'. One can find even stronger reasons of the same kind. For instance, take the reason R'' to believe P: 'agents $\{i_1, ..., i_n\}$, constituting an overwhelming majority of the community of experts for P's subject matter, believe P'. Notice that neither R', nor R'' offer any kind of direct support for the truth of P. At most, we might say that they offer indirect support to the truth of P, via the assumption that if i is an expert, and i believes P, then i themselves must have a (direct) reason for believing P.

Note that these direct reasons might not be easily communicable to the other agent. Scientific communication is, after all, a difficult process. The schooling, scientific techniques and intuitions based on years that the expert has and that lead them believe P, are not as easily transferable to the agent as a fact like "the chance it will rain in Amsterdam tomorrow is 80%". Thus, even in the presence of all the facts, it might still be rational for the agent to defer to the expert's analysis. All of this is to say that, quite contrary to the picture sketched by Cartesian epistemology, "rationality sometimes consists in refusing to think for one-self" [45, p. 336].

One need not even look for cases as extreme as complete deference to another agent to observe the manifestation of epistemic dependence. Let us say that an agent has a reason R to believe P, which is based upon a set of pieces of evidence $\{e_1, ..., e_n\}$, which could be experiments, forensic evidence, direct observations, or other forms of evidence. Let us further say that R is an argument constructed by the agent themselves on the basis of this evidence. If the evidence has not been collected entirely by the agent, can we say that the agent's epistemic standing is not causally dependent on other persons? It may well be that one piece of evidence in the set which is critical for the argument is a faulty result from a lab that did not meet cleanliness standards. Could we possibly blame the agent themselves, and ascribe them a fault of reasoning? We think not.

This suggests a fault in the above definition. Indeed, the definition is considered somewhat deficient in the social epistemology literature, for it cannot distinguish between individualist and anti-individualist stances on epistemology [42, p. 116]. The definition over-attributes dependence, since any person is, in practice, epistemically dependent upon others in this sense. In other words, using this definition, we cannot make a difference between epistemologies that are *truly* anti-individualist, i.e. that claim that there are parts of epistemology that are irreducible to individual epistemology together with communication between agents, and individualist epistemologies, that claim that this reduction is possible. [42] contrasts "mere reliance" with "vulnerable dependence", the former involving knowledge of the dependability of those depended on, and the latter the absence of this.

But we enter here into rather thorny territory, because it seems like requiring *knowledge* of the dependability of those depended on makes almost all cases of dependence vulnerable dependence. Even if we relaxed the requirement for knowledge, the distinction between mere reliance and vulnerable dependence is not so easy to parse. After all, if I commit myself to a doxastic or epistemic attitude on the basis of an argument that uses a piece of evidence received from another agent, I am making myself vulnerable to failure if it turns out that this agent's method of evidence collection was faulty. As such, in this thesis, we will consider most cases of dependence to be vulnerable dependence. This is partly because the framework we present in Section 3 ought to act as a toy model for scientific networks, and often whist working in the context of a scientific network, one places complete trust that other scientists' results that you base your work upon is well done. As a matter of fact, our model will reflect this in a big way, since agents will be at the risk of *losing knowledge* on the basis of the social connections they choose to form. But more on this later.

2.3 A Formal Reason for Social Epistemology

Besides the historical cases and epistemic dependence, a formal result highlights the need for social epistemology. The independence thesis states that "rational individuals can form irrational groups; and rational groups might be composed of irrational individuals" [62]. The result is deeply unintuitive, given that it seems to point out that we call agents rational when there is a context in which they do not act rationally. Furthermore, this misalignment carries over to knowledge creation, where the best course to the creation of individual knowledge might not necessarily be the best course to the creation of group knowledge. This is precisely what makes the independence thesis a strong argument for the existence of independent standards of rationality for individuals and groups.

It seems like drawing a distinction between two levels of analysis teases out the place where this divergence happens. First there is the 'micro' level — an agent's individual reaction to and management of new facts or data coming in. Second there is the 'macro' level, studying, after some assumptions about the epistemic properties of the individuals within, what "community structures best serve the epistemic aim of that community" [92, p. 574]. Based on this divergence, in a series of very well-appreciated papers [92, 93, 94], Zollman sets out to provide a model of scientific networks using earlier work by Bala and Goyal [4, 3]. His purpose in doing this is to tease out general principles that may underlie the functioning of (and therefore aid in optimizing) actual scientific networks. In summary, his results state the following:

> First, in some contexts, a community of scientists is, as a whole, more reliable when its members are less aware of their colleagues' experimental results. Second, there is a robust trade-off between the reliability of a community and the speed with which it reaches a correct conclusion. [92, p. 574]

The first part is motivated by the fact that a less connected network experiences less change in strategy due to random strings of misleading evidence remaining more local, and not spreading to the entire network [92, pp. 582–583]. The reason for the second part is simply that there is more communication in the more connected network.

To start with, Zollman studies three special networks: the total network (which he calls "the complete graph"), the cycle, and the wheel. These are illustrated in figure 2.3, for six agents.

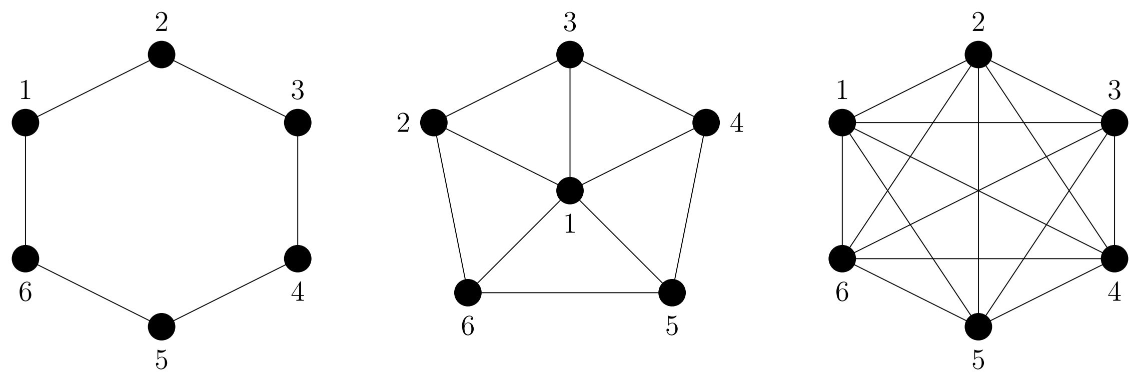

A cycle *A wheel for 1* *The total network*

Figure 1: The cycle, the wheel, and the total network for six agents

In this model, presented by [92, 93, 94] and originally formulated in [4], there are two states of the world, φ_1 and φ_2, and two actions available to each agent, A_1 and A_2. The expected utility of A_1 is the same irrespective of the world we are in, whereas the expected utility of A_2 is lower in φ_1 and higher in φ_2. The expected utilities are common knowledge among the agents, but which world is the actual one is unknown to any agent. However, they do form beliefs regarding this matter, and at each round, they take the action with the highest expected payoff. The players observe the outcomes of their own action and of the actions of their neighbors.

Running this game 10,000 times on a computer obtains Zollman the following results. For reliability, meaning the probability of learning the actual world, we have the plot in Figure 2.3. The x axis is the number of agents in the network, and the y axis represents the percentage out of the 10,000 runs where the agents successfully learned the actual world.

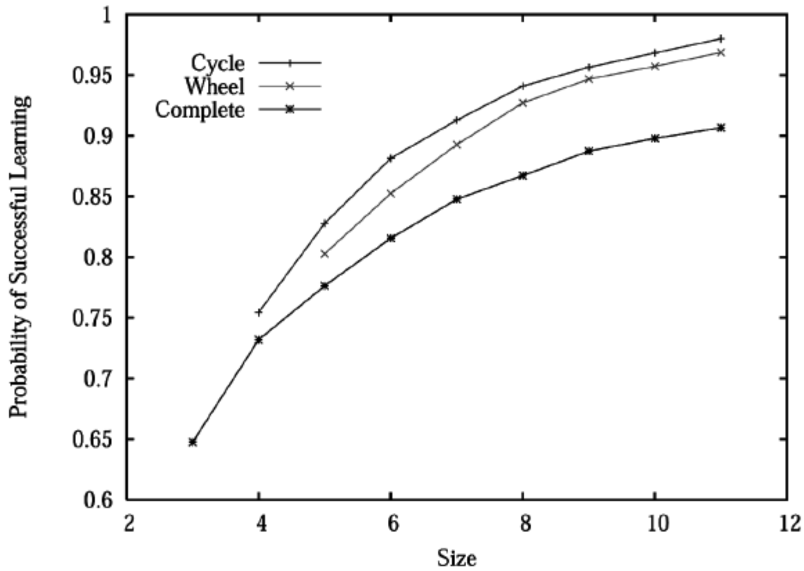

Figure 2: Reliability results from Zollman's simulations. Plot from [92, p. 580]

The original formulation is clearly observed on this plot: the sparser the network, the more reliable it is. As for speed, the plot in Figure 3 shows that the more densely connected a network, the faster it reaches the desired result.

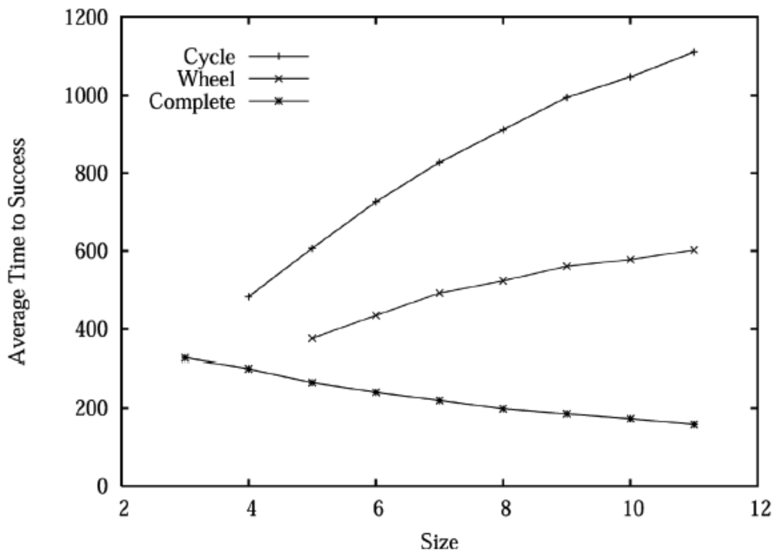

Figure 3: Speed results from Zollman's simulations. Plot from [92, p. 581]

Zollman then finds that the trade-off between speed and reliability generalizes by performing an analysis of all the networks between size 3 and 6 [92, p. 581]. Thus, it might seem that this is a universal principle that governs social networks of rational agents. However, the models of Bala and Goyal make critical assumptions about the agents' cognitive and computational limitations [4, 3] and rely on computer simulations and statistical methods to derive results [92, 93, 94].

In order to test the broader applicability of these results, it is of paramount importance to develop a parallel framework in which such questions can get their answer. This thesis aims to undertake precisely this task, and does so by abstracting away from the setting that Zollman uses for his quantitative analysis. We instead focus on the direct mechanisms of information flow in social-epistemic contexts in which the kind of reasoning being done is evidence-based reasoning. This is done by providing a logical framework for the treatment of the evidence-based individual and group knowledge and belief for agents that are placed within the context of a social network.

3 A Model for Evidential Networks

We argue that logic is the appropriate setting for the treatment of the issues outlined in the previous framework. At the end of this section, it will also become clear that the developed framework fits naturally within the logical literature on social epistemology. Social networks have long been formally studied with the tools of graph theory [29, Ch. 2] and, likewise, the notion of knowledge, together with its dynamics, have received ample formal treatment in the discipline of modal logic [47, 28, 15]. It is only recently that frameworks integrating the two have begun to emerge in order to study problems with regard to the logic of social network creation [77], the epistemic dimension of networks [12], informational cascades [11], rational groups [86] etc. A similarly recent advancement in epistemic logic is the use of topology [51], and more specifically topological evidence models [70, 69, 9], which allow for a much more nuanced analysis of knowledge, encoding a variety of notions as evidence, arguments and justifications [69, Ch. 5]. In what follows, we put these ingredients together and develop multi-agent topological models for the diffusion of pieces of evidence within a social network.

First, we introduce some preliminaries in the form of topological definitions and their epistemic interpretations. Second, we introduce the setting of single-agent topological evidence models. Third, we present partitional models, which we defend as the natural multi-agent extension to topological evidence models. Fourth, we introduce threshold models for diffusion and their logical setting. Finally, we put these ingredients together and present our main contribution: evidence diffusion models.

3.1 Preliminaries

3.1.1 Topological Preliminaries

We start by introducing the basic topological notions that will be used throughout the formal part of the thesis. For a more in-depth introduction to topology, we refer the reader to any standard textbook of the topic (such as [65, 31]). Throughout this section, we refer to a fixed set X as the domain of a topology.

Definition 1 (Topological Subbasis and Basis). *A family of sets $\mathcal{X} \subseteq \wp(X)$ is a subbasis if and only if for every $x \in X$ there exists a set $U \in \mathcal{X}$ such that $x \in U$. Equivalently, $\bigcup \mathcal{X} = X$. We call $\mathcal{X}$ a basis if and only if, in addition to being a subbasis, $\mathcal{X}$ satisfies: for every $U, U' \in \mathcal{X}$ and every $x \in U \cap U'$ there exists U'' such that $U'' \subseteq U \cap U'$ and $x \in U''$.*

It is easy to see that the closure of any subbasis under finite intersections is a basis. Topological bases are often introduced before topologies, because they *underlie* topologies and their properties are often easier to prove. That is helpful because often properties of the bases transfer neatly to topologies themselves. This is because there is a notion of a basis *generating* a topology, which we introduce now.

Definition 2 (Topology Generated by Basis). *Let $\mathcal{X}$ be a basis. Then the family of sets τ obtained by requiring that if $\mathcal{U} \subseteq \mathcal{X}$, then $\bigcup \mathcal{U} \in \tau$ is called the topology generated by $\mathcal{X}$.*

The relationship between subbases, bases and topologies is neatly captured in the diagram below ([5, Lecture 1.1, Slide 32]):

$$\text{Subbasis} \xrightarrow[\text{finite intersections}]{} \text{Basis} \xrightarrow[\text{arbitrary unions}]{} \text{Topology}$$

Figure 4: Relationship between subbases, bases and topologies

This gives a hint of the general definition of a topology: it is a basis that is closed under *arbitrary* unions. We can now formulate a definition of a topology starting from the original set X itself.

Definition 3 (Topology)**.** *A family $\tau \subseteq \wp(X)$ is called a topology on X if and only if*

(i)	$X \in \tau$	*— it contains the whole set;*
(ii)	*if* $\mathcal{U} \subseteq \tau$ *with* $\mathcal{U}$ *finite, then* $\bigcap \mathcal{U} \in \tau$	*— it is closed under finite intersections;*
(iii)	*if* $\mathcal{U} \subseteq \tau$*, then* $\bigcup \mathcal{U} \in \tau$	*— it is closed under arbitrary unions.*

The pair $\langle X, \tau \rangle$ is called a *topological space*. We call elements $x \in X$ *points* and elements $U \in \tau$ *open sets* or *opens*. The complements of the open sets are called *closed sets*. Note that openness and closedness are not mutually exclusive properties: for example the set X is both open and closed, since its complement, $\emptyset$ is open by virtue of it being the vacuous union. The sets that are both open and closed are called *clopen* sets.

Definition 4 (Open Neighborhood)**.** *Let τ be a topology on X. For an $x \in X$ and a $U \in \tau$, we call U an open neighborhood of x if and only if $x \in U$.*

With this definition in hand, we go on to define two notions that will be important for topological modal logic: the interior and closure of a set.

Definition 5 (Interior, Closure)**.** *Let $\langle X, \tau \rangle$ be a topological space, $A \subseteq X$ and $x \in X$. Then x is an interior point of A if and only if there exists an open neighborhood of x, U, such that $U \subseteq A$. Then we define the interior of A as $Int(A) = \{x \mid x \text{ is an interior point of } A\}$. Equivalently, $Int(A) = \bigcup\{U \subseteq A \mid U \in \tau\}$. Thus, the interior of A is the greatest open contained in A. The closure of A is $Cl(A) = \bigcap\{C \supseteq A \mid C = X \backslash U \text{ for some } U \in \tau\}$. Thus, the closure of A is the smallest closed set containing A.*

Much like the diamond and box operators of modal logic, these notions are dual to each other. That is, $\mathrm{Cl}(A) = X \setminus \mathrm{Int}(X \setminus A)$. It is also not difficult to see from the definitions above that Int and Cl can be defined as operators on $\langle X, \tau \rangle$. These operators satisfy the following axioms, which are known as the Kuratowski axioms [31, pp. 14–15]. For any $A, B \subseteq X$,

(I1) $\mathrm{Int}(X) = X$	(C1) $\mathrm{Cl}(\emptyset) = \emptyset$
(I2) $\mathrm{Int}(A) \subseteq A$	(C2) $A \subseteq \mathrm{Cl}(A)$
(I3) $\mathrm{Int}(A \cap B) = \mathrm{Int}(A) \cap \mathrm{Int}(B)$	(C3) $\mathrm{Cl}(A \cup B) = \mathrm{Cl}(A) \cup \mathrm{Cl}(B)$
(I4) $\mathrm{Int}(\mathrm{Int}(A)) = \mathrm{Int}(A)$	(C4) $\mathrm{Cl}(\mathrm{Cl}(A)) = \mathrm{Cl}(A)$

Finally, we define the following notions:

Definition 6 (Dense Sets, Nowhere Dense Sets)**.** *$A \subseteq X$ is dense in X if and only if $Cl(A) = X$, or equivalently if for all $U \in \tau \setminus \{\emptyset\}$, $A \cap U \neq \emptyset$. It is called nowhere dense if and only if $Int(Cl(A)) = \emptyset$. For any $B \subseteq X$, A is dense in B if and only if $B \subseteq Cl(A \cap B)$.*

For our setting, it will be important that we have a way to 'move' from a topological space to a subpart of that space. It is important to do this in a way that preserves the topological structure we are interested in — this is where the notion of subspace topology comes in.

Definition 7 (Subspace Topology)**.** *For the topological space $\langle X, \tau \rangle$ and for $A \subseteq X$, we define $\tau\big|_A = \{A \cap U \mid U \in \tau\}$ and say that $\langle A, \tau\big|_A \rangle$ is a subspace of the original space.*

In the same vein, we will need an operation of putting multiple topologies together in a structure-preserving way. This will be achieved by using the so-called join topology, which we define now.

Definition 8 (Join Topology)**.** *Let $Ag = \{1, ..., n\}$ and for each $i \in Ag$, let τ_i be a topology on X. Then the join topology, denoted by $\bigvee_{i \in Ag} \tau_i$ is the smallest topology containing τ_i for each $i \in Ag$. Formally, $\bigcup_{i \in Ag} \tau_i$ constitutes a subbasis and $\bigvee_{i \in Ag} \tau_i$ is the topology generated by this subbasis.*

3.1.2 Epistemic Interpretation of Topology

Epistemic logic, for which there is a long-standing tradition, pioneered by [47], allows for an axiomatic analysis of knowledge and belief, taken as primitive notions. In practice, however there are other pieces of information that help us derive or arrive at these beliefs or at knowledge. This granularity should be taken into consideration in an analysis of epistemic and doxastic dynamics and so we make the assumption, together with others in the literature [9, 84, 83, 82], that rational beliefs and knowledge are based on the primitive notion of *evidence*. We define evidence broadly, as the informational resources that an agent uses in their process of reasoning; or in other words that which founds their beliefs [52]. This definition is, in fact, very broad — it includes false or misleading pieces of information that *seem* (to the agent) to be true, priors, biases and *a priori* knowledge [9, p. 2].

In order to capture these notions formally, topology is a natural setting, for "topological spaces emerge naturally as information structures that can provide a deeper insight into the evidence-based justification of knowledge and belief" [9, p. 2]. In fact, topological spaces have a broad spectrum of applications in formal epistemology. Tracing back to intuitionistic languages, where open sets were interpreted as pieces of evidence or observable properties [80], to topological evidence models [9, 6, 69, 10, 7], and formal learning theory [13, 14, 51, 50, 36, 37].

Having $\langle X, \tau \rangle$ as a topological space, we think of the points $x \in X$ as possible worlds. The parallel with Kripke models [17] becomes immediately apparent: we replace the (epistemic or doxastic) accessibility relation with a topology. But if the accessibility relation in Kripke models meant epistemic indistinguishability, or doxastic priority, what are the open sets of the topology to mean? Following [9, 69], we think of the sets in the subbasis $\mathcal{S}$ as *directly observable* pieces of evidence, of those in the basis $\mathcal{B}$ as pieces of *combined* evidence (as they are obtained by closing $\mathcal{S}$ under finite intersection, which is like conjunction) and of the sets in the topology τ simply as evidence (as they are obtained by closing $\mathcal{B}$ under arbitrary unions, a stronger version of disjunction). This will allow us to define the evidence based notions we are after.

3.2 Topological Evidence Models

We shall now present the setting that underlies our contribution, in its single-agent variant. This framework provides a topological semantics for the epistemic notions of evidence, argument, justification, belief and knowledge, and allows for the exploration of their interplay [9]. It has been explored in [9, 69, 8, 70], and this section more closely follows [69, Ch.5][1]. The setting is that of topological evidence models (or for short, topo-e models). First, we give the language we will be using and then the definition for topo-e models proper.

Definition 9. *The language $\mathcal{L}_s$ of single-agent logic for evidence, belief, and knowledge is defined via the recursive rule below. We take $p \in \mathbb{P}$, where $\mathbb{P}$ is a countable set of propositional atoms.*

$$p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid E\varphi \mid \widetilde{E}\varphi \mid \Box\varphi \mid \widetilde{\Box}\varphi \mid [\forall]\varphi \mid K\varphi \mid B\varphi$$

The modality $E\varphi$ is read as 'the agent has a basic piece of evidence for φ' and the modality $\Box\varphi$ is read as 'the agent has a factive basic piece of evidence for φ'. The modalities $\widetilde{E}$ and $\widetilde{\Box}$ are the same, except they make reference to combined evidence, and not to basic evidence. The modality $[\forall]\varphi$ is read as 'the agent infallibly knows φ'. K and B are the knowledge and belief operators, and they are read as 'the agent defeasibly knows (respectively, believes) φ'. We follow the standard conventions in omitting brackets when there is no ambiguity.

Definition 10 (Single-agent Topo-e Model)**.** *A topo-e model is a triple $\mathfrak{M} = \langle X, \mathcal{E}, [\![\cdot]\!]\rangle$, where X is a set of possible worlds, taken to be nonempty, $\mathcal{E} \subseteq \wp(X)$ is a set of basic pieces of evidence, satisfying $X \in \mathcal{E}$ and $\emptyset \notin \mathcal{E}$, and $[\![\cdot]\!] : \mathbb{P} \mapsto \wp(X)$ is a valuation function.*

Note that our definition is actually that of uniform evidence models, after [84], since [9, 69] include the topology of evidence as part of the model. Yet, since van Benthem and Pacuit [84] do not explicitly make use of the topological aspect, and [9, 69], together with us which follow them, do, and furthermore the topology of evidence can be generated from $\mathcal{E}$, we take the liberty to call these models topo-e-models.

Pieces of evidence are represented here as subsets of the set of possible worlds. The set $\mathcal{E}$ of basic pieces of evidence comes with some constraints, however. Namely, it must contain the set of all possible worlds, X, and not contain the empty set. These two formal requirements correspond to the following conceptual requirements: tautologies are always evidence and, respectively, contradictions are never evidence [84].

Now, as stated before, we are using a broad and inclusive notion of evidence: evidence may be false (in the sense that a set in $\mathcal{E}$ might not contain the actual world), or inconsistent (in the sense that intersections of the subsets of $\mathcal{E}$ may be empty). Despite this, we want our agent to be able to put these pieces of evidence together finitely and consistently, so that, ultimately, they use informational resources that are in some sense 'good' to arrive at their epistemic or doxastic attitudes. This is where the notion of *body of evidence* comes in. In what follows, for any two sets A and B, we use $A \subseteq_{\text{fin}} B$ to mean that $A \subseteq B$ and A is finite.

[1]Note, however, that our thesis makes use of different notation. In [69], notions that have to do with basic evidence are subscripted with 0 ($\mathcal{E}_0, E_0, \Box_0$), whereas we use no subscript. This is because we will later index these sets with agents. For combined evidence [69] uses $\mathcal{E}, E, \Box$, but we use $\widetilde{\mathcal{E}}, \widetilde{E}, \widetilde{\Box}$.

Definition 11 (Body of Evidence, Combined Evidence, Topology of Evidence)**.** *Given a basic evidence set $\mathcal{E}$, a subset $F \subseteq \mathcal{E}$ is a body of evidence if and only if $F \neq \emptyset$ and for all $F' \subseteq_{fin} F$, if $F' \neq \emptyset$, then $\bigcap F' \neq \emptyset$. When F is finite we call it a finite body of evidence. Let $\mathcal{F} = \{F \mid F \text{ is a body of evidence}\}$ and let $\mathcal{F}^{fin} = \{F \mid F \text{ is a finite body of evidence}\}$. We define $\widetilde{\mathcal{E}} = \{\bigcap F \mid F \in \mathcal{F}^{fin}\}$ to be the family of combined evidence. Naturally, we call any $e \in \widetilde{\mathcal{E}}$ a piece of combined evidence. Note that $\widetilde{\mathcal{E}}$ is a topological basis, and the topology of evidence τ is generated from $\widetilde{\mathcal{E}}$ by closing it under arbitrary unions.*

Before we proceed it would be worth it to introduce a few of the semantic notions we will be working with. Given a topo-e model, $\mathfrak{M} = \langle X, \mathcal{E}, [\![\cdot]\!]\rangle$, and a world $x \in X$, we define satisfaction for the propositional variables and Boolean operators in the usual way:

- $\mathfrak{M}, x \vDash p$ if and only if $x \in [\![p]\!]$;
- $\mathfrak{M}, x \vDash \neg\varphi$ if and only if it is not the case that $\mathfrak{M}, x \vDash \varphi$;
- $\mathfrak{M}, x \vDash \varphi \wedge \psi$ if and only if $\mathfrak{M}, x \vDash \varphi$ and $\mathfrak{M}, x \vDash \psi$.

The other Boolean operators are taken to be short-hand notations for the following: for any $\varphi, \psi \in \mathcal{L}_s$, $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, $\varphi \to \psi := \neg(\varphi \wedge \neg\psi)$ and $\varphi \leftrightarrow \psi := (\varphi \to \psi) \wedge (\psi \to \varphi)$. We also introduce the following short-hands for the possibility operators corresponding to our modalities: $\Diamond\varphi := \neg\Box\neg\varphi$, $\widetilde{\Diamond}\varphi := \neg\widetilde{\Box}\neg\varphi$, $\langle K\rangle\varphi := \neg K\neg\varphi$, $\langle B\rangle\varphi := \neg B\neg\varphi$. When it is not the case that $\mathfrak{M}, x \vDash \varphi$, we will write $\mathfrak{M}, x \nvDash \varphi$.

Now, we know by definition that for any propositional letter $p \in \mathbb{P}$, $[\![p]\!] \subseteq X$. This is the set of worlds at which p is true, which we will call the truth set of p, or the propositional atom's corresponding proposition. The function $[\![\cdot]\!]$ can be extended so that its domain is the entirety of $\mathcal{L}_s$, and not just $\mathbb{P}$. For the Boolean operators, we do this in the standard way, and we leave the discussion of truth-sets for the other operators for later in this section.

Propositions are convenient to work with in this topological setting, because many of the logical operations we employ correspond very naturally to set-theoretical operations that are relevant for topology. In this vein, it is easy to observe that for any $\varphi, \psi \in \mathcal{L}_s$, $[\![\neg\varphi]\!] = X \setminus [\![\varphi]\!]$, $[\![\varphi \wedge \psi]\!] = [\![\varphi]\!] \cap [\![\psi]\!]$, $[\![\varphi \vee \psi]\!] = [\![\varphi]\!] \cup [\![\psi]\!]$ etc. Oftentimes, when we work with an arbitrary proposition, i.e. an arbitrary subset of the set of worlds X, we will, for simplicity's sake, use the uppercase letter P instead of using the truth set notation $[\![p]\!]$. We move on with the presentation of the evidential notions in the framework.

Definition 12 (Evidential support)**.** *A body of evidence F supports a proposition $P \subseteq X$ if and only if $\bigcap F \subseteq P$.*

In words, the definition above amounts to a body of evidence supporting P if P is true in every state satisfying all the evidence in F [69, p. 53]. Said otherwise, if we conjoin the evidence in body F, and obtain that the world is one of several states (the states consistent with the intersection $\bigcap F$), and in all of those states P is true, then our body of evidence supports P. The definition of a basic or combined piece of evidence supporting a proposition P is the obvious adaptation of the one above.

Given this definition, and as in the real world, it is entirely possible that there are multiple evidence sets supporting P. A body of evidence that supports P is merely one evidential path to the proposition P. Oftentimes, more than one such evidential path is required to support a proposition in a satisfactory way. We will now define an object that brings together a multiplicity of evidential paths to some proposition P: an argument for P.

Definition 13 (Argument for P)**.** *An argument A for a proposition P is defined as $A = \bigcup \widetilde{\mathcal{E}}'$, where $\widetilde{\mathcal{E}}' \subseteq \widetilde{\mathcal{E}}$ and $\widetilde{\mathcal{E}}' \neq \emptyset$ and for all $e \in \widetilde{\mathcal{E}}'$, $e \subseteq P$ or equivalently $A \subseteq P$. We call A factive at $x \in X$ if and only if $x \in A$.*

Very conveniently, from a topological point of view, an argument for a proposition P corresponds to an open set contained in P — a set $U \in \tau \setminus \{\emptyset\}$ such that $U \subseteq P$. Thus, $\mathrm{Int}(P)$ is the weakest (by virtue of restricting the possible worlds the least) argument for P.

Definition 14 (Justification for P)**.** *A justification J for P is an argument such that for all $e \in \widetilde{\mathcal{E}}$, $J \cap e \neq \emptyset$. We call J factive at $x \in X$ if and only if $x \in J$.*

A justification is a stronger notion than that of argument. In fact, it is an argument that is, in some sense, undefeated, or consistent with all the evidence. Topologically, the fact that it has to have a non-empty intersection with all available evidence amounts to it being dense: a justification for P is a dense open subset of P. Now that we have defined these notions, we move on to give semantics to all the modal operators in $\mathcal{L}_s$. Given a topo-e-model $\mathfrak{M} = \langle X, \mathcal{E}, [\![\cdot]\!]\rangle$ and a world $x \in X$,

- $\mathfrak{M}, x \vDash [\forall]\varphi$ if and only if $[\![\varphi]\!] = X$;
- $\mathfrak{M}, x \vDash E\varphi$ if and only if there exists $e \in \mathcal{E}$ such that $e \subseteq [\![\varphi]\!]$;
- $\mathfrak{M}, x \vDash \Box\varphi$ if and only if there exists $e \in \mathcal{E}$ such that $e \subseteq [\![\varphi]\!]$ and $x \in e$;
- $\mathfrak{M}, x \vDash \widetilde{E}\varphi$ if and only if there exists $e \in \widetilde{\mathcal{E}}$ such that $e \subseteq [\![\varphi]\!]$;
- $\mathfrak{M}, x \vDash \widetilde{\Box}\varphi$ if and only if there exists $e \in \widetilde{\mathcal{E}}$ such that $e \subseteq [\![\varphi]\!]$ and $x \in e$.

The first is the global modality, $[\forall]$, which encodes the strongest notion of knowledge in this framework. When $[\forall]\varphi$ holds, φ is absolutely certain information for the agent: the alternative is not even considered possible. It holds when all the worlds in X satisfy φ. Importantly, when it holds, it is true at all possible worlds. This will be a feature of E and $\widetilde{E}$ as well, since their clauses make no reference to the world of evaluation itself. The truth set of $[\![[\forall]\varphi]\!]$ is then defined in the following way:

$$[\![[\forall]\varphi]\!] = \begin{cases} X & \text{if } [\![\varphi]\!] = X; \\ \emptyset & \text{otherwise.} \end{cases}$$

The next modality, $E\varphi$, amounts to 'having a basic piece of evidence for φ' and $\widetilde{E}\varphi$ to 'having a combined piece of evidence for φ'. $E\varphi$ holds when there is a basic piece of evidence supporting φ, and $\widetilde{E}\varphi$ when there is a combined piece of evidence supporting φ. Their definitions are straight-forward, and so too are their truth sets.

$$[\![E\varphi]\!] = \begin{cases} X & \text{if there exists } e \in \mathcal{E} \\ & \text{such that } e \subseteq [\![\varphi]\!]; \\ \emptyset & \text{otherwise.} \end{cases} \qquad [\![\widetilde{E}\varphi]\!] = \begin{cases} X & \text{if there exists } e \in \widetilde{\mathcal{E}} \\ & \text{such that } e \subseteq [\![\varphi]\!]; \\ \emptyset & \text{otherwise.} \end{cases}$$

The corresponding box modalities are $\Box$ and $\widetilde{\Box}$, which can be read as 'there is a factive basic (respectively combined) piece of evidence for φ'. They are defined in much the same way, except they make reference to the world of evaluation, which is precisely what encodes the notion of factivity we are interested in. This is a local notion of factivity — instead of there being a designated unique 'actual world', we take a piece of evidence to be factive *from the point of view of the world of evaluation*, if the piece of evidence contains that world. In other words, 'actual' is, as far as our definitions are concerned, indexical with respect to the world of evaluation. The truth sets of formulas using these modalities are defined as follows:

$$[\![\Box\varphi]\!] = \{x \in X \mid \text{there exists } e \in \mathcal{E} \text{ such that } e \subseteq [\![\varphi]\!] \text{ and } x \in e\}$$
$$[\![\widetilde{\Box}\varphi]\!] = \{x \in X \mid \text{there exists } e \in \widetilde{\mathcal{E}} \text{ such that } e \subseteq [\![\varphi]\!] \text{ and } x \in e\}$$

With these notions in hand, we can define evidence-based belief and defeasible knowledge, in the sense of the defeasibility theory of knowledge developed in [58, 57, 55, 56], in a natural way. First, we give simple satisfaction clauses for the B and K operators. Then, we discuss in more detail how these clauses correspond to our prior definitions and to other topological notions. For a topo-e model $\mathfrak{M} = \langle X, \mathcal{E}, [\![\cdot]\!]\rangle$,

- $\mathfrak{M}, x \vDash B\varphi$ if and only if there exists a justification J for φ;
- $\mathfrak{M}, x \vDash K\varphi$ if and only if there exists a factive justification J for φ.

So, we cash out our earlier definitions of evidential concepts by defining notions of belief and knowledge that are, as we desired, evidentially justified. As before, we provide the truth sets of $B\varphi$ and $K\varphi$.

$$[\![B\varphi]\!] = \begin{cases} X & \text{if there exists a justification } J \text{ for } \varphi; \\ \emptyset & \text{otherwise.} \end{cases}$$
$$[\![K\varphi]\!] = \{x \in X \mid \text{there exists a justification } J \text{ for } \varphi \text{ and } x \in J.\}$$

Not only are K and B defined in terms of the evidential concepts, but it turns out that these notions are also purely *topological* in nature [69, §5.3.2. for discussion], unlike in prior accounts [84, 83, 82]. They are topological in the sense that, as the next proposition shows, our doxastic (and later, epistemic) attitudes correspond exactly to topological concepts. The following proposition, due to [69, p. 62, 9, Proposition 2], showcases the various equivalent definitions of belief that we obtain through the notions defined earlier.

Proposition 1 ([69, Proposition 5.3.4])**.** *The following are equivalent:*

1. *$B\varphi$ holds at any state, or equivalently, there is a justification for φ (there exists $U \in \tau$ such that $U \subseteq [\![\varphi]\!]$ and for all $e \in \widetilde{\mathcal{E}}$, $U \cap e \neq \emptyset$);*
2. *φ is supported by all "sufficiently strong" evidence (for all $F \in \mathcal{F}^{fin}$ there exists $F' \in \mathcal{F}^{fin}$ such that $F \subseteq F'$ and $\bigcap F' \subseteq [\![\varphi]\!]$);*
3. *Every piece of evidence can be strengthened to some evidence supporting φ (for all $e \in \widetilde{\mathcal{E}}$ there exists $e' \in \widetilde{\mathcal{E}}$ such that $e' \subseteq e \cap [\![\varphi]\!]$);*
4. *Every argument can be strengthened to an argument for φ (for all $U \in \tau \setminus \{\emptyset\}$ there exists $U' \in \tau \setminus \{\emptyset\}$ such that $U' \subseteq U \cap [\![\varphi]\!]$);*

5. *$[\![\varphi]\!]$ includes some dense open set*
   (there exists $U \in \tau$ such that $U \subseteq [\![\varphi]\!]$ and $Cl(U) = X$);

6. *$Int([\![\varphi]\!])$ is dense in τ*
   ($Cl(Int([\![\varphi]\!])) = X$);

7. *$[\forall]\widetilde{\Diamond}\widetilde{\Box}\varphi$ holds at any state.*

Point (1) is the basic definition for belief, and it makes use of the concept of justification. The fact that justifications are open encode the fact that they are based on the available evidence, and the density condition encodes the fact that beliefs must be coherent with all the available evidence — these properties appear explicitly in (5). Thus, what we defined is a notion of *evidentially justified belief*. Justifications were in turn defined using the more primitive notions of bodies of evidence, combined evidence and arguments. (2), (3) and (4) show that there are equivalent definitions of belief that make use only of these primitives. (6) is interesting from a topological point of view because it makes use of the customary closure and interior operators. (7) shows that B is definable in terms of $[\forall]$ and $\widetilde{\Box}$. A similar proposition obtains for knowledge as well [69, p. 69, 9, Proposition 4].

Proposition 2 ([69, Proposition 5.5.1]). *The following are equivalent:*

1. *There is a factive justification for φ at x*
   (there exists $U \in \tau$ such that $U \subseteq [\![\varphi]\!]$, $x \in U$, and for all $e \in \widetilde{\mathcal{E}}$, $U \cap e \neq \emptyset$);

2. *$K\varphi$ holds at x.*
   (there exists $U \in \tau$ such that $U \subseteq [\![\varphi]\!]$);

3. *$Int([\![\varphi]\!])$ contains x and is dense in τ;*

4. *$B\varphi \wedge \widetilde{\Box}\varphi$ holds at x.*

As has become evident from the above, this is a notion of knowledge as *correctly justified belief* or *belief based on true justification*. It is, however, not just 'justified true belief', as that would fall prey to Gettier counterexamples [35]: justified true belief is too easily lost, and consistent with having wrong justifications for a true conclusion [69, p. 49]. The knowledge defined here (captured by the modal operator K) is also not absolutely certain (like that captured by modal operator $[\forall]$) — it is amenable to revision, or *defeasible*. In terms of strength, it is stronger than the knowledge characterized by the "no false lemma" of Clark [25], since our concept of justification requires not only the use of no falsehood, but also consistency with all available (combined) evidence [9, p. 25], and weaker than that of the defeasibility theory of knowledge [58, 57, 55, 56]. For an ample discussion of both of these assessments, we refer the reader to [9, pp. 25–29]

So far, we have presented the evidence-based belief- and knowledge-forming process as somewhat of a static process, where all the information that the agent has is fixed and readily-available. But, of course, it would be more realistic if the model had a way to account for changes in information, such as the agent learning a fact with absolute certainty (what we will call truthful public announcements) or gaining a new piece of evidence (evidence addition). [69] presents the topological versions of these evidence dynamics initially formalized by [84] as model transformers.

Formally, we can think of model transformers as operations on models, so we shall say 'model $\mathfrak{M}'$ is the result of applying operation f to model $\mathfrak{M}$'. For what follows, let us fix a topo-e model $\mathfrak{M} = \langle X, \mathcal{E}, [\![\cdot]\!]\rangle$. We present truthful public announcements and evidence addition here, and later on, when defining evidence diffusion models, we will define model transformers that are more closely related to our own modelling setting. Topologically, modelling the effect of a truthful public announcement with φ amounts to moving to the φ-consistent subspace of the original space.

Definition 15 (Truthful Public Announcement)**.** *The result of a truthful public announcement of φ in $\mathfrak{M}$ is the model $\mathfrak{M}^{!\varphi} = \langle X^{!\varphi}, \mathcal{E}^{!\varphi}, [\![\cdot]\!]^{!\varphi}\rangle$, where:*

- $X^{!\varphi} = [\![\varphi]\!]$;
- $\mathcal{E}^{!\varphi} = \{e \cap [\![\varphi]\!] \mid e \in \mathcal{E} \text{ with } e \cap [\![\varphi]\!] \neq \emptyset\}$;
- $[\![p]\!]^{!\varphi} = [\![p]\!] \cap [\![\varphi]\!]$ *for every* $p \in \mathbb{P}$.

Evidence addition is defined in the obvious way: by letting all other objects in the model be the same, and just adding the newly acquired piece of evidence to the basic evidence set.

Definition 16 (Evidence Addition)**.** *The result of applying evidence addition with evidence $e \subseteq X$ to $\mathfrak{M}$ is the model $\mathfrak{M}^{+e} = \langle X, \mathcal{E}^{+e}, [\![\cdot]\!]\rangle$, where, $\mathcal{E}^{+e} = \mathcal{E} \cup \{e\}$.*

As one might expect, the effect of a truthful public announcement with φ on a model is that that sentence becomes true at all possible worlds in that model (given that we have eliminated all φ-inconsistent worlds). Of course, then, it becomes known (both fallibly and defeasibly) and believed by all agents — in fact it becomes common knowledge. As for evidence addition, it is difficult to say in general what effects it will have on the agents' knowledge. Adding a piece of evidence might make an agent acquire more knowledge, or it might make the agent lose knowledge. We shall explore this later on in this chapter. Now that we have defined informational dynamics for topo-e models as well, let us consider an example model in this setting, for clarity's sake.

Example 1. Our agent, Alice, a pest-management entomologist, has been hired by a large houseplant producer to assess and resolve a pest infestation on their Monstera Deliciosa farm. The producer is scared about some kind of insect that they saw. She knows with absolute certainty that the only[2] insects that are seen on or around Monsteras are: spider mites, mealybugs, scale insects, thrips, soil mites and springtails. In reality, and quite unfortunately, thrips are infesting the Monsteras. After a day of looking at the plants on the farm and at the insects that infest them, she has gathered the following evidence: there is damage to the leaves, the insect is dark in color. On the second day, after bringing a loop to more closely see the insect, she gathers the evidence that the insect is thin as well.

Formally, we model the situation on the first day using $\mathfrak{M} = \langle X, \mathcal{E}, [\![\cdot]\!]\rangle$, where $X = \{$Thr(ips), Spr(ingtails), Sca(le), Spi(der mites), Mea(lybugs), Soi(l mites)$\}$. Let us use the propositional letters l, d, t to mean 'there is damage to the *l*eaves', 'the insect is *d*ark in color', and 'the insect is *t*hin' respectively; and for convenience, let us also designate $thr, spr, mea, spi, sca, soi$ as propositional atoms, each standing for one of the insects. We set $[\![l]\!] = \{$Thr, Spi, Mea, Sca$\}$ (green in 5), $[\![d]\!] = \{$Thr, Spi, Sca$\}$ (blue in 5), $[\![t]\!] = \{$Thr, Spr$\}$ (red in 5) and the truth sets of the other propositions that were mentioned are

[2]There might be more insects that thrive around Monsteras that I do not know about, but for the sake of this example, let us say that these are the only ones.

assigned in the obvious way. Of course, the set of basic pieces of evidence is $\mathcal{E} = \{[\![l]\!], [\![d]\!]\}$. The model on the second day will be $\mathfrak{M}^{+\{\text{Thr, Spr}\}}$ We can represent the two models and the change between them in a diagram such as the one below (see Figure 5). Above, we said that the infestation is actually due to thrips, so we are fixing the usually indexical notion of actuality to the world Thr, which we mark by making the node have a double outline.

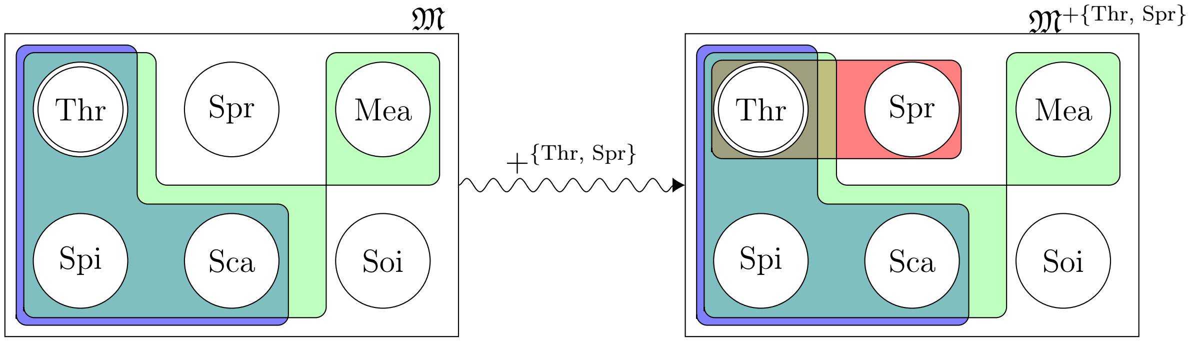

Figure 5: Single-agent Topo-e Model and Evidence Addition

Let us illustrate some of the concepts introduced by looking at the second model. In this model, we have that Ed, Et and El are true, since $[\![d]\!], [\![t]\!], [\![l]\!] \in \mathcal{E}^{+\{Thr,Spr\}}$. By virtue of the fact that $\widetilde{\mathcal{E}}^{+\{Thr,Spr\}}$ is the closure of $\mathcal{E}^{+\{Thr,Spr\}}$ under finite unions, we also have $[\![d]\!] \cap [\![t]\!] = [\![thr]\!] \in \widetilde{\mathcal{E}}^{+\{Thr,Spr\}}$, so we obtain $\widetilde{E}thr$, whereas we do not have $Ethr$ — there is no direct observation pointing to thrips. The same applies to the modalities $\Box$ and $\widetilde{\Box}$, as all evidence sets mentioned thus far contain the actual world. Notice now that the set $\{Thr\}$ is open in the topology of evidence $\tau^{+\{Thr,Spr\}} = \{X, [\![d]\!], [\![t]\!], [\![l]\!], [\![thr]\!], [\![\neg soi]\!], \emptyset\}$ obtained by closing $\widetilde{\mathcal{E}}^{+\{Thr,Spr\}}$ under arbitrary unions. Furthermore, it is dense by virtue of the fact that it has non-empty intersection with all other non-empty open sets. So, $\{Thr\}$ is, and therefore contains, a dense open set, so $Kthr$ holds at Thr — Alice defeasibly knows that the infestation she was called for is a thrips infestation.

Now that the single-agent setting is defined and exemplified, we want to aim our attention towards the principles that guide the interactions between the knowledge and beliefs of multiple agents, and so we move on to defining the multi-agent version of topo-e models.

3.3 Partitional Models as Multi-Agent Topo-e Models

A natural next step for any single-agent setting in epistemic logic is to extend it to be able to a multi-agent one. As discussed in [9, 1, 33], in the case of topological evidence models, the extension is neither immediate, nor obvious. In this section, based closely on [33], we present a multi-agent extension to topo-e models, as well as the conceptual issues that motivate some of the design choices of this setting. We might think a simple way to define the multi-agent setting would be to specify a set of basic pieces of evidence for each agent, as in $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, V\rangle$ where Ag is the set of agents. This would give rise to natural agent-particular definitions for knowledge and belief in terms of openness and density in the topologies generated from the appropriate basic evidence sets. For example, we might propose a satisfaction clause for the belief operator for agent i as: $\mathfrak{M}, x \vDash B_i\varphi$ if and only if $[\![\varphi]\!]$ has a subset that is open and dense in τ_i, the topology generated by $\mathcal{E}_i$.

However, as [33] points out, this approach has some issues. Firstly, we are making the undesirable assumption that all agents have the same 'hard' information, because they all consider the same set of worlds possible. Note that this is particularly visible in the equivalence between the density definition for knowledge — φ is known if $[\![\varphi]\!]$ includes some dense open set containing the actual world — and the formula $[\forall]\widetilde{\Diamond}_i\widetilde{\Box}_iP$ (for any $i \in Ag$). In particular, the outermost operator is the global modality, which makes no reference to any agent. Secondly, the natural semantics borne out of these models validates some interaction principles that are undesirable. For example, for any two agents i, j we have that $\langle K_i\rangle K_i\varphi \rightarrow K_j\langle K_i\rangle K_i\varphi$ is valid. In words, this means that if an agent considers it epistemically possible that they know φ, then any agent knows this. Once again, this is an undesirable principle.

The principal cause of these problems seems to be that we did not adopt *fully* local notions of knowledge and belief. In order to have true locality, we must give up the assumption that all agents possess the same hard information. To do so, we have to change the type of models we work with. The framework presented in this section is the one originally introduced in [33].

Firstly, the set X now represents the set of worlds that *any* agent considers possible. Secondly, a particular agent's 'horizon' of epistemic possibility is no longer the whole set of possible worlds X, but is encoded by a partition of X: the worlds an agent considers possible at a world x will be those in the equivalence class of x. We formally define partitions and introduce some notation that will be useful in later sections.

Definition 17 (Partition, Equivalence Class and Relation)**.** *A family of sets* $\Pi \subseteq \wp(X) \setminus \{\emptyset\}$ *is a partition of* X *if and only if it covers* X ($\bigcup_{\pi\in\Pi}\pi = X$) *and its elements are pairwise disjoint* (for any $\pi, \pi' \in \Pi$, $\pi \neq \pi'$ entails $\pi \cap \pi' = \emptyset$)*. We call* $\pi \in \Pi$ *an equivalence class, and for any* $x \in X$*, we denote the unique* $\pi \in \Pi$ *such that* $x \in \pi$ *by* $\Pi(x)$ *and call it the equivalence class of* x*. The unique relation* $\sim \subseteq X \times X$ *such that for any* $x, y \in X$*,* $x \sim y$ *if and only if* $x, y \in \pi$ *for some* $\pi \in \Pi$ *is an equivalence relation on* X*. We call it the equivalence relation induced by* Π *and we call* Π *the partition induced by* $\sim$*.*

There are many ways to partition a set. Some partitions contain more equivalence classes than others, and beyond the number of equivalence classes, sometimes every equivalence class in a partition is a subset of some equivalence class of another. It turns out that, because of these facts among others, a natural order arises between the partitions of a set. In fact, the partitions of a set form a lattice structure, for which we direct the interested reader to [16]. For our purposes, it will be useful to introduce a few of these notions that aid in comparing partitions.

Definition 18 (Partition Order, Coarsest Common Refinement)**.** *Let* Π_1, Π_2 *be partitions of a set* X*. Then we say* Π_2 *is finer than* Π_1 *(and* Π_1 *coarser than* Π_2*) if and only if for every* $\pi_2 \in \Pi_2$ *there exists* $\pi_1 \in \Pi_1$ *such that* $\pi_2 \subseteq \pi_1$*. For partitions* $\Pi_1, ..., \Pi_n$ *of* X*, we define the coarsest common refinement* $\bigwedge_{i=1}^n \Pi_i$ *in Equation 1.*

$$\bigwedge_{i=1}^{n} \Pi_i = \{\bigcap_{i=1}^{n} \pi_i \mid \pi_i \in \Pi_i \textit{ and } \bigcap_{i=1}^{n} \pi_i \neq \emptyset\}. \tag{1}$$

These notions have natural epistemic interpretations. The finer a partition, i.e. the more equivalence classes it has, the more an agent can distinguish between possible worlds and therefore the more hard information they have. The coarsest common refinement is going to be very useful in the next section, for it can encode the operation whereby agents share all their hard information and combine it — as if they openly shared all of their infallible knowledge.

Wrapping up the discussion of how we change topo-e models to obtain partitional models, each agent gets their own set of basic pieces of evidence, which, much like the single-agent basic evidence sets, require that hard information is evidence too. In the context of these partitions, though, we not only require that $X \in \mathcal{E}_i$, but also that $\Pi_i \subseteq \mathcal{E}_i$. The formal definition of a partitional model is below.

Definition 19 (Partitional Models). *Given a fixed, finite and non-empty set of agents Ag, $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, [\![\cdot]\!]\rangle$ is a partitional model, where X is a nonempty set of possible worlds, $\mathcal{E}_i \subseteq \wp(X)$ is a set of basic pieces of evidence for agent i, Π_i is agent i's partition, and $[\![\cdot]\!]$ is the valuation function. For every $i \in Ag$, we require $X \in \mathcal{E}_i$ and $\Pi_i \subseteq \mathcal{E}_i$.*

Note that after this point, we continue to use finite and non-empty sets of agents, even if we do not explicitly mention this. Furthermore, as we 'localized' hard information and evidence for each agent, it is only natural that we will require a local notion of density, which we introduce below.

Definition 20 (i-local Density). *Given a partitional model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, [\![\cdot]\!]\rangle$, we let τ_i be the topology generated from $\mathcal{E}_i$. For $U \subseteq X$ we say it is i-locally dense in $\pi \in \Pi_i$ if and only if U is dense in $\langle \pi, \tau_i|_{\pi}\rangle$. We call U i-locally dense if and only if it is i-locally dense for every π in Π_i.*

We now define the language of multi-agent logic for evidence, belief, and knowledge. As one can see, it is basically the same language as the single-agent $\mathcal{L}_s$, but each modality is relativized to an agent. The modality $[\Pi_i]$ ought to capture a notion of infallible knowledge, just like $[\forall]$ did in the single-agent setting.

Definition 21 (Language). *Let Ag be a set of agents and let $\mathbb{P}$ be a countable set of propositional atoms. The language $\mathcal{L}_m$ of multi-agent logic for evidence, belief, and knowledge is defined by the recursive rule below. We take $p \in \mathbb{P}$ to be a propositional atom.*

$$\varphi := p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid E_i\varphi \mid \widetilde{E}_i\varphi \mid \Box_i\varphi \mid \widetilde{\Box}_i\varphi \mid [\Pi_i]\varphi \mid K_i\varphi \mid B_i\varphi$$

The reading of these operators is the same as for the single-agent topological evidence models, for which we refer the reader to Definition 9, except we are now indexing the modalities by the agents. For example $E_i\varphi$ is read as 'agent i has a basic piece of evidence for φ'. The reading of $[\Pi_i]\varphi$ is 'agent i knows that φ infallibly'. As for the semantics, for the sake of brevity, we will resort to giving the satisfaction clauses for each operator instead of providing the truth sets as well, except for one illustrative case. The truth sets are easily derived from the definitions provided. For a partitional model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, [\![\cdot]\!]\rangle$ and a world $x \in X$,

- $\mathfrak{M}, x \models p$ iff $x \in [\![p]\!]$;
- $\mathfrak{M}, x \models \neg\varphi$ iff $\mathfrak{M}, x \not\models \varphi$;

- $\mathfrak{M}, x \vDash \varphi \land \psi$ iff $\mathfrak{M}, x \vDash \varphi$ and $\mathfrak{M}, x \vDash \psi$;
- $\mathfrak{M}, x \vDash E_i\varphi$ iff there exists $e \in \mathcal{E}_i$ such that $e \cap \Pi_i(x) \subseteq [\![\varphi]\!]$;
- $\mathfrak{M}, x \vDash \widetilde{E}_i\varphi$ iff there exists $e \in \widetilde{\mathcal{E}}_i$ such that $e \cap \Pi_i(x) \subseteq [\![\varphi]\!]$;
- $\mathfrak{M}, x \vDash \Box_i\varphi$ iff there exists $e \in \mathcal{E}_i$ such that $e \cap \Pi_i(x) \subseteq [\![\varphi]\!]$ and $x \in e$;
- $\mathfrak{M}, x \vDash \widetilde{\Box}_i\varphi$ iff there exists $e \in \widetilde{\mathcal{E}}_i$ such that $e \cap \Pi_i(x) \subseteq [\![\varphi]\!]$ and $x \in e$;
- $\mathfrak{M}, x \vDash [\Pi_i]\varphi$ iff $\Pi_i(x) \subseteq [\![\varphi]\!]$;
- $\mathfrak{M}, x \vDash B_i\varphi$ iff there exists $U \subseteq [\![\varphi]\!]$ such that U is open in τ_i and i-locally dense in $\Pi_i(x)$;
- $\mathfrak{M}, x \vDash K_i\varphi$ iff there exists $U \subseteq [\![\varphi]\!]$ such that U is open in τ_i and i-locally dense in $\Pi_i(x)$ and $x \in U$.

Just like before, the clauses for the modalities E_i, $\widetilde{E}_i$, and $[\Pi_i]$ make no reference to the world of evaluation, except for $\Pi_i(x)$ to implement the locality to the relevant partition. In the single-agent setting, this resulted in these notions being 'global'. But since we have localized our definitions with respect to the equivalence classes of partitions, in effect, from the point of view of a world x within equivalence class π, the new 'global' is π itself, not X. The truth sets for these modalities will reflect this: if a world x in an equivalence class satisfies the clause for, say, $E_i\varphi$, then all the worlds in $\Pi_i(x)$ will satisfy it too. As an example truth set, for E_i, we have

$$[\![E_i\varphi]\!] = \bigcup_{\pi \in \Pi_i} \pi \quad \text{where } \pi \text{ is such that there exists } e \in \mathcal{E}_i \text{ such that } e \cap \pi \subseteq [\![\varphi]\!].$$

At the beginning of this section, we highlighted two problems with the naive approach to multi-agent topo-e models: the conceptual issue of hard information, and the unwanted validity $\langle K_i \rangle K_i\varphi \to K_j\langle K_i \rangle K_i\varphi$. It is quite obvious that the first issue is solved by partitional models — the fact that each agent has their own infallible information is now encoded in the models themselves. As for the second issue, we point to the result in [33, Ch. 3], showing that the logic of n-agent partitional models is the fusion logic combining S4.2K_i for each $i \in \{1, ..., n\}$. In other words, the logic of the knowledge-only fragment of these models is exactly the sum of n many single-agent knowledge-only logics. From this, it is clear that no strange interaction principles such as the one above will be present in our semantics.

Example 2. As an example model in this framework, let us enrich the story from Example 1. Bob is the owner of the farm. He doesn't know much about Monsteras, let alone about pests. But, he cares about the profitability of his farm, so when he heard from his gardeners that there might be an infestation, he took matters into his own hands and went to check the plants out. He is sure that he saw no flying insects on the farm. He is not absolutely convinced, but he thinks that there was a relatively big, elongated insect on one plant — he mistook a bit of dirt for an insect. After some searching the internet, he (mistakenly) thinks that the insect should be one of eight: the six from before, and additionally caterpillars or fruit flies. He also trusts his eyes, and so he takes it that he gathered the following evidence: the insect is dark in color and it is large and elongated (which he takes to mean that it should be either an agglomeration of scale, or a caterpillar).

Formally, we model this using $\mathfrak{M} = \langle X, \{\mathcal{E}_a, \mathcal{E}_b\}, \{\Pi_a, \Pi_b\}, [\![\cdot]\!]\rangle$. X is now the set of worlds that any agent considers possible, so we have to add fruit flies and caterpillars: $X = \{$Thr(ips), Spr(ingtails), Sca(le), Spi(der mites), Mea(lybugs), Soi(l mites), (Fruit) Fly, Cat(erpillar)$\}$. As before, we use propositional letters l, d, t, e to mean 'there is damage to the *l*eaves', 'the insect is *d*ark in color', 'the insect is *t*hin', and 'the insect appears *e*longated' respectively; and for convenience, let us also designate $thr, spr, mea, spi, sca, soi, fly, cat$ as propositional atoms, each standing for one of the insects. We set $[\![l]\!] = \{$Thr, Spi, Mea, Sca$\}$ (green), $[\![d]\!] = \{$Thr, Spi, Sca, Cat$\}$ (blue), $[\![t]\!] = \{$Thr, Spr$\}$ (red), $[\![e]\!] = \{$Sca, Cat$\}$ (purple) and the truth sets of the other propositions that were mentioned in the obvious way.

Now, Alice's partition is $\Pi_a = \{\pi_1, \pi_2\}$, with $\pi_1 = X \backslash \pi_2$ and $\pi_2 = \{$Fly, Cat$\}$, encoding what in Example 1 was implicit in the model's set of worlds: she knows with certainty that fruit flies and caterpillars do not affect monsteras. As for Bob's partition, we have $\Pi_b = \{\pi_3, \pi_4\}$ with $\pi_3 = X \setminus \pi_4$ and $\pi_4 = \{$Fly$\}$, for he is certain that the insect does not fly. As for their evidence, we now have $\mathcal{E}_a = \{X, \pi_1, \pi_2, l, d, t\}$ and $\mathcal{E}_b = \{X, \pi_3, \pi_4, d, e\}$. The diagram for the model[3] is shown below, see Figure 6.

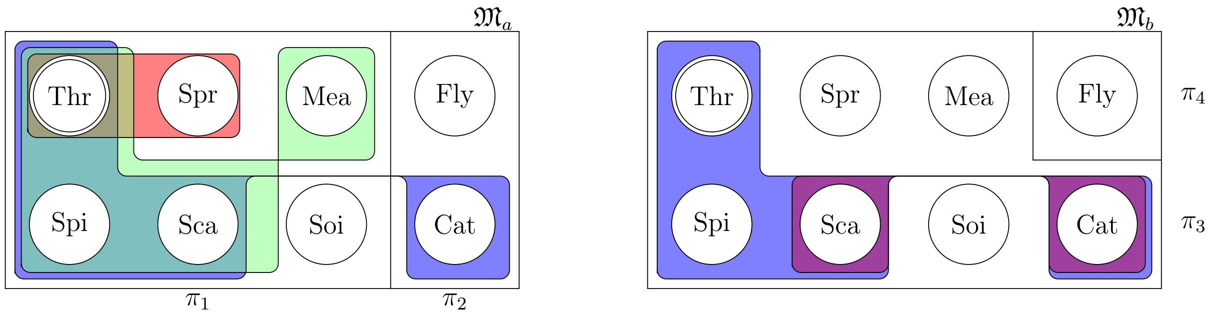

Figure 6: Multi-agent topo-e model

Now let us point out some facts about this model that illustrate the concepts introduced in this section. From here on, when we say φ is true, we mean $\mathfrak{M}, thr \vDash \varphi$. Firstly, notice that we have $E_a(thr \vee spr \vee mea \vee spi \vee sca \vee soi)$, whereas this is not true for Bob. This comes from the condition that $\Pi_i \subseteq \mathcal{E}_i$ for every $i \in Ag$: more specifically, from $\pi_1 \in \mathcal{E}_a$. Further, we have that $E_a(thr \vee spi \vee sca)$ is true, but $E_b(thr \vee spi \vee sca)$ is false. This is because we require the existence of $e \in E_b$ such that $e \cap \pi_3 \subseteq [\![thr \vee spi \vee sca]\!]$, and there is no such e for Bob. What explains this difference between the two agents is the difference in their partitions: $\pi_1 \cap [\![d]\!] = \{Thr, Spi, Sca\}$, whereas $\pi_3 \cap [\![d]\!] = [\![d]\!]$. For what Alice knows and believes, we point out that $\langle X, \tau_a|_{\pi_1}\rangle$ is exactly the topological space of the model in Example 1, so many of her beliefs and knowledge will be the same as in that model, except that she now has beliefs and knowledge about Bob's beliefs and knowledge, and beliefs and knowledge that include the new possibilities fly and cat. Bob, on the other hand, has a partition with less hard information at the actual world, and a piece of evidence that is false, so his strongest piece of knowledge is $thr \vee spi \vee sca \vee cat$. On the other hand, his strongest, i.e. most restrictive in terms of the number of states he believes the real world could be, belief is $B_b(sca \vee cat)$ — a false belief. This is because $[\![sca \vee cat]\!]$ is dense in $\langle X, \tau_b|_{\pi_3}\rangle$.

Now note that Alice knows that Bob believes $sca \vee cat$. This is because $[\![B_b(sca \vee cat)]\!] = \pi_3$ and for $K_a B_b(sca \vee cat)$ we require $U \subseteq \pi_3$ that is open and locally dense in π_1:

[3] We split $\mathfrak{M}$ in $\mathfrak{M}_a = \langle X, \mathcal{E}_a, \Pi_a, [\![\cdot]\!]\rangle$ and $\mathfrak{M}_b = \langle X, \mathcal{E}_b, \Pi_b, [\![\cdot]\!]\rangle$ for clarity.

$\{thr, spi, sca\} \in \tau_a$ is just such a set. Bob, on the other hand, neither knows nor believes that Alice knows thr. He doesn't even believe that she believes it! This is because $[\![B_a thr]\!] = \pi_1$, so for either $B_b B_a thr$ or $K_b B_a thr$ we would need $U \subseteq \pi_1$ which is open and b-locally dense in π_3 and there is no such set. Actually, because all of Bob's pieces of evidence contain the world cat, which is outside π_1, Alice's beliefs and knowledge are absolutely opaque to Bob, for he doesn't have an open that is a subset of π_1. So we can see that now our setting contains not only the means to reason about evidence-based knowledge and belief, but also about what agents know or believe about what other agents know or believe.

A keen reader might have already wondered — what if Alice, the expert, shares her evidence with Bob? And what if Bob told Alice about his mistaken observation when she got on the farm, and the latter took it for granted? If Bob had all the evidence Alice possesses, then he would know $thr \vee sca$, which is certainly an improvement, since he reduced the possible suspect-pests from four to two. If Alice had the evidence $[\![sca \vee cat]\!]$ in her set, $[\![thr]\!]$ would no longer be dense, and Alice would no longer know thr. She would not even believe it! This illustrates the fact that communication between agents is an important part of doxastic-epistemic dynamics, which has the potential to both aid and impair agents' efforts to find the truth. This makes it an especially fascinating avenue to explore in the context of this thesis.

Some work in the classic paradigm of Dynamic Epistemic Logic has been done in this direction. For example, in [69], as well as in the previous section of this thesis, topological versions of the model transformers from [84] are presented. These formalize notions like public announcements, evidence addition, and others, but only for the single-agent case. Apart from these operators in the classical dynamic paradigm, we will be interested in studying a different kind of communication protocol — diffusion. Diffusion via communication is generally defined as the passive, sequential process by means of which information spreads across a network, as long as some conditions hold. For the purposes of this thesis, this type of communication protocol is interesting because it enables us to study the 'long-term' effects of social influence on beliefs and knowledge, as well as the effects of network structure in scientific communities.

3.4 Threshold Models for Diffusion

The dynamics of diffusion under threshold-limited influence have received a lot of scientific attention, gaining early traction through [41, 74, 73], and continuing with [29, 39, 53, 61, 64]. There is also a rich body of literature within the logic community that studies precisely this kind of effect [21, 22, 24, 23, 60, 72, 75, 91]. In this section, closely based on [12], we present threshold models, which are used to study the spread of behaviors, opinions, fashions etc. across a network of agents.

The dissemination of information is formalized via a rule that mimics the effects of social pressure: an agent will adopt, say, a behavior, when the proportion of agents in their social network who have adopted that behavior to the total number of agents in their network is over some threshold. [12] presents a logical setting to reason about threshold-limited influence and its dynamics. We have adapted the notation in that paper to be closer to that used in this thesis. The first notion we introduce is the notion of (social) network.

Definition 22 (Network)**.** *A network is a pair $\langle Ag, \mathcal{N} \rangle$, where Ag is a non-empty finite set of agents and $\mathcal{N} : Ag \mapsto \wp(Ag)$ is a function such that for every $i, j \in Ag$ we have:*

$i \notin \mathcal{N}(i)$ (irreflexivity); $j \in \mathcal{N}(i)$ if and only if $i \in \mathcal{N}(j)$ (symmetry); for every $i \neq j$ there exist $k_0, ..., k_n$ agents such that $k_0 = i, k_n = j$ and for every $\ell \in \{1, ..., n\}$, $k_\ell \in \mathcal{N}(k_{\ell-1})$ (connectedness).

As is visible, there are some pre-imposed constraints on the structure of the network. These are supposed to encode a notion of symmetric contact, or 'friendship'. The irreflexivity requirement is there to ensure that we are not counting any person as their own friend. Symmetry is there to encode the bilateral nature of friendship. Connectedness is enforced to ensure that there are no junctions or isolated agents in the network. This is because we want to study one network at a time, and not multiple.

We will now introduce the simplest kind of models to study threshold-limited influence. These encode one behavior that an agent in the network may or may not have, which we take to be extensional. That is to say that the adopted behavior, denoted by B, is represented by the set of agents that have adopted it. Furthermore, the adoption threshold is uniform in these models.

Definition 23 (Threshold Model)**.** *A threshold model is a tuple $\mathfrak{M} = \langle Ag, \mathcal{N}, B, \theta \rangle$, where $B \subseteq Ag$ is a behavior and $\theta \in [0, 1]$ is a uniform adoption threshold.*

Threshold models consist of a network, together with a behavior and a uniform adoption threshold. Now that we have defined these models, in order to study threshold-limited influence, a dynamic process, we must define an update rule that transforms one model into another based on the network, behavior, and threshold, whilst taking into account our target mechanism. The rule defined below is the way to do this.

Definition 24 (Threshold Model Update)**.** *Let $\mathfrak{M} = \langle Ag, \mathcal{N}, B, \theta \rangle$ be a threshold model. Then the model $\mathfrak{M}' = \langle Ag, \mathcal{N}, B', \theta \rangle$ is called its update, where*

$$B' = B \cup \{i \in Ag \mid \frac{|\mathcal{N}(i) \cap B|}{|\mathcal{N}(i)|} \geq \theta\}$$

Observe that the new set B' of adopters contains the old set (so that $B \subseteq B'$) and those agents for whom the proportion of adopters in their network ($\mathcal{N}(i) \cap B$) to all their neighbors ($\mathcal{N}(i)$) is at least as large as the threshold (θ). This is the standard definition of adoption rule in the literature [12, 29].

The above update rule can be applied infinitely often, and what obtains from this is an infinite sequence of models — we call this a diffusion sequence, and give the formal definition below. Effectively, this allows us to have a 'history' of the spread of the behavior we are investigating.

Definition 25 (Diffusion Sequence)**.** *Let $\mathfrak{M} = \langle Ag, \mathcal{N}, B, \theta \rangle$ be a threshold model. The tuple $\mathcal{S}_{\mathfrak{M}} = \langle \mathfrak{M}_0, \mathfrak{M}_1, ... \rangle$ is a diffusion sequence, where for every $n \in \mathbb{N}$, $\mathfrak{M}_n = \langle Ag, \mathcal{N}, B_n, \theta \rangle$ such that B_n is given by: $B_0 = B$ and $B_{n+1} = B_n'$.*

Even though these sequences are infinite by definition, a neat property of our definition of threshold model update is the fact that a fixed point is always reached. In other words, for any diffusion sequence, there exists a certain point after which no more changes occur.

Proposition 3. *Let $\mathcal{S}_{\mathfrak{M}}$ be a diffusion sequence. For some $n \in \mathbb{N} < |Ag|$, we reach a fixed point such that $\mathfrak{M}_n = \mathfrak{M}_{n+1}$.*

The interested reader can consult a proof of this in [12, Proposition 2.1]. We will prove a similar proposition for our new system in the next section in more detail, see Proposition 12. Now we will define the language used to describe the situation in a threshold model at a given point.

Definition 26 (Language)**.** *Let Ag be a finite set and let atoms be given by $\Phi = \{N_{ij} \mid i, j \in Ag\} \cup \{\beta_i \mid i \in Ag\}$. The language $\mathcal{L}_{[]}$ is given by the following recursive rule:*

$$\varphi := N_{ij} \mid \beta_i \mid \neg\varphi \mid \varphi \wedge \varphi \mid [adopt]\varphi$$

This language has a static component, consisting of all sentences that do not involve the modality $[adopt]$, and a dynamic component, consisting of those sentences that do involve the modality. The atom β_i encodes that an agent has adopted in the model, and the atom N_{ij} encodes that agent i and j are connected. The dynamic modality $[adopt]$ encodes the sentences that will be true after the next update. The next definition provides the truth conditions for these linguistic objects.

Definition 27 (Truth Clauses for $\mathcal{L}_{[]}$)**.** *Given a model $\mathfrak{M} = \langle Ag, \mathcal{N}, B, \theta \rangle$, $N_{ij}, \beta_i \in \Phi$ and $\varphi, \psi \in \mathcal{L}_{[]}$:*

- $\mathfrak{M} \models \beta_i$ *iff* $i \in B$;
- $\mathfrak{M} \models N_{ij}$ *iff* $j \in \mathcal{N}(i)$;
- $\mathfrak{M} \models \neg\varphi$ *iff* $\mathfrak{M} \not\models \varphi$;
- $\mathfrak{M} \models \varphi \wedge \psi$ *iff* $\mathfrak{M} \models \varphi$ *and* $\mathfrak{M} \models \psi$;
- $\mathfrak{M} \models [adopt]\varphi$ *iff* $\mathfrak{M}' \models \varphi$, *where* $\mathfrak{M}'$ *is the updated threshold model (Definition 24).*

It is not a difficult observation to make that threshold models pre-encode the next rounds, in the sense that all the information required to compute the updates is already in the model. This is reflected in our language as well, in the sense that we can construct sentences which, in some sense, 'predict' what is going to happen at the next rounds. We therefore define the following abbreviations, for which we credit [12].

Abbreviation. The abbreviation $[adopt]^n\varphi$, defined recursively as

$$[adopt]^0\varphi := \varphi$$
$$[adopt]^{n+1}\varphi := [adopt][adopt]^n\varphi$$

Abbreviation.

$$\beta_{\mathcal{N}(i)\geq\theta} := \bigvee_{\{G\subseteq F\subseteq Ag \mid \frac{|G|}{|F|}\geq\theta\}} \left(\bigwedge_{j\in F} N_{ij} \wedge \bigwedge_{j\notin F} \neg N_{ij} \wedge \bigwedge_{j\in G} \beta_j \right)$$

The first abbreviation introduces a neater way to write sentences that talk about what happens at later rounds, so that instead of writing the $[adopt]$ modality many times, we can just write the number of the future round we are talking about as a superscript to the adopt modality. The second abbreviation encodes the fact that agent i will adopt at the next round. Note that the set of agents is finite, so we can read $\bigvee$ as an existential quantifier and $\bigwedge$ as a universal quantifier. Then, $\beta_{\mathcal{N}(i)\geq\theta}$ is true if and only if there exist two sets: F such that every agent in F is connected to i and none of the agents outside

F are — so exactly i's network — and G, a subset of F such that the proportion of agents in the group to all of i's friends is over the threshold *and* all the agents in this group have adopted.

The reason why we are interested in this kind of diffusion dynamic is because of the ease with which cascades can be studied in this context. A cascade is like a 'chain reaction' in the diffusion sequence: an agent adopting causes another agent to adopt, which causes another agent to adopt, and so on. A cascade is called *full* (or complete [12]) when the behavior is adopted by all agents. Note here that, even though we previously proved that every diffusion sequence reaches a fixed point, this does not mean that every diffusion sequence reaches a full cascade. This is, in fact, not the case in this setting, as the counterexample below shows. From here on, if we use the terms 'cascade' and 'full cascade' interchangeably, and always mean full cascade.

Example 3. Let us say that each node represents an agent, that the connecting lines represent the network, and that the agents in red are the ones that adopted. Let us also say that the threshold is $\theta = \frac{3}{4}$. Agent 3 will never adopt, because exactly half of his network has adopted, and exactly half of his network has not adopted. Thus, a cascade is never reached.

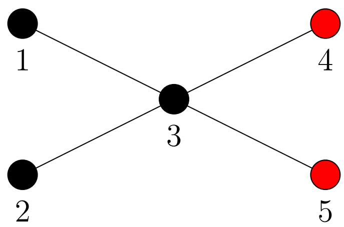

Figure 7: Threshold model reached fixed point, but not cascade.

Our language is already expressive enough to capture whether or not a model will *eventually* form a full cascade. This is achieved via the abbreviation below, which, says that in $|Ag| - 1$ updates, all agents have adopted. We make a claim about the $(|Ag| - 1)^{\text{th}}$ update because we know, via Proposition 3 that the update sequence has certainly reached a fixed point at that point.

Abbreviation.

$$cascade := [adopt]^{|Ag|-1} \bigwedge_{i \in Ag} \beta_i$$

In what follows, we look at the necessary and sufficient conditions for cascades to form. For some intuition, it will be worth looking at the example above and considering what prevented the formation of a full cascade. Literally speaking, we have already done this in the explanation above: agent 3 will never adopt. However, we can also put things in a different light: the group of agents $\{1, 2, 3\}$ *as a unit* will never adopt, because they form a group that, *together*, is quite resilient to adopting. And why are they resilient? Because every agent in that group has enough non-adopting neighbors to never adopt the behavior. This "enough" is a notion of density that we can capture formally. We introduce the concept of a cluster of density d in the network now, after [12, 29].

Definition 28 (Cluster of Density d)**.** *Given a network $\langle Ag, \mathcal{N} \rangle$ and a subset of the set of agents $C \subseteq Ag$, we say it is a cluster of density d if and only if for all $i \in C$,*

$$\frac{|\mathcal{N}(i) \cap C|}{|\mathcal{N}(i)|} \geq d$$

The subgroup $\{1, 2, 3\}$ of non-adopting agents is a cluster of density 0.5. A classic result in the study of threshold limited influence [64, 29, Ch. 19.3] states that this is precisely the condition that prevents cascades from forming, and more precisely that the cluster is of density greater than $1 - \theta = 1 - \frac{3}{4} = \frac{1}{4} = 0.25 < 0.5$. We formulate the result for threshold models here, as in [12].

Theorem 1 (The Cluster Theorem)**.** Given a threshold model $\mathfrak{M}$ with threshold $\theta \neq 0$ and a set $B \subset Ag$ who have adopted, all agents will eventually adopt if and only if there exists no cluster of density greater than $1 - \theta$ in $Ag \setminus B$.

As before, it turns out that we can express that a group C is a cluster of density d in our language. The proof of this Proposition is given in [12, Proposition 2.3].

Proposition 4. *A group C is a cluster of density d in $\langle Ag, \mathcal{N} \rangle$ iff $\mathfrak{M} = \langle Ag, \mathcal{N}, B, \theta \rangle$ satisfies*

$$\bigwedge_{i \in C} \bigvee_{\{G \subseteq F \subseteq Ag \mid \frac{|G \cap C|}{|F|} \geq d\}} \left(\bigwedge_{j \in F} N_{ij} \wedge \bigwedge_{j \notin F} \neg N_{ij} \right)$$

And further, we can use the fact that our set of agents is finite, and use disjunction as a kind of existential quantifier to express the existence of a cluster of a certain density. Below, we showcase a sentence presented by [12] that encodes the existence of a cluster of density d of agents that have not adopted the behavior.

Abbreviation.

$$\exists C_{\geq d} \neg \beta := \bigvee_{C \subseteq Ag} \bigwedge_{i \in C} \bigvee_{\{G \subseteq F \subseteq Ag \mid \frac{|G \cap C|}{|F|} \geq d\}} \left(\bigwedge_{j \in F} N_{ij} \wedge \bigwedge_{j \notin F} \neg N_{ij} \wedge \bigwedge_{j \in G} \neg \beta_j \right)$$

With this, as it has become quite clear, we can express the cluster theorem within our language as well. We follow [12] in the formulation. For $\mathfrak{M} = \langle Ag, \mathcal{N}, B, \theta \rangle$ with $\theta \neq 0$:

$$\mathfrak{M} \vDash \textit{cascade} \leftrightarrow \neg \exists C_{>1-\theta} \neg \beta$$

Threshold models are a great tool to study the effects of network structure on social influence. Their simplicity also makes it easy to implement in other settings as well, which is exactly what we will proceed to do in the following section. We will study the effects of social network structure on the diffusion of pieces of evidence in networks where agents form evidence-based beliefs and knowledge.

3.5 Evidence Diffusion Models

The main phenomenon we would like to investigate in our framework, which combines Partitional Models with Threshold Models, is the diffusion of pieces of evidence in the context of a scientific social network. Concretely, we can think of a group of scientists that have all performed experiments so that each has a specific set of (possibly inconsistent) pieces of evidence, as well as some scientific contacts which they often talk to or collaborate with. As time goes on, they might perform more experiments and obtain new pieces of evidence, either alone, or in collaboration with other scientists. They might also often talk to their scientific contacts about various pieces of evidence that one or the other has not heard about, and spread awareness about new results. These are the kinds of phenomena we would like to model in this section.

3.5.1 Semantic Notions

As the reader might expect, the formalization of evidence, evidence-based beliefs and (defeasible) knowledge will be taken care of by the (multi-agent) topo-e models we have presented in the previous sections. As is the case in partitional models, any agent has the ability to reason about the evidence, beliefs, and knowledge of any other agent. However, we would not want, in the present modelling setting, to unrealistically assume that communication and collaboration also happen in this unfettered way, for this is very rarely how things happen in reality. For this reason, in our setting, two agents being scientific contacts of each other boils down to the existence of an avenue to share pieces of evidence with each other — a scientific network. We also separate scientific collaboration from scientific communication. We understand the former to be the joined effort of multiple agents to create new evidence, for instance through experimentation. The latter we understand to be the sharing and diffusion of certain information or of evidence from agents who possess it to agents who do not.

We model the spread of evidence across the network as a communication process which is *social*, *local* [81, 44], and *threshold-limited* [29, 76]. It is social in the sense that you adopt pieces of evidence not on the basis of some individual rational process, but on the basis of the evidence others are using. It is local in the sense that you take into account not what everyone is doing, but precisely what your close contacts are doing. It is threshold-limited in the sense that a specific proportion of your scientific contacts will need to possess a piece of evidence for you to adopt it. A keen reader might already anticipate the way we shall formalize this: an agent i will adopt a piece of evidence e if out of his network of scientific contacts, $\mathcal{N}(i)$, the proportion of agents who have this piece of evidence in their evidence sets to all the agents in $\mathcal{N}(i)$ is over some threshold θ.

Definition 29 (Evidence Diffusion Model). *Given a fixed, finite and non-empty set of agents Ag and a threshold $\theta \in [0,1]$ that is fixed for all agents and pieces of evidence, $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ is an evidence diffusion model. For every $i \in Ag$, we require $X \neq \emptyset$, $X \in \mathcal{E}_i$, $\emptyset \notin \mathcal{E}_i$, and $\Pi_i \subseteq \mathcal{E}_i$. Finally, $\mathcal{N} : Ag \mapsto \wp(Ag)$ is a neighbor function, satisfying: irreflexivity ($a \notin \mathcal{N}(a)$), symmetry ($a \in \mathcal{N}(b)$ implies $b \in \mathcal{N}(a)$) and connectedness (for every $i \neq j$ there exist $k_0, ..., k_n$ agents such that $k_0 = i, k_n = j$ and for every $\ell \in \{1, ..., n\}$, $k_\ell \in \mathcal{N}(k_{\ell-1})$).*

The first ingredients of our modelling setting are the set of agents and the so-called diffusion threshold. In this particular setting, we are letting the diffusion threshold be constant for all agents and all possible pieces of evidence. Against this background, our Evidence Diffusion models consist of a set of possible worlds, a set of basic pieces of evidence for each agent, an information partition for each agent, a valuation, and a neighbor function specifying the social network.

We constrain the network in three significant ways, all conceptually motivated by the phenomenon we are trying to model. Firstly, we take scientific connections to be symmetric — if I am open to sharing my evidence with you, then so are you with me. Secondly, we impose the constraint of connectedness — this eliminates the possibility of agents that have no connections, which is desirable since our focus is on scientific networks *proper*. Finally, we take the network to be irreflexive — this is because we want the threshold θ to reflect the proportion of genuine scientific connections required for evidence adoption, and counting oneself as a connection doesn't make much sense.

Before we move on with the formal presentation of the language and semantics we will be using, it would be worth discussing the point at which the analogy between our intended interpretation of evidence as ‘experiments’ and real-world experiments breaks down. In our framework, a piece of evidence is represented as the set of possible worlds consistent with that evidence. We may put this, equivalently, as the piece of evidence being represented as the strongest (most restrictive) hypothesis it supports, as propositions can also be taken as the set of worlds at which they are true. This latter equivalence highlights a simplifying assumption of our model and an incongruence between the framework and the way the concept of ‘evidence’ is normally understood [52, 2]. Colloquially, even if two different experiments e_1 and e_2 support the exact same hypothesis H, they will not be understood as equivalent, but as two separate avenues of support for H. Not so in our framework. Thus, when we say ‘experiment’, we really use this *semantic* notion of experiment, and when we say ‘experimental validation’, we really mean ‘experimental validation of a hypothesis H by any experiment e such that H is the maximally restrictive hypothesis supported by e’. With this in mind, we go on to define the language of evidence diffusion.

Definition 30 (Language)**.** *Let Ag be a finite set of agents and let $\mathbb{P}$ be a countable set of propositional atoms. Then we define the language of evidence diffusion for individual knowledge and belief $\mathcal{L}_d$ by the following recursive rule, where $p \in \mathbb{P}$ and $i, j \in Ag$:*

$$\varphi := p \mid N_{ij} \mid \neg\varphi \mid \varphi \wedge \varphi \mid E_i\varphi \mid [\Pi_i]\varphi \mid K_i\varphi \mid B_i\varphi$$

This language is different from that of $\mathcal{L}_m$, the language of multi-agent evidence, belief, and knowledge in two major ways. Firstly, we have added the set of network atoms $\{N_{ij} \mid i, j \in Ag\}$, which give us the linguistic means to capture network properties in the language itself. Secondly, we replaced the four evidence modalities from before with just one evidence modality, E, which, though using the same notation, will be interpreted differently than in the previous section. We read N_{ij} as ‘j a neighbor of i’ and $E_i\varphi$ as ‘the proposition φ is part of i’s basic evidence set’. Note that this latter reading is different from ‘i has a basic piece of evidence for φ’ — the difference will become clearer after we provide a semantics for this language. Later on, we will add several new modalities pertaining to group knowledge and dynamics. We provide satisfaction clauses for the notions introduced so far.

Definition 31 (Satisfaction)**.** *Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be an evidence diffusion model, $x \in X$ be a world, $p \in \mathbb{P}$ a propositional atom, and $i, j \in Ag$ agents. We define the semantics as follows:*

- $\mathfrak{M}, x \vDash p$ *iff* $x \in [\![p]\!]$;
- $\mathfrak{M}, x \vDash N_{ij}$ *iff* $j \in \mathcal{N}(i)$;
- $\mathfrak{M}, x \vDash \neg\varphi$ *iff* $\mathfrak{M}, x \nvDash \varphi$;
- $\mathfrak{M}, x \vDash \varphi \wedge \psi$ *iff* $\mathfrak{M}, x \vDash \varphi$ *and* $\mathfrak{M}, x \vDash \psi$;
- $\mathfrak{M}, x \vDash E_i\varphi$ *iff* $[\![\varphi]\!] \in \mathcal{E}_i$;
- $\mathfrak{M}, x \vDash [\Pi_i]\varphi$ *iff* $\Pi_i(x) \subseteq [\![\varphi]\!]$;
- $\mathfrak{M}, x \vDash B_i\varphi$ *iff there exists* $U \subseteq [\![\varphi]\!]$ *such that* U *is open in* τ_i *and* i*-locally dense in* $\Pi_i(x)$;
- $\mathfrak{M}, x \vDash K_i\varphi$ *iff there exists* $U \subseteq [\![\varphi]\!]$ *such that* U *is open in* τ_i *and* i*-locally dense in* $\Pi_i(x)$ *and* $x \in U$.

Notably, the modality E requires that φ's truth set itself be a part of i's set of basic evidence. This is very different from the way we interpreted E in the context of partitional models, which only required the existence of a piece of evidence whose intersection with the equivalence class of the world of evaluation is a subset of φ. Having this modality which more precisely picks out evidence sets is desirable for us because it will allow us to encode relevant properties of diffusion sequences, much like the atom β_i does in the case of threshold models. Moreover, it matches the diffusion sequence's 'blindness' to agents' partitions in the sense that a piece of evidence may get transmitted, which counts as a piece of evidence for some proposition φ for an agent, but not for another, like in Example 2.

Since later on in this section we will be proving some properties of the notions introduced above, we define entailment and validity for evidence diffusion models. Note that we are universally quantifying over thresholds, so that if a proposition is valid, it is valid for all thresholds.

Definition 32 (Entailment, Validity)**.** *For any $\Gamma \subseteq \mathcal{L}_d$ and any $\varphi \in \mathcal{L}_d$ we write $\Gamma \vDash \varphi$ if and only if for any threshold θ, for any evidence diffusion model $\mathfrak{M}$ for θ, for any world $x \in X$, if for every $\gamma \in \Gamma$, $\mathfrak{M}, x \vDash \gamma$, then $\mathfrak{M}, x \vDash \varphi$ and say that Γ entails φ. If $\emptyset \vDash \varphi$ we write $\vDash \varphi$ and call φ valid or a validity. We write $\mathfrak{M} \vDash \varphi$ if for every world $x \in X$, $\mathfrak{M}, x \vDash \varphi$.*

With the semantic notions pertaining to evidence diffusion models introduced, we can move on to the discussion of knowledge and evidence dynamics on these models.

3.5.2 Dynamics

It is quite clear that the scientific process is dynamic and continuous. As such, a model of scientific networks ought to capture at least a part of this dynamicity. We think that collaboration and communication, which, we remind the reader, are different notions for us, are two essential epistemic actions that occur within the context of actual scientific networks. In this subsection, we introduce three model transformers and corresponding dynamic operators that capture these two processes. We capture collaboration by means of arbitrary evidence addition, which encodes the discovery of a piece of evidence by a subgroup of agents. We also formalize two types of communication: the sharing of hard information, or certain knowledge, and that of soft information, or evidence. Law establishment is the way we encode the sharing of certainties: it is the operation whereby all agents come together and combine their infallible knowledge to obtain the 'laws' of their field. Evidence diffusion is how we encode the sharing of evidence: and we take it to be the process by means of which soft information spreads across the network. We start by introducing evidence diffusion, which is the main focus of our modelling setting.

We think of evidence-sharing as taking place in 'diffusion rounds'. At each such round, the agents accept from their neighbours the basic pieces of evidence that are at or above the threshold, incorporating this new piece of evidence into their own set of basic pieces of evidence. We restrict attention to the basic pieces of evidence because, from a modelling point of view, they much more closely fit the phenomenon at hand. This is because the combined pieces of evidence and the open sets in the topology of evidence represent the result of agents' reasoning with the basic pieces of evidence, and we only want the 'empirical data' to be transmitted, so agents on the receiving end can do their own logical reasoning on the basis of it. This communication procedure is borne out in the update rule defined below.

Definition 33 (Evidence Diffusion). *The result of applying evidence diffusion to the evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *for* θ *is the evidence diffusion model* $\mathfrak{M}^{\because} = \langle X, \{\mathcal{E}_i^{\because}\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *for* θ, *where for any* $i \in Ag$,

$$\mathcal{E}_i^{\because} = \mathcal{E}_i \cup \{e \in \bigcup_{j\in Ag} \mathcal{E}_j \mid \frac{|\{j \mid j \in \mathcal{N}(i) \text{ and } e \in \mathcal{E}_j\}|}{|\{j \mid j \in \mathcal{N}(i)\}|} \geq \theta\}.$$

Note that neither this nor any other opdate changes the threshold for diffusion. As one might imagine, this model transformer can be applied repeatedly, creating a sequence of models, where each model is the result of applying evidence diffusion to the previous one. We call such a sequence a diffusion sequence, and give the definition as well as some notational conventions regarding these sequences below.

Definition 34 (Diffusion Sequence). *Let* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *be an evidence diffusion model. We define its diffusion sequence to be the tuple* $\mathcal{S}_{\mathfrak{M}} = \langle \mathfrak{M}_0, \mathfrak{M}_1, ...\rangle$ *where* $\mathfrak{M}_0 = \mathfrak{M}$ *and for any* $n \in \mathbb{N}$, $\mathfrak{M}_{n+1} = \mathfrak{M}_n^{\because}$. *When it is convenient and clear within the context, we will refer to the basic evidence set of agent* i *in the model* $\mathfrak{M}_k$ *for* $k \in \mathbb{N}$ *by* $\mathcal{E}_i^k$ *and similarly by* $\widetilde{\mathcal{E}}_i^k$ *for their combined evidence set and* τ_i^k *for their topology of evidence.*

The way we capture the sharing of hard information will have to do with the partitions which we used to encode each agent's infallible knowledge. Formally, the result of all the agents putting together their infallible knowledge will be that they all end up with the same partition of the set of possible worlds. This partition will be the coarsest common refinement of all of their partitions — the least refined partition that takes into account all of the hard information in each agent's partition. Note that the change to the agents' basic evidence sets is required to make $\mathfrak{M}^{\bigwedge}$ an evidence diffusion model.

Definition 35 (Law Establishment). *The result of applying law establishment to* $\mathfrak{M}$ *is the model* $\mathfrak{M}^{\bigwedge} = \langle X, \{\mathcal{E}_i^{\bigwedge}\}_{i\in Ag}, \{\Pi_i^{\bigwedge}\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$, *where for every* $i \in Ag$, $\Pi_i^{\bigwedge} = \bigwedge_{j\in Ag} \Pi_j$ *(recall Definition 18) and* $\mathcal{E}_i^{\bigwedge} = \mathcal{E}_i \cup \Pi_i^{\bigwedge}$.

Arbitrary evidence addition is defined in the obvious way: by keeping all the elements of the model the same, but adding to the basic evidence set of each agent in the desired subgroup the piece of evidence in question. For us, arbitrary evidence addition will play the role of formalizing the discovery of new pieces of evidence by direct observation or experimentation.

Definition 36 (Arbitrary Evidence Addition). *The result of applying arbitrary evidence addition with evidence* $e \subseteq X$ *for subgroup* $G \subseteq A$ *to* $\mathfrak{M}$ *is the model* $\mathfrak{M}^{+_G e} = \langle X, \{\mathcal{E}_i^{+_G e}\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$, *where for every* $i \in Ag$ *we have*

$$\mathcal{E}_i^{+_G e} = \begin{cases} \mathcal{E}_i \cup \{e\} & \textit{if } i \in G; \\ \mathcal{E}_i & \textit{otherwise.} \end{cases}$$

.

Having these model transformers, we can naturally extend our language with dynamic modalities corresponding to these and obtain easy semantic clauses for them.

Definition 37 (Language). *Let Ag be a finite set of agents and let $\mathbb{P}$ be a countable set of propositional atoms. Then we define the dynamic language of evidence diffusion for individual knowledge and belief $\mathcal{L}_d^{\square}$ by the following recursive rule, where $p \in \mathbb{P}$ and $i, j \in Ag$:*

$$\varphi := p \mid N_{ij} \mid \neg\varphi \mid \varphi \wedge \varphi \mid E_i\varphi \mid [\Pi_i]\varphi \mid K_i\varphi \mid B_i\varphi \mid [\because]\varphi \mid [\bigwedge]\varphi \mid [+_G e]\varphi$$

For an evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ and a world $x \in X$, we have the following new semantic clauses, the rest staying the same as before.

- $\mathfrak{M}, x \models [\because]\varphi$ if and only if $\mathfrak{M}^{\because}, x \models \varphi$;
- $\mathfrak{M}, x \models [\bigwedge]\varphi$ if and only if $\mathfrak{M}^{\bigwedge}, x \models \varphi$;
- $\mathfrak{M}, x \models [+_G e]\varphi$ if and only if $\mathfrak{M}^{+_G e}, x \models \varphi$.

We now move on to the presentation of the group knowledge modalities.

3.5.3 Group Knowledge Modalities

As the tension between individual and social epistemology is a focus of our thesis, it is only natural that our formal setting would encode various notions of group knowledge. By exploring the formal relationships between our definitions of group knowledge and individual evidence-based knowledge and belief, we tease out possible relationships between individual and group rationality. By having the linguistic and semantic means to reason about the structure of the agent network and the diffusion sequence, our setting can place these relationships in the context of the evidential dynamics within scientific networks. And by doing so, our results provide insights into how the structure of the agent network affects various epistemic factors: the robustness of (defeasible) knowledge, truth-tracking capabilities, speed of evidence diffusion, speed of knowledge-formation etc.

There are multiple possible avenues for defining group knowledge in our setting. We are mostly interested here in notions of knowledge that capture what the agents *would* know if they were to communicate with each other. As it turns out, there are multiple notions that could fit this description in this framework. One of them, which is the one introduced in Definition 39, we will call 'distributed evidence', and we denote it by the modal operator ƎD. The other will be called 'distributed knowledge', and indeed it will correspond to the traditional notion of distributed knowledge [43, 32, 46, 87, 71]. We will denote it by ꓘD. We augument the language further, by adding the clauses ƎDφ | ꓘDφ to our initial definition (to which, recall, we also added the dynamic modalities).

Definition 38 (Language). *Let Ag be a finite set of agents and let $\mathbb{P}$ be a countable set of propositional atoms. Then we define the dynamic language of evidence diffusion for group knowledge $\mathcal{L}_{dg}^{\square}$ by the following recursive rule, where $p \in \mathbb{P}$ and $i, j \in Ag$:*

$$\varphi := p \mid N_{ij} \mid \neg\varphi \mid \varphi \wedge \varphi \mid E_i\varphi \mid [\Pi_i]\varphi \mid K_i\varphi \mid B_i\varphi \mid [\because]\varphi \mid [\bigwedge]\varphi \mid [+_G e]\varphi \mid \text{ƎD}\varphi \mid \text{ꓘD}\varphi$$

Definition 39 (Distributed Evidence). $\mathfrak{M}, x \models$ ƎDφ *if and only if for every $i \in Ag$, there exists $U_i \in \tau_i$ such that:*

1. $Cl_{\vee\tau_{\cap}}(\bigcap_{i\in Ag} U_i) = \bigcap_{i\in Ag} \Pi_i(x)$, *where* $\vee\tau_{\cap} = \bigvee_{i\in Ag} \tau_i\Big|_{\bigcap_{i\in Ag}\Pi_i(x)}$[4].

   (If the agents communicated all their information, hard and soft, without boundaries, they would have a justification...)

2. $x \in \bigcap_{i\in Ag} U_i$;

   (...which is factive...)

3. $\bigcap_{i\in Ag} U_i \subseteq [\![\varphi]\!]$;

   (... and which entails φ.)

What this notion encodes is the knowledge the agents would come to possess if they put all of their evidence and infallible knowledge together and reasoned as one agent. It was formally defined by [1] and it was explored by [33] whether it is the appropriate notion of distributed knowledge for Partitional Models (which was answered in the negative, and for which the notion in Definition 40 was provided). In our framework, since we have actual encodings of inter-agent communication, we shall find that this will coincide with the knowledge of an agent given certain conditions, which we prove in Proposition 14.

As [69] points out, this notion of knowledge is not monotonic, in the sense that some proposition φ might be individual knowledge (there is an agent i such that $K_i\varphi$ holds at the actual world) but not hold under distributed evidence (the actual world doesn't make $\mathrm{ED}\varphi$ true). This should be pretty clear, as we have seen in Example 2 that sometimes, adding evidence to an agent's set does not confer them more knowledge, but less.

Definition 40 (Distributed Knowledge). $\mathfrak{M}, x \vDash \mathbf{KD}\varphi$ *if and only if for every* $i \in Ag$, *there exists* $U_i \in \tau_i$ *such that:*

1. U_i *is* $\Pi_i(x)$*-locally dense;*

   (All agents have justifications...)

2. $x \in \bigcap_{i\in Ag} U_i$;

   (... which are factive ...)

3. $\bigcap_{i\in Ag} U_i \subseteq [\![\varphi]\!]$.

   (... and when put together, these justifications entail φ.)

This notion is more akin to the regular understanding of distributed knowledge [43, 32, 46, 87, 71], coinciding with the one on Kripke models in the sense that distributed knowledge is that which is entailed by what the agents know individually. Despite its conceptual appeal, as we shall see in the next section, this notion is not robust whatsoever. That is to say that knowledge is very easily lost under evidence diffusion.

To illustrate how these notions work in concert, we give an example, much like in the previous sections.

[4]For the definition of the join topology, $\bigvee_{i\in Ag} \tau_i$, see Definition 8.

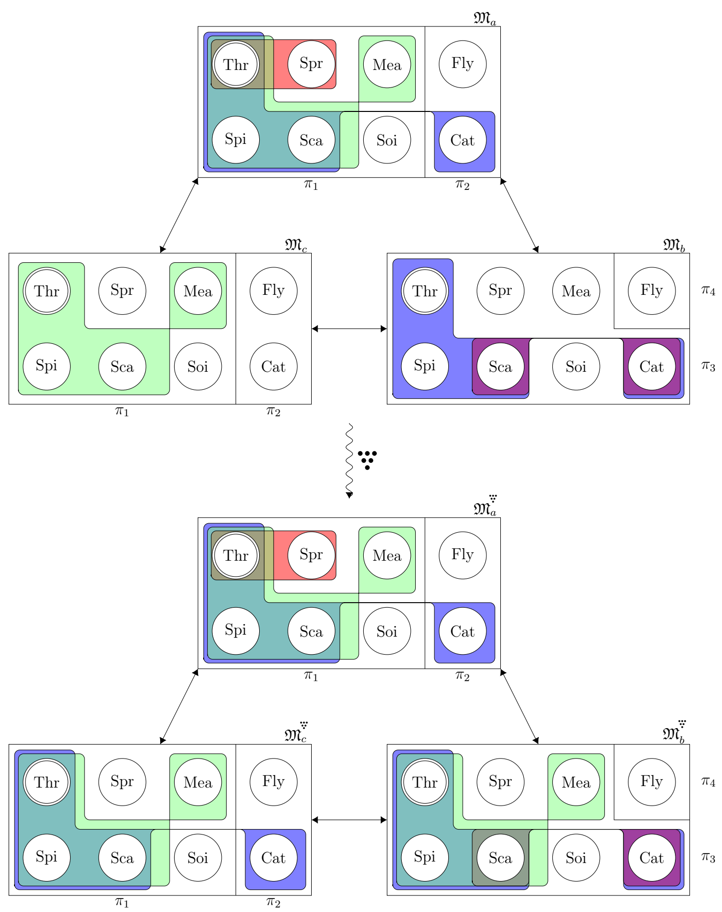

Figure 8: Evidence Diffusion Model and Update

Example 4. We can consider a similar story to the one in Example 2, but with a few twists. We will let Alice and Bob be in the same position, and have the same pieces of evidence as before. We are adding one additional agent, Clarice, who is a gardener on the Monstera farm. Clarice waters the plants every day, and so she, like Alice, noticed that many leaves look damaged. She also knows with certainty, from an old gardener's handbook which we take to be infallible, that caterpillars and fruit flies do not affect Monsteras. We shall say that all three agents communicate openly the day after Alice's arrival on the farm (so after Alice gathered all her evidence).

All of this information is encoded in the model on the top. Note that we set the threshold to $\theta = 0.51$, so that one agent in your network having a piece of evidence will be insufficient to receive that piece of evidence. Using the straight arrows, we represent the network: everyone is connected to everyone else. The model on the bottom, $\mathfrak{M}^{\because}$ is the result of applying the evidence diffusion rule to $\mathfrak{M}$.

After one round of evidence diffusion, Alice's evidence set remains the same, because any evidence she might receive from other agents is either already in her evidence set, or has too few adopting agents to be adopted. This means that her knowledge (excluding knowledge about other agents' knowledge and beliefs) will also remain the same. Clarice gains the piece of evidence {Thr, Spi, Sca, Cat}, as both Bob and Alice possess it, and Bob receives the piece of evidence {Thr, Spi, Sca, Mea}, since both Alice and Clarice have it. This diffusion round is epistemically good for both Bob and Clarice, as they now both know $thr \vee spi \vee sca$, more than they knew before.

Note that ƎDthr does not hold at the actual world in either model. To see why, recall that every agent is required to have a set in their topology of evidence such that that set is dense in $\vee\tau_{\cap} = \bigvee_{i \in Ag} \tau_i\Big|_{\bigcap_{i \in Ag} \Pi_i(x)}$. Since $\vee\tau_{\cap}$ contains the set {Sca}, we see that it is impossible to select sets from the three agents such that the requirements of ƎDthr are met. At the same time, ⅄Dthr holds at the actual world in both models, since every agent has a $\Pi_i(x)$ locally dense set ({Thr, Spr} for Alice, {Thr, Spi, Sca, Cat} for Bob and Clarice) such that they all contain *thr* and their intersection is a subset of {Thr}. This is a happy example in which Distributed Knowledge is not lost after an update, but as we shall see, this is certainly no rule, but an exception.

In this chapter, we introduced several formal frameworks, closing with the contribution of this thesis, evidence diffusion models in Section 3.5. This framework provides the means of expressing and analyzing different notions of group knowledge in the context of a social network subject to evidence diffusion phenomena. For this reason, it serves as a (toy) model for scientific networks, which means that we can interpret formal facts about the model as good indications for the validity of the epistemological concepts introduced in Chapter 2. We therefore proceed to the next chapter, which consists of this analysis of formal facts about the system in the light of formal social epistemology.

4 Epistemic and Evidential Phenomena on Evidence Diffusion Models

In this chapter we analyse some of the formal properties of evidence diffusion models. First, we look at properties of agent networks, and introduce several network topologies that we will pay special attention to. Second, we focus on diffusion as a model transformer, and assess the stability (or invariance) of the modalities introduced in the previous chapter under this transformation. Third, we take networks and diffusion together, and we look at the interplay between the two. Fourth, we focus in on the special networks introduced in the first section of this chapter and investigate possible measures of speed and reliability for these. Finally, we compare our results with Zollman's claims about the same networks.

4.1 Networks

An essential ingredient to our modelling setting are the social networks that the agents form. Not only are they the base for diffusion phenomena, but since what spreads through the network are pieces of evidence, and the notions of knowledge and belief we have defined are evidence-based, networks critically affect the epistemic and doxastic attitudes that agents hold as well.

Our networks are specified by graphs, which are mathematical objects that have the double advantage of being intuitive and easy to illustrate, as well as supported by a rich literature, including literature that focuses on social networks specifically [29, 88, 41, 40, 19, 90, 89]. We now introduce some graph-theoretic notions, including measures on graphs that we will interpret as measures of connectivity in Section 4.5. In doing this, we use definitions from [88]. For the following, take a non-empty finite set of agents Ag and a network $\mathcal{N} : Ag \mapsto \wp(Ag)$ satisfying: irreflexivity ($a \notin \mathcal{N}(a)$), symmetry ($a \in \mathcal{N}(b)$ implies $b \in \mathcal{N}(a)$) and connectedness (for every $i \neq j$ there exist $k_0, ..., k_n$ agents such that $k_0 = i, k_n = j$ and for every $\ell \in \{1, ..., n\}$, $k_\ell \in \mathcal{N}(k_{\ell-1})$).

Definition 41 (Path). *A path between two agents $i, j \in Ag$ is a sequence of agents which starts at i, ends at j, and repeats no vertices. Formally, a path between i and j is a tuple $\langle a_1, ..., a_n \rangle$, where $a_1 = i, a_n = j$, for every $k, \ell \in \{1, ..., n\}$, $a_{k+1} \in \mathcal{N}(a_k)$, and if $k \neq \ell$, then $a_k \neq a_\ell$. The length of a path is the number of edges that appear in it, or equivalently one less than the number of vertices.*

An important measure for social networks is how 'close' or 'far' two agents are from each other — we will call this the distance between two agents and define it as the smallest path connecting them. Another measure we will also look at is the distance between the two agents that are farthest from each other — we will call this the diameter of the network and define it as the maximum distance between any two agents.

Definition 42 (Distance, Diameter). *Let $\mathcal{P}_{ij} = \{P \mid P \text{ is a path between } i \text{ and } j\}$. We denote the distance between two agents i and j by $d(i, j) = min_{P \in \mathcal{P}_{ij}}(|P|)$. The diameter of the network is specified by $max_{i \in Ag}(max_{j \in Ag}(d(i, j)))$.*

We might also wonder: given a network and a random agent in that network, how far can we expect them to be from all other agents? This is another measure of network inter-connectivity, which we will call average node distance.

Definition 43 (Average Node Distance)**.** *Let* $Pairs = \{(i,j) \mid i,j \in Ag, (j,i) \notin Pairs\}$*. We define the average node distance of the network as*

$$Avg := \frac{\sum_{(i,j)\in Pairs} d(i,j)}{|Pairs|}$$

Note that the set $Pairs$ is just the set of pairs of agents where order does not matter, or the combination of the set of agents taken 2 at a time, so that $|Pairs| = C^2_{|Ag|} = \frac{|Ag|\cdot(|Ag|-1)}{2}$.

In principle, as long as the constraints of Definition 29 are met, the network in an evidence diffusion model could be arranged in any way. This means that the number of possible networks is very large, even given a relatively small number of agents. For this reason, and for the sake of comparing our formal setting with earlier work that uses these [92, 93, 94], we will restrict our attention to specific classes of networks in what follows.

Notation. We will use $\mathcal{T}$ to denote the network $\mathcal{T} : Ag \mapsto \wp(Ag)$ such that for every $i \in Ag$, $\mathcal{T}(i) = Ag \setminus \{i\}$. We will denote by $\bigstar_i$ the network $\bigstar_i : Ag \mapsto \wp(Ag)$ such that $\bigstar_i(i) = Ag$ and for every $i \neq j \in Ag$, $\bigstar_i(j) = \{i\}$. We will denote by $\bigcirc$ the network $\bigcirc : Ag \mapsto \wp(Ag)$ such that for every i, there exist distinct j,k such that $j,k \in \mathcal{N}(i)$ and for every $\ell \in Ag \setminus \{i,j,k\}, \ell \notin \mathcal{N}(i)$. We will denote by $\mathcal{W}_i$ the network $\mathcal{W}_i : Ag \mapsto \wp(Ag)$ such that for every j, $j \in \mathcal{N}(i)$ and for every j, there exist distinct $k,\ell \in Ag \setminus \{i,j\}$ such that $k,\ell \in \mathcal{N}(j)$ and for every $m \in Ag \setminus \{i,j,k,\ell\}, m \notin \mathcal{N}(j)$.

In Figure 9 we illustrate the networks $\mathcal{T}$, the total network, $\bigstar_i$, the star network, where i is the star, $\bigcirc$, the cycle network, and $\mathcal{W}_i$, the wheel network, for a six agent set.

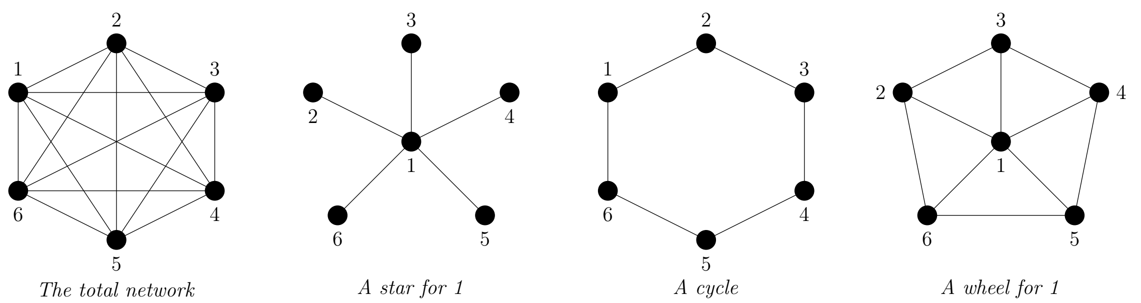

Figure 9: $\mathcal{T}, \bigstar_i, \bigcirc, \mathcal{W}_i$ for $|Ag| = 6$

In the following table we apply the measures defined earlier for the four classes of networks just introduced.

Proposition 5. *For any set of agents with*[5] $|Ag| \geq 4$ *and for any network* $\mathcal{N} : Ag \mapsto \wp(Ag)$ *that satisfies irreflexivity* ($a \notin \mathcal{N}(a)$)*, symmetry* ($a \in \mathcal{N}(b)$ *implies* $b \in \mathcal{N}(a)$) *and connectedness* (*for every* $i \neq j$ *there exist* $k_0, ..., k_n$ *agents such that* $k_0 = i, k_n = j$ *and for every* $\ell \in \{1, ..., n\}$, $k_\ell \in \mathcal{N}(k_{\ell-1})$)*, if* $\mathcal{N}$ *is of the type specified in the first column, then the measures in the next columns are satisfied.*

[5] We impose this constraint because of the way we defined the wheel network — more specifically because we imposed that for any non-central j, there exist $k \neq \ell$ that are also non-central such that $k, \ell \in \mathcal{N}(j)$.

Type	Diameter	Average Node Distance
Total	1	1
Star	2	$\frac{2(\lvert Ag\rvert-1)}{\lvert Ag\rvert}$
Cycle	$\lfloor\frac{\lvert Ag\rvert}{2}\rfloor$	$\frac{\sum_{k=1}^{\lfloor\frac{\lvert Ag\rvert}{2}\rfloor} 2k-((1-rem(\frac{\lvert Ag\rvert}{2}))\cdot\lfloor\frac{\lvert Ag\rvert}{2}\rfloor)}{\lvert Ag\rvert-1}$
Wheel	$\begin{cases} 1 & \text{if } \lvert Ag\rvert = 4 \\ 2 & \text{otherwise} \end{cases}$	$\frac{2\cdot(\lvert Ag\rvert-2)}{\lvert Ag\rvert}$

Proof. For the total network, note that for every $i \in Ag$, $j \in \mathcal{N}(i)$ for every $j \in Ag \setminus \{i\}$. Thus, the shortest path from an agent to another is simply their connection, which has length 1. This immediately entails the fact that the average node distance is also 1.

Now, let us define the average distance between one agent and the other nodes in the network. For any $i \in Ag$, let

$$Avg_i := \frac{\sum_{j\in Ag\setminus\{i\}} d(i,j)}{|Ag| - 1}$$

Observe the following equalities:

$$\begin{aligned} Avg &= \frac{\sum_{(i,j)\in Pairs} d(i,j)}{|Pairs|} \\ &= \frac{2\cdot\sum_{(i,j)\in Pairs} d(i,j)}{2\cdot|Pairs|} \\ &= \frac{\sum_{(i,j)\in \mathcal{T}} d(i,j)}{|\mathcal{T}|} \\ &= \frac{\sum_{i\in Ag}\sum_{j\in Ag\setminus\{i\}} d(i,j)}{|Ag|\cdot(|Ag|-1)} \\ &= \frac{\sum_{i\in Ag} Avg_i}{|Ag|} \end{aligned}$$

So the average distance between the nodes of the network is equal the average of the average distance between one agent and the other nodes.

For $\bigstar_i$, it is clear that i is at distance 1 from every other agent. For any $j \neq i \in Ag$, and for every other agent $j' \neq i \neq j \in Ag$, we have the path $\langle j, i, j'\rangle$ of length 2. Thus, the diameter of the network is 2. For i, the average distance to another agent is 1. For any $j \neq i$, the average distance is clearly $\frac{1+2\cdot(|Ag|-2)}{|Ag|-1}$. Thus the average distance for this network is

$$\frac{1+(|Ag|-1)\cdot\frac{1+2\cdot(|Ag|-2)}{|Ag|-1}}{|Ag|} = \frac{1+1+2\cdot(|Ag|-2)}{|Ag|} = \frac{2+2\cdot(|Ag|-1)-2}{|Ag|} = \frac{2(|Ag|-1)}{|Ag|}$$

For the diameter of the cycle, observe first that every agent is in the same position with respect to their distance to their network — they have two agents at distance 1, two at distance 2, and so on. Note then that from any agent to another, there are two paths, and that generally if the first path is of length x, the second is of length $|Ag| - x$. It is clear then, that the maximal distance between two agents is obtained by maximising the length of the paths on both sides. So, for even sets of agents, the diameter is $\frac{|Ag|}{2}$. For odd sets of agents, observe that there are two 'furthest' agents, and the paths to either are of length $\lfloor \frac{|Ag|}{2} \rfloor$ and $\lceil \frac{|Ag|}{2} \rceil$. Since the distance between two agents is defined as the shortest path, we have established that the diameter of the cycle in general is $\lfloor \frac{|Ag|}{2} \rfloor$.

As for the average length, as mentioned above, there are two agents at each path length, up to $\lfloor \frac{|Ag|}{2} \rfloor - 1$. At path length $\lfloor \frac{|Ag|}{2} \rfloor - 1$, there is one agent on even networks, and there are two agents on odd networks. Since this is for every agent, this will be the overall average as well. So if we denote by $rem(\frac{x}{y})$ the remainder of the division between x and y, for every $i \in Ag$, we have:

$$Avg = \frac{\sum_{k=1}^{\lfloor \frac{|Ag|}{2} \rfloor} 2k - ((1 - rem(\frac{|Ag|}{2})) \cdot \lfloor \frac{|Ag|}{2} \rfloor)}{|Ag| - 1}$$

For the wheel for i network, if $|Ag| = 4$, then note that the network is total, so the diameter is 1. Otherwise, similarly to the star, the diameter is 2, as every agent can be reached by every other agent in at most 2 steps through i. For the average distance, like before, i's distance from every agent is 1, and so the average distance will also be 1. For every $j \neq i$, there are exactly 3 agents they are connected to, and hence at distance 1, and the rest $|Ag| - 3$ of the agents are reachable through i via a path of length 2. So for any $j \neq i$, $Avg_j = \frac{3+2\cdot(|Ag|-3)}{|Ag|-1}$. And thus,

$$Avg = \frac{1 + (|Ag| - 1) \cdot \frac{3+(|Ag|-3)\cdot 2}{|Ag|-1}}{|Ag|} = \frac{1 + 3 + (|Ag| - 3) \cdot 2}{|Ag|} = \frac{2 \cdot (|Ag| - 2)}{|Ag|}$$

□

Now that these measures are set in place, we can investigate what aspects of the network structure are expressible in our logical language, and what the agents know about the network. By making use of the finiteness of the set of agents and the atom N_{ij} in the language, we can introduce abbreviations that express that the network is one of these four. After each abbreviation, we prove that the propositions correspond exactly to the network structure. The first abbreviation we introduce is that for the total network, which states that for every agent, and for every different agent, they are connected by the network.

Abbreviation. The network is total:

$$tot := \bigwedge_{i \in Ag} \bigwedge_{j \in Ag \setminus \{i\}} N_{ij}.$$

Proposition 6. *For any evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!] \rangle$, $\mathcal{N} = \mathcal{T}$ *if and only if* $\mathfrak{M} \vDash tot$.

Proof. Pick an arbitrary evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!] \rangle$ and an arbitrary world $x \in X$. $\mathcal{N} = \mathcal{T}$ is by definition equivalent to the property that for every $i \in Ag$, $\mathcal{N}(i) = Ag \setminus \{i\}$. This is, again by definition, equivalent to saying that for $i \neq j \in Ag$, $\mathfrak{M}, x \vDash N_{ij}$. Since the set of agents is finite, we can transform this universal quantification into a conjunction, which is exactly what *tot* is. Since we reasoned with equivalences, and started from an arbitrary evidence diffusion model and world, the proposition holds. □

For the star network, the following abbreviation captures the fact that the agent i is connected to every other agent, and that no agent is connected to anybody but i.

Abbreviation. The network is a star for i:

$$star_i := \bigwedge_{j \in Ag \setminus \{i\}} (N_{ij} \wedge \neg \bigvee_{k \in Ag \setminus \{i,j\}} N_{jk}).$$

Proposition 7. *For any evidence diffusion model* $\mathfrak{M} = \langle X,\ \{\mathcal{E}_i\}_{i \in Ag},\ \{\Pi_i\}_{i \in Ag},\ \mathcal{N},\ [\![\cdot]\!] \rangle$, $\mathcal{N} = \bigstar_i$ *if and only if* $\mathfrak{M} \vDash star_i$.

Proof. Let $\mathfrak{M} = \langle X,\ \{\mathcal{E}_i\}_{i \in Ag},\ \{\Pi_i\}_{i \in Ag},\ \mathcal{N},\ [\![\cdot]\!] \rangle$ be an arbitrary evidence diffusion model and x an arbitrary world. Note that the set of agents is finite, so a conjunction spanning all agents acts as universal quantification, and disjunction as an existential quantification. Thus $\mathfrak{M}, x \vDash star_i$ if and only if for all agents j except i, $j \in \mathcal{N}(i)$ and it is not the case that there exists some k other than i, j such that $k \in \mathcal{N}(j)$. In other words, i is connected to every agent, and every other agent is connected only to i. This is precisely the definition of the star network. Since we worked with equivalences and used an arbitrary world and model, our proposition is proven. □

To obtain the cycle, given our already existing constraints, we only need to ensure that every agent is connected to exactly two other agents.

Abbreviation.

$$cycle := \bigwedge_{i \in Ag} \bigvee_{j,k \in Ag \setminus \{i\}}^{j \neq k} (N_{ij} \wedge N_{ik} \wedge \bigwedge_{\ell \in Ag \setminus \{i,j,k\}} \neg N_{i\ell})$$

Proposition 8. *For any evidence diffusion model* $\mathfrak{M} = \langle X,\ \{\mathcal{E}_i\}_{i \in Ag},\ \{\Pi_i\}_{i \in Ag},\ \mathcal{N},\ [\![\cdot]\!] \rangle$, $\mathcal{N} = \bigcirc$ *if and only if* $\mathfrak{M} \vDash cycle$.

Proof. The proof is analogous to the previous proof. □

To obtain the wheel for i, what is required is that i is connected to every other agent, and that every agent other than i is connected to exactly three agents.

Abbreviation. The network is a wheel for i:

$$wheel_i := \bigwedge_{j \in Ag \setminus \{i\}} (N_{ij} \wedge \bigvee_{k,\ell \in Ag \setminus \{i,j\}}^{k \neq \ell} (N_{jk} \wedge N_{j\ell} \wedge \bigwedge_{m \in Ag \setminus \{i,j,k,\ell\}} \neg N_{jm}))$$

Proposition 9. *For any evidence diffusion model* $\mathfrak{M} = \langle X,\ \{\mathcal{E}_i\}_{i \in Ag},\ \{\Pi_i\}_{i \in Ag},\ \mathcal{N},\ [\![\cdot]\!] \rangle$, $\mathcal{N} = \mathcal{W}_i$ *if and only if* $\mathfrak{M} \vDash wheel_i$.

Proof. The proof is analogous to the previous proof. □

As we have stressed earlier, and as we will go on to prove, the network topology will affect the epistemic and doxastic attitudes in a major way. As such, a natural question to ask is — do the agents know the network? As it turns out, they do. The following propositions show that the agents have both infallible knowledge, defeasible knowledge, and belief about the network.

Proposition 10. *Every agent infallibly knows the network. Formally, for every* $i, j, k \in Ag$, $\vDash N_{jk} \leftrightarrow [\Pi_i]N_{jk}$.

Proof. Take an arbitrary threshold θ, an arbitrary model $\mathfrak{M}$, world x and agents i, j, k. For the left to right direction, suppose $\mathfrak{M}, x \vDash N_{jk}$. Note that the truth clause for the network atoms make no reference to the world of evaluation, so $[\![N_{jk}]\!] = X$, which immediately means that for all $y \in \Pi_i(x)$, $\mathfrak{M}, y \vDash N_{jk}$, giving $\mathfrak{M}, x \vDash [\Pi_i]N_{jk}$. For the right to left direction, suppose $\mathfrak{M}, x \vDash [\Pi_i]N_{jk}$. This means that every world in $\Pi_i(x)$ makes N_{jk} true, and since $x \in \Pi_i(x)$, $\mathfrak{M}, x \vDash N_{jk}$. Since $\theta, \mathfrak{M}, x$ and i, j, k were picked arbitrarily, our proposition holds. □

The same result easily obtains for defeasible knowledge, belief, and the evidence modalities. The first is formulated in Proposition 11, whereas the latter are left out, since the proofs are analogous to the given proof.

Proposition 11. *Every agent defeasibly knows the network. Formally, for every* $i, j, k \in Ag$, $\vDash N_{jk} \leftrightarrow K_i N_{jk}$.

Proof. Take an arbitrary threshold θ, an arbitrary model $\mathfrak{M}$, world x and agents i, j, k. For the left to right direction, suppose $\mathfrak{M}, x \vDash N_{jk}$, meaning just like above that $[\![N_{jk}]\!] = X$. Clearly, $\Pi_i(x)$ is dense in $\Pi_i(x)$ and $\Pi_i(x) \subseteq X$, so $\mathfrak{M}, x \vDash K_i N_{jk}$. Now for the right to left direction, suppose $\mathfrak{M}, x \vDash K_i N_{jk}$. This means that $[\![N_{jk}]\!]$ contains a $\Pi_i(x)$ dense open set, which means $[\![N_{jk}]\!] \neq \emptyset$, immediately entailing $[\![N_{jk}]\!] = X$ and $\mathfrak{M}, x \vDash N_{jk}$. Since $\theta, \mathfrak{M}, x$ and i, j, k were picked arbitrarily, the proposition holds. □

4.2 Diffusion

We now turn to a formal investigation of evidence diffusion, in the general sense. Diffusion is, once again, defined as an operation on models. We are interested in the invariance of the epistemic (and doxastic) operators we have defined under this operation, both from a mathematical point of view, and from a philosophical point of view. Mathematically, the invariance of the modalities under an operation is revelatory for the functioning of both the former and the latter. Philosohpically, whether or not individual or group knowledge is preserved under diffusion gives insights into evidence-based reasoning in a dynamic environment where agents send and receive pieces of evidence.

First, let us establish a basic fact about diffusion sequences. Just like in [12], any diffusion sequence reaches a fixed point bounded by the number of agents. The fact that multiple evidence pieces spread does not affect this property, as the pieces of evidence do not interact with each other, as far as diffusion is concerned.

Notation. For an evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!] \rangle$, we let $Ag_e = \{i \in Ag \mid e \in \mathcal{E}_i\}$. For any $k \in \mathbb{N}$, given the diffusion sequence $\mathcal{S}_{\mathfrak{M}} = \langle \mathfrak{M}, \mathfrak{M}_1, ... \rangle$, we let $Ag_e^k = \{i \in Ag \mid e \in \mathcal{E}_i^k\}$, where $\mathcal{E}_i^k \in \mathfrak{M}_k$.

Proposition 12. *For any evidence diffusion model* $\mathfrak{M}$*, there exists some* $n \in \mathbb{N} \leq |Ag| - 1$ *where we reach a fixed point in* $\mathcal{S}_{\mathfrak{M}}$*, i.e. for any* $\mathfrak{M}_m \in \mathcal{S}_{\mathfrak{M}}$ *with* $m \geq n$, $\mathfrak{M}_m = \mathfrak{M}_n$*. For clarity, we will often denote the upper bound* $|Ag| - 1$ *by* fix *(for fixpoint).*

Proof. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!] \rangle$ be an arbitrary evidence diffusion model. By definition, for any $i \in Ag$ and any $k \in \mathbb{N}$, $\mathcal{E}_i^k \subseteq \mathcal{E}_i^{k+1}$. Thus, for any $e \in \bigcup_{i \in Ag} \mathcal{E}_i$, $Ag_e^k \subseteq Ag_e^{k+1} \subseteq Ag$. Note that the set of agents Ag is finite, and so for every $e \in \bigcup_{i \in Ag} \mathcal{E}_i$, there exists $n \in \mathbb{N}$ such that $Ag_e^n = Ag_e^{n+1}$. Clearly, the slowest diffusion scenario for some $e \in \bigcup_{i \in Ag} \mathcal{E}_i$ is that in which $|Ag_e| = 1$ and for every $k \in \mathbb{N}$, $|Ag_e^{k+1}| = |Ag_e^k| + 1$, which is clearly bounded by $|Ag| - 1$. This concludes the proof. $\square$

We define the following abbreviation for iterations of the dynamic diffusion operator, which will be convenient as we begin to prove properties about specific points in the diffusion sequence.

Abbreviation. We define the formula $[\because]^n\varphi$ as the abbreviation defined recursively by

$$[\because]^0\varphi := \varphi;$$
$$[\because]^{n+1}\varphi := [\because][\because]^n\varphi.$$

Since all of the epistemic and doxastic attitudes we have defined are evidence-based, and since agents receive pieces of evidence through the diffusion process, it will be useful to have a short-hand way of referring to when a piece of evidence will be received by an agent at the next round. Furthermore, we define when a piece of evidence will be received by an agent at *some* round.

Definition 44 (Freeness of Evidence)**.** *We call a piece of evidence $e \in \bigcup_{i \in Ag} \mathcal{E}_i$ free for i if and only if there exists $\mathcal{G} \subseteq \mathcal{N}(i)$ such that $e \in \mathcal{E}_j$ for every $j \in \mathcal{G}$ and $\frac{|\mathcal{G}|}{|\mathcal{N}(i)|} \geq \theta$. We say that e is free for i at some point in the diffusion sequence if and only if there exists some point $k \in \mathbb{N}$ in the diffusion sequence $\mathcal{S}_{\mathfrak{M}} = \langle \mathfrak{M}, \mathfrak{M}_1, ... \rangle$ such that there exists $\mathcal{G} \subseteq \mathcal{N}(i)$ with $e \in \mathcal{E}_i^k$ for every $j \in \mathcal{G}$ and $\frac{|\mathcal{G}|}{|\mathcal{N}(i)|} \geq \theta$.*

As in [12], and as shown in Section 3.4, we can define an abbreviation for when a certain piece of evidence is free for an agent at the next round (and with the help of $[\because]$, at further rounds as well). We make essential use of the E_i modality for this purpose. The abbreviation states: there exists a subset of the group of agents, F, and a subset of F, G, such that $\frac{|G|}{|F|} \geq \theta$ and such that F contains exactly the agents that are in i's network, and every agent in G has $[\![\varphi]\!]$ in their basic evidence set.

Abbreviation. We define the abbreviation $\phi_i\varphi$, read 'φ is ɸree (free) for i', as

$$\phi_i\varphi := \bigvee_{\{G \subseteq F \subseteq Ag \mid \frac{|G|}{|F|} \geq \theta\}} \left(\bigwedge_{j \in F} N_{ij} \wedge \bigwedge_{j \notin F} \neg N_{ij} \wedge \bigwedge_{j \in G} E_j\varphi\right).$$

We can also express the fact that $[\![\varphi]\!]$ is free at some point in the diffusion sequence for i: $\bigvee_{k \in \{1,\ldots,fix\}} [\because]^{k-1}\phi_i\varphi$. The negation of this disjunction says that $[\![\varphi]\!]$ is *never* free for i, so that i will never have this piece of evidence in their set of basic pieces of evidence. Note that they might still have this piece of evidence in their topology of evidence by some finite intersection or arbitrary union.

The following proposition states that an agent has a piece of evidence $[\![\varphi]\!]$ in their set of basic pieces of evidence at the next round if and only if they had it at this round, or it is free for i.

Proposition 13. *For any $i \in Ag$, $\models [\because]E_i\varphi \leftrightarrow \phi_i\varphi \vee E_i\varphi$.*

Proof. Let $\mathfrak{M} = \langle X,\ \{\mathcal{E}_i\}_{i\in Ag},\ \{\Pi_i\}_{i\in Ag},\ \mathcal{N},\ [\![\cdot]\!]\rangle$ be an arbitrary diffusion model, $x \in X$ an arbitrarily picked world and $i \in Ag$ an arbitrary agent. $\mathfrak{M}, x \vDash [\because]E_i\varphi$ means that $\mathfrak{M}^{\because}, x \vDash E_i\varphi$, which by definition means that $[\![\varphi]\!] \in \mathcal{E}_i^{\because}$. Also by definition, either $[\![\varphi]\!] \in \mathcal{E}_i$, in which case $\mathfrak{M}, x \vDash E_i\varphi$ and by extension $\mathfrak{M}, x \vDash \phi_i\varphi \vee E_i\varphi$, or $[\![\varphi]\!]$ was received via evidence diffusion, in which case there must exist $\mathcal{G} \subseteq \mathcal{N}(i)$ such that $\frac{|\mathcal{G}|}{|\mathcal{N}(i)|} \geq \theta$ and for every $j \in \mathcal{G}$, $[\![\varphi]\!] \in \mathcal{E}_j$, which entails by definition that $\mathfrak{M}, x \vDash \phi_i\varphi$, and by extension $\mathfrak{M}, x \vDash \phi_i\varphi \vee E_i\varphi$. Since we reasoned with equivalences, we proved both directions at once, and so $\mathfrak{M}, x \vDash [\because]E_i\varphi \leftrightarrow \phi_i\varphi \vee E_i\varphi$, and since the agent, world, and model were picked arbitrarily, we obtain the desired result. □

An important fact to point out is that, unlike in [12], where the atom β could be used to fully specify the set of adopting agents B, we cannot always fully specify an agent's set of basic pieces of evidence using the modality E_i. On some models, there are pieces of evidence in some agent's set of evidence $e \in \mathcal{E}_i$ that are not captured by the valuation, so that there is no $\varphi \in \mathcal{L}_d$ with $[\![\varphi]\!] = e$.

We begin our analysis of the modalities by proving a statement that was made in the previous chapter. Namely, that defeasible evidence is the knowledge of a theoretical agent that has all of the hard knowledge ('laws') and soft knowledge (evidence) of all of the agents combined. Formally, we prove that distributed evidence holds of a proposition if and only if that proposition is known at the end of the diffusion sequence by an agent who is set to receive all available pieces of evidence and all hard information.

Proposition 14. *For any evidence diffusion model* $\mathfrak{M} = \langle X,\ \{\mathcal{E}_i\}_{i\in Ag},\ \{\Pi_i\}_{i\in Ag},\ \mathcal{N},\ [\![\cdot]\!]\rangle$, *if for every* $e \in \bigcup_{j\in Ag}\mathcal{E}_j$, *e is free for i at some point in the diffusion sequence* $\mathcal{S}_{\mathfrak{M}}$, *then* $\mathfrak{M} \vDash \mathfrak{D}\varphi \leftrightarrow [\bigwedge][\because]^{fix}K_i\varphi$.

Proof. Let $\mathfrak{M} = \langle X,\ \{\mathcal{E}_i\}_{i\in Ag},\ \{\Pi_i\}_{i\in Ag},\ \mathcal{N},\ [\![\cdot]\!]\rangle$ be an arbitrarily picked evidence diffusion model, x an arbitrarily picked world and suppose that for every $e \in \bigcup_{j\in Ag}\mathcal{E}_j$, e is free for $i \in Ag$ at some point in the diffusion sequence. Now we want to show $\mathfrak{M}, x \vDash \mathfrak{D}\varphi \leftrightarrow [\bigwedge][\because]^{fix}K_i\varphi$.

We proceed via equivalences and prove both directions at the same time. $\mathfrak{M}, x \vDash \mathfrak{D}\varphi$ means that for every $i \in Ag$, there exists $U_i \in \tau_i$ such that: 1) $\mathrm{Cl}_{\vee\tau_\cap}(\bigcap_{i\in Ag} U_i) = \bigcap_{i\in Ag}\Pi_i(x)$, where $\vee\tau_\cap = \bigvee_{i\in Ag}\tau_i\Big|_{\bigcap_{i\in Ag}\Pi_i(x)}$; 2) $x \in \bigcap_{i\in Ag} U_i$; 3) $\bigcap_{i\in Ag} U_i \subseteq [\![\varphi]\!]$. Now let us look at the model $\mathfrak{M}^{\bigwedge fix}$ (the fixpoint of $\mathfrak{M}^{\bigwedge}$'s diffusion sequence). By supposition, $\mathcal{E}_i^{\bigwedge fix} = \bigcup_{j\in Ag}\mathcal{E}_j$. Note that the closure of $\bigcup_{j\in Ag}\mathcal{E}_j$ under finite intersections and arbitrary unions (i.e. $\tau_i^{\bigwedge fix}$) contains both the closure under finite intersections and arbitrary unions of sets *within* the sets of basic evidence for each agent, but also *between* the basic evidence sets of different agents. Thus, this closure contains just the same sets as the closure of $\bigcup_{j\in Ag}\tau_j$ under arbitrary unions and finite intersections (i.e. $\bigvee_{j\in Ag}\tau_j$ by definition). As such, $\tau_i^{\bigwedge fix} = \bigvee_{j\in Ag}\tau_j$. This means that all the sets which the definition $\mathfrak{D}$ establishes the existence of are in i's topology. Since topologies are closed under finite intersections, and there are finitely many agents, the intersection of these sets $\bigcap_{j\in Ag} U_j \in \tau_i^{\bigwedge fix}$. Furthermore, since this is a model obtained after law distribution, we have $\Pi_i^{\bigwedge fix}(x) = \bigcap_{j\in Ag}\Pi_i(x)$. By 1) we get that $\bigcap_{j\in Ag} U_j$ is $\Pi_i^{\bigwedge fix}(x)$-locally dense, by

2) we get that it is factive, and by 3) that it is a subset of $[\![\varphi]\!]$. Thus, $\mathfrak{M}^{\bigwedge fix}, x \vDash K_i\varphi$, which entails $\mathfrak{M}^{\bigwedge}, x \vDash [\because]^{fix} K_i\varphi$, which in turn yields $\mathfrak{M}, x \vDash [\bigwedge][\because]^{fix} K_i\varphi$.

Since $\mathfrak{M}$ and x were chosen arbitrarily, the proposition is established. □

Note that this agent need not exist for distributed evidence to hold of a proposition, it is the knowledge of a *hypothetical* agent that is in possession of all hard and soft information. As we shall see, the fact that distributed evidence makes a hypothetical statement as opposed to a statement about the current state of affairs is important, as it is the critical factor that makes it *robust*, i.e. invariant under diffusion, as we show in the next proposition.

Proposition 15. $\vDash \mathcal{E}D\varphi \leftrightarrow [\because]\mathcal{E}D\varphi$

Proof. Let $\theta \in [0,1]$ be arbitrarily given and let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ and $x \in X$ be arbitrary as well. For the left to right direction, suppose $\mathfrak{M}, x \vDash \mathcal{E}D\varphi$ for some arbitrary φ. We want to show $\mathfrak{M}, x \vDash [\because]\mathcal{E}D\varphi$, which, by the definition of $[\because]$, is equivalent to showing that $\mathfrak{M}^{\because}, x \vDash \mathcal{E}D\varphi$. By definition, because $\mathfrak{M}, x \vDash \mathcal{E}D\varphi$, we have that for every $i \in Ag$ there exists $U_i \in \tau_i$ such that:

1. $\bigcap_{i\in Ag} U_i \subseteq [\![\varphi]\!]$. But by the definition of $\mathfrak{M}^{\because}$, we know that for every $i \in Ag$, $\tau_i \subseteq \tau_i^{\because}$, the topology generated by the subbasis $\mathcal{E}_i^{\because}$, so that clearly for all $i \in Ag$ there exists $U_i \in \tau_i^{\because}$ such that $\bigcap_{i\in Ag} U_i \subseteq [\![\varphi]\!]$;
2. $x \in \bigcap_{i\in Ag} U_i$ similarly entails that for every $i \in Ag$ there exists $U_i \in \tau_i^{\because}$ such that $x \in \bigcap_{i\in Ag} U_i$;
3. $\mathrm{Cl}_{\vee\tau_\cap}(\bigcap U_i) = \bigcap \Pi_i(x)$, where $\vee\tau_\cap = \bigvee_{i\in Ag} \tau_i\Big|_{\bigcap_{i\in Ag}\Pi_i(x)}$. But since $\bigvee_{i\in Ag}\tau_i$ is the topology generated by $\bigcup_{i\in Ag}\tau_i$ and by Definition 33, updates always introduce pieces of evidence that were already in some agent's basic evidence set, $\bigvee_{i\in Ag}\tau_i^{\because} = \bigvee_{i\in Ag}\tau_i$. Also, the evidence update rule leaves the partitions unchanged. So, $\bigvee_{i\in Ag}\tau_i^{\because}\Big|_{\bigcap_{i\in Ag}\Pi_i(x)} = \vee\tau_\cap^{\because} = \vee\tau_\cap$ and thus, for every $i \in Ag$, there exists $U_i \in \tau_i^{\because}$ such that $\mathrm{Cl}_{\vee\tau_\cap^{\because}}(\bigcap_{i\in Ag} U_i) = \bigcap_{i\in Ag}\Pi_i(x)$.

So $\mathfrak{M}^{\because}, x \vDash \mathcal{E}D\varphi$, which entails that $\mathfrak{M}, x \vDash [\because]\mathcal{E}D\varphi$. Thus, $\mathfrak{M}, x \vDash \mathcal{E}D\varphi \to [\because]\mathcal{E}D\varphi$.

For the right to left direction, similar arguments pertaining to the fact that diffusion does not affect points 1,2, and 3 of the definition of $\mathcal{E}D$ can be made. Thus, we infer that $\mathfrak{M}, x \vDash [\because]\mathcal{E}D\varphi \to \mathcal{E}D\varphi$. This means that $\mathfrak{M}, x \vDash \mathcal{E}D\varphi \leftrightarrow [\because]\mathcal{E}D\varphi$.

And since we picked an arbitrary world, model and threshold, we obtain that $\mathcal{E}D\varphi \leftrightarrow [\because]\mathcal{E}D\varphi$ is valid on evidence diffusion models.

□

By contrast, distributed knowledge, which makes a statement about the current knowledge of the agents, is not robust, as the next proposition indicates. Recall that a proposition is distributed knowledge if the agents have factive justifications that together entail

φ. Since some agent's contribution to this factive justification might get defeated by incoming evidence, i.e. lose density, it stands to reason that distributed knowledge might also be lost in this way. This is what the counterexample exploits.

Proposition 16. $\not\models \mathbb{D}\varphi \to [\because]\mathbb{D}\varphi$

Proof. We prove the proposition by providing a counterexample. We require a threshold θ, a proposition φ, a model and a world at which $\mathbb{D}\varphi$ is true, but at which $[\because]\mathbb{D}\varphi$ is false. Consider, for threshold $\theta = 0.5$, the model in Figure 10, pictured on the left — call this model $\mathfrak{M}$. Formally, we have two agents, $Ag = \{\text{Alice}, \text{Bob}\}$, the set of worlds $X = \{1, 2, 3, 4\}$, and we consider two propositions, $p, q \in \mathbb{P}$ with $[\![p]\!] = P = \{1, 2\}$ (green in Figure 10) and $[\![q]\!] = Q = \{1, 3\}$ (blue in 10) and the rest of the propositions being false at all worlds. Alice's partition is the same as Bob's: $\Pi_{\text{Alice}} = \Pi_{\text{Bob}} = \{X\}$. Alice's evidence set is $\mathcal{E}_{\text{Alice}} = \{X, P, Q\}$ and Bob's evidence set is $\mathcal{E}_{\text{Bob}} = \{X, P, \neg(P \leftrightarrow Q)\}$.

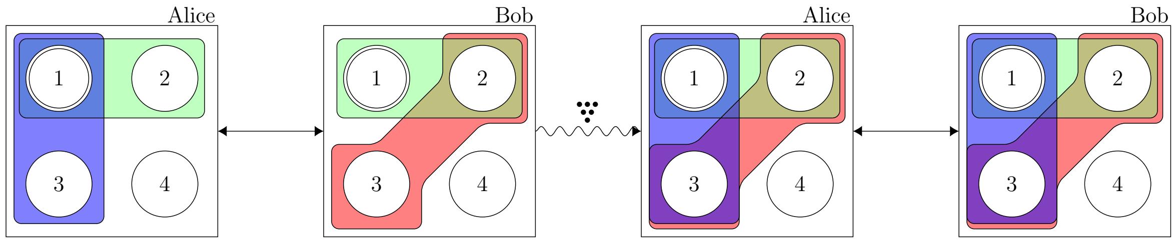

Figure 10: Distributed knowledge loss

$\mathfrak{M}, 1 \models \mathbb{D}p$, since: a) P is $\Pi_{\text{Alice}}(1)$-locally dense (X-locally dense) by virtue of having non-empty intersection with all opens in τ_{Alice} and $\Pi_{\text{Bob}}(1)$-locally dense (X locally dense) by virtue of having non-empty intersection with all opens in τ_{Bob}; b) $x \in P$, as defined; c) $P \cap P \subseteq P$, evidently.

$\mathfrak{M}, 1 \not\models [\because]\mathbb{D}p$. By the definition of $[\because]$, this is equivalent to showing that $\mathfrak{M}^{\because}, 1 \not\models \mathbb{D}p$. This is indeed the case, because for neither of the agents do we have a $\Pi_{\text{Alice}}(1)$-locally dense (X-locally dense) set. After the agents put together their evidence the set $\{3\} = \{1, 3\} \cap \{2, 3\}$ is part of both of their topologies, and so any dense set will have to have non-empty intersection with this set, which is not the case.

This shows that $\mathbb{D}\varphi \to [\because]\mathbb{D}\varphi$ is not valid on evidence diffusion models. □

The next proposition shows that distributed knowledge is not robust under strong premises either — even if all the agents already know φ individually, it is possible that after a round of diffusion, they don't even know φ together! Note that this is even stronger than saying that now none of them know φ, for if one of them knew φ, it would automatically become distributed knowledge as well. What this shows is that distributed knowledge is extremely fragile. As before, this is because they might have sets of evidence that are inconsistent with each other (which is what happens in the counterexample).

Proposition 17. $\not\models (\mathbb{D}\varphi \wedge \bigwedge_{i \in Ag} K_i\varphi) \to [\because]\mathbb{D}\varphi$

Proof. Consult the counterexample in Figure 10 and observe that P is dense in τ_{Alice} in the initial model, $\mathfrak{M}$, and thus $\mathfrak{M}, 1 \models K_{\text{Alice}}p$, and likewise for Bob, $\mathfrak{M}, 1 \models K_{\text{Bob}}p$. So we have that $\mathfrak{M}, 1 \models \bigwedge_{i \in Ag} K_i p$, and thus $\mathfrak{M}, 1 \models (\mathbb{D}p \wedge \bigwedge_{i \in Ag} K_i\varphi)$, yet, as discussed in the previous proof, $\mathfrak{M}, 1 \not\models [\because]\mathbb{D}p$. Thus, $(\mathbb{D}\varphi \wedge \bigwedge_{i \in Ag} K_i\varphi) \to [\because]\mathbb{D}\varphi$ is not valid on evidence diffusion models. □

A natural question at this point would be: what are the conditions under which distributed knowledge *is* preserved? And what precisely does it mean for the basic evidence sets of two agents to be inconsistent? For these reasons, we introduce the notion of misleading evidence, customary to the literature on topological evidence models [69, 9]. The notion introduced here is, of course, indexed by the agent, so that a piece of evidence can be misleading for an agent i, but not for another agent, j. Intuitively, a piece of evidence e is misleading for agent i in model $\mathfrak{M}$ when in the model $\mathfrak{M}^{+_{\{i\}}e}$, i.e. when i receives piece of evidence e in their basic evidence set, a false piece of evidence is created by intersection with some piece of evidence that i already had.

Definition 45. *We call a piece of evidence e i-misleading at $x \in X$ if and only if there is some $e' \in \widetilde{\mathcal{E}}_i$ such that $x \notin (e \cap e')$ and $(e \cap e') \notin \widetilde{\mathcal{E}}_i \cup \{\emptyset\}$.*

The condition that no agent i receives an i-misleading piece of evidence is sufficient to ensure that distributed knowledge is preserved at the next diffusion round, as the next proposition shows.

Proposition 18. *If $\mathfrak{M}, x \models \mathbb{D}\varphi$ and for every i, there is no i-misleading piece of evidence that is free for i, then $\mathfrak{M}, x \models [\because]\mathbb{D}\varphi$.*

Proof. We suppose the antecedent and aim to show the consequent. Since $\mathfrak{M}, x \models \mathbb{D}\varphi$, for every $i \in Ag$ there exists $U_i \in \tau_i$ such that: 1) U_i is $\Pi_i(x)$-locally dense; 2) $x \in \bigcap U_i$; 3) $\bigcap U_i \subseteq [\![\varphi]\!]$. From 2) and 3) we obtain that $x \in [\![\varphi]\!]$. By 1), for every $i \in Ag$, U_i is $\Pi_i(x)$-locally dense, meaning that for every $V \in \tau_i \setminus \{\emptyset\}$, $V \cap \Pi_i(x) \neq \emptyset$ entails $V \cap U_i \neq \emptyset$. From the fact that no piece of evidence that is i-misleading is free for i, we deduce the fact that for every $i \in Ag$ and for every $e \in \widetilde{\mathcal{E}}_i^{\because}$, $e \notin \widetilde{\mathcal{E}}_i$ entails $x \in e$. This immediately means that the U_i that were invoked for $\mathfrak{M}, x \models \mathbb{D}\varphi$ are still dense in $\tau_i^{\because}$, since for every $e \in \widetilde{\mathcal{E}}_i^{\because}$, $x \in (U_i \cap e)$. Thus we can use the same sets as before and claim that $\mathfrak{M}^{\because}, x \models \mathbb{D}\varphi$. □

This condition that no agent i receives an i-misleading piece of evidence is, however, not necessary for the preservation of distributed knowledge. Namely, there are models and diffusion sequences where some agent receives an i-misleading piece of evidence, yet still retains their knowledge. Consider the example below.

Example 5. In the initial model (Figure 11, below), Alice's strongest knowledge is $1 \vee 2$. In the diffusion sequence, she receives the piece of evidence $\{2\}$, which does not contain the actual world, 1, so it is misleading. Yet, since $\{2\}$ is a subset of $\{1, 2\}$ Alice does not lose her strongest piece of knowledge.

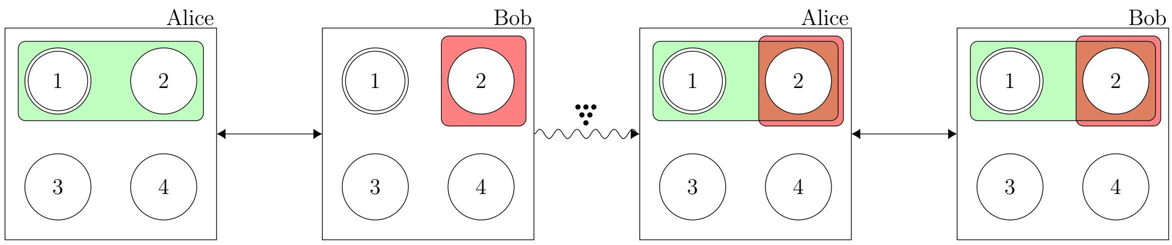

Figure 11: Alice receives misleading evidence, but still knows $1 \vee 2$

In order to find necessary and sufficient conditions for the preservation of these operators under diffusion, we must take into account the critical factors contributing to the loss of knowledge that we have been discussing. Both individual and distributed knowledge require the existence of factive justifications, i.e. dense sets containing the actual world. Many of our counterexamples have relied on newly received pieces of evidence breaking the density of all the agent's justifications, leading to knowledge loss. Thus, in order to find the conditions for invariance, this situation must be prevented. Tools from network analysis give us exactly the means to do this, and we turn to this next.

4.3 Diffusion on Networks in General

Before we proceed, recall the notion of Cluster of density d outlined in Definition 28. The definition neatly transfers to the networks that underlie evidence diffusion models. Furthermore, since we have the atoms N_{ij} and the modality E_i, we can define parallel abbreviations to the ones in Section 3.4, with the same caveat as before — the modality E_i can only refer to pieces of evidence expressible by propositions. The parallel extends to a more significant result, which we present now — a cluster theorem for pieces of evidence.

Theorem 2 (Evidence Cluster Theorem)**.** For every $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ with $\theta \neq 0$, for every piece of evidence $e \in \bigcup_{i\in Ag} \mathcal{E}_i$, all agents will eventually adopt e if and only if there does not exist a cluster of density greater than $1-\theta$ in $Ag \setminus Ag_e$.

Proof. We follow the same steps required to prove the original cluster theorem [64, 29]. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be an arbitrary evidence diffusion model with $\theta \neq 0$ and let $e \in \bigcup_{i\in Ag} \mathcal{E}_i$ be an arbitrarily picked piece of evidence.

($\Rightarrow$) We proceed contrapositively by supposing that there exists a cluster C of density greater than $1-\theta$ in $Ag \setminus Ag_e$ and showing that no agent in C will adopt e. Assume for a contradiction that some $i \in C$ does adopt e, and without loss of generality let us say he is (among) the first adopter(s) from C. Then there exists a point k in the diffusion sequence $\mathcal{S}_{\mathfrak{M}}$ such that $e \in \mathcal{E}_i^k \in \mathfrak{M}_k$ but $e \notin \mathcal{E}_i^{k-1} \in \mathfrak{M}_{k-1}$. In order to adopt e at k, we must have $\frac{|\mathcal{N}(i)\cap Ag_e^{k-1}|}{|\mathcal{N}(i)|} \geq \theta$. Since $C \cap Ag_e^{k-1} = \emptyset$, and since C has density greater than $1-\theta$, i.e. $\frac{|\mathcal{N}(i)\cap C|}{|\mathcal{N}(i)|} > 1-\theta$, it must be that $\frac{|\mathcal{N}(i)\cap Ag_e^{k-1}|}{|\mathcal{N}(i)|} < \theta$, contradiction.

As such, we have established that all agents eventually adopting e entails that there is no cluster of density greater than $1-\theta$ in $Ag \setminus Ag_e$.

($\Leftarrow$) We also prove this direction by contraposition, so suppose that there exists an $i \in Ag$ that has not adopted, i.e. for which $e \notin \mathcal{E}_i^{fix}$. We aim to show that the set $C = Ag \setminus Ag_e^{fix}$ is a cluster of density greater than $1-\theta$. For this, take an arbitrary $j \in C$. Since $j \in C$, he is part of the non-adopting agents, and thus $\frac{|\mathcal{N}(j)\cap Ag_e^{fix}|}{|\mathcal{N}(j)|} < \theta$. This entails that $\frac{|\mathcal{N}(j)\cap C|}{|\mathcal{N}(j)|} > 1-\theta$. And since j was picked arbitrarily, this means that for all $j \in C$, $\frac{|\mathcal{N}(j)\cap C|}{|\mathcal{N}(j)|} > 1-\theta$, which is equivalent to C being a cluster of density greater than $1-\theta$.

So, we have established if there exists no cluster of density greater than $1-\theta$ in $Ag \setminus Ag_e$, then all agents will eventually adopt e.

And since e was picked arbitrarily, this holds for all $e \in \bigcup_{i\in Ag} \mathcal{E}_i$. And since $\mathfrak{M}$ was also picked arbitrarily, this holds for any evidence diffusion model with $\theta \neq 0$. The result is hence established. □

This may not be very surprising — after all, there is not much of a difference between threshold models (with multiple non-interacting behaviors) and our pieces of evidence.

Furthermore, whether or not a full cascade forms for a piece of evidence might not immediately seem very interesting from an epistemological point of view. The following Corollary shows that, in fact, this Theorem gives quite a bit of insight into the epistemic proclivities of the agents in this model.

Corollary 1. For every $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ with $\theta \neq 0$, for every piece of evidence $e \in \bigcup_{i\in Ag} \mathcal{E}_i$, for every agent $i \in Ag$, $e \in \mathcal{E}_i^{fix}$ if and only if for every cluster C of density greater than $1-\theta$ with $i \in C$, there exists an agent $j \in C$ such that $e \in \mathcal{E}_j$.

Proof. Take an arbitrary evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$, an arbitrary piece of evidence $e \in \bigcup_{i\in Ag} \mathcal{E}_i$ and an arbitrary agent $i \in Ag$.

For the left to right direction, we go contrapositively and suppose there exists a cluster of density greater than $1-\theta$ containing i where for all $j \in C, e \notin \mathcal{E}_j$. We assume for a contradiction that $e \in \mathcal{E}_i^{fix}$. At some point k in the diffusion sequence $\mathcal{S}_{\mathfrak{M}}$, we must have had $\frac{|\mathcal{N}(i)\cap Ag_e^k|}{|\mathcal{N}(i)|} \geq \theta$. Now we reason by cases. In the first case, $\mathcal{N}(i)\cap Ag_e^k \cap C = \emptyset$, meaning that all adopting agents in $\mathcal{N}(i)$ are outside of the cluster. But then i cannot have adopted e, since, as explained in the proof of the theorem, it must be that $\frac{|\mathcal{N}(i)\cap Ag_e^k|}{|\mathcal{N}(i)|} < \theta$, so this is a contradiction. In the second case, $\mathcal{N}(i) \cap Ag_e^k \cap C \neq \emptyset$. But there must have been $\ell \leq k$ such that $C \cap Ag_e^\ell \neq \emptyset$, but $C \cap Ag_e^{\ell-1} = \emptyset$, i.e. there must have been a point in the diffusion sequence where some agent(s) in C adopt e, where none have done so before. In this case, reason analogously to obtain a contradiction. Since we obtained a contradiction in both cases, we have, contrapositively, the left to right direction of this corollary.

For the right to left direction, we again aim to go contrapositively, and proceed exactly as in the right to left direction of the evidence cluster theorem. Observe that i is part of the set C constructed in that proof and that for all $j \in C$, $e \notin \mathcal{E}_j$, and that thus, C is exactly the set required for this proof as well.

Since we worked with an arbitrary model, piece of evidence, and agent, the corollary is established. □

So, an agent will end up adopting a piece of evidence if and only if there is an 'early adopter' [29, Ch. 19] of that piece of evidence in *every* Cluster of density greater than $1-\theta$ that they are a part of. If we take cluster density as a measure of group cohesion, we can rephrase this as saying: agents refrain from adopting pieces of evidence if they have any contact group that is 'close enough' in which none of the members have adopted that piece of evidence. Put in these terms, it certainly seems like the process of adopting pieces of evidence from diffusion is not a rational process, as we would like to think that scientists have the ability to individually and independently evaluate groundbreaking science when they come across it. In many real cases of collective failure such as the case of Ignaz Semmelweis outlined in Chapter 2 or the case of the Einstein-de Haas effect [85], though, it certainly seems like precisely this kind of 'groupthink' is at play.

This Corollary not only provides the essential ingredient to establishing the necessary and sufficient conditions for individual and distributed knowledge to be preserved in diffusion, but also to be attained at the end of a diffusion sequence. We now go on to prove theorems to this extent, the first pertaining to individual knowledge.

Theorem 3 (Knowledge Cluster Theorem)**.** For any $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ with $\theta \neq 0$, $\mathfrak{M}, x \vDash [\because]^{fix} K_i\varphi$ if and only there exists $E \subseteq \bigcup_{j\in Ag} \mathcal{E}_j$ such that:

- There exists $U \in \tau_i^{+E}$, the topology generated from $\mathcal{E}_i \cup E$, such that $x \in U \cap \Pi_i(x) \subseteq [\![\varphi]\!]$; and

- For every cluster C of density greater than $1-\theta$ containing i:
  - for every $e \in E$ there is some $j \in C$ such that $e \in \mathcal{E}_j$; and
  - for all $E' \subseteq \bigcup_{j \in Ag} \mathcal{E}_j$, if there exists $J \subseteq C$ such that for every $e \in E'$ there is $j \in J$ with $e \in \mathcal{E}_j$, then $\bigcap E' \cap (U \cap \Pi_i(x)) \neq \emptyset$.

For this theorem, it will be worth it to explain the conditions on a more intuitive level before going on with the proof. Before anything else, bear in mind that there are two main parts to this theorem: one we will call 'receiving', the second we call 'blocking'. For the receiving part, notice that we state that an agent attains individual knowledge of a proposition φ at the end of the diffusion sequence if and only if there exists a subset of the totality of pieces of evidence that all agents possess that satisfies certain conditions. We can think of this subset as an 'information package' that the agent must have if he is to know φ. The first condition is that, adding this package to the agent's set of basic pieces of evidence ensures the existence of a set U in the resulting topology that, when intersected with the partition of the world of evaluation $\Pi_i(x)$, is a subset of $[\![\varphi]\!]$. In other words, adding this package to the agent's evidence generates evidence *supporting* φ. The next condition also belongs to the 'receiving' part, as it ensures, according to the Evidence Cluster Theorem, that all the pieces of evidence from E end up being part of the agent's set of basic evidence pieces at the end of the diffusion sequence. So, in sum, up to this point, this information package exists and is received. For the 'blocking' part, we have one condition left. Observe first the quantification and the part after 'then'. We are quantifying over all subsets of the total evidence that, if received, *could* make the set U not dense. The conditions in between are once more taken from the Evidence Cluster Theorem. So, in sum, the 'blocking' condition states: if there is a subset of pieces of evidence that i already has, or will receive at some point, then it does not break U's density. Putting all of these things together, we obtain a dense subset of $[\![\varphi]\!]$: exactly what is needed for i to know φ. What follows is the proof of this theorem.

Proof. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be an arbitrary evidence diffusion model with $\theta \neq 0$.

($\Rightarrow$) Suppose that $\mathfrak{M}, x \vDash [\because]^{fix} K_i\varphi$. Thus, we have by the definition of $[\because]$ that $\mathfrak{M}_{fix}, x \vDash K_i\varphi$, which in turn means that there exists $U \in \tau_i^{fix}$, the topology of evidence generated from $\mathcal{E}_i^{fix}$, such that $U \cap \Pi_i(x) \subseteq [\![\varphi]\!]$ and U is i-locally dense in $\Pi_i(x)$.

Now, for the set E required by the consequent, take $E = \mathcal{E}_i^{fix} \setminus \mathcal{E}_i$. For the subset of φ required, we take U mentioned above. Note that the topology generated by $\mathcal{E}_i \cup E = \mathcal{E}_i \cup (\mathcal{E}_i^{fix} \setminus \mathcal{E}_i) = \mathcal{E}_i^{fix}$ is exactly the τ_i^{fix}. And so, we know by assumption that $U \in \tau_i^{fix}$ and $x \in U \cap \Pi_i(x) \subseteq [\![\varphi]\!]$, so the first bullet point is satisfied.

Take any $e \in E$. Suppose towards a contradiction that there is a cluster C of density greater than $1-\theta$ containing i with all $j \in C$ having $e \notin \mathcal{E}_j$. By the corollary to the evidence cluster theorem, i never adopts e. But we know by assumption that $e \in \mathcal{E}_i^{fix}$, contradiction. So the first requirement is met.

For the second requirement, take an arbitrary $E' \in \bigcup_{j \in Ag} \mathcal{E}_j$ and suppose there exists $J \subseteq C$ such that for every $e \in E'$ there is $j \in J$ with $e \in \mathcal{E}_j$. By the corollary to the evidence cluster theorem, we have that for every $e \in E'$, $e \in \mathcal{E}_i^{fix}$, so that $\bigcap E' \in \tau_i^{fix}$. But note that, by assumption, U is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix}. Thus, we cannot have that $\bigcap E' \cap (U \cap \Pi_i(x)) = \emptyset$. And since E' was picked arbitrarily, the second requirement is met as well.

This proves the consequent, and so the left-to-right direction of this theorem is established.

($\Leftarrow$) Now let us suppose the consequent and try to prove the antecedent, namely $\mathfrak{M}, x \vDash [\because]^{fix} K_i\varphi$. This is the case if and only if $\mathfrak{M}_{fix}, x \vDash K_i\varphi$, which in turn is equivalent to the existence of a subset that is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix}.

We know that there exists a set $E \subseteq \bigcup_{j \in Ag}$ such that we have $U \in \tau_i^{+E}$, the topology generated by $\mathcal{E}_i \cup E$ with $x \in U \cap \Pi_i(x) \subseteq [\![\varphi]\!]$. By assumption, for every cluster C of density greater than $1-\theta$ containing i and for every $e \in E$, there is some $j \in C$ such that $e \in \mathcal{E}_j$. Thus, by the corollary to the evidence cluster theorem, for every $e \in E$, $e \in \mathcal{E}_i^{fix}$, and indeed $\tau_i^{+E} \subseteq \tau_i^{fix}$, entailing $U \in \tau_i^{fix}$.

Now to show that U is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix}, suppose the contrary. This means that there exists $A \in \tau_i^{fix}$ such that $(U \cap \Pi_i(x)) \cap (A \cap \Pi_i(x)) = \emptyset$. From this we can infer that there exists $e \in \widetilde{\mathcal{E}}_i^{fix}$ such that $(U \cap \Pi_i(x)) \cap (e \cap \Pi_i(x)) = \emptyset$. Further, we can obtain that there exists $E' \subseteq \mathcal{E}_i^{fix}$ such that $(U \cap \Pi_i(x)) \cap (\bigcap E' \cap \Pi_i(x)) = \emptyset$, which by the distributivity and idempotence of $\cap$ yields $\bigcap E' \cap (U \cap \Pi_i(x)) = \emptyset$. But now note that for every $e' \in E'$, $e' \in \mathcal{E}_i^{fix}$, and so by the evidence cluster theorem we have that for every cluster C of density greater than $1-\theta$ there must exist $j \in C$ such that $e' \in \mathcal{E}_j$. So we have found $E' \subseteq \bigcup_{j \in Ag}$ such that there exists $J \subseteq C$ with every $e' \in E'$ having a corresponding $j \in J$ with $e' \in \mathcal{E}_j$, but $\bigcap E' \cap (U \cap \Pi_i(x)) = \emptyset$, contradiction. So U is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix}.

As such, $[\![\varphi]\!]$ contains a factive τ_i^{fix}-open that is i-locally dense in $\Pi_i(x)$. By definition, we have that $\mathfrak{M}_{fix}, x \vDash K_i\varphi$ and in turn $\mathfrak{M}, x \vDash [\because]^{fix} K_i\varphi$. This establishes the right to left direction of the theorem.

Both directions of the equivalence have been shown, so the proof is concluded and the theorem holds. □

Of course, from this theorem, we immediately obtain that if all agents are to know something, that is, if the group is to have what is known as *mutual* knowledge [63, 32, 68, 27] of a proposition φ, the conditions must hold for all of the agents. The following Corollary states this formally.

Corollary 2. For any evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ with $\theta \neq 0$, $\mathfrak{M}, x \vDash [\because]^{fix} \bigwedge_{i \in Ag} K_i\varphi$ if and only if the conditions in the previous theorem are met for every agent.

Proof. Take an arbitrary evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ with $\theta \neq 0$. For the left to right direction, suppose $\mathfrak{M}, x \vDash [\because]^{fix} \bigwedge_{j \in Ag} K_j\varphi$ and pick an arbitrary $i \in Ag$. We can infer that $\mathfrak{M}_{fix}, x \vDash \bigwedge_{j \in Ag} K_j\varphi$ and therefore $\mathfrak{M}, x \vDash K_j\varphi$. From the knowledge cluster theorem we thereby obtain that the conditions hold for i, and since i was picked arbitrarily, they hold for all agents. Now for the right to left direction suppose that the conditions hold for every agent. Applying the knowledge cluster theorem for an arbitrary agent i yields $\mathfrak{M}, x \vDash [\because]^{fix} K_i\varphi$. Thus, $\mathfrak{M}_{fix}, x \vDash K_i\varphi$. And since i was arbitrary, for all $i \in Ag$, $\mathfrak{M}_{fix}, x \vDash K_i\varphi$. Since the set Ag is finite, we can form the conjunction and claim $\mathfrak{M}_{fix}, x \vDash \bigwedge_{i \in Ag} K_i\varphi$. And by the definition of $[\because]$, we have $\mathfrak{M}, x \vDash [\because]^{fix} \bigwedge_{i \in Ag} K_i\varphi$, as required. Since both directions were shown, the corollary holds. □

For distributed knowledge, the conditions are similar, but they, of course, have to mimic the definition of distributed knowledge instead of individual knowledge. What we mean by this is that instead of requiring a factive dense open subset of φ, we require that each agent has an 'information package' which gives them a factive dense set, and that all of these sets *taken together* entail φ (i.e. their intersection is a subset of $[\![\varphi]\!]$). The proof is similar to the previous one.

Theorem 4 (Distributed Knowledge Cluster Theorem)**.** For any evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ with $\theta \neq 0$, $\mathfrak{M}, x \vDash [\because]^{fix}\reflectbox{$\mathbf{K}$}\mathbf{D}\varphi$ if and only if for every $i \in Ag$ there exists $E \subseteq \bigcup_{j\in Ag}\mathcal{E}_j$ such that there exists $U_i \in \tau_i^{+E}$, the topology generated by $\mathcal{E}_i \cup E$, with $x \in U_i$ such that:

- for every cluster C of density greater than $1 - \theta$ containing i, we have:
  - for every $e \in E$, there exists $j \in C$ such that $e \in \mathcal{E}_j$; and
  - for every $E' \subseteq \bigcup_{j\in Ag}\mathcal{E}_j$, if there exists $J \subseteq C$ such that for all $j \in J$, $e \in \mathcal{E}_j$, then $\bigcap E' \cap (U_i \cap \Pi_i(x)) \neq \emptyset$;
- $\bigcap_{i\in Ag} U_i \subseteq [\![\varphi]\!]$.

Proof. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be an arbitrary evidence diffusion model with $\theta \neq 0$.

($\Rightarrow$) Suppose that $\mathfrak{M}, x \vDash [\because]^{fix}\reflectbox{$\mathbf{K}$}\mathbf{D}\varphi$. Thus, we have by the definition of $[\because]$ that $\mathfrak{M}_{fix}, x \vDash \reflectbox{$\mathbf{K}$}\mathbf{D}\varphi$, which in turn means that for every $i \in Ag$ there exists $U_i \in \tau_i^{fix}$ such that: 1) U_i is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix}; 2) $x \in \bigcap_{i\in Ag} U_i$; 3) $\bigcap_{i\in Ag} U_i \subseteq [\![\varphi]\!]$. We proceed in a similar fashion to how we did in the proof of the knowledge cluster theorem.

Let $i \in Ag$ be arbitrarily picked. Now, for the set E required by the consequent, take $E = \mathcal{E}_i^{fix} \setminus \mathcal{E}_i$. For the subset of φ required, we take the U_i invoked above. Note that the topology generated by $\mathcal{E}_i \cup E = \mathcal{E}_i \cup (\mathcal{E}_i^{fix} \setminus \mathcal{E}_i) = \mathcal{E}_i^{fix}$ is exactly the τ_i^{fix}. And so, we know by assumption that $U \in \tau_i^{fix}$ and $x \in U_i$.

Take any $e \in E$. Suppose towards a contradiction that there is a cluster C of density greater than $1 - \theta$ containing i with all $j \in C$ having $e \notin \mathcal{E}_j$. By the corollary to the evidence cluster theorem, i never adopts e. But we know by assumption that $e \in \mathcal{E}_i^{fix}$, contradiction. So the first requirement is met.

For the second requirement, take an arbitrary $E' \in \bigcup_{j\in Ag}\mathcal{E}_j$ and suppose there exists $J \subseteq C$ such that for every $e \in E'$ there is $j \in J$ with $e \in \mathcal{E}_j$. By the corollary to the evidence cluster theorem, we have that for every $e \in E'$, $e \in \mathcal{E}_i^{fix}$, so that $\bigcap E' \in \tau_i^{fix}$. But note that, by assumption, U_i is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix}. Thus, we cannot have that $\bigcap E' \cap (U_i \cap \Pi_i(x)) = \emptyset$. And since E' was picked arbitrarily, the second requirement is met as well.

Now note that we have by assumption that $\bigcap_{j\in Ag} U_j \subseteq [\![\varphi]\!]$, and so the final bullet point is also satisfied.

This proves the consequent, and so the left-to-right direction of this theorem is established.

($\Leftarrow$) Now let us suppose the consequent and try to prove the antecedent, namely $\mathfrak{M}, x \vDash [\because]^{fix}K_i\varphi$. This is the case if and only if 1),2), and 3) mentioned above are satisfied.

Pick an arbitrary agent $i \in Ag$. We know that there exists a set $E \subseteq \bigcup_{j \in Ag}$ such that we have $U_i \in \tau_i^{+E}$, the topology generated by $\mathcal{E}_i \cup E$ with $x \in U_i$. By assumption, for every cluster C of density greater than $1 - \theta$ containing i and for every $e \in E$, there is some $j \in C$ such that $e \in \mathcal{E}_j$. Thus, by the corollary to the evidence cluster theorem, for every $e \in E$, $e \in \mathcal{E}_i^{fix}$, and indeed $\tau_i^{+E} \subseteq \tau_i^{fix}$, entailing $U_i \in \tau_i^{fix}$.

Now to show that U_i is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix}, suppose the contrary. This means that there exists $A \in \tau_i^{fix}$ such that $(U_i \cap \Pi_i(x)) \cap (A \cap \Pi_i(x)) = \emptyset$. From this we can infer that there exists $e \in \widetilde{\mathcal{E}}_i^{fix}$ such that $(U_i \cap \Pi_i(x)) \cap (e \cap \Pi_i(x)) = \emptyset$. Further, we can obtain that there exists $E' \subseteq \mathcal{E}_i^{fix}$ such that $(U_i \cap \Pi_i(x)) \cap (\bigcap E' \cap \Pi_i(x)) = \emptyset$, which by the distributivity and idempotence of $\cap$ yields $\bigcap E' \cap (U \cap \Pi_i(x)) = \emptyset$. But now note that for every $e' \in E'$, $e' \in \mathcal{E}_i^{fix}$, and so by the evidence cluster theorem we have that for every cluster C of density greater than $1 - \theta$ there must exist $j \in C$ such that $e' \in \mathcal{E}_j$. So we have found $E' \subseteq \bigcup_{j \in Ag}$ such that there exists $J \subseteq C$ with every $e' \in E'$ having a corresponding $j \in J$ with $e' \in \mathcal{E}_j$, but $\bigcap E' \cap (U \cap \Pi_i(x)) = \emptyset$, contradiction. So U_i is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix}.

Since i was arbitrarily chosen, we can infer that for every agent there exists $U_i \in \tau_i^{fix}$ such that U_i is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix} and $x \in \bigcap_{i \in Ag} U_i$. Since we are assuming $\bigcap_{i \in Ag} U_i \subseteq [\![\varphi]\!]$, this requirement is met as well. So $\mathfrak{M}_{fix}, x \vDash \mathsf{ᗡ}\mathbf{D}\varphi$, which means $\mathfrak{M}, x \vDash [\because]^{fix} \mathsf{ᗡ}\mathbf{D}\varphi$, as needed.

Both directions of the equivalence have been shown, so the proof is concluded and the theorem holds. $\square$

As we have seen earlier, namely in Proposition 15, distributed evidence ($\mathsf{Ǝ}\mathbf{D}$) is invariant under evidence diffusion. For this reason, a theorem similar to the ones we have proven so far does not make sense for this modality, as it is unaffected by clusters, and more generally by diffusion.

These results are stated in a general form, for arbitrary evidence diffusion models — this is part of what makes the conditions so complex. In Section 4.1, we introduced four special types of network: the total network, the star, the cycle, and the wheel. We would therefore now be interested to see what special behavior these networks elicit, as well as effects the extra assumptions imposed by these special networks have on the conditions required for individual and distributed knowledge at the end of a diffusion sequence. Furthermore, we would like to define measures of speed and reliability for these networks, so that later on we can compare the results in our model with those of Zollman [92, 93, 94].

4.4 Diffusion on the Total, Star, Cycle, and Wheel Networks

Investigating the principles that guide diffusion on these special networks seems to be a clear enough goal — we will see how the extra assumptions we are making about the networks allow us to modify the cluster conditions from the previous theorems. It is not yet clear, however, what speed and reliability are in this setting — for we need different metrics compared to those of Zollman (see Chapter 2) — and there are many possible avenues of definition, so these notions will need to be clarified before we proceed with our analysis. In general, these measures will have to be tied to the behavior of evidence diffusion on the four networks, since even if we look at the speed or reliability of accession to knowledge (or learning), the epistemic and doxastic attitudes we have defined are evidence-based.

In this framework, given the step-wise nature of evidence diffusion, a natural measure of speed is the number of rounds it takes for a specific outcome to be reached. For example, the speed of a cascade for a specific piece of evidence is the number of diffusion rounds it takes for every agent to acquire that piece of evidence through diffusion. Similarly, the speed of the accession of an agent i to knowledge of a proposition φ is defined as the number of diffusion rounds that have to pass before the agent robustly knows φ. Formally, this would be the least number $k \in \mathbb{N}$ such that $\mathfrak{M}_k, x \vDash K_i\varphi$ and for every $k' \geq k$, $\mathfrak{M}_{k'}, x \vDash K_i\varphi$. The definition for distributed knowledge is analogous, and the definition for distributed evidence need not be provided, as $\boxplus$ and diffusion do not interact.

Throughout this section, we will mostly be looking at the speed with which these networks reach the fixpoint in the diffusion sequence, as it would be difficult to tease out the speed of accession to knowledge without further assumptions — a fact which will become clear after our discussion about reliability.

Indeed, reliability on evidence diffusion models is a very difficult characteristic to associate to the networks directly. As opposed to speed, where we might be interested in the speed of evidence diffusion in general, and not just the speed of accession to knowledge, reliability is much more closely related to the success or failure of the network to reach a specific, 'good' epistemic state. This will inevitably depend on the *seed sets* [61] of the pieces of evidence. By seed set for a piece of evidence $e \in \bigcup_{i \in Ag} \mathcal{E}_i$, we mean the *initial* set of adopters Ag_e, or those agents who have adopted the piece of evidence before any diffusion rounds have occurred. In what follows, we outline the three main factors that contribute to the reliability of evidence diffusion models in general, and discuss what conclusions can be drawn about networks after that. We tease these out from the Knowledge Cluster Theorem.

Firstly, the Knowledge Cluster Theorem requires the existence of a subset of $\bigcup_{i \in Ag} \mathcal{E}_i$ with certain properties. Earlier, we called this subset the 'information package' that the agent must receive in order to have a factive (combined) piece of evidence for φ. Of course, such an information package must exist in the first place for an agent (or the group) to attain (distributed) knowledge, and this existence requirement is wholly dependent on the set $\bigcup_{i \in Ag} \mathcal{E}_i$. This does not depend on the network, but rather on the pieces of evidence that the agents have to begin with.

Secondly, the agent must be in a position to receive such an information package. The agent's receptivity for these 'good' pieces of evidence would be a measure that is calculated *across* evidence diffusion models, not *within* an evidence diffusion model, since whether or not an agent receives such an information package in a specific model is a yes-or-no question, and does not lie on a spectrum. There are many possible ways to specify such a measure, but this lies beyond the scope of this thesis. Finally, the agent must be in a position to fend off all of the possible pieces of evidence that would break the density of their justification for φ. Much like the previous measure, this non-receptivity for 'bad' pieces of evidence would be a cross-model measure. We bring this subject up again in the next, concluding chapter, which traces limitations and future lines of research.

What further complicates the analysis of reliability is that it seems like most of the network-related factors that would make an agent more receptive to 'good' evidence would make them more receptive to 'bad' evidence as well. For example, networks that will form full cascades with relatively low thresholds (such as the Cycle, see below) put the agent in the best spot to receive the good information packages, but at the same time in the most vulnerable position to receive 'bad' pieces of evidence. Thus, even out of the select few networks we have highlighted, it will be difficult to formally capture which is more reliable than which with the tools we have been using so far. As well, we conjecture that the size of the agent set, as well as the size of the set of possible worlds, both affect reliability, the former by making it likelier that the 'bad' pieces of evidence will be further from the agent, and the latter by increasing the number of possible pieces of evidence that break a justification's density. Nevertheless, we proceed with our analysis of these special networks.

4.4.1 The Total Network

As one may imagine, the total network makes talk of clusters somewhat obsolete, since all agents are connected to all other agents. In practice, this intuition is materialized in the following proposition, which is the Evidence Cluster Theorem (henceforth ECT) adapted to the stronger assumption of a total network. What is noteworthy is that what is required now for a full cascade with a piece of evidence is simply that there exist a numerous enough subset of adopters, where 'numerous enough' means $\frac{|Ag_e|}{|Ag|-1} \geq \theta$.

Proposition 19 (ECT for Total Network)**.** *For any evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *,* $e \in \mathcal{E}_i^{fix}$ *if and only if* $e \in \mathcal{E}_i$ *or* $\frac{|Ag_e|}{|Ag|-1} \geq \theta$*.*

Proof. For the left to right direction, note that $\mathcal{E}_i^{fix}$ is the basic evidence set of i at the end of the diffusion sequence, which means that $e \in \mathcal{E}_i$ or e was received by i through diffusion. In case $e \notin \mathcal{E}_i$ and it was received through diffusion it must be that $\frac{|\mathcal{N}(i)\cap Ag_e|}{|\mathcal{N}(i)|} \geq \theta$. But since the network is total, $\mathcal{N}(i) = Ag \setminus \{i\}$, so $|\mathcal{N}(i)| = |Ag| - 1$. We know $i \notin Ag_e$, and $\mathcal{N}(i) = Ag \setminus \{i\}$, so $\mathcal{N}(i) \cap Ag_e = Ag(e)$. Thus, $\frac{|\mathcal{N}(i)\cap Ag_e|}{|\mathcal{N}(i)|} \geq \theta$ if and only if $\frac{|Ag_e|}{|Ag|-1} \geq \theta$, so the disjunction holds and the left to right direction is established. For the right to left direction, $e \in \mathcal{E}_i$ or $\frac{|Ag_e|}{|Ag|-1} \geq \theta$. If the former is the case, note that $\mathcal{E}_i \subseteq \mathcal{E}_i^{fix}$ and we are done. If the latter is the case, by arguments similar to those for the previous direction, e is received by i through diffusion, so $e \in \mathcal{E}_i^{fix}$. The right to left direction is also established, which proves the equivalence. □

This has a strong influence on the number of diffusion rounds that are required to reach the fixpoint. Since no non-adopting agent will become an adopter unless there is a numerous enough set of adopters, and all of them will if there is such a set, all diffusion of evidence will happen in one round — for every piece of evidence, it either gets transmitted to the entire network, or it stays in the seed set. Thus, a fixpoint is reached for this kind of network in one round, as the following proposition encodes. This means that the total network is maximal in terms of speed.

Proposition 20. $\vDash tot \to ([\because]^1\varphi \to [\because]^2\varphi)$. *In words, if the model's network is total, then the diffusion sequence reaches its fixpoint after one round.*

Proof. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be an arbitrary evidence diffusion model and $x \in X$ an arbitrary world. We want to show that $\mathfrak{M}, x \vDash tot \to ([\because]\varphi \to [\because][\because]\varphi)$, so we suppose $\mathfrak{M}, x \vDash tot$ and $\mathfrak{M}, x \vDash [\because]\varphi$ and aim to show that $\mathfrak{M}, x \vDash [\because][\because]\varphi$. $\mathfrak{M}, x \vDash tot$, so by Proposition 6, we know that $\mathcal{N} = \mathcal{T}$ and so for every $i \in Ag$, $\mathcal{N}(i) = Ag \setminus \{i\}$. For any piece of evidence e and agent $i \notin Ag_e$, because of the totality of the network, e is adopted at the next diffusion round by i if and only if $\frac{|A_e|}{|Ag|-1}$. Since this is the case for every non-adopting agent, they will act in unison — either all of them adopt, or non of them adopt. If they adopt, then $Ag_e^{\because} = Ag$, and if they do not, then $Ag_e^{\because} = Ag_e$. In the former case, nothing regarding e can change anymore, as a full cascade has already been reached, and in the latter case, again nothing can change after an additional diffusion round, since the set $Ag_e^{\because}$ will remain the same for the same reasons it did not change during the first round. Since this holds for every piece of evidence, the evidence sets of all agents will remain fixed after the first diffusion round. This clearly entails that $\mathfrak{M}_2, x \vDash \varphi$, and thus $\mathfrak{M}, x \vDash [\because][\because]\varphi$. Since we reasoned with an arbitrary model and world, our proposition is proved: $tot \to ([\because]\varphi \to [\because][\because]\varphi)$ is valid on evidence diffusion models. □

We can now replace the cluster-dependent conditions in the Knowledge Cluster Theorem(s) (henceforth KCT) using Proposition 19, where the following obtains.

Proposition 21 (KCT for Total Network)**.** *For any* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *, if* $\mathcal{N} = \mathcal{T}$*, then* $\mathfrak{M}, x \vDash [\because]^{fix} K_i\varphi$ *if and only if:*

- *there exists* $E \subseteq \bigcup_{j \in Ag} \mathcal{E}_j$ *such that* $E \subseteq \mathcal{E}_i$ *or for every* $e \in E$*,* $\frac{|Ag_e|}{|Ag|-1} \geq \theta$*, and there exists* $U \in \tau_i^{+E}$ *such that* $x \in U \subseteq [\![\varphi]\!]$*;*
- *For every* $E' \subseteq \bigcup_{j \in Ag} \mathcal{E}_j$*, if* $E' \subseteq \mathcal{E}_i$ *or for every* $e \in E'$*,* $\frac{|Ag_e|}{|Ag|-1} \geq \theta$*, then* $\bigcap E' \cap (U \cap \Pi_i(x)) \neq \emptyset$*.*

Proof. The proposition obtains immediately by substituting the conditions in the KCT with those obtained in Proposition 19. □

From the point of view of epistemic reliability, a few comments can be made about this network. A positive aspect of it, which we obtain from the second condition, is that it seems like it could be quite robust to negative results. If a false piece of evidence somehow makes its way into an agent's set of basic pieces of evidence — we may imagine a faulty experiment, or an artefact in the data — then it would need to be independently reproduced by enough agents before it could make its way to all other agents, and we know that faulty experiments rarely reproduce. A negative aspect, which we obtain from the first condition, would be that, if an agent doesn't have the necessary information package in their set of basic pieces of evidence, then many (depending on the threshold) other agents would have to have this package for it to transmit, so that agents are needlessly resistant to accepting pieces of evidence from which they have a lot of epistemic advantages to gain. Both aspects seem to point to a kind of conservativity to non-well-vetted evidence. Depending on the overall evidence available, and on the seed sets for every piece of evidence, this could yield both good and bad results.

4.4.2 The Star Network

We move on to the star (for i) network. As one would expect, our analysis will find that the agent i will have a lot of authority in this network, in the sense that their evidence will be transmitted without any hurdles, and they will act as a barrier to the diffusion of much of the other evidence. The following proposition encapsulates the first part of what was just mentioned.

Proposition 22. *For any* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *, if* $\mathcal{N} = \bigstar_i$*, then* $\mathcal{E}_i \subseteq \mathcal{E}_j^{\because}$ *for every* $j \in Ag$*.*

Proof. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be arbitrarily picked, and take an arbitrary world x. Suppose $\mathcal{N} = \bigstar_i$ and suppose $e \in \mathcal{E}_i$. We have $\mathcal{N}(i) = Ag \setminus \{i\}$ and $\mathcal{N}(j) = \{i\}$ for every $j \neq i \in Ag$. Thus, it is clear that for every $e \in \mathcal{E}_i$ and for every $j \neq i \in Ag$, $\frac{|\mathcal{N}(j)\cap Ag_e|}{|\mathcal{N}(j)|} = \frac{|\{i\}|}{|\{i\}|} = 1 \geq \theta$. As such, for every $j \in Ag$, for every $e \in \mathcal{E}_i$, we have $e \in \mathcal{E}_j^{\because}$. And since $\mathfrak{M}$ and x were arbitrarily picked, the proposition holds. □

It is clear now that i's evidence gets transmitted to everybody else after one round, but what of the evidence of the other agents? We capture what happens with that evidence by investigating the conditions in which i does not have a piece of evidence initially, but receives it after one diffusion round. In short, the requirement from the total network shows up again: there must be a group of adopters numerous enough for i to adopt.

Proposition 23. *For any* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *, if* $\mathcal{N} = \bigstar_i$ *and* $e \notin \mathcal{E}_i$*, then* $e \in \mathcal{E}_i^{\because}$ *if and only if* $\frac{|Ag_e|}{|Ag|-1} \geq \theta$*.*

Proof. Let us take an arbitrary evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ with $\mathcal{N} = \bigstar_i$. Suppose some piece of evidence $e \notin \mathcal{E}_i$. Now for the left to right direction suppose $e \in \mathcal{E}_i^{\because}$. This means that e was received through diffusion, i.e. $\frac{|\mathcal{N}(i)\cap Ag_e|}{|\mathcal{N}(i)|} \geq \theta$. But since $\mathcal{N}(i) = Ag\setminus\{i\}$ and $i \notin Ag_e$, we have that the previous inequality is equivalent to $\frac{|Ag_e|}{|Ag|-1} \geq \theta$. For the right to left direction, by a similar argument, i will receive e at the next round, so that $e \in \mathcal{E}_i^{\because}$. The equivalence and thus the proposition as a whole are proven. □

Note that, if i receives a piece of evidence as the previous proposition describes, then we may apply Proposition 22 once more, leading to everyone else receiving that piece of evidence. Thus, the two previous results together give us the necessary and sufficient conditions for an agent to have a piece of evidence at the end of the diffusion sequence.

Proposition 24 (ECT on Star Network)**.** *For any evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *, if* $\mathcal{N} = \bigstar_i$*, then* $e \in \mathcal{E}_j^{fix}$ *if and only if* $e \in \mathcal{E}_j$ *or* $e \in \mathcal{E}_i$ *or* $\frac{|Ag_e|}{|Ag|-1} \geq \theta$*.*

Proof. Take an arbitrary evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ with $\mathcal{N} = \bigstar_i$. For the left to right direction, note that $\mathcal{E}_j^{fix}$ is the basic evidence set of j at the end of the diffusion sequence, which means that $e \in \mathcal{E}_j$ or e was received by j through diffusion. In case $e \notin \mathcal{E}_j$ and it was received through diffusion, let us draw a case distinction. If $e \in \mathcal{E}_i$, then by Proposition 22, $e \in \mathcal{E}_j^{\because}$, whereby $e \in \mathcal{E}_j^{fix}$. If $e \notin \mathcal{E}_i$, then since i is j's only connection, by Proposition 23 and then Proposition 22, $e \in \mathcal{E}_i^2$. So the disjunction holds. For the right to left direction, we argue in much the same way, using the previous propositions. This proves the equivalence.

□

In short, for an agent to end up with a piece of evidence at the fixpoint, they must have either started with it, received it from i during the first round, or received it from i during the second round. An analogue of Proposition 20 can be proven for star networks. Star networks are thus slower than total networks.

Proposition 25. $\vDash star_i \to ([\because]^2\varphi \to [\because]^3\varphi)$. *In words, if the model's network is a star for i, then we reach a fixpoint at diffusion round 2.*

Proof. Analogous to Proposition 20. □

As before, we can substitute the cluster conditions for the ones that we derived with the extra assumption of the star network.

Proposition 26 (KCT for Star Network)**.** *for any* $\mathfrak{M} = \langle X,\ \{\mathcal{E}_i\}_{i\in Ag},\ \{\Pi_i\}_{i\in Ag},\ \mathcal{N},\ [\![\cdot]\!]\rangle$ *, if* $\mathcal{N} = \bigstar_i$*, then* $\mathfrak{M}, x \vDash [\because]^{fix}K_j\varphi$ *if and only if:*

- *there exists* $E \subseteq \bigcup_{k\in Ag}\mathcal{E}_k$ *such that* $E \subseteq \mathcal{E}_j$ *or* $E \subseteq \mathcal{E}_i$ *or for every* $e \in E$, $\frac{|Ag_e|}{|Ag|-1} \geq \theta$*, and there exists* $U \in \tau_i^{+E}$ *such that* $x \in U \subseteq [\![\varphi]\!]$*;*
- *For every* $E' \subseteq \bigcup_{k\in Ag}\mathcal{E}_k$*, if* $E' \subseteq \mathcal{E}_j$ *or* $E' \subseteq \mathcal{E}_i$ *or for every* $e \in E'$, $\frac{|Ag_e|}{|Ag|-1} \geq \theta$*, then* $\bigcap E' \cap (U \cap \Pi_i(x)) \neq \emptyset$*.*

Proof. The proposition obtains immediately by substituting the conditions in the KCT with those obtained in Proposition 24. □

From the point of view of reliability, it seems hard to put this network above the total network. The same condition for a full cascade — that enough agents over-all are in possession of a piece of evidence — is present here, as it was in the case of the previous network. But additionally, we see that the central authority is able to spread their evidence without hurdle or rebuttal. Quite obviously, if the star of the network has false or misleading pieces of evidence, the entire network has to pay the price. Even if the central authority is known to be reliable, the process of verification that the total network encapsulates seems to be more epistemically healthy.

4.4.3 The Cycle Network

For the cycle network, it is harder to obtain general conditions, like we have been doing so far. Thankfully, however, there is a property that this network has which still makes it tenable to analysis. Namely, all clusters except the whole set of agents and the singleton sets have a density of $\frac{1}{2}$.

Proposition 27. *For any* $\mathfrak{M} = \langle X,\ \{\mathcal{E}_i\}_{i\in Ag},\ \{\Pi_i\}_{i\in Ag},\ \mathcal{N},\ [\![\cdot]\!]\rangle$ *, If* $\mathcal{N} = \bigcirc$*, for every* $C \subsetneq Ag$*, where for every* $i \in C$ *there exists* $j \in C$ *with* $j \in \mathcal{N}(i)$*, C is a cluster of density* $\frac{1}{2}$*.*

Proof. Take any $C \subsetneq Ag$ such that for every $i \in C$ there exists $j \in C$ with $j \in \mathcal{N}(i)$. Since $\mathcal{N} = \bigcirc$, the sets that we are excluding with the aforementioned conditions are Ag and the singletons $\{i\}$ for any $i \in Ag$. There are finitely many such sets, since the set of agents is finite. If $|C| = 2$, note that each agent in C has a neighbor in C, and a neighbor outside C, meaning $\frac{|\mathcal{N}(i)\cap C|}{|\mathcal{N}(i)|} = \frac{1}{2}$ for both of these agents. This carries over to $|C| > 2$, for the agents that are on the edges of the set — and there are always such agents because we are excluding the set Ag. For the agent in the middle we have $\frac{|\mathcal{N}(i)\cap C|}{|\mathcal{N}(i)|} = 1$, as we know from the properties of the cycle that every agent is connected exclusively to two neighbors, and both are in the set by assumption. This proves the proposition. □

Given this, we can analyse what happens in case the threshold is less than or equal to, or greater than $\frac{1}{2}$. Starting with the former, it should be clear that, if a piece of evidence e is in some agent's basic evidence set, there will be no clusters of density strictly greater than $\frac{1}{2}$ in $Ag \setminus Ag_e$, which means that there will be a full cascade with that piece of evidence. And since this happens with all pieces of evidence, the agents fully and openly share all pieces of evidence. This is captured and proven in the following proposition.

Proposition 28. *For every* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\mathcal{N} = \bigcirc$*, if* $\theta \leq \frac{1}{2}$*, then for every* $i \in Ag$*,* $\mathcal{E}_i^{fix} = \bigcup_{j\in Ag} \mathcal{E}_j$*.*

Proof. Evidently, for any $e \in \bigcup_{i\in Ag} \mathcal{E}_i$, if $Ag_e = Ag$, then $Ag_e^{fix} = Ag$. From Proposition 27, we can derive that for any $e \in \bigcup_{i\in Ag} \mathcal{E}_i$ with $Ag_e \neq Ag$, there is no cluster of density greater than $1 - \theta$ in $Ag \setminus Ag_e$, which, by the ECT, means that $Ag_e^{fix} = Ag$. Putting this all together, we get exactly that for every $i \in Ag$, $\mathcal{E}_i^{fix} = \bigcup_{j\in Ag} \mathcal{E}_j$, as required. □

An immediate consequence of this is that all agents have, as their topology of evidence, the join topology $\bigvee_{i\in Ag} \tau_i$. Thus, we can simplify the KCT(s) and make no reference to information packages or clusters, since we know directly what topology the agents will obtain at the end of the diffusion sequence.

Proposition 29 (KCT for Cycle with $\theta \leq \frac{1}{2}$)**.** *For any* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\mathcal{N} = \bigcirc$ *and* $\theta \leq \frac{1}{2}$*,* $\mathfrak{M}, x \vDash [\because]^{fix} K_i\varphi$ *if and only if there exists* $U \in \bigvee_{j\in Ag} \tau_j$ *such that* $x \in U \subseteq [\![\varphi]\!]$ *and* U *is* i*-locally dense in* $\Pi_i(x)$*.*

Proof. Take an arbitrary evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ for $\theta \leq \frac{1}{2}$ and with $\mathcal{N} = \bigcirc$. By Proposition 28, for any $i \in Ag$, $\mathcal{E}_i^{fix} = \bigcup_{j\in Ag} \mathcal{E}_j$. Thus, as we have shown before in the proof of Proposition 14, $\mathcal{E}_i^{fix} = \bigcup_{j\in Ag} \mathcal{E}_j$ entails $\tau_i^{fix} = \bigvee_{j\in Ag} \tau_j$. So, the necessary and sufficient condition for $[\because]^{fix} K_i\varphi$ is the existence of a factive $\bigvee_{j\in Ag} \tau_j$-open subset of $[\![\varphi]\!]$ that is i-locally dense in $\Pi_i(x)$. □

The Distributed KCT is modified in a similar way, in the sense that every agent i must have a factive $\bigvee_{j\in Ag} \tau_j$-open that is i-locally dense in $\Pi_i(x)$, and the intersection of these has to be a subset of $[\![\varphi]\!]$. Note that even though the agents do have the same topology, they do not have the same partitions, and so not everybody knows the same propositions, and distributed knowledge does not coincide with any particular agent's knowledge in a necessary manner. It should be clear, though, that if we apply law distribution to the model, then since all agents will have the same topology and partition, all forms of knowledge will coincide.

We now prove an upper bound for the fixpoint of a diffusion sequence based on a model with this network. Unlike before, we cannot say exactly how many diffusion rounds there will be before the fixpoint is reached. This is because, on this network, the number of diffusion rounds depends on the *seed set* [61] for each piece of evidence. For example, a model in which there is only one piece of evidence $e \subsetneq X$, and it has only been adopted by one agent. Even though it will form a cascade according to our previous results, the number of rounds it will take to propagate is $\lfloor \frac{|Ag|}{2} \rfloor$. On the other hand, if $Ag_e = Ag \setminus \{i\}$ for some i, it will only take one round before we reach a full cascade with the one and only piece of evidence, therefore reaching the fixpoint. For these parameters, this is clearly the slowest network we have seen so far, and the slowest over all.

Proposition 30. *For any evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\theta \leq \frac{1}{2}$, $\mathfrak{M} \vDash cycle \rightarrow ([\because]^{\lfloor\frac{|Ag|}{2}\rfloor}\varphi \rightarrow [\because]^{\lfloor\frac{|Ag|}{2}\rfloor+1}\varphi)$

Proof. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be an arbitrary evidence diffusion model with $\theta \leq \frac{1}{2}$, take an arbitrary world x and suppose $\mathfrak{M}, x \vDash cycle$ and $\mathfrak{M}, x \vDash [\because]^{\lfloor\frac{|Ag|}{2}\rfloor}\varphi$. Take an arbitrary piece of evidence e and suppose it has the minimal seed set $Ag_e = \{i\}$ for some i. Then, for i's neighbors, as $\theta \leq \frac{1}{2}$, they will adopt, and so too for their neighbors, and so on. The set of agents is finite, and so we reach a full cascade in $\lfloor\frac{|Ag|}{2}\rfloor$. Since this was the minimal seed set, diffusion with evidence that has more numerous seed sets can only happen faster. So, for every $i \in Ag$, $\mathcal{E}_i^{\lfloor\frac{|Ag|}{2}\rfloor} = \bigcup_{j\in Ag}\mathcal{E}_j$. This clearly entails that $\mathfrak{M}, x \vDash [\because]^{\lfloor\frac{|Ag|}{2}\rfloor+1}\varphi$, as needed. □

For $\theta \leq \frac{1}{2}$, we might say that this network is the least reliable so far. The unfettered diffusion of all pieces of evidence means that *all* bad results will spread to *all* agents. As always, the evidence available over-all is important, but with these parameters, the cycle has no way to fend against bad evidence.

We move on now to the other case, where $\theta > \frac{1}{2}$. Intuitively speaking, since any agent is only connected to two other agents, and in order to obtain a piece of evidence more than half of the agents need to possess that piece of evidence, the only case where diffusion happens is when we have adopter-non-adopter-adopter configurations. That is, when a non-adopting agent has neighbors who are both adopters. This intuition is made precise with the following proposition: the agent's set of basic pieces of evidence will consist of their own initial set, together with the pieces of evidence that *both* his neighbors have.

Proposition 31 (ECT for Cycle with $\theta > \frac{1}{2}$)**.** *For any evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\mathcal{N} = \bigcirc$ *and* $\theta > \frac{1}{2}$, *for any* $i \in Ag$ *and for every* $e \in \bigcup_{j\in Ag}\mathcal{E}_j$, $e \in \mathcal{E}_i^{fix}$ *if and only if* $e \in \mathcal{E}_i \cup (\mathcal{E}_j \cap \mathcal{E}_k)$ *where* $\mathcal{N}(i) = \{j, k\}$.

Proof. Take an arbitrary evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ for $\theta > \frac{1}{2}$ with $\mathcal{N} = \bigcirc$. For the left to right direction, note that $\mathcal{E}_i^{fix}$ is i's basic evidence set at the end of the diffusion sequence, so let us draw the following case distinction for an arbitrary piece of evidence e. If $e \in \mathcal{E}_i$, then clearly $e \in \mathcal{E}_i \cup (\mathcal{E}_j \cap \mathcal{E}_k)$ where $\mathcal{N}(i) = \{j, k\}$. If $e \notin \mathcal{E}_i$, then at some round, e was received through diffusion, which requires $\frac{|\mathcal{N}(i)\cap Ag_e|}{|\mathcal{N}(i)|} \geq \theta > \frac{1}{2}$, and since i is only connected to j and k, this means that both must have e at some round. Suppose towards a contradiction that $e \notin \mathcal{E}_j \cap \mathcal{E}_k$, i.e. either j or k doesn't have it in the initial model. Then, since they both have it at *some* later round, one of them, without loss of generality say j, must have received it at some later round. Analogously, for j to receive e at a later round, i must have had e at an earlier round, which we know is false. So $e \in \mathcal{E}_j \cap \mathcal{E}_k$, which establishes the left to right direction. For the right to left direction, we easily observe that if $e \in \mathcal{E}_i$, then $e \in \mathcal{E}_i^{fix}$, and if $e \in \mathcal{E}_j \cap \mathcal{E}_k$, then e is received by i through diffusion at the next round, so $e \in \mathcal{E}_i^{fix}$. The proposition is proven. □

Similarly to the previous case, this proposition gives us precisely the topology that the agent will have after diffusion, so we can omit the mention of clusters and specify the condition for knowledge after diffusion directly.

Proposition 32 (KCT for Cycle with $\theta > \frac{1}{2}$)**.** *For any* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\mathcal{N} = \bigcirc$, *if* $\theta > \frac{1}{2}$ *then* $\mathfrak{M}, x \vDash [\because]^{fix}K_i\varphi$ *if and only if there exists* $U \in \tau_i'$, *the* *topology generated by* $\mathcal{E}_i \cup (\mathcal{E}_j \cap \mathcal{E}_k)$*, where* $\mathcal{N}(i) = \{j, k\}$*, such that* $x \in U \subseteq [\![\varphi]\!]$ *and* U *is i-locally dense in* $\Pi_i(x)$

Proof. It is clear from Proposition 31 that $\mathcal{E}_i \cup (\mathcal{E}_j \cap \mathcal{E}_k)$, where $\mathcal{N}(i) = \{j, k\}$, is the subbasis for τ_i^{fix}, so $[\because]^{fix} K_i\varphi$ holds by definition if and only if, in the fixpoint model, there is a factive τ_i^{fix}-open subset of $[\![\varphi]\!]$ that is i-locally dense in $\Pi_i(x)$. □

As we have stated before, the only cases where diffusion happens on the cycle network with thresholds greater than $\frac{1}{2}$ are adopter-non-adopter-adopter configurations, which turn into adopter-adopter-adopter configurations in one round. This means that a fixpoint is reached after exactly one round of diffusion. For these parameters, the cycle matches the speed of the total network.

Proposition 33. *For any evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\theta > \frac{1}{2}$, $\mathfrak{M} \vDash cycle \to ([\because]^1\varphi \to [\because]^2\varphi)$

Proof. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be an arbitrary evidence diffusion model with $\theta > \frac{1}{2}$, take an arbitrary world x and suppose $\mathfrak{M}, x \vDash cycle$ and $\mathfrak{M}, x \vDash [\because]^1\varphi$. Take an arbitrary agent $i \in Ag$. Either $\mathcal{E}_i^1 = \bigcup_{j\in Ag}\mathcal{E}_j$, in which case i cannot receive any further pieces of evidence, or $\mathcal{E}_i^1 \neq \bigcup_{j\in Ag}\mathcal{E}_j$. In this latter case, for any piece of evidence $e \in \bigcup_{j\in Ag}\mathcal{E}_j \setminus \mathcal{E}_i^1$, to receive it at the next round, it must be the case that $e \in \mathcal{E}_j^1 \cap \mathcal{E}_k^1$ where $\{j, k\} = \mathcal{N}(i)$. Since $e \notin \mathcal{E}_i^1$, we know $e \notin \mathcal{E}_j$ or $e \notin \mathcal{E}_k$. Without loss of generality, say $e \notin \mathcal{E}_j$ but $e \in \mathcal{E}_k^1$. Then for $e \in \mathcal{E}_i^2$, $e \in \mathcal{E}_j^1$, which requires $e \in \mathcal{E}_i^1$. We know this is not the case, so $e \notin \mathcal{E}_i^2$, and since e and i were arbitrary, this means that we have reached a fixpoint in diffusion, and thus, $\mathfrak{M}, x \vDash [\because]^2\varphi$, proving the proposition. □

As for reliability, these parameters seem to better poise cycle networks against false evidence. That being said, the fact that what an agent adds to their evidence set during diffusion is wholly dictated by what their immediate neighbors agree on seems to be a principle of poor epistemic standing.

4.4.4 The Wheel Network

As with the cycle network, we can obtain an upper bound on the density of the clusters appearing in wheel networks. As expected, this upper bound is higher than the previous one, since the wheel is more densely connected than the cycle.

Proposition 34. *For every* $C \subsetneq Ag$ *where for every* $j \in C$ *there exists* $k \in C$ *with* $k \in \mathcal{N}(j)$*,* C *is a cluster of density at most* $\frac{2}{3}$.

Proof. Take any $C \subsetneq Ag$ such that for every $j \in C$ there exists $k \in C$ with $k \in \mathcal{N}(j)$. Since $\mathcal{N} = \mathcal{W}_i$, the sets that we are excluding with the aforementioned conditions are Ag and the singletons $\{j\}$ for any $j \in Ag$. There are finitely many such sets, since the set of agents is finite. If $i \notin C$, then if $|C| = 2$, note that each agent in C has a neighbor in C, and two neighbors outside C, meaning $\frac{|\mathcal{N}(i)\cap C|}{|\mathcal{N}(i)|} = \frac{1}{3}$ for both of these agents. This carries over to $|C| > 2$, for the agents that are on the edges of the set — and there are always such agents because we are excluding the set Ag. For the agent in the middle we have $\frac{|\mathcal{N}(i)\cap C|}{|\mathcal{N}(i)} = \frac{2}{3}$, as we know from the properties of the cycle that every agent is connected exclusively to two neighbors and to i. If $i \in C$, then every non-i agent that is not on the edges of the set, has $\frac{|\mathcal{N}(i)\cap C|}{|\mathcal{N}(i)}=1$, those on the edges have $\frac{|\mathcal{N}(i)\cap C|}{|\mathcal{N}(i)}=\frac{2}{3}$ and i has $\frac{|\mathcal{N}(i)\cap C|}{|\mathcal{N}(i)}=\frac{|C|-1}{|Ag|-1}$. If this latter fraction is lower than $\frac{2}{3}$, that becomes the density of C, but if it is larger $\frac{2}{3}$ is the density by definition. These were all the possible cases, so the highest density of such a cluster is $\frac{2}{3}$, as stated. □

What we immediately obtain as a consequence of the ECT is that, for thresholds at least as low as $1-\frac{2}{3}$, a full cascade forms for all pieces of evidence.

Proposition 35. *For every* $\mathfrak{M}=\langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\mathcal{N}=\mathcal{W}_i$, *if* $\theta\leq\frac{1}{3}$, *then for every* $j\in Ag$, $\mathcal{E}_j^{fix}=\bigcup_{k\in Ag}\mathcal{E}_k$.

Proof. Evidently, for any $e\in\bigcup_{i\in Ag}\mathcal{E}_i$, if $Ag_e=Ag$, then $Ag_e^{fix}=Ag$. From Proposition 34, we can derive that for any $e\in\bigcup_{i\in Ag}\mathcal{E}_i$ with $Ag_e\neq Ag$, there is no cluster of density greater than $1-\theta$ in $Ag\setminus Ag_e$, which, by the ECT, means that $Ag_e^{fix}=Ag$. Putting this all together, we get exactly that for every $i\in Ag$, $\mathcal{E}_i^{fix}=\bigcup_{j\in Ag}\mathcal{E}_j$, as required. □

In terms of speed, for these parameters, we prove the following proposition.

Proposition 36. *For every* $\mathfrak{M}=\langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$, *if* $\theta\leq\frac{1}{3}$,

$$\mathfrak{M}\vDash wheel_i\rightarrow([\because]^{\left\lceil\frac{\frac{\lceil|Ag|-1\rceil}{3}-1}{2}\right\rceil+1}\varphi\rightarrow[\because]^{\left\lceil\frac{\frac{\lceil|Ag|-1\rceil}{3}-1}{2}\right\rceil+2}\varphi)$$

Proof. Take an arbitrary $\mathfrak{M}=\langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ for $\theta\leq\frac{1}{3}$ and suppose $\mathfrak{M}\vDash wheel_i$ and $\mathfrak{M},x\vDash[\because]^{\left\lceil\frac{\frac{\lceil|Ag|-1\rceil}{3}-1}{2}\right\rceil+1}\varphi$. Then we can infer that $\mathcal{N}=\mathcal{W}_i$. Let $k=\left\lceil\frac{\frac{\lceil|Ag|-1\rceil}{3}-1}{2}\right\rceil+1$. Now note that $\mathfrak{M}_k$ is the model after $\left\lceil\frac{\frac{\lceil|Ag|-1\rceil}{3}-1}{2}\right\rceil+1$ diffusion rounds. Observe that $\lceil\frac{|Ag|-1}{3}\rceil$ is the upper bound of the number of agents needed to make i adopt, for we know that i is connected to every other agent, and that $\theta\leq\frac{1}{3}$. Then note that, in the slowest diffusion scenario, i.e. when there is one non-i initial adopter, two agents adopt per round, so that at round t there are $1+2t$ adopters of this piece of evidence. Thus to obtain our upper bound we simply solve $\lceil\frac{|Ag|-1}{3}\rceil\geq1+2t$, getting $\frac{\frac{\lceil|Ag|-1\rceil}{3}-1}{2}\geq t$. Putting $\left\lceil\frac{\frac{\lceil|Ag|-1\rceil}{3}-1}{2}\right\rceil$ ensures that i will adopt, so the round after that, $k+1$, is an upper bound for the fixpoint. This proves the proposition □

Next up, we substitute the cluster conditions in the KCT. This case is exactly analogous to the cycle with $\theta\leq\frac{1}{2}$.

Proposition 37 (KCT for Wheel with $\theta\leq\frac{1}{3}$)**.** *For any* $\mathfrak{M}=\langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\mathcal{N}=\mathcal{W}_i$ *and* $\theta\leq\frac{1}{3}$, $\mathfrak{M},x\vDash[\because]^{fix}K_i\varphi$ *if and only if there exists* $U\in\bigvee_{j\in Ag}\tau_j$ *such that* $x\in U\subseteq[\![\varphi]\!]$ *and* U *is i-locally dense in* $\Pi_i(x)$.

Proof. Analogous to Proposition 29. □

For thresholds where full cascades are not guaranteed, we observe properties similar to those of the star network and of the cycle, which the wheel is a combination of. Much like in the star network, the agent in the center of the wheel has a high level of authority, but it is somewhat blunted when compared to the star network. On the star network, the evidence of the central agent was diffused to all other agents with no caveat, but here, we observe that the central agent must have at least one other agent to support his case for a piece of evidence, so to speak, for that evidence to reach a full cascade. This is what the next proposition shows.

Proposition 38. *For every* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\mathcal{N} = \mathcal{W}_i$, *if* $\theta \in (\frac{1}{3}; \frac{2}{3}]$, *then for every* $e \in \bigcup_{j \in Ag} \mathcal{E}_j$, *if* $i \in Ag_e$ *and there is some* $j \neq i$ *with* $j \in Ag_e$, *then* $Ag_e^{fix} = Ag$.

Proof. Take an arbitrary evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ with $\mathcal{N} = \mathcal{W}_i$ and $\theta \in (\frac{1}{3}; \frac{2}{3}]$. Take an arbitrary $e \in \bigcup_{j \in Ag} \mathcal{E}_j$ and suppose that $i, j \in Ag_e$ with $i \neq j$. By the definition of $\mathcal{W}_i$, we have that there exist $k, \ell \in \mathcal{N}(j)$ with k, ℓ, i all distinct. Observe that for k, since $i, j \in Ag_e$, we have $\frac{|\mathcal{N}(k) \cap Ag_e|}{|\mathcal{N}(k)|} = \frac{|\{i,j,m\} \cap \{i,j\}|}{|\{i,j,m\}|} = \frac{2}{3} \geq x$ for any $x \in (\frac{1}{3}; \frac{2}{3}]$, so that k adopts e at the next round. Similarly for ℓ. This argument can be repeated, until we reach a full cascade, which we are assured to do, since Ag is finite. Since we picked an arbitrary piece of evidence, this happens for all pieces of evidence. The proposition is proven. □

For $\theta > \frac{1}{3}$, the necessary and sufficient conditions for a full cascade to form for a piece of evidence are not as clear cut as in the other cases, so we do not provide modified versions of the ECT and KCT. We present this as a reliability advantage at the end of the section.

Proposition 39. *For any evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\theta \in (\frac{1}{3}; \frac{2}{3}]$, $\mathfrak{M} \vDash wheel_i \to ([\because]^{\lfloor \frac{|Ag|-1}{2} \rfloor}\varphi \to [\because]^{\lfloor \frac{|Ag|-1}{2} \rfloor + 1}\varphi)$

Proof. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be an arbitrary evidence diffusion model with $\theta \leq \frac{1}{2}$, take an arbitrary world x and suppose $\mathfrak{M}, x \vDash wheel_i$ and $\mathfrak{M}, x \vDash [\because]^{\lfloor \frac{|Ag|-1}{2} \rfloor}\varphi$. Take an arbitrary piece of evidence e and suppose it has the seed set $Ag_e = \{i, j\}$ for some $j \neq i$. Then, for j's neighbors, as $\theta \in (\frac{1}{3}; \frac{2}{3}]$, they will adopt, and so too for their neighbors, and so on. The set of agents is finite, and so we reach a full cascade in $\lfloor \frac{|Ag|-1}{2} \rfloor$, by analogy with the cycle network. Note that if the seed set contains i and is greater, then diffusion can only happen faster, and that if it it does not contain i, but a cascade forms, then necessarily there have to be enough agents to make i adopt, after which the same pattern of diffusion happens, so that it is still faster than the seed set $\{i, j\}$. This provides a clear upper bound to the speed of diffusion, so $\mathfrak{M}, x \vDash [\because]^{\lfloor \frac{|Ag|-1}{2} \rfloor + 1}\varphi$, as needed. □

So, for $\theta \in (\frac{1}{3}; \frac{2}{3}]$, the wheel network is slightly faster than the cycle network, by virtue of being one round faster for some sets of agents.

The following example shows that even for high thresholds, an agent's set of basic pieces of evidence isn't completely determined by their immediate neighbors. We can see this as a consequence of the central agent acting as a 'bridge' between all the other agents.

Example 6. Let $Ag = \{1, ..., 6\}$ and let $\mathcal{N} = \mathcal{W}_1$. Consider an evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ for threshold $\theta > \frac{2}{3}$ and let us focus on a piece of evidence e where $Ag_e = \{2, 4, 5, 6\}$. Below we depict the network of agents, with the red nodes being the agents that are in possession of the piece of evidence e at the depicted round. After one diffusion round, since $\frac{4}{5} > \frac{2}{3}$, 1 receives e in their basic evidence set. Note that at this round, 3 does not receive e since only two out of three connections possess e.

At the next round, however, all of 3's contacts are in possession of e, and so 3 receives it too. So we observe that, unlike in the case of the cycle network, $\mathcal{E}_3^{fix} \neq \mathcal{E}_3 \cup (\mathcal{E}_2 \cap \mathcal{E}_4 \cap \mathcal{E}_1)$, since $e \notin \mathcal{E}_1$, $e \notin \mathcal{E}_3$, but $e \in \mathcal{E}_3^{fix}$.

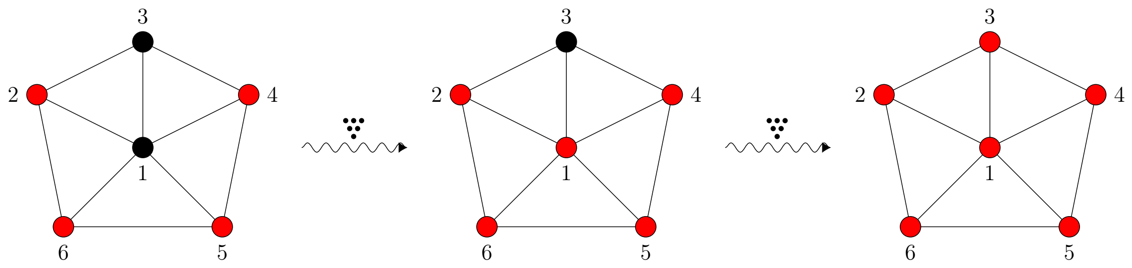

In fact, this example is a great illustration of the upper bound on the fixpoint for diffusion for the wheel network for $\theta > \frac{2}{3}$. The only scenario in which diffusion happens for a non-central agent is when all their neighbors are adopters and they are not. The only scenario in which diffusion happens for a central agent is when $\frac{|Ag_e|}{|Ag|-1} \geq \theta$. A sequence of these two is the maximal diffusion sequence for the wheel network with such a high threshold. Thus, the speed with which the fixpoint is reached is as good as the star network.

Proposition 40. *For any evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ *with* $\theta > \frac{2}{3}$, $\mathfrak{M} \vDash wheel_i \to ([\because]^2\varphi \to [\because]^3\varphi)$

Proof. Let $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!]\rangle$ be an arbitrary evidence diffusion model with $\theta > \frac{2}{3}$, take an arbitrary world x and suppose $\mathfrak{M}, x \vDash wheel_i$ and $\mathfrak{M}, x \vDash [\because]^2\varphi$. Take an arbitrary piece of evidence e. Either $i \in Ag_e$ or $i \notin Ag_e$. If $i \in Ag_e$, then for every j such that $k, \ell \in \mathcal{N}(j)$ with k, ℓ, i all distinct and $k, \ell \in Ag_e$, $j \in Ag_e^1$. After this, it is clear that $Ag_e^1 = Ag_e^2$, since no non-adopting agent j has $k, \ell \in \mathcal{N}(j)$ with k, ℓ, i all distinct and $k, \ell \in Ag_e$. If $i \notin Ag_e$, then either $\frac{|Ag_e|}{|Ag|-1} \geq \theta$, in which case $i \in Ag_e^1$ and the situation from the previous case repeats, or otherwise $i \notin Ag_e$ and so since i is a neighbor of all agents, for all $j \in Ag$, $\frac{|\mathcal{N}(j) \cap Ag_e|}{|\mathcal{N}(j)|} < \theta$. As just demonstrated, the fixpoint is at most 2 in every case. So $\mathfrak{M}, x \vDash [\because]^3\varphi$, as needed. □

Now let us discuss the reliability of this network. This network seems to be the most reliable so far, as it manages to strike a great balance between speed and reliability. A first advantage is that it is for a less broad spectrum of thresholds than for the cycle networks that we encounter the negative scenario in which all pieces of evidence cascade. The second advantage has to do with its relationship with the star and the cycle. For higher thresholds, we did not provide substitutions for the cluster conditions in the ECT and KCT. Though this might initially seem like a negative aspect of the wheel, it is a sign of the diversity of diffusion behavior given different seed sets. This we interpret to be a strength of the network, because it shows a form of rational balance between the authority of the central agent in the star network and the rigidity of evidence diffusion (for high thresholds) in the cycle network.

4.5 Zollman's Findings in This Model

Before we begin this comparison, it is important to note a major caveat: the radically different nature of the two analyses makes it close to impossible to draw any practical conclusions regarding which model is 'better'. As outlined in Chapter 2, in the model Zollman uses there are two states of the world, φ_1 and φ_2, and two actions available to each agent, A_1 and A_2. The agents have expected utilities for these actions, and they form beliefs about which is the actual world given the utilities they receive after acting. Learning means learning which is the actual world In our model, there are possibly infinitely many worlds, and the agents have pieces of evidence they form beliefs and knowledge on the basis of. Learning in our context is broader, as any proposition can be learned. We therefore present our results and this comparison as a *companion* to the work of Zollman [92, 93, 94], and not as a competitor.

Still, it is important to note the strengths and weaknesses of both approaches to the issue. Zollman ran 10,000 simulations on each network [92, p. 579]. Furthermore, for a limited number of agents, he explored the properties of all possible networks [92, §3.2]. As things stand currently, this kind of quantitative analysis is beyond the scope of this thesis. What we have done in this thesis, however, is to provide propositions and theorems, i.e. mathematical facts, about the networks, which is a more difficult task for a quantitative setting. With this in mind, we proceed with the analysis.

In order to compare the results in our framework with those in Zollman's [92, 93, 94], we first discuss the ordering of the networks in terms of connectivity. Earlier in this chapter, we introduced two measures for connectivity: the diameter of the network, and the average node distance. In the table below, we can observe several values for the general functions from Proposition 5, where the diameter is in the first position, and the average node distance is in the second. As is quite obvious, the order of the networks from most connected to least connected is as follows: Total, Wheel, Star, Cycle.

$\lvert Ag \rvert$ / Network	4	5	6	...
Total	(1,1)	(1,1)	(1,1)	...
Star	(2, 1.5)	(2,1.6)	(2,1.66667)	...
Cycle	(2,1.333...)	(2,1.5)	(3,1.8)	...
Wheel	(1,1)	(2,1.2)	(2,1.333...)	...

To measure connectivity, Zollman used different graph metrics, namely density, the percentage of possible connections that actually appear in the graph, and the clustering coefficient, the degree to which one's neighbors are connected to each other [92, 66]. For the Total, Cycle and Wheel networks, which are the ones we will be focusing on due to Zollman's focus on them, these give the same ordering.

If we are to compare the speed with which the fixpoint in the model is reached with the speed of learning in Zollman's model, our results differ quite drastically. For Zollman, the more highly connected a network is, the faster it converges. In our framework, a comparable trend is that the total network is the fastest, arriving at the fixpoint in one round, and the cycle is potentially the slowest (for $\theta \leq \frac{1}{2}$), with the maximum diffusion speed being $\lfloor \frac{|Ag|}{2} \rfloor$. That being said, the wheel is more connected than the star, but is at *most* as fast as it. Likewise, the cycle is as fast as the very fastest network, the total network, for a large part of the possible thresholds, i.e. for every $\theta > \frac{1}{2}$, despite being the least well connected.

There is a point to be made here about the validity of the comparison. After all, the fixpoint of a diffusion sequence is merely that stage of the diffusion sequence after which nothing more changes. This is quite dissimilar to the way Zollman sets the problem up, where agents have to choose between which of two actions to pursue and update their beliefs and actions based on the payoffs [92, pp. 578–579]. However, recall that our model offers no guarantee of learning, and that in fact the conditions for learning are based on multiple, rather complex factors. And then, if we are assured that a proposition will be learned, then the speed with which the learning occurs would *still* depend on the speed of diffusion, for which speed of reaching the fixpoint is a proxy.

As for the reliability, Zollman finds that more sparsely connected networks more reliably track the truth [92, p. 579]. As discussed above, reliability is a concept that is difficult to model in our framework. For this reason, we will limit ourselves to the reliability analysis we made in terms of the epistemic properties of different networks. In this arena as well, we find that our results differ from Zollman's. The most sparsely connected network we found to be the least reliable, as, for low thresholds it guaranteed the diffusion of misleading evidence to all agents, and for higher thresholds it exhibited a large amount of rigidity, in the sense that an agent received only the pieces of evidence that were possessed by both the agent's neighbors. The next we found to be the total network, where the agents showed an interesting amount of conservativeness to new pieces of evidence. Given a small extra assumption that false pieces of evidence appear less often than true ones (assuming a fixed actual world), we may see this as a good property to have, and one that increases reliability. The most reliable network was the wheel, where full cascades with all pieces of evidence happen on a less broad spectrum of thresholds, and where, for the rest of the spectrum of thresholds, we neither ran into the issue of the star network, where the star agent had too much authority, nor into the issue of the cycle, where an agent rigidly accepts only that which their neighbors can agree on.

One noteworthy aspect of the previous discussion is that the ranking is based on interpretations of the formal properties we obtained and the comparison between our work and Zollman's relies on the different modeling choices that were made. These choices are crucial to take into account when reflecting on the comparison. In future work, other such choices or deviations, for example exploring alternative epistemic principles to rank the networks, or performing a quantitative analysis, would be fascinating to see.

5 Limitations, Conclusion, and Further Research

This thesis has explored the complex dynamics of evidence diffusion in networks of evidence-based reasoners. By bridging the gap between topological evidence models and threshold models, we have developed evidence diffusion models for multi-agent systems, providing insights into the effects of network structure on both individual and group knowledge.

In Chapter 2, we laid the philosophical groundwork for our contributions by providing three arguments for the necessity of Social Epistemology; that is, for its irreducibility to Individual Epistemology with communication. Through the historical case study of Ignaz Semmelweis, we demonstrated the potentially deadly effects of poor scientific network organization. Through the concept of epistemic dependence, we highlighted the fact that agents come to believe and know things by more than individual means. By outlining the Independence Thesis and Zollman's results pertaining to the potential trade-offs between network speed and reliability, we showed that there is ample formal support for Social Epistemology. We furthermore set our thesis up as a continuation of this formal support.

Chapter 3 introduced the formal model that is at the core of this thesis' contribution. In Section 3.1 we outlined the topological preliminaries that we need for our models and results and we outlined their epistemic interpretations. We then presented Topological Evidence Models for single-agent evidence-based belief and knowledge (Section 3.2), Partitional Models as the natural multi-agent extension of Topological Evidence Models (Section 3.3) and Threshold Models for diffusion (Section 3.4) as the parent-frameworks for our own. The chapter culminated with the development of Evidence Diffusion Models in Section 3.5, the main contribution of this thesis. These models allow for individual and group notions of knowledge and belief to be studied in the context of a network of evidence-based reasoners in which pieces of evidence are being communicated via threshold-limited diffusion.

In Chapter 4, we applied the developed models to analyse the dynamics of information flow in the context of evidential networks. In Section 4.1, graph-theoretic network measures were outlined that helped us analyse network connectivity's effect on knowledge. We also highlighted the four special networks to which Section 4.4 was dedicated. Furthermore, we demonstrated our language's power to express the properties of these networks and we showcased the agents' knowledge regarding the networks they find themselves in. In Section 4.2, we focused on the invariance of the individual and group knowledge modalities under the evidence diffusion model transformer. We also provided epistemic interpretations of these formal results, pertaining to the robustness of Distributed Evidence and the fragility of Defeasible Knowledge and (Defeasible) Distributed Knowledge. In Section 4.3, we obtained sweeping results regarding evidence cascades and individual and group learning. These were formulated as 'Cluster Theorems', whereby we can learn, from the available evidence in the network and the network structure in the initial model alone, what evidence agents will have, as well as what they will know individually and in a group at the end of a diffusion sequence. In Section 4.4 we provided speed and reliability results for the four special networks that we defined earlier. Finally, in Section 4.5, we compare these results with Zollman's results on the speed and reliability of the same networks and indeed point out some stark differences between our results. We also point out that these divergences are more than likely caused by the stark differences in modeling paradigms, and that further comparative work is needed.

It is important to point out the limitations of the model at its current stage of development. We presently see three potential interconnected limitations. The first is that no axiomatization was presented, or is currently known, for Evidence Diffusion Models. The second is that our work so far would be difficult to explore quantitatively as well. Second, once computer implementation for Evidence Diffusion Models becomes feasible, simulations could be run so that results more similar to Zollman's in nature could be obtained, and a more faithful comparison could be drawn. Finally, a more rational diffusion process, which involves checking and cross-validation could be implemented to model a more realistic, *softer* social influence. For instance, agents would become *aware* of a piece of evidence once enough of their friends are as well, and they *check it* when this condition is met. Checking a piece of evidence could involve probabilities, whereby true evidence 'replicates' with a very high probability, and false evidence 'replicates' with a very low probability.

We see many avenues for further work beyond the ones mentioned just now. First, one might explore variations of this model where the thresholds are not uniform across agents, i.e. each agent might have their own threshold. This would be a good setting to model the difference between agents that are more or less skeptical of new results. The doxastic and epistemic effects of skepticism in a scientific network of evidence-based reasoners would be fascinating to see.

Second, one might explore a variation of this model where each piece of evidence has its own threshold for adoption. This would seem like a good modelling setting for the plausibility of evidence pieces. In this setting, we can imagine that some evidence pieces that are 'obvious' or 'very plausible' would reach all agents, whereas 'not so plausible' evidence pieces might be limited to only a few clusters. Of course, combinations of the two aforementioned extensions would be highly complex, but also highly interesting.

A third avenue for further research could be the exploration of qualitative expert deference principles. The principle by which agents choose to adopt pieces of evidence in our model is quantitative, since it is based on a proportion, but in Chapter 2 we outlined an expert deference principle that is not based on anything numerical. There are several ways this could look in an extension of our model. For example, the agents in the network might be labelled as 'experts' or 'non-experts' and choose what to believe based on the following principle: non-expert agent i believes φ if and only if there exists $j \in \mathcal{N}(i)$ such that j is an expert and believes φ.

A fourth possible extension of our model would be the addition and exploration of other group modalities. Our line of work mostly focused on potential knowledge that can be obtained by communication, i.e. distributed knowledge (and evidence). A parallel line of work is that which focuses on common knowledge — that which everybody knows, and everybody knows that everybody knows, and so on. Much like the distinction between distributed knowledge and distributed evidence, one might define the analogues of common knowledge and common evidence, and explore their properties under evidence diffusion.

A final avenue for further research would be to probe the Independence Thesis in our model in a more systematic way. This could be done by exploring the differences in individual and group knowledge that the combination of properties of the family of sets of evidence $\{\mathcal{E}_i\}_{i \in Ag}$ and the network $\mathcal{N}$ yield. For example, one might imagine it beneficial for both individual agents and for groups if no agent possesses any false evidence, but if false evidence is relegated to a cluster of the network, then false evidence might be bad individually, but neutral for the distributed knowledge of the group, since the agents outside this false-evidence cluster could still reason with true evidence only.

References

[1] Aldo Iván Ramírez Abarca. "Topological Models for Group Knowlede and Belief". MSc Thesis. Amsterdam: University of Amsterdam, 2015.

[2] Peter Achinstein. *The Book of Evidence*. New York: Oxford University Press, 2001.

[3] Venkatesh Bala and Sanjeev Goyal. "A Noncooperative Model of Network Formation". In: *Econometrica* 68.5 (2000), pp. 1181–1229.

[4] Venkatesh Bala and Sanjeev Goyal. "Learning from Neighbours". In: *The Review of Economic Studies* 65.3 (1998), pp. 595–621.

[5] Alexandru Baltag. *Topology, Logic and Learning Lecture 1.1*. University of Amsterdam, 2023.

[6] Alexandru Baltag, Nick Bezhanishvili, and David Fernández-Duque. "The Topology of Surprise". In: *Proceedings of the Nineteenth International Conference on Principles of Knowledge Representation and Reasoning*. Haifa, Israel: International Joint Conferences on Artificial Intelligence Organization, 2022, pp. 33–42.

[7] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. "A Topological Approach to Full Belief". In: *Journal of Philosophical Logic* 48.2 (2019), pp. 205–244.

[8] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. "Justified Belief and the Topology of Evidence". In: *Logic, Language, Information, and Computation*. Ed. by Jouko Väänänen, Åsa Hirvonen, and Ruy De Queiroz. Vol. 9803. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 83–103.

[9] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. "Justified Belief, Knowledge, and the Topology of Evidence". In: *Synthese* 200.6 (2022), p. 512.

[10] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. "The Topology of Belief, Belief Revision and Defeasible Knowledge". In: *Logic, Rationality, and Interaction*. Ed. by David Hutchison et al. Vol. 8196. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 27–40.

[11] Alexandru Baltag, Zoé Christoff, Jens Ulrik Hansen, and Sonja Smets. "Logical Models of Informational Cascades". In: *Logic across the University: Foundations and Applications: Proceedings of the Tsinghua Logic Conference, Beijing*. Ed. by Johan van Benthem and Fenrong Liu. London: College Publications, 2013.

[12] Alexandru Baltag, Zoé Christoff, Rasmus K. Rendsvig, and Sonja Smets. "Dynamic Epistemic Logics of Diffusion and Prediction in Social Networks". In: *Studia Logica* 107.3 (2019), pp. 489–531.

[13] Alexandru Baltag, Nina Gierasimczuk, and Sonja Smets. "Belief Revision as a Truth-Tracking Process". In: *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*. Groningen The Netherlands: ACM, 2011, pp. 187–190.

[14] Alexandru Baltag, Nina Gierasimczuk, and Sonja Smets. "On the Solvability of Inductive Problems: A Study in Epistemic Topology". In: *Electronic Proceedings in Theoretical Computer Science* 215 (2016), pp. 81–98.

[15] Alexandru Baltag and Bryan Renne. "Dynamic Epistemic Logic". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Metaphysics Research Lab, Stanford University, 2016.

[16] Garrett Birkhoff. *Lattice Theory*. Providence, RI: American Mathematical Society, 1995.

[17] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge: Cambridge University Press, 2010.

[18] Kay Codell Carter and Barbara R. Carter. *Childbed Fever: A Scientific Biography of Ignaz Semmelweis, with a New Introduction by the Authors*. New Brunswick, N.J: Transaction Publishers, 2005.

[19] Damon Centola and Michael Macy. "Complex Contagions and the Weakness of Long Ties". In: *American Journal of Sociology* 113.3 (2007), pp. 702–734.

[20] Tom Chivers. "What's next for Psychology's Embattled Field of Social Priming". In: *Nature* 576.7786 (2019), pp. 200–202.

[21] Zoé Christoff. "Dynamic Logics of Networks: Information Flow and the Spread of Opinion". PhD thesis. University of Amsterdam, 2016.

[22] Zoé Christoff and Davide Grossi. "Stability in Binary Opinion Diffusion". In: *Logic, Rationality, and Interaction*. Ed. by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada. Vol. 10455. Berlin, Heidelberg: Springer Berlin Heidelberg, 2017, pp. 166–180.

[23] Zoé Christoff and Jens Ulrik Hansen. "A Logic for Diffusion in Social Networks". In: *Journal of Applied Logic* 13.1 (2015), pp. 48–77.

[24] Zoé Christoff and Jens Ulrik Hansen. "A Two-Tiered Formalization of Social Influence". In: *Logic, Rationality, and Interaction*. Ed. by David Hutchison et al. Vol. 8196. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 68–81.

[25] M. Clark. "Knowledge and Grounds: A Comment on Mr. Gettier's Paper". In: *Analysis* 24.2 (1963), pp. 46–48.

[26] C. A. J. Coady. *Testimony: A Philosophical Study*. New York: Oxford University Press, 1992.

[27] Hans van Ditmarsch, Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. "Knowledge and Simplicial Complexes". In: *Philosophy of Computing*. Ed. by Björn Lundgren and Nancy Abigail Nuñez Hernández. Cham: Springer International Publishing, 2022, pp. 1–50.

[28] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Pieter Kooi. *Dynamic Epistemic Logic*. v. 337. Dordrecht: Springer, 2007.

[29] David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge: Cambridge University Press, 2010.

[30] Adam Elga. "Reflection and Disagreement". In: *Noûs* 41.3 (2007), pp. 478–502.

[31] Ryszard Engelking. *General topology*. Berlin: Heldermann, 1989.

[32] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Vardi. *Reasoning About Knowledge*. The MIT Press, 1995.

[33] Saúl Fernández González. "Generic Models for Topological Evidence Logics". MSc Thesis. University of Amsterdam, 2018.

[34] Viktor Ya Frenkel. "On the History of the Einstein–de Haas Effect". In: *Soviet Physics Uspekhi* 22.7 (1979), p. 580.

[35] Edmund L. Gettier. “Is Justified True Belief Knowledge?” In: *Analysis* 23.6 (1963), pp. 121–123.

[36] Nina Gierasimczuk. “Knowing One’s Limits: Logical Analysis of Inductive Inference”. PhD thesis. Amsterdam: University of Amsterdam, 2010.

[37] Nina Gierasimczuk, Vincent F. Hendricks, and Dick De Jongh. “Logic and Learning”. In: *Johan van Benthem on Logic and Information Dynamics*. Ed. by Alexandru Baltag and Sonja Smets. Vol. 5. Cham: Springer International Publishing, 2014, pp. 267–288.

[38] Alvin I. Goldman. *Knowledge in a Social World*. Oxford : New York: Oxford University Press, 1999.

[39] Sanjeev Goyal. “Interaction Structure and Social Change”. In: *Journal of Institutional and Theoretical Economics (JITE) / Zeitschrift für die gesamte Staatswissenschaft* 152.3 (1996), pp. 472–494.

[40] Mark Granovetter. “The Strength of Weak Ties: A Network Theory Revisited”. In: *Sociological Theory* 1 (1983), p. 201.

[41] Mark Granovetter. “Threshold Models of Collective Behavior”. In: *American Journal of Sociology* 83.6 (1978), pp. 1420–1443.

[42] John Greco. “What Is Social Epistemic Dependence?” In: *Philosophical Topics* 49.2 (2021), pp. 113–132.

[43] Joseph Y. Halpern and Yoram Moses. “Knowledge and Common Knowledge in a Distributed Environment”. In: *Journal of the ACM* 37.3 (1990), pp. 549–587.

[44] Pelle G. Hansen, Vincent F. Hendricks, and Rasmus K. Rendsvig. “Infostorms”. In: *Metaphilosophy* 44.3 (2013), pp. 301–326.

[45] John Hardwig. “Epistemic Dependence”. In: *The Journal of Philosophy* 82.7 (1985), pp. 335–349.

[46] F. A. Hayek. “The Use of Knowledge in Society”. In: *The American Economic Review* 35.4 (1945), pp. 519–530.

[47] Jaakko Hintikka. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ed. by Vincent F. Hendricks and John Symons. London: King’s College London Publications, 2005.

[48] Jonathan Jenkins Ichikawa and Matthias Steup. “The Analysis of Knowledge”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Metaphysics Research Lab, Stanford University, 2018.

[49] Daniel Kahneman. *Thinking, Fast and Slow*. London: Penguin Books, 2012.

[50] Kevin T. Kelly. “A Computational Learning Semantics for Inductive Empirical Knowledge”. In: *Johan van Benthem on Logic and Information Dynamics*. Ed. by Alexandru Baltag and Sonja Smets. Vol. 5. Cham: Springer International Publishing, 2014, pp. 289–337.

[51] Kevin T. Kelly. *The Logic of Reliable Inquiry*. New York: Oxford University Press, 1996.

[52] Thomas Kelly. “Evidence”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Metaphysics Research Lab, Stanford University, 2016.

[53] David Kempe, Jon Kleinberg, and Éva Tardos. "Maximizing the Spread of Influence through a Social Network". In: *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Washington, D.C.: ACM, 2003, pp. 137–146.

[54] Philip Kitcher. *The Advancement of Science: Science Without Legend, Objectivity Without Illusions*. New York: Oxford University Press, 1993.

[55] Peter D. Klein. "A Proposed Definition of Propositional Knowledge". In: *The Journal of Philosophy* 68.16 (1971), pp. 471–482.

[56] Peter David Klein. *Certainty, a Refutation of Scepticism*. Minneapolis: University of Minnesota Press, 1981.

[57] Keith Lehrer. *Theory of Knowledge*. Boulder, Colo.: Westview Press, 1990.

[58] Keith Lehrer and Thomas Paxson. "Knowledge: Undefeated Justified True Belief". In: *The Journal of Philosophy* 66.8 (1969), pp. 225–237.

[59] Kevin Leyton-Brown and Yoav Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. San Rafael, California: Morgan & Claypool Publishers, 2008.

[60] Fenrong Liu, Jeremy Seligman, and Patrick Girard. "Logical Dynamics of Belief Change in the Community". In: *Synthese* 191.11 (2014), pp. 2403–2431.

[61] Cheng Long and Raymond Chi-Wing Wong. "Minimizing Seed Set for Viral Marketing". In: *2011 IEEE 11th International Conference on Data Mining*. Vancouver, BC, Canada: IEEE, 2011, pp. 427–436.

[62] Conor Mayo-Wilson, Kevin J. S. Zollman, and David Danks. "The Independence Thesis: When Individual and Social Epistemology Diverge". In: *Philosophy of Science* 78.4 (2011), pp. 653–677.

[63] John-Jules Ch Meyer and Wiebe van der Hoek. *Epistemic Logic for AI and Computer Science*. 41. Cambridge: Cambridge Univ. Press, 2004.

[64] Stephen Morris. "Contagion". In: *The Review of Economic Studies* 67.1 (2000), pp. 57–78.

[65] James Raymond Munkres. *Topology*. Upper Saddle River, NJ: Prentice Hall, 2000.

[66] M. E. J. Newman, D. J. Watts, and S. H. Strogatz. "Random Graph Models of Social Networks". In: *Proceedings of the National Academy of Sciences* 99.suppl. 1 (2002), pp. 2566–2572.

[67] Cailin O'Connor, Sanford Goldberg, and Alvin Goldman. "Social Epistemology". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Metaphysics Research Lab, Stanford University, 2023.

[68] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. Cambridge, Mass: MIT Press, 1994.

[69] Aybüke Özgün. "Evidence in Epistemic Logic : A Topological Perspective". PhD thesis. Amsterdam: University of Amsterdam, 2017.

[70] Aybüke Özgün. "Topological Models for Belief and Belief Revision". MSc Thesis. Amsterdam: University of Amsterdam, 2013.

[71] Floris Roelofsen. "Distributed Knowledge". In: *Journal of Applied Non-Classical Logics* 17.2 (2007), pp. 255–273.

[72] Ji Ruan and Michael Thielscher. "A Logic for Knowledge Flow in Social Networks". In: *AI 2011: Advances in Artificial Intelligence*. Ed. by Dianhui Wang and Mark Reynolds. Vol. 7106. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 511–520.

[73] Thomas C. Schelling. *Micromotives and Macrobehavior*. New York, NY: Norton, 2006.

[74] Thomas C. Schelling. "Models of Segregation". In: *The American Economic Review* 59.2 (1969), pp. 488–493.

[75] Jeremy Seligman, Fenrong Liu, and Patrick Girard. "Logic in the Community". In: *Logic and Its Applications*. Ed. by Mohua Banerjee and Anil Seth. Vol. 6521. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 178–188.

[76] P. Singh, S. Sreenivasan, B. K. Szymanski, and G. Korniss. "Threshold-Limited Spreading in Social Networks with Multiple Initiators". In: *Scientific Reports* 3.1 (2013), p. 2330.

[77] Sonja Smets and Fernando R. Velázquez-Quesada. "How to Make Friends: A Logical Approach to Social Group Creation". In: *Logic, Rationality, and Interaction*. Ed. by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada. Vol. 10455. Berlin, Heidelberg: Springer Berlin Heidelberg, 2017, pp. 377–390.

[78] Matthias Steup and Ram Neta. "Epistemology". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Metaphysics Research Lab, Stanford University, 2024.

[79] Michael G. Titelbaum. *Fundamentals of Bayesian Epistemology*. New York: Oxford University Press, 2022.

[80] Anne Sjerp Troelstra and Dirk van Dalen. *Constructivism in Mathematics: An Introduction*. Vol. 1 and 2. Amsterdam; New-York; Oxford.: North Holland, 1988.

[81] Thomas W. Valente. "Social Network Thresholds in the Diffusion of Innovations". In: *Social Networks* 18.1 (1996), pp. 69–89.

[82] Johan van Benthem, David Fernández-Duque, and Eric Pacuit. "Evidence and Plausibility in Neighborhood Structures". In: *Annals of Pure and Applied Logic* 165.1 (2014), pp. 106–133.

[83] Johan van Benthem, David Fernández-Duque, and Eric Pacuit. "Evidence Logic: A New Look at Neighborhood Structures". In: *Advances in Modal Logic*. Vol. 9. King's College Press, 2012, pp. 97–118.

[84] Johan van Benthem and Eric Pacuit. "Dynamic Logics of Evidence-Based Beliefs". In: *Studia Logica: An International Journal for Symbolic Logic* 99.1/3 (2011), pp. 61–92.

[85] Hanna Sofie Van Lee. "The Reliability of Scientific Communities: A Logical Analysis". MSc Thesis. Amsterdam: University of Amsterdam, 2015.

[86] Hanna Sofie Van Lee and Sonja Smets. "The Logic of Observation and Belief Revision in Scientific Communities". In: *Journal for General Philosophy of Science* 51.2 (2020), pp. 243–266.

[87] Yì N. Wáng and Thomas Ågotnes. "Public Announcement Logic with Distributed Knowledge: Expressivity, Completeness and Complexity". In: *Synthese* 190.S1 (2013), pp. 135–162.

[88] Stanley Wasserman and Katherine Faust. *Social Network Analysis: Methods and Applications*. 8. Cambridge ; New York: Cambridge University Press, 1994.

[89] Duncan J. Watts, Peter Sheridan Dodds, and John Deighton served as editor and Tulin Erdem served as associate editor for this article. “Influentials, Networks, and Public Opinion Formation”. In: *Journal of Consumer Research* 34.4 (2007), pp. 441–458.

[90] Duncan J. Watts and Steven H. Strogatz. “Collective Dynamics of ‘Small-World’ Networks”. In: *Nature* 393.6684 (1998), pp. 440–442.

[91] Liang Zhen and Jeremy Seligman. “A Logical Model of the Dynamics of Peer Pressure”. In: *Electronic Notes in Theoretical Computer Science* 278 (2011), pp. 275–288.

[92] Kevin J. S. Zollman. “The Communication Structure of Epistemic Communities”. In: *Philosophy of Science* 74.5 (2007), pp. 574–587.

[93] Kevin J. S. Zollman. “The Epistemic Benefit of Transient Diversity”. In: *Erkenntnis* 72.1 (2010), pp. 17–35.

[94] Kevin J.S. Zollman. “Network Epistemology: Communication in Epistemic Communities”. In: *Philosophy Compass* 8.1 (2013), pp. 15–27.