UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Evidence Diffusion in Social Networks: A Topological Perspective

Özgün, A.; Smets, S.; Zotescu, Teodor-Ștefan

DOI
10.1007/978-981-95-2481-5_8

Publication date
2026

Document Version
Final published version

Published in
Logic, Rationality, and Interaction

License
Article 25fa Dutch Copyright Act (https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment)

Link to publication

Citation for published version (APA):
Özgün, A., Smets, S., & Zotescu, T.-Ș. (2026). Evidence Diffusion in Social Networks: A Topological Perspective. In V. Goranko, C. Shi, & W. Wang (Eds.), *Logic, Rationality, and Interaction: 10th International Conference on Logic, Rationality and Interaction, LORI 2025, Xi'an, China, October 16–19, 2025 : proceedings* (pp. 110-123). (Lecture Notes in Computer Science; Vol. 16010), (FoLLI Publications on Logic, Language and Information). Springer. https://doi.org/10.1007/978-981-95-2481-5_8

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:18 Nov 2026

Evidence Diffusion in Social Networks: a Topological Perspective

Aybüke Özgün(✉), Sonja Smets, and Teodor-Ștefan Zotescu

ILLC, University of Amsterdam, Amsterdam, The Netherlands
{a.ozgun,s.j.l.smets}@uva.nl

Abstract. This paper explores the effect of social structures on the beliefs and knowledge of agents who reason in an evidence-based manner. We introduce Evidence Diffusion Models for the formal analysis of multi-agent evidence-based reasoning, focusing on the context of a social network where evidence pieces are being communicated between agents via threshold-limited diffusion. Our models bridge the gap between multi-agent Topological Evidence Models (Partitional Models) and Threshold Models for Diffusion. Firstly, we show that, in our setting, network structures are expressible and known both defeasibly and infallibly by the agents. Then, we prove that defeasible knowledge and (defeasible) 'evidence-based distributed knowledge' are easily lost under the diffusion of pieces of evidence in a network, whereas 'group knowledge with distributed evidence' is strongly robust. Finally, 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.

Keywords: Dynamic Epistemic Logic · Topological Evidence Models · Logics for Social Networks · Distributed Knowledge · Formal Epistemology

1 Introduction

We position this paper within the recent trend of using logic for the formal analysis of social phenomena [2,7,10,14,19–21,24,27]. More specifically we adopt Dynamic Epistemic Logic [3,4,9] to model evidence-based reasoning in social networks. Here we focus on agents in a multi-agent system whose epistemic and doxastic attitudes are directly tied to their available evidence (aligned with earlier work in [3,18]). These agents are able to share their evidence with their neighbors in the network, allowing for evidence to spread in this network via a process of diffusion. Our leading example is the spread of evidence within a scientific community. In a diffusion round, individual agents (e.g. scientists) share evidence, which will then have effects on the epistemic states of other agents and thereby on the epistemic states of groups of agents in the network. In this paper, we are particularly interested in the stability of these epistemic attitudes

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2026
V. Goranko et al. (Eds.): LORI 2025, LNCS 16010, pp. 110–123, 2026.
https://doi.org/10.1007/978-981-95-2481-5_8

under the diffusion process. In the context of evidence-based reasoning, we will distinguish different group attitudes in this paper. We investigate the notions of 'Group Knowledge with Distributed Evidence' and 'Evidence-based Distributed Knowledge', which, as we shall see, behave quite differently. We provide the definitions of these attitudes and analyze their stability or robustness under the process of evidence diffusion. Further, we zoom in on so-called Cluster Theorems, which characterize the evidential and network conditions for evidence cascades to form, and for individuals and groups to obtain a certain epistemic state in the diffusion process.

This paper is organized as follows. In Sect. 2 we provide the preliminary topological notions for evidence-based knowledge, introducing our logic for evidence diffusion (both the language and semantics). In Sect. 3 we continue with an analysis of the dynamics of our models to describe the processes of diffusion, distinguishing evidence diffusion and law distribution. In Sect. 4 we analyze the robustness of our epistemic notions under the diffusion process and state our main cluster theorems. We end with a brief conclusion in Sect. 5.[1]

2 Evidence Models for Diffusion

We introduce now the preliminary topological notions we will be using throughout the paper. For a more in-depth introduction to topology, we refer the reader to a standard topology textbook, e.g. [17].

Topological Preliminaries: A *topological space* is a pair $\mathcal{X} = (X, \tau)$, where X is a non-empty set and $\tau \subseteq \mathcal{P}(X)$ such that (1) $X, \emptyset \in \tau$, (2) for all finite $\mathcal{U} \subseteq \tau$ $\bigcap \mathcal{U} \in \tau$, and (3) for all $\mathcal{U} \subseteq \tau$, $\bigcup \mathcal{U} \in \tau$. Elements of τ are called *open sets* (or *opens*) and their complements are called *closed sets*.

Given a family $\mathcal{E} \subseteq \mathcal{P}(X)$ of subsets of X, there exists a unique, smallest topology $\tau_{\mathcal{E}}$ with $\mathcal{E} \subseteq \tau_{\mathcal{E}}$. The family $\tau_{\mathcal{E}}$ consists of $\emptyset$, X, all finite intersections of the elements of $\mathcal{E}$, and all arbitrary unions of these finite intersections. $\mathcal{E}$ is called a *subbasis* for $\tau_{\mathcal{E}}$, and $\tau_{\mathcal{E}}$ is said to be *generated* by $\mathcal{E}$. The set $\overline{\mathcal{E}}$ of finite intersections of members of $\mathcal{E}$ forms a basis for $\tau_{\mathcal{E}}$.

Given a topological space $\mathcal{X} = (X, \tau)$ and $A \subseteq X$, we define its interior as $\mathrm{Int}(A) = \bigcup\{U \subseteq A \mid U \in \tau\}$, i.e. the greatest open contained in A. The closure of A is $\mathrm{Cl}(A) = \bigcap\{C \supseteq A \mid C = X \setminus U \text{ for some } U \in \tau\}$, i.e. the smallest closed set containing A. We call a set $A \subseteq X$ dense if $\mathrm{Cl}(A) = X$, or equivalently if for all $U \in \tau \setminus \{\emptyset\}$, $A \cap U \neq \emptyset$.

We now introduce evidence diffusion models. Throughout the paper, we work with a non-empty, finite set of agents Ag.

Definition 1 (Evidence Diffusion Model). *An* evidence diffusion model *is a tuple* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!], \theta \rangle$, *where* X *is a non-empty set of*

[1] The results in this paper are based on the work in [27]. Due to the page-limit of this paper, several proofs are omitted from the text and can be found in the appendix that accompanies this paper, which is available here.

possible worlds; $\mathcal{E}_i$ is agent i's set of basic pieces of evidence such that $X \in \mathcal{E}_i$ and $\emptyset \notin \mathcal{E}_i$; Π_i is a partition of X, representing agent i's hard information/infallible knowledge, such that $\Pi_i \subseteq \mathcal{E}_i$; $[\![\cdot]\!] : \mathbb{P} \to X$ is a standardly defined valuation function; and $\theta \in [0,1]$ is a diffusion threshold (fixed for all agents and pieces of evidence). 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 $\ell \in \{1, ..., n\}$, $k_\ell \in \mathcal{N}(k_{\ell-1})$).

This paper fits within the approach to epistemic logic which investigates knowledge not as a primitive notion, but as derived from other pieces of information, called broadly *evidence* [23]. We say broadly because we include in our notion of evidence false or misleading pieces of information, as well as priors, biases and *a priori* knowledge [3,18]. Topological evidence models [3,18], which our setting is based on, replace the usual accessibility relation of Kripke models, which represents epistemic indistinguishability, with a set of basic pieces of evidence for each agent, $\mathcal{E}_i$, the elements of which represent *directly observable* pieces of evidence. Topologically, this is a subbasis, and the elements of the topology generated by this subbasis, τ_i, represent simply evidence, which we understand as any piece of information derived by taking finite conjunctions (intersection) and arbitrary disjunctions (union) of basic evidence. We will also make use of $\overline{\mathcal{E}}_i$, the basis generated by closing $\mathcal{E}_i$ under finite intersections.

We take the same approach to multi-agent topological evidence models as [12], who introduces partitions Π_i of the set of possible worlds for each agent. These partitions represent the agents' hard information, which might not be the same for each agent. Given this, however, we need to *localize* the usual topological and epistemic notions [3,18] to the agents' hard information. Effectively, we introduce a notion of density restricted to the subspace induced by a specific partition cell, which will represent what the agent defeasibly knows, given what she infallibly knows. Let x be a world in X. Then $\Pi_i(x)$ is the unique $\pi \in \Pi_i$ such that $x \in \pi$. We denote by $\tau\big|_{\Pi_i(x)} = \{U \cap \Pi_i(x) \mid U \in \tau\}$, the subspace topology obtained by restricting τ to $\Pi_i(x)$. A set $U \subseteq X$ is i-locally dense in $\Pi_i(x)$ if and only if U is dense in $\tau\big|_{\Pi_i(x)}$.

Finally, we impose three constraints on the network, all motivated by the modeled phenomenon: (1) symmetry, because scientific connections are mutual; (2) connectedness, because excluding isolated agents ensures a proper network; and (3) irreflexivity, so self-connections are excluded and the threshold θ reflects genuine scientific influence.

The *language $\mathcal{L}$ of evidence diffusion for group knowledge* is given by the following grammar, where $p \in \mathbb{P}$, a countable set of propositional atoms, and $i, j \in Ag$:

$$\varphi := p \mid N_{ij} \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid E_i\varphi \mid [\Pi_i]\varphi \mid K_i\varphi \mid \text{∃D}\varphi \mid \text{∀D}\varphi$$

We read N_{ij} as 'j is i's neighbor' and $E_i\varphi$ as 'φ is a basic piece of evidence of agent i'. For the individual knowledge operators, we have $[\Pi_i]\varphi$ read as 'i

infallibly knows that φ', $K_i\varphi$ read as 'i defeasibly knows that φ'. Finally, we read the group knowledge operators as follows: $\mathrm{E}\!\!\!\!-\mathrm{D}\varphi$ is read as 'there is group knowledge with distributed evidence that φ', and $\wr\mathrm{D}\varphi$ as 'there is evidence-based distributed knowledge that φ'. In the following, $[\![\varphi]\!]_{\mathfrak{M}} = \{x \in X \mid \mathfrak{M}, x \vDash \varphi\}$ is the *truth set* of φ in a model $\mathfrak{M}$. We omit the subscript $\mathfrak{M}$ when the model is contextually clear. We follow the usual rules for elimination of the parenthesis in the language.

Definition 2 (Semantics for $\mathcal{L}$). *Given an evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$ and $x \in X$, we interpret $\mathcal{L}$ recursively as:*

$$
\begin{array}{ll}
\mathfrak{M}, x \vDash p & \textit{iff } x \in [\![p]\!] \\
\mathfrak{M}, x \vDash N_{ij} & \textit{iff } j \in \mathcal{N}(i) \\
\mathfrak{M}, x \vDash \neg\varphi & \textit{iff not } \mathfrak{M}, x \vDash \varphi \\
\mathfrak{M}, x \vDash \varphi \wedge \psi & \textit{iff } \mathfrak{M}, x \vDash \varphi \textit{ and } \mathfrak{M}, x \vDash \psi \\
\mathfrak{M}, x \vDash E_i\varphi & \textit{iff } [\![\varphi]\!] \in \mathcal{E}_i \\
\mathfrak{M}, x \vDash [\Pi_i]\varphi & \textit{iff } \Pi_i(x) \subseteq [\![\varphi]\!] \\
\mathfrak{M}, x \vDash K_i\varphi & \textit{iff there exists } U \in \tau_i \textit{ such that } x \in U \subseteq [\![\varphi]\!] \\
 & \quad \textit{and } U \textit{ is } i\textit{-locally dense in } \Pi_i(x)
\end{array}
$$

Note that our E_i operator is different from the usual E_i operator on topological evidence models [3,12,18], where it means 'agent i has a basic piece of evidence *for* φ', which is interpreted as $\exists e \in \mathcal{E}_i(e \subseteq [\![\varphi]\!])$.

For the distributed operators we write the definitions separately, as they are more complex. We begin with a topological concept which is necessary for group knowledge with distributed evidence, the join topology, which is the smallest topology containing all the agents' topologies.

Definition 3 (Join Topology). *Let X be a set of possible worlds and 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.*

Definition 4 (Group Knowledge with Distributed Evidence). $\mathfrak{M}, x \vDash \mathrm{E}\!\!\!\!-\mathrm{D}\varphi$ *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* $\vee\tau_\cap = \bigvee_{i\in Ag} \tau_i\Big|_{\bigcap_{i\in Ag} \Pi_i(x)}$;
   (If the agents communicated all their information, both hard and soft, 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 we concur with [12] that it is not the appropriate notion of evidence-based distributed knowledge. For a more detailed analysis of why this is, consult [12,27].

In our framework, since we encode inter-agent communication, we show in Proposition 5 that this coincides with the knowledge of an agent given certain communication conditions. As [18] 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 group knowledge with distributed evidence (the actual world doesn't make $\mathbb{D}\varphi$ true).

Definition 5 (Evidence-based Distributed Knowledge). $\mathfrak{M}, x \vDash \mathbb{D}\varphi$ *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 [11,13], 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 shown in Proposition 7, this notion is not robust: it is very easily lost under evidence diffusion.

The definition of *logical consequence* is standard [5]: for any $\Gamma \subseteq \mathcal{L}$ and $\varphi \in \mathcal{L}$, Γ *logically entails* φ, $\Gamma \vDash \varphi$, iff for all models $\mathfrak{M}$ and all $x \in X$: if $\mathfrak{M}, x \vDash \psi$ for all $\psi \in \Gamma$, then $\mathfrak{M}, x \vDash \varphi$. We write $\mathfrak{M} \vDash \varphi$ if for every world $x \in X$, $\mathfrak{M}, x \vDash \varphi$; and $\vDash \varphi$ means that $\mathfrak{M} \vDash \varphi$ for every evidence diffusion model $\mathfrak{M}$, that is, φ is *valid*.

3 Evidential Dynamics

We now go on to define two dynamic operators that capture, on the one hand the sharing of soft information, or evidence, and on the other hand the sharing of hard information, or certain knowledge. We call the former Evidence Diffusion and the latter Law Distribution.

We model the spread of evidence across the network as a communication process which is *social*, *local* [22], and *threshold-limited* [10]. An agent i will adopt a piece of evidence e if out of her 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 the threshold θ. We restrict attention to the basic pieces of evidence 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.

Definition 6 (Evidence Diffusion). *The result of applying evidence diffusion to the evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$ *is the evidence diffusion model* $\mathfrak{M}^{\sim} = \langle X, \{\mathcal{E}_i^{\sim}\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$*, where for any* $i \in Ag$*,*

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

Note that evidence diffusion does not change the threshold. When an agent receives a new piece of evidence through diffusion, we say that agent 'adopted' that piece of evidence. The diffusion rule could easily be adapted so that combined evidence or opens get transmitted as well. Unlike other models of threshold-limited influence [6,7,10,22], what is subjected to diffusion in our model is not the belief directly, but the 'data' for belief. We believe that this strikes a middle ground between agents who are completely swayed by those around them and those whose epistemic states are completely independent of the evidence of other agents in their network; acknowledging both the effects of social influence on belief and the rational, evidence-based nature of belief.

We call a countably infinite sequence of applications of this transformation a diffusion sequence.

Definition 7 (Diffusion Sequence). *Given an evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$*, 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* $n \in \mathbb{N}$*,* $\mathfrak{M}_{n+1} = \mathfrak{M}_n^{\sim}$*.*

Given a diffusion sequence $\mathcal{S}_{\mathfrak{M}} = \langle \mathfrak{M}_0, \mathfrak{M}_1, ...\rangle$, we denote the basic evidence set of agent i in the model $\mathfrak{M}_k$ for $k \in \mathbb{N}$ by $\mathcal{E}_i^k$. Similarly $\overline{\mathcal{E}}_i^{\,k}$ denotes the associated combined evidence set, τ_i^k their topology of evidence, and $[\![\varphi]\!]^k$ the truth set of proposition φ at diffusion step k.

The sharing of hard information is captured using the partitions that encode each agent's infallible knowledge. When all agents combine their infallible knowledge, they arrive at a common partition of possible worlds: the coarsest common refinement of their individual partitions. This is the least refined partition that incorporates all their hard information. Note that the change to the agents' basic evidence sets is required to make $\mathfrak{M}^{\bigwedge}$ an evidence diffusion model.

Definition 8 (Coarsest Common Refinement). *Given finitely many partitions* $\Pi_1, ..., \Pi_n$ *of* X*, the* coarsest common refinement*,* $\bigwedge_{i=1}^n \Pi_i$*, of these partitions is defined as:*

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

Definition 9 (Law Establishment). *Given an evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$*, the result of applying law establishment to* $\mathfrak{M}$ *is the model* $\mathfrak{M}^{\bigwedge} = \langle X, \{\mathcal{E}_i^{\bigwedge}\}_{i\in Ag}, \{\Pi_i^{\bigwedge}\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$*, where for every* $i \in Ag$*,* $\Pi_i^{\bigwedge} = \bigwedge_{j\in Ag} \Pi_j$ *and* $\mathcal{E}_i^{\bigwedge} = \mathcal{E}_i \cup \Pi_i^{\bigwedge}$*.*

Having defined these model transformations, we extend our language with corresponding dynamic modalities and obtain natural semantic clauses for them.

The *dynamic language* $\mathcal{L}^{[]}$ *of evidence diffusion for group knowledge* is obtained by extending the static language $\mathcal{L}$ with the dynamic modalities $[\sim]\varphi$ (read 'adopt φ') and $[\bigwedge]\varphi$ (read 'laws φ') for evidence diffusion and law establishment, respectively.

Given an evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$ and $x \in X$, we interpret the new dynamic operators as follows:

$$\begin{array}{l} \mathfrak{M}, x \vDash [\sim]\varphi \ \text{ iff } \ \mathfrak{M}^{\sim}, x \vDash \varphi \\ \mathfrak{M}, x \vDash [\bigwedge]\varphi \ \text{ iff } \ \mathfrak{M}^{\bigwedge}, x \vDash \varphi. \end{array}$$

4 Robustness Under Diffusion and Cluster Theorems

One characteristic of this framework, highlighted by the validity in Proposition 1, is that the network structure is transparent to all agents, that is, all agents infallibly (and also defeasibly), know who is connected to whom in the network.

Proposition 1. *For every* $i, j, k \in Ag$, $\vDash N_{jk} \leftrightarrow [\Pi_i]N_{jk}$ *and* $\vDash N_{jk} \leftrightarrow K_i N_{jk}$.

Proof. Follows from the fact that the semantic clause for N_{jk} is world-independent, that is $[\![N_{jk}]\!] = X$ or $[\![N_{jk}]\!] = \emptyset$. Moreover, $x \in \Pi_i(x) \subseteq Cl_{\tau_i}(\Pi_i(x)) \subseteq X$.

Even though we defined diffusion sequences as countably infinite, the following proposition establishes that a fixpoint is always reached at some step within the sequence. In other words, diffusion always 'ends' at some point and the model no longer changes.

Notation 1. *For an evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$ *, we let* $Ag_e = \{i \in Ag \mid e \in \mathcal{E}_i\}$, the set of all agents who have e as a basic piece of evidence. *Given the diffusion sequence* $\mathcal{S}_{\mathfrak{M}} = \langle \mathfrak{M}, \mathfrak{M}_1, ...\rangle$ *and* $k \in \mathbb{N}$, *we let* $Ag_e^k = \{i \in Ag \mid e \in \mathcal{E}_i^k\}$, the set of all agents who have e as a basic piece of evidence at the k^{th} diffusion.

Proposition 2. *For any evidence diffusion model* $\mathfrak{M}$, *there exists some* $n \in \mathbb{N} \leq |Ag| - 1$ *where the diffusion sequence* $\mathcal{S}_{\mathfrak{M}}$ *reaches a fixed point, that is, 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 $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$ be an evidence diffusion model. By the definition of $\mathcal{S}_{\mathfrak{M}}$, for any $i \in Ag$ and $k \in \mathbb{N}$, $\mathcal{E}_i^k \subseteq \mathcal{E}_i^{k+1}$. Thus, for all $e \in \bigcup_{i\in Ag} \mathcal{E}_i$, $Ag_e^k \subseteq Ag_e^{k+1}$. 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 $|Ag| - 1$.

Using the dynamic language of evidence diffusion for group knowledge, we can capture whether or not an agent receives a piece of evidence at the next diffusion round.

Definition 10 (Free Evidence for agent i). *Given an evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$, we call a piece of evidence $e \in \bigcup_{i\in Ag} \mathcal{E}_i$* free for i in $\mathfrak{M}$ *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 $\frac{|\mathcal{G}|}{|\mathcal{N}(i)|} \geq \theta$.*

Abbreviation 1. *We define the formula $[\sim]^n\varphi$ as the abbreviation defined recursively by $[\sim]^0\varphi := \varphi$ and $[\sim]^{n+1}\varphi := [\sim][\sim]^n\varphi$.*

Making use of this abbreviation, we can express that some piece of evidence φ is free for some agent i.

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

Before we proceed with our analysis of robustness results, we want to restrict attention to Boolean formulas. Our agents share their direct observations and infallible first-order knowledge, not what they know about other agents' evidence or knowledge, so it stands to reason that we focus on first-order group and individual knowledge and evidence. The following proposition establishes that the truth sets of Boolean formulas does not change under our model transformations. We denote by $\mathcal{L}_{CPL}$ the language of classical propositional logic.

Proposition 3. *Given an evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$ and $\varphi \in \mathcal{L}_{CPL}$, $[\![\varphi]\!]_{\mathfrak{M}} = [\![\varphi]\!]_{\mathfrak{M}^\sim} = [\![\varphi]\!]_{\mathfrak{M}^\wedge}$.*

The following validity describes that the only way an agent i comes to possess evidence after diffusion is if they already had that piece of evidence, or if it was free for i in the initial model.

Proposition 4. *For any $i \in Ag$ and any $\varphi \in \mathcal{L}_{CPL}$, $\vDash [\sim]E_i\varphi \leftrightarrow \Phi_i\varphi \vee E_i\varphi$.*

The following proposition shows that defeasible evidence is the knowledge of a hypothetical agent that has all of the hard knowledge ('laws') and soft knowledge (evidence) of all of the agents combined. The definition of group knowledge with distributed evidence (Definition 4) already hinted at this fact.

Proposition 5. *Given an evidence diffusion model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$, if for every $e \in \bigcup_{j\in Ag} \mathcal{E}_j$, e is free for i in some model in $\mathcal{S}_{\mathfrak{M}}$, then for any $\varphi \in \mathcal{L}_{CPL}$, $\mathfrak{M} \vDash \text{ED}\varphi \leftrightarrow [\bigwedge][\sim]^{fix}K_i\varphi$.*

Since distributed evidence is the knowledge of this *hypothetical* agent, it stands to reason that the *actual* diffusion pattern of a model does not affect it—distributed evidence is robust (in fact, invariant) under diffusion.

Proposition 6. *For any* φ *in* $\mathcal{L}_{CPL}$, $\vDash \exists\!\!\!\!\mathbf{D}\varphi \leftrightarrow [\sim]\exists\!\!\!\!\mathbf{D}\varphi$

In stark contrast, our notion of evidence-based distributed knowledge is not robust under evidence diffusion. It is so easily lost under evidence diffusion that we cannot assure its preservation even if a proposition is known both distributively and individually by all the agents before evidence diffusion.

Proposition 7. $\nvDash (\not\exists\!\!\!\!\mathbf{D}\varphi \wedge \bigwedge_{i\in Ag} K_i\varphi) \rightarrow [\sim]\not\exists\!\!\!\!\mathbf{D}\varphi$

Proof. Consider the evidence diffusion model $\mathfrak{M} = \langle\{1,2,3,4\}, \{\mathcal{E}_{\text{Alice}}, \mathcal{E}_{\text{Bob}}\}, \{\Pi_{Alice}, \Pi_{Bob}\}, \mathcal{N}, [\![\cdot]\!], 0.5\rangle$ for two agents $Ag = \{Alice, Bob\}$, where $[\![p]\!] = P = \{1,2\}$ (green in Fig. 1) and $[\![q]\!] = Q = \{1,3\}$ (blue in Fig. 1); $\mathcal{E}_{\text{Alice}} = \{X, P, Q\}$ and $\mathcal{E}_{\text{Bob}} = \{X, P, \neg(P \leftrightarrow Q)\}$; $\Pi_{\text{Alice}} = \Pi_{\text{Bob}} = \{X\}$, and Alice and Bob are connected (as depicted by the arrows in Fig. 1). We stipulate that 1 is the actual world. Since the partitions of both agents is the whole space, local density with respect to the agents' partitions and global density coincide.

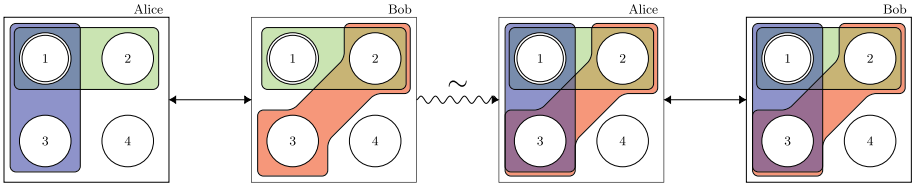

Fig. 1. Distributed knowledge loss

$\mathfrak{M}, 1 \vDash \not\exists\!\!\!\!\mathbf{D}p$ since $[\![p]\!] = P$ itself is a dense open set both in τ_{Alice} and τ_{Bob} and $1 \in P$. For the same reasons, we also have $\mathfrak{M}, 1 \vDash K_{\text{Alice}}p$ and $\mathfrak{M}, 1 \vDash K_{\text{Bob}}p$, thus, $\mathfrak{M}, 1 \vDash \bigwedge_{i\in Ag} K_i p$. However, $\mathfrak{M}, 1 \nvDash [\sim]\not\exists\!\!\!\!\mathbf{D}p$: after evidence diffusion, we obtain that $\{3\} \in \tau^{\sim}_{Alice}$ and $\{3\} \in \tau^{\sim}_{Bob}$. So, any dense open set $U_1 \in \tau^{\sim}_{Alice}$ and $U_2 \in \tau^{\sim}_{Bob}$ is such that $3 \in U_1 \cap U_2 \not\subseteq P$.

One condition that is sufficient for the preservation of evidence-based distributed knowledge under evidence diffusion is that no agent receives so-called misleading evidence (defined in [3,18]) through the diffusion process.

Definition 11 (Misleading evidence). *We call a piece of evidence* e *i-misleading at* $x \in X$ *if and only if there is some* $e' \in \overline{\mathcal{E}}_i$ *such that* $x \notin (e \cap e')$ *and* $(e \cap e') \notin \overline{\mathcal{E}}_i \cup \{\emptyset\}$.

Proposition 8. *Given an evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$ *such that for all* i*, there is no* i*-misleading piece of evidence that is free for* i*, we have that for any* $\varphi \in \mathcal{L}_{CPL}$, $\mathfrak{M} \vDash \not\exists\!\!\!\!\mathbf{D}\varphi \rightarrow [\sim]\not\exists\!\!\!\!\mathbf{D}\varphi$.

The condition that no agent i receives an i-misleading piece of evidence is, however, not a necessary condition for the preservation of evidence-based distributed knowledge. For a simple example, consider a variation $\mathfrak{M}'$ of the

model presented in the proof of Proposition 7, where $\mathcal{E}'_{\text{Alice}} = \{X, P\}$ and $\mathcal{E}'_{\text{Bob}} = \{X, \{2\}\}$. In this case, $\mathfrak{M}', 1 \vDash \langle\!\rangle \mathbf{D}p$, and $\{2\} \in \mathcal{E}_{\text{Bob}}$ is misleading for Alice. Yet, $\mathfrak{M}', 1 \vDash [\sim]\langle\!\rangle \mathbf{D}p$.

The following theorem specifies the conditions under which a so-called *cascade* [6,10] is formed for a piece of evidence, i.e. when it propagates to all agents. It makes use of the notion of clusters, which we introduce now, following [6,10].

Definition 12 (Cluster of Density d). *Given a set of agents Ag and network $\langle Ag, \mathcal{N}\rangle$ defined on Ag, a subset of the set of agents $C \subseteq Ag$ is called a* cluster of density d *if and only if for all $i \in C$, $\frac{|\mathcal{N}(i) \cap C|}{|\mathcal{N}(i)|} \geq d$.*

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

Corollary 1. *For every model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$ with $\theta \neq 0$, $e \in \bigcup_{i\in Ag} \mathcal{E}_i$, and $i \in Ag$, $e \in \mathcal{E}_i^{fix}$ 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 $e \in \mathcal{E}_j$.*

From this corollary, an agent will end up adopting a piece of evidence if and only if there is an 'early adopter' [10, Ch. 19] of that piece of evidence in every cluster of density greater than $1 - \theta$ that she is 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 though, such as the case of Ignaz Semmelweis [27, Ch. 2], it certainly seems like precisely this kind of 'groupthink' is at play.

Corollary 1 not only provides the essential ingredient to establishing the necessary and sufficient conditions for evidence-based 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 2 (Knowledge Cluster Theorem). *For any model $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i\in Ag}, \{\Pi_i\}_{i\in Ag}, \mathcal{N}, [\![\cdot]\!], \theta\rangle$ with $\theta \neq 0$, $x \in X$ and $\varphi \in \mathcal{L}_{CPL}$, $\mathfrak{M}, x \vDash [\sim]^{fix} K_i\varphi$ if and only if there exists $E \subseteq \bigcup_{j\in Ag} \mathcal{E}_j$ such that:*

1. *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*
2. *For every cluster C of density greater than $1 - \theta$ containing i:*
   (a) for every $e \in E$ there is some $j \in C$ such that $e \in \mathcal{E}_j$; and

(b) for all $E' \subseteq \bigcup_{j\in Ag} \mathcal{E}_j$ *and* $\bigcap E' \neq \emptyset$*, 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* $\bigcap E' \cap (U \cap \Pi_i(x)) \neq \emptyset$*.*

Before presenting the proof, we outline the theorem's key conditions. It consists of two parts: **receiving** and **blocking**.

In the receiving part, an agent attains individual knowledge of φ at the end of the diffusion sequence if and only if there exists a subset of the total evidence satisfying certain conditions. This subset functions as an 'information package' the agent must acquire to know φ. First, adding this package to the agent's basic evidence must ensure the existence of a set U in the resulting topology such that $U \cap \Pi_i(x) \subseteq [\![\varphi]\!]$, meaning the evidence *supports* φ. Second, by Theorem 1, all elements of this package must become part of the agent's evidence by the end of the diffusion sequence. Thus, the package both exists and is received.

The blocking part ensures that no subset of evidence the agent already has or will receive disrupts the density of U. Again, this follows from the Evidence Cluster Theorem (Theorem 1). Together, these conditions establish a dense subset of $[\![\varphi]\!]$, ensuring that the agent defeasibly knows φ. We now proceed to the proof.

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

($\Rightarrow$) Suppose that $\mathfrak{M}, x \vDash [\sim]^{fix} K_i\varphi$. This means, by the definition of $[\sim]$, that $\mathfrak{M}_{fix}, x \vDash K_i\varphi$, that is, there exists $U \in \tau_i^{fix}$, the topology of evidence generated from $\mathcal{E}_i^{fix}$, such that $U \cap \Pi_i(x) \subseteq [\![\varphi]\!]^{fix} = [\![\varphi]\!]$ (by Proposition 3) 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. 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 τ_i^{fix}, so we know by assumption that $U \in \tau_i^{fix}$ and $x \in U \cap \Pi_i(x) \subseteq [\![\varphi]\!]$, so (1) is proven.

For (2), let C be a cluster of density greater than $1 - \theta$ such that $i \in C$. To prove (2a), suppose further, toward contradiction, that there is $e \in E$ such that $e \notin \mathcal{E}_j$ for all $j \in C$. Then, by Corollary 1, $e \notin \mathcal{E}_i^{fix}$, contradicting $e \in E \subseteq \mathcal{E}_i^{fix}$.

To prove (2b), take an arbitrary $E' \subseteq \bigcup_{j\in Ag} \mathcal{E}_j$ 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 by Corollary 1, 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.

($\Leftarrow$) Now suppose that there exists $E \subseteq \bigcup_{j\in Ag} \mathcal{E}_j$ that satisfies items (1) and (2) and prove that $\mathfrak{M}, x \vDash [\sim]^{fix} K_i\varphi$, that is, there is $U \in \tau_i^{fix}$ such that $x \in U$ and U is i-locally dense in $\Pi_i(x)$ with respect to τ_i^{fix}.

By the assumption, we know that there exists a set $E \subseteq \bigcup_{j\in Ag} \mathcal{E}_j$ such that there exists $U \in \tau_i^{+E}$ 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 Corollary 1, 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 $(U \cap \Pi_i(x)) \cap (A \cap \Pi_i(x)) = \emptyset$. That is, $A \cap (U \cap \Pi_i(x)) = \emptyset$. From this we can infer that there exists $e \in \overline{\mathcal{E}}_i^{fix}$ such that $e \cap (U \cap \Pi_i(x)) = \emptyset$. Moreover, there exists $E' \subseteq \mathcal{E}_i^{fix}$ such that $\bigcap E' \cap (U \cap \Pi_i(x)) = \emptyset$. 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 $E' \subseteq \bigcup_{j \in Ag} \mathcal{E}_j$ 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}.

The following corollary establishes the conditions for what is known in the literature as 'mutual knowledge' [11,16], which amounts to 'everybody knows'.

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

For evidence-based distributed knowledge, the conditions are similar, but they, of course, have to mimic the definition of evidence-based 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 3 (Evidence-Based Distributed Knowledge Cluster Theorem). *For any evidence diffusion model* $\mathfrak{M} = \langle X, \{\mathcal{E}_i\}_{i \in Ag}, \{\Pi_i\}_{i \in Ag}, \mathcal{N}, [\![\cdot]\!], \theta \rangle$ *with* $\theta \neq 0$, $x \in X$ *and* $\varphi \in \mathcal{L}_{CPL}$, $\mathfrak{M}, x \vDash [\sim]^{fix} D\varphi$ *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:*

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

5 Further Work

There are several directions in which this research could be taken further, which includes a continued logical investigation of our formal setting as well as the study of its implications and relevance for debates in philosophy of science, social epistemology and related philosophical studies. Axiomatizing evidence diffusion

models is a natural next step for this line of research. Computer implementations would enable simulations, allowing for quantitative analyses of evidence diffusion on specific models, and in turn more direct comparisons with related work, specifically that of Zollman [25,26]. Another avenue is to explore communication protocols that differ from evidence diffusion, especially those that allow for various social-epistemic factors at play in more fine-grained representations of scientific networks. For instance, one could imagine a diffusion protocol which more closely mimics the replication of results in scientific practice. That is to say, before an agent adopts a piece of evidence from a neighbor, the agent 're-runs the experiment', and only adopts that evidence if the experiment is validated. Whether or not this evidence gets validated can be modeled probabilistically: if the evidence is true (contains the real world), then it replicates with, say, 99% probability, the remaining 1% representing false negatives; and if the evidence is false, it replicates with 5% probability, representing false negatives. Another possible communication protocol worth exploring would be one involving deliberation[2], where agents have to persuade their neighbors to adopt a piece of evidence, similar to [8].

Beyond these refinements, extending the model to allow agent-specific thresholds could better capture variations in skepticism, while evidence-specific thresholds could reflect differences in plausibility. Further, incorporating qualitative expert deference principles—where non-expert agents adopt beliefs based on expert testimony—would offer an alternative to purely quantitative diffusion mechanisms.

Finally, investigating common knowledge and common evidence within the diffusion framework, alongside a systematic exploration of the Independence Thesis [15] in this model, could provide deeper insights into the epistemic trade-offs between network structure and collective reasoning.

Acknowledgement. The research priority area Human(e) AI at the University of Amsterdam has contributed to sponsoring the authors' participation at the conference. We thank the anonymous reviewers of LORI 2025 for their valuable feedback.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abarca, A.I.R.: Topological models for group knowledge and belief. MSc Thesis, University of Amsterdam, Amsterdam (2015)
2. Baltag, A., Bezhanishvili, N., Özgün, A., Smets, S.: A topological approach to full belief. J. Philos. Log. **48**(2), 205–244 (2019)
3. Baltag, A., Bezhanishvili, N., Özgün, A., Smets, S.: Justified belief, knowledge, and the topology of evidence. Synthese **200**(6), 512 (2022)
4. Baltag, A., Renne, B.: Dynamic Epistemic Logic. In: Stanford Encyclopedia of Philosophy, Winter 2016 edn. Metaphysics Research Lab, Stanford University (2016)

[2] We thank one of our anonymous reviewers for this suggestion.

5. Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic. Cambridge University Press (2010)
6. Christoff, Z.: Dynamic logics of networks: information flow and the spread of opinion. Ph.D. thesis, University of Amsterdam (2016)
7. Christoff, Z., Hansen, J.U., Proietti, C.: Reflecting on social influence in networks. J. Log. Lang. Inf. **25**(3–4), 299–333 (2016)
8. Ding, H., Pivato, M.: Deliberation and epistemic democracy. J. Econ. Behav. Organ. **185**, 138–167 (2021)
9. van Ditmarsch, H., van der Hoek, W., Kooi, B.P.: Dynamic Epistemic Logic. Synthese Library, vol. 337. Springer, Dordrecht (2007)
10. Easley, D., Kleinberg, J.: Networks, Crowds, and Markets: Reasoning about a Highly Connected World. Cambridge University Press, Cambridge (2010)
11. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.: Reasoning About Knowledge. The MIT Press (1995)
12. Fernández González, S.: Generic models for topological evidence logics. MSc thesis, University of Amsterdam (2018)
13. Halpern, J.Y., Moses, Y.: Knowledge and common knowledge in a distributed environment. J. ACM **37**(3), 549–587 (1990)
14. Liu, F., Seligman, J., Girard, P.: Logical dynamics of belief change in the community. Synthese **191**(11), 2403–2431 (2014). https://doi.org/10.1007/s11229-014-0432-3
15. Mayo-Wilson, C., Zollman, K.J.S., Danks, D.: The independence thesis: when individual and social epistemology diverge. Philos. Sci. **78**(4), 653–677 (2011)
16. Meyer, J.J.C., van der Hoek, W.: Epistemic Logic for AI and Computer Science, vol. 41. Cambridge University Press, Cambridge (2004)
17. Munkres, J.R.: Topology, 2nd edn. Prentice Hall, Upper Saddle River (2000)
18. Özgün, A.: Evidence in epistemic logic: a topological perspective. Ph.D. thesis, University of Amsterdam, Amsterdam (2017)
19. Seligman, J., Liu, F., Girard, P.: Logic in the community. In: Banerjee, M., Seth, A. (eds.) ICLA 2011. LNCS (LNAI), vol. 6521, pp. 178–188. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-18026-2_15
20. Smets, S., Velázquez-Quesada, F.R.: A logical study of group-size based social network creation. J. Log. Alegbr. Methods Program **106**, 117–140 (2019)
21. Smets, S., Velázquez-Quesada, F.R.: A closeness- and priority-based logical study of social network creation. J. Logic Lang. Inform. **29**(1), 21–51 (2020). https://doi.org/10.1007/s10849-019-09311-5
22. Valente, T.W.: Social network thresholds in the diffusion of innovations. Soc. Netw. **18**(1), 69–89 (1996)
23. van Benthem, J., Fernández-Duque, D., Pacuit, E.: Evidence and plausibility in neighborhood structures. Ann. Pure Appl. Logic **165**(1), 106–133 (2014)
24. Zhen, L., Seligman, J.: A logical model of the dynamics of peer pressure. Electron. Notes Theor. Comput. Sci. **278**, 275–288 (2011)
25. Zollman, K.J.S.: The communication structure of epistemic communities. Philos. Sci. **74**(5), 574–587 (2007)
26. Zollman, K.J.: Network epistemology: communication in epistemic communities. Philos Compass **8**(1), 15–27 (2013)
27. Zotescu, T.Ş.: Multi-agent topological models for evidence diffusion. MSc thesis, University of Amsterdam (2024)