Dynamic Epistemic Logics of Diffusion and Prediction in Social Networks

Alexandru Baltag[1], Zoé Christoff[1], Rasmus K. Rendsvig[2], and Sonja Smets[1]

[1] Institute for Logic, Language and Computation, University of Amsterdam
A.Baltag@uva.nl, zoe.christoff@gmail.com, S.J.L.Smets@uva.nl
[2] Theoretical Philosophy, Lund University
Center for Information and Bubble Studies, University of Copenhagen
rendsvig@gmail.com

Abstract. *We take a logical approach to* threshold models*, used to study the diffusion of opinions, new technologies, infections, or behaviors in social networks. Threshold models consist of a network graph of agents connected by a social relationship and a threshold value which regulates the diffusion process. Agents adopt a new behavior/product/opinion when the proportion of their neighbors who have already adopted it meets the threshold. Under this adoption policy, threshold models develop dynamically towards a guaranteed fixed point. We construct a minimal dynamic propositional logic to describe the threshold dynamics and show that the logic is sound and complete. We then extend this framework with an epistemic dimension and investigate how information about more distant neighbors' behavior allows agents to anticipate changes in behavior of their closer neighbors. Overall, our logical formalism captures the interplay between the epistemic and social dimensions in social networks.*

Keywords: social network theory, threshold models, diffusion in networks, social epistemology, formal epistemology, dynamic epistemic logic, opinion dynamics, opinion dynamics under uncertainty

1 Introduction

An individual's actions or opinions are often influenced by the actions of people around her. The way a new product or fashion gets adopted by a population depends on how agents are influenced by others, which in turn depends both on the way the population is structured and on how influenceable agents are.

This paper focuses on one particular account of social influence, "threshold-limited influence", as presented in e.g. [10,26], relying on an imitation or conformity pressure effect: *agents adopt a behavior/product/opinion/fashion whenever a critical fraction of their neighbors in the network have adopted it already*. In this sense, diffusion in social networks can be seen as a study of local influence, triggering agents to adopt a similar behavior/opinion/product as their neighbors [27,13]. The so-called *threshold models*, first introduced by [12,22], are used precisely to represent the dynamics of diffusion under threshold-limited influence. This type of models has received a lot of attention in the recent literature [10,15,19,25,1,11,17,18].

This paper has two goals. Our first goal is to propose logics for reasoning about threshold models and their dynamics. Our second goal is to investigate how the agents' knowledge affects such dynamics.

After recalling standard threshold models in Subsection 2.1, a dynamic logic for modeling threshold influence within social networks is introduced in Subsection 2.2. While conceptually in line with [24,29,21,23,7,8,20] in using logic to model social influence effects within network structures, our new framework distinguishes itself by avoiding the use of static modalities or hybrid logic tools. In this sense, the logical setting we introduce is "minimal": propositional logic is used to specify both the network structure and the agents behavior, and a single dynamic modality is used to represent the threshold-limited influence. Moreover, while [24,29,23,7,8] focus on the limit thresholds of 100% (all neighbors) and non-0% (at least one neighbor), we allow here for any (uniform) adoption threshold, as is standard within the literature on threshold models. Subsection 2.3 shows how the logic captures the relationship between clusters and diffusion of a behavior to the whole network.

In Section 3 we introduce *epistemic* threshold models. These models come equipped with a specific knowledge-dependent update procedure, called "informed adoption", where agents must possess sufficient information about their surroundings before they adopt. This is a conceptual jump from the initial minimal modeling of influence from Section 2 to a more sophisticated (information dependent) diffusion policy. Instead of modelling the kind of agents who adopt a behavior whenever enough of their neighbors have adopted it already, we focus in secion 2 on agents who adopt whenever they *know* that enough of their neighbors have already adopted. We then relate these two adoption policies by showing under which epistemic conditions their diffusion dynamics is step-wise identical. The section is concluded by extending the logic to a sound and complete dynamic epistemic logic for the epistemic threshold models and the informed update procedure.

We further notice an interesting feature of the informed update procedure. Even though the "informed update" requires that agents have *enough* information to be influenced, the update does not require them to use *all* their available information when making their choices. Hence, if we consider threshold models as representing reflecting agents who are driven by a coordination goal, the new knowledge dependent update procedure makes our agents choose an action even when *they know they could do better*. To overcome this shortcoming, in Section 4, we introduce a third adoption policy, a "prediction update", where agents utilize *all* the available information to *predict the future behavior of other agents in the network*, and act upon their predictions. In other words, they anticipate, and it is common knowledge that they do. We show that the agents' reasoning about other predicting agents always reaches a fixed point and that making adoption dependent on this very fixed point captures the best response of agents trying to coordinate to the best of their knowledge. We give an example illustrating how knowledge about the network and about the behavior of other agents can be interpreted as an "accelerator" of diffusion dynamics, under this last prediction policy: the fixed point of the diffusion process under the prediction update is the same as under the informed update, but it can be reached faster if agents know more about the network around them.

Finally, Section 5 discusses the in-built assumptions of the introduced updates as well as several alternative diffusion policies and Section 6 gives some directions for further research.

2 Threshold Models and their Dynamic Logic

This section introduces the notion of threshold models and designs a logic to capture their dynamics. Subsection 2.1 first reminds the reader of the standard definition of threshold models.

2.1 Threshold Models for Social Influence

A social network may be seen as a graph, where nodes represent agents and edges represent a binary social relationship among them. This paper restricts itself to finite and undirected graphs without self-loops, that is, to symmetric and irreflexive social relationships, e.g. being neighbors or friends. Moreover, we impose that each agent has at least one neighbor in the network, as isolated agents are irrelevant to a discussion of social influence:

Definition 1 (Network). *A network is a pair* $(\mathscr{A}, N)$ *where* $\mathscr{A}$ *is a non-empty finite set of* agents *and the function* $N : \mathscr{A} \to \mathscr{P}(\mathscr{A})$ *assigns a set* $N(a)$ *to each* $a \in \mathscr{A}$*, such that*
– $a \notin N(a)$ *(Irreflexivity),*
– $b \in N(a)$ *if and only if* $a \in N(b)$ *(Symmetry).*
– $N(a) \neq \emptyset$ *(Seriality).*

The simplest type of threshold model consists of such a network together with a unique behavior B (or opinion, fashion, product, or “like-able item”) distributed over $\mathscr{A}$ and a fixed uniform adoption threshold θ. A threshold model thus represents the current spread of behavior B throughout the network, while containing the adoption threshold which prescribes how this spread will evolve.

Definition 2 (Threshold Model). *A threshold model is a tuple* $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ *where* $(\mathscr{A}, N)$ *is a network,* $B \subseteq \mathscr{A}$ *is a* behavior *and* $\theta \in [0,1]$ *is a uniform* adoption threshold.

It is assumed throughout this paper that both the network structure and the adoption threshold stay constant under updates. Thus, the spread of the behavior (i.e., the extension of B) at ensuing time steps may be calculated using the fixed threshold and network structure as follows:

Definition 3 (Threshold Model Update). *The update of threshold model* $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ *is the threshold model* $\mathscr{M}' = (\mathscr{A}, N, B', \theta)$*, where* B' *is given by*

$$B' = B \cup \{a \in \mathscr{A} : \frac{|N(a) \cap B|}{|N(a)|} \geq \theta\}.$$

This definition captures the idea that the new set of agents who adopted the behavior B' (in the new updated model $\mathscr{M}'$) does include the set of agents B who had already adopted the behavior before and it includes those agents who have enough influental neighbors (given by the number θ) that have adoped already. This definition is set in line with the standard approach on adopt rules in the literature [10].

By repeatedly applying this update rule in an initial threshold model, we obtain a unique sequence of threshold models, which we call a diffusion sequence:

Definition 4 (Diffusion Sequence). *Let* $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ *be a threshold model. The diffusion sequence* $\mathscr{S}_{\mathscr{M}}$ *is the sequence of threshold models* $\langle \mathscr{M}_0, \mathscr{M}_1, \mathscr{M}_2, ..., \mathscr{M}_n, \mathscr{M}_{n+1}, ...\rangle$ *such that, for any* $n \in \mathbb{N}$, $\mathscr{M}_n = (\mathscr{A}, N, B_n, \theta)$ *where* B_n *is given by:*

$$B_0 = B \text{ and } B_{n+1} = B'_n.$$

Note that this diffusion process always reaches a fixed point, and that the number of agents in the model gives an upper bound on the number of updates that can be performed before reaching the fixed point:

Proposition 1. *Let* $\mathscr{S}_{\mathscr{M}}$ *be a diffusion sequence. For some* $n \in \mathbb{N} < |\mathscr{A}|$*, we reach a fixed point* $\mathscr{M}_n = \mathscr{M}_{n+1}$ *in the sequence* $\mathscr{S}_{\mathscr{M}}$.

Proof. The fact that there is a $n \in \mathbb{N}$ such that $\mathscr{M}_n = \mathscr{M}_{n+1}$ follows immediately from the fact that $\mathscr{A}$ is finite and $B_n \subseteq B_{n+1}$ for all $n \in \mathbb{N}$. The fact that $n < |\mathscr{A}|$ is given by considering the slowest possible diffusion scenario, i.e. where $|B_0| = 1$ and only one agent adopts per round, i.e. for each $m < n \in \mathbb{N}$, $|B_m| = m + 1$. In this case $\left|B_{|\mathscr{A}|-1}\right| = |\mathscr{A}|$. □

Interpretation. Threshold models and their dynamics may be interpreted in two ways. One interpretation assumes that agents are mere automata and that their behavior is forced upon them by their environment. This interpretation suits the models that are used in e.g. epidemiology: viral infection "just happens" to agents. Alternatively, agents may be interpreted as rational beings aiming towards coordination with their neighbors. In fact, the above update rule also corresponds to the best response dynamics of an associated coordination game [19], under the assumption that there is a 'seed' set of players who always, possibly irrationally, play B [10].

Numerous variations of threshold models exist in the literature, including infinite networks [19], networks with non-inflating behavior adoption [19], agent-specific thresholds [15], weighted links [15] and multiple behaviors [1]. For simplicity, and to fit most examples in the literature, we will stick to the above simpler notion of finite threshold models. The next subsection proposes a logical framework to reason about them.

2.2 The Logic of Threshold-Limited Influence

This section introduces a minimal logic to express the standard notion of threshold-limited influence introduced in the section above. To describe the *situation* of a social network at a given moment, the static language needs to capture two things: who is related to whom and who is displaying the contagious behavior B. In this paper, both features will be encoded using propositional variables. To describe the *change* of situation of a social network, the language includes a dynamic modality. This modality represents how agents adopt the behavior of their neighbors, whenever the given adoption threshold is reached, i.e., whenever enough neighbors have adopted.

Definition 5 (Languages $\mathscr{L}_{[]}$ and $\mathscr{L}$). *Let $\mathscr{A}$ be a finite set and let atoms be given by $\Phi = \{N_{ab} : a, b \in \mathscr{A}\} \cup \{\beta_a : a \in \mathscr{A}\}$. The language $\mathscr{L}_{[]}$ is then given by:*

$$\varphi := N_{ab} \mid \beta_a \mid \neg\varphi \mid \varphi \wedge \varphi \mid [adopt]\varphi$$

The formulas of $\mathscr{L}$ are those of $\mathscr{L}_{[]}$ that do not involve the $[adopt]$-modality.

Disjunction and material implication are defined in the standard way. $\mathscr{L}_{[]}$ is an extension of propositional logic with a unary dynamic modality, denoted $[adopt]$. The language is interpreted over threshold models, using the behavior set and the social network to determine the extension of the atomic formulas. The $[adopt]$ modality is interpreted as is standard in dynamic epistemic logic[3] [3,5,28,6]: intuitively, we evaluate $[adopt]\varphi$ as true "today" if and only if φ is true "tomorrow". Here, "tomorrow" is given by the threshold update of Definition 3.

Definition 6 (Truth Clauses for $\mathscr{L}_{[]}$). *Given a model $\mathscr{M} = (\mathscr{A}, N, B, \theta)$, $N_{ab}, \beta_a \in \Phi$, and $\varphi, \psi \in \mathscr{L}_{[]}$:*

$\mathscr{M} \vDash \beta_a$	*iff*	$a \in B$
$\mathscr{M} \vDash N_{ab}$	*iff*	$b \in N(a)$
$\mathscr{M} \vDash \neg\varphi$	*iff*	$\mathscr{M} \nvDash \varphi$
$\mathscr{M} \vDash \varphi \wedge \psi$	*iff*	$\mathscr{M} \vDash \varphi$ *and* $\mathscr{M} \vDash \psi$
$\mathscr{M} \vDash [adopt]\varphi$	*iff*	$\mathscr{M}' \vDash \varphi$*, where* $\mathscr{M}'$ *is the updated threshold model (Definition 3).*

[3] The dynamic operators in Dynamic Epistemic Logic are taken to be model transformers, they transform a given model into a new model.

Let us also introduce some abbreviations:

Abbreviation. We introduce the formula $[adopt]^n\varphi$ as an abbreviation which is defined recursively:

$$[adopt]^0\varphi := \varphi$$

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

Abbreviation. We introduce the following abbreviation:

$$\beta_{N(a)\geq\theta} := \bigvee_{\{\mathscr{G}\subseteq\mathscr{N}\subseteq\mathscr{A}:\frac{|\mathscr{G}|}{|\mathscr{N}|}\geq\theta\}} (\bigwedge_{b\in\mathscr{N}} N_{ab} \wedge \bigwedge_{b\notin\mathscr{N}} \neg N_{ab} \wedge \bigwedge_{b\in\mathscr{G}} \beta_b)$$

This formula $\beta_{N(a)} \geq \theta$ expresses that the proportion of agent a's neighbors who have adopted is equal to or above the threshold θ.

The following proposition captures within our language the fact (as noted in Prop. 1) that all diffusion sequences stabilize after some finite number of updates, illustrating how our language allows for capturing features of threshold model dynamics, such as stability and stabilization of the diffusion sequence:

Proposition 2. *Let $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ be a threshold model. There exists $n \in \mathbb{N} < |\mathscr{A}|$ such that, for any $\varphi \in \mathscr{L}_{[]}$:*

$$[adopt]^n\varphi \leftrightarrow [adopt]^{n+1}\varphi$$

Proof. As noted in the proof of Proposition 1, in the diffusion sequence $\mathscr{S}_{\mathscr{M}}$, for some $n \in \mathbb{N} < |\mathscr{A}|$, $\mathscr{M}_n = \mathscr{M}_{n+1}$. Hence $\mathscr{M}_n$ and $\mathscr{M}_{n+1}$ are guaranteed to satisfy the same formulas, whereby $\mathscr{M} \models [adopt]^n\varphi \leftrightarrow [adopt]^{n+1}\varphi$. □

Axiomatization. We obtain an axiomatization of the logic for threshold models and their update dynamics by using the standard method of reduction rules from dynamic epistemic logic [3,28,5,6].

Definition 7 (The Logic of Threshold-Limited Influence, L_θ). *The logic L_θ is comprised of any axiomatization of the propositional calculus and of the axioms and derivation rules of Table 1, for a given threshold $\theta \in [0, 1]$.*

The *static* logic consists of the axioms of propositional logic, the network axioms of Table 1 and the rule of Modus Ponens. These capture the constraints imposed on the networks. In the *dynamic* part of the logic, we define rules that reduce formulas that contain the $[adopt]$ modality to formulas without it. This is possible as the update procedure is deterministic: all the information required to determine the update threshold model is present in the current model. Hence the next state is "pre-encoded" in the present state.

As the $[adopt]$ modality only affects the extension of B, the reduction axioms are trivial in all cases except those involving β_a. The corresponding reduction axiom, Red.Ax.β, relies on the mentioned pre-encoding. The axiom Red.Ax.β states that a has adopted B after the update just in case 1) she had already adopted it before the update or 2) the proportion of her neighbors who had already adopted it before the update was above threshold θ.

Definition 8 ($\mathscr{C}_\theta$). *Let the threshold $\theta \in [0, 1]$ be given. The class of threshold models $\mathscr{C}_\theta$ contains all and only models with the same threshold θ.*

Network Axioms	
$\neg N_a a$	Irreflexivity
$N_{ab} \leftrightarrow N_b a$	Symmetry
$\bigvee_{b \in \mathscr{A}} N_{ab}$	Seriality
Reduction Axioms	
$[adopt]N_{ab} \leftrightarrow N_{ab}$	Red.Ax.N
$[adopt]\neg\varphi \leftrightarrow \neg[adopt]\varphi$	Red.Ax.$\neg$
$[adopt]\varphi \wedge \psi \leftrightarrow [adopt]\varphi \wedge [adopt]\psi$	Red.Ax.$\wedge$
$[adopt]\beta_a \leftrightarrow \beta_a \vee \beta_{N(a)} \geq \theta$	Red.Ax.β
Inference Rules	
From φ and $\varphi \to \psi$, infer ψ	Modus Ponens
From φ, infer $[adopt]\varphi$	$Nec_{[adopt]}$

Table 1. Hilbert-style proof system L_θ.

For any given threshold $\theta \in [0,1]$, the minimal logic L_θ is sound and complete with respect to the corresponding class of models $\mathscr{C}_\theta$: [4]

Theorem 1 (Completeness). *Let $\theta \in [0,1]$. For any $\varphi \in \mathscr{L}$,*

$$\models_{\mathscr{C}_\theta} \varphi \text{ iff } \vdash_{L_\theta} \varphi$$

Proof. Soundness: Let $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ be an arbitrary threshold model with $a, b \in \mathscr{A}$. Then $\mathscr{M}$ satisfies Irreflexivity (Symmetry/seriality) directly by the semantics and the assumption of irreflexivity (symmetry/seriality) of the network. $\mathscr{M} \models [adopt]N_{ab} \leftrightarrow N_{ab}$ as the adoption operation never alters the network. Soundness of Red.Ax.$\neg$ and Red.Ax.$\wedge$ may be shown straightforwardly using induction on the length of formulas.

To see that $\mathscr{M}$ satisfies Red.Ax.β, let $\mathscr{M}'$ be the adoption update of $\mathscr{M}$. Then $\mathscr{M} \models [adopt]\beta_a$ iff $\mathscr{M}' \models \beta_a$ iff $a \in B' = B \cup \left\{b \in \mathscr{A} : \frac{N(b)\cap B}{N(b)} \geq \theta\right\}$ iff $\mathscr{M} \models \beta_a$ or $a \in \left\{b \in \mathscr{A} : \frac{N(b)\cap B}{N(b)} \geq \theta\right\}$. A syntactic decoding following Definition 2.2 of the large, right-hand disjunct of Red.Ax.β (called $\beta_{N(a)\geq\theta}$) shows that it is satisfied iff $a \in \left\{b \in \mathscr{A} : \frac{N(b)\cap B}{N(b)} \geq \theta\right\}$: The outer disjunction requires/ensures the existence of two sets of agents, $\mathscr{G}$ and $\mathscr{N}$, such that $\mathscr{G} \subseteq \mathscr{N}$ and $\frac{|\mathscr{G}|}{|\mathscr{N}|} \geq \theta$. The inner conjunction in Definition 2.2 is satisfied iff $\mathscr{N} = N(a)$ and $\mathscr{G} \subseteq B$. Hence φ is satisfied iff $\exists \mathscr{G} \subseteq N(a) \cap B : \frac{|\mathscr{G}|}{|N(a)|} \geq \theta$ iff $\frac{|N(a)\cap B|}{|N(a)|} \geq \theta$ iff $a \in \left\{b \in \mathscr{A} : \frac{N(b)\cap B}{N(b)} \geq \theta\right\}$. Hence $\mathscr{M} \models [adopt]\beta_a$ iff $\mathscr{M} \models \beta_a$ or $\mathscr{M} \models \beta_{N(a)\geq\theta}$.

Completeness: The proof goes via translation of the dynamic language into the static part of the language, in the usual way (see for instance [28, Ch. 7]). □

[4] The proof system and model class are further parametrized by the set of agents $\mathscr{A}$ used to define the corresponding language.

2.3 Clusters and Cascades

An agent adopting a new behavior may influence some of her neighbors to adopt it at the next moment, which in turn may cause further agents to adopt it, and so on. Such a chain reaction is termed a *cascade* in the literature (see e.g. [10, Ch. 19]), and a cascade is said to be *complete* when it results into a state where *all* agents have adopted the new behavior. Because the above given updates of threshold models always reach a fixed point, any cascade will eventually stop. However, a cascade may stop before all agents have adopted, i.e. without being complete. The following recalls a known result about how cascading effects are constrained by the network structure and shows how the suitable constraint may be captured by the minimal logic L_θ.

First of all, our language can express that a diffusion sequence will reach a complete cascade, given the upper bound on the number of updates before stabilization of the diffusion process noted in Proposition 1:

Definition 9. *The sentence abbreviated by '$cascade$' expresses that all agents will have adopted eventually:*

$$cascade := [adopt]^{|\mathscr{A}|-1} \bigwedge_{a \in \mathscr{A}} \beta_a$$

Some parts of a network structure may be more "dense" than others. Strongly connected groups of agents are more resilient to external influence. E.g., a tightly knit group may be hard to convert to a particular opinion if all its members support one another in disagreeing with the opinion. Tightly connected components of a network might therefore block the diffusion of a behavior when it stems from outside this component. Briefly put, dense components of a network may prevent complete cascades and the denser a group, the better it resists change induced from the outside. The required precise notion of a "dense" group is that of a *d-cohesive set* [19], also referred to as a *cluster of density d* [10]. A cluster of density d is a set of agents such that for each agent in the set, the proportion of her neighbors which are also in the group is at least d.

Definition 10 (Cluster of density d). *Given a network $(\mathscr{A}, N)$, a cluster of density d is any group $C \subseteq \mathscr{A}$ such that for all $a \in C$,*

$$\frac{|N(a) \cap C|}{|N(a)|} \geq d.$$

Notice that any network will contain at least one cluster of density 1, namely the group $\mathscr{A}$, and that each singleton $\{a\} \subseteq \mathscr{A}$ is a cluster of density 0 (by irreflexivity).

Example: Clusters. Let model $\mathscr{M}$ given as illustrated below, with $B = \{d\}$. In this model, $C = \{a, b, c\}$ is a cluster of density $\frac{2}{3}$, in which no member belongs to B.

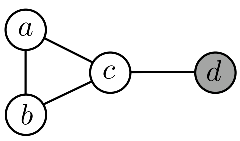

Fig. 1. A social network with a cluster of density $\frac{2}{3}$.

The language $\mathscr{L}$ can express the existence of a cluster: if C is a cluster of density d then for each a in C, there is a big enough subset of C which are a's neighbors.

Proposition 3. *The group C is a cluster of density d in $(\mathscr{A}, N)$ iff $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ satisfies*

$$\bigwedge_{a \in C} \bigvee_{\left\{\mathscr{G} \subseteq \mathscr{N} \subseteq \mathscr{A} : \frac{|\mathscr{G} \cap C|}{|\mathscr{N}|} \geq d\right\}} \left(\bigwedge_{b \in \mathscr{N}} N_{ab} \wedge \bigwedge_{b \notin \mathscr{N}} \neg N_{ab} \right) \tag{1}$$

Proof. Left to right: Let $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ and assume C is a cluster of density d in $(\mathscr{A}, N)$. Then by definition, for all $a \in C$, $\frac{|N(a) \cap C|}{|N(a)|} \geq d$. As $\mathscr{M}$ is based on $(\mathscr{A}, N)$, $\{b : \mathscr{M} \models N_{ab}\} = N(a)$ for all $a \in \mathscr{A}$. Let a be given and pick $\mathscr{N} = N(a)$ and $\mathscr{G} = N(a) \cap C$. Then $\frac{|\mathscr{G}|}{|\mathscr{N}|} \geq d$. Given the choice of $\mathscr{N}$, $\mathscr{M} \models \bigwedge_{b \in \mathscr{N}} N_{ab} \wedge \bigwedge_{b \notin \mathscr{N}} \neg N_{ab}$. So $\mathscr{M}$ satisfies (1).

Right to left: Assume that $\mathscr{M}$ satisfies (1) for some $C \subseteq \mathscr{A}$ and some $d \in [0,1]$. Then for each $a \in C$, there is are sets $\mathscr{G}$ and $\mathscr{N}$ with $\mathscr{G} \subseteq \mathscr{N}$ and $\frac{|\mathscr{G} \cap C|}{|\mathscr{N}|} \geq d$, such that $\mathscr{N} = \{\mathscr{M} \models N_{ab}\} = N(a)$. Hence $\frac{|\mathscr{G} \cap C|}{|N(a)|} = \frac{|\mathscr{G} \cap C|}{|\mathscr{N}|} \geq d$. As $\mathscr{G} \cap C \subseteq \mathscr{N} = N(a)$, $\frac{|N(a) \cap C|}{|N(a)|} \geq d$. As a was arbitrary from C, C is indeed a cluster of density d in $(\mathscr{A}, N)$. □

Given Proposition 3, it is easy to see that the sentence below characterizes the existence of a cluster of density d among agents *who have not adopted* (abbreviated $\exists C_{\geq d} \neg \beta$):

$$\exists C_{\geq d} \neg \beta := \bigvee_{C \subseteq \mathscr{A}} \bigwedge_{a \in C} \bigvee_{\left\{\mathscr{G} \subseteq \mathscr{N} \subseteq \mathscr{A} : \frac{|\mathscr{G} \cap C|}{|\mathscr{N}|} \geq d\right\}} \left(\bigwedge_{b \in \mathscr{N}} N_{ab} \wedge \bigwedge_{b \notin \mathscr{N}} \neg N_{ab} \wedge \bigwedge_{b \in \mathscr{G}} \neg \beta_b \right)$$

Note that we can express in the same way that there is a cluster of density greater than d, by replacing $\geq$ by the strict $>$ in the formula (abbreviated $\exists C_{>d} \neg \beta$).

Example: Clusters, cont.. The model illustrated in Fig. 1 contains a cluster $C = \{a, b, c\}$ of density $\frac{2}{3}$, such that no agent in C has adopted. Hence, the model should satisfy $\exists C_{\frac{2}{3}} \neg \beta$:

$$\bigvee_{C \subseteq \mathscr{A}} \bigwedge_{a \in C} \bigvee_{\left\{\mathscr{G} \subseteq \mathscr{N} \subseteq \mathscr{A} : \frac{|\mathscr{G} \cap C|}{|\mathscr{N}|} \geq \frac{2}{3}\right\}} \left(\bigwedge_{b \in \mathscr{N}} N_{ab} \wedge \bigwedge_{b \notin \mathscr{N}} \neg N_{ab} \wedge \bigwedge_{b \in \mathscr{G}} \neg \beta_b \right). \tag{2}$$

To verify this, assume C is a group that satisfies the outmost disjunction. Then for each $a \in C$ there is must a $\mathscr{G}$ and $\mathscr{N}$ such that $\frac{|\mathscr{G} \cap C|}{|\mathscr{N}|} \geq \frac{2}{3}$ for which $\mathscr{M}$ satisfies

$$\bigwedge_{b \in \mathscr{N}} N_{ab} \wedge \bigwedge_{b \notin \mathscr{N}} \neg N_{ab} \wedge \bigwedge_{b \in \mathscr{G}} \neg \beta_b. \tag{3}$$

To see that $\mathscr{M}$ satisfies (3), regard first agent c, for whom the appropriate $\mathscr{N}$ is $\{a, b, d\}$. As $|\mathscr{N}| = 3$, we must identify a group $\mathscr{G} \subseteq C$ with $|\mathscr{G}| \geq 2$ such that for all $b \in \mathscr{G}$, $\mathscr{M} \models N_{cb}$. Such a $\mathscr{G}$ exists, being $\{a, b\}$. Finally, indeed $\mathscr{M} \models \neg \beta_a \wedge \neg \beta_b$, and hence the conjunct for c is satisfied. Similar reasoning shows that the conjuncts for a and b also hold. This gives us (2).

The Cluster Theorem. The following theorem from [19],[10, Ch.19.3] characterizes the possibility of a complete adoption cascade in a network:

> Given a threshold model $\mathscr{M}$ with threshold $\theta \neq 0$ and a set $B \subset \mathscr{A}$ of agents who have adopted, all agents will eventually adopt *if, and only if* there does not exist a cluster of density greater than $1 - \theta$ in $\mathscr{A} \backslash B$.

As both the complete cascade and the existence of the relevant clusters are expressible in $\mathscr{L}_{[]}$, the cluster theorem can also be encoded in our setting, in the following way:

Let $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ with $\theta \neq 0$. Then

$$\mathscr{M} \models cascade \leftrightarrow \neg\exists C_{>1-\theta}\neg\beta.$$

2.4 Logics for Generalizations of Threshold Models

So far, we have considered the "simplest" possible network structures: the networks are finite, symmetric, irreflexive and serial. The constraints of symmetry and irreflexivity could easily be relaxed in the initial definition of threshold models (Def. 2) to generalize the logics to different types of social relationships (for instance a hierarchical network).

For simplicity, we work with uniform thresholds. Obtaining logics for settings without this uniformity constraint is unproblematic: 1) define θ not as a constant but as a function assigning a particular threshold to each agent; i.e., set $\theta : \mathscr{A} \to [0,1]$ in the definition of threshold models (Def. 2); 2) replace θ by $\theta(a)$ in the definition of the update (Def. 3) and in the reduction axiom Red.Ax.β (in Table 1). This will generate a logic for each such function θ, that is, for each distribution of thresholds among agents.

The logical setting may also be generalized to capture the spread of *several* behaviors and their interaction. This amounts to: 1) modify the definition of threshold models (Def. 2) to let $\mathscr{B}$ be a finite *set* of behaviors ($\mathscr{B} = \{B_1, B_2, ...B_n\}$) and define $\theta : \mathscr{A} \times \mathscr{B} \to [0,1]$; 2) Relativize the definition of the update to each behavior B_i; 3) extend our set of atomic propositions: $\Phi = \{N_{ab} : a, b \in \mathscr{A}\} \cup \{\beta_{ia} : a \in \mathscr{A}, i \in 1, ...n\}$; 4) relativize the semantic clause in the obvious way: $\mathscr{M} \vDash \beta_{ia}$ iff $a \in B_i$, and replace the reduction axiom $Red.Ax.\beta$ by $Red.Ax.\beta_i$ accordingly. The "signature" of the resulting logic will then be given by $[\theta, \mathscr{A}, \mathscr{B}]$. Such a logic allows reasoning about the diffusion of a fixed number of behaviors, given a specific distribution of thresholds for each behavior to each agent, for any particular network structure.

Furthermore, we consider the *proportion* of neighbors who have adopted as the only relevant factor for decision making. This makes every neighbor as influential as any other. To generalize, weighted links representing different "degrees of influence" could be used instead. The condition for being influenced into adopting would become: the *weighted sum* of my neighbors which have adopted is at least θ. Alternatively, we could fix an ordering of neighbors of each agent a with $b \geq_a c$ stating that agent b influences agent a at least as much as agent c does. Based on such an ordering, one possible update policy would be that a adopts when a given proportion of $\geq_a$-maximal agents have adopted.

Additional alternative policies will be discussed in Section 5. These will also involve epistemic considerations, the topic to which we turn next.

3 Epistemic Threshold Models and Their Dynamic Logic

By the definition of the above given update on threshold models, agents *react to their environment*: they are always influenced by the actual behavior of their direct neighbors. In many situations, this "nomothetic" update style seems to pose unrealistic requirements. The update requires that agents act in accordance with the *facts* of others' behavior, even in the face of uncertainty. Hence, the above threshold model update may require of agents that they act in accordance with information that they do not actually possess. For an example, see Fig. 2.

To accommodate this shortcoming, we extend the standard threshold models with an epistemic dimension and define a refined adoption policy where agents' behavior change depends on their knowl-

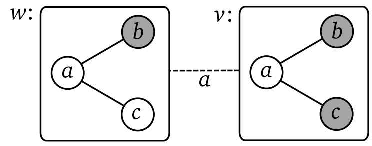

Fig. 2. A situation of uncertainty. Agent a cannot tell whether world w or world v is the actual one, as indicated by the dashed line (when representing indistinguishability relations we omit reflexive and transitive links). Hence, a does not know whether c has adopted or not. Assume that the threshold is $\theta > 1/2$ and that v is the actual world. Then, according to the 'threshold model update', a should adopt – but a does not know that!

edge of others' behavior. We moreover define a logical system suitable to reason about *epistemic threshold models* and their dynamics.

To add an epistemic dimension to threshold models, we add for each agent a subjective epistemic indistinguishability relation, as illustrated in Fig. 2, in the standard way since [14]. Or equivalently, following [2], each agent is given an "information partition" over a given set of possible worlds. Each information cell in this partition indicates the uncertainty of the agent: i.e. the things she cannot tell apart. This modeling of uncertainty is commonplace in logic, economics and computer science.

3.1 Epistemic Threshold Models

The most general version of threshold models with an epistemic dimension that we will work with in this paper is the following:

Definition 11 (Epistemic Threshold Model (ETM)). *An epistemic threshold model (ETM) is a tuple*

$$\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a\in\mathscr{A}})$$

where:

$\mathscr{W}$	*is a finite, non-empty set of possible worlds (or states),*
$\mathscr{A}$	*is a finite non-empty set of agents,*
$\sim_a \subseteq \mathscr{W} \times \mathscr{W}$	*is an equivalence relation, for each agent $a \in \mathscr{A}$,*
$N : \mathscr{W} \to (\mathscr{A} \to \mathscr{P}(\mathscr{A}))$	*assigns a neighborhood $N(w)(a)$ to each $a \in \mathscr{A}$ in each $w \in \mathscr{W}$, such that:*
	$a \notin N(w)(a)$ *(Irreflexivity)*
	$b \in N(w)(a) \Leftrightarrow a \in N(w)(b)$ *(Symmetry)*
	$N(w)(a) \neq \emptyset$ *(Seriality)*
$B : \mathscr{W} \to \mathscr{P}(\mathscr{A})$	*assigns to each $w \in \mathscr{W}$ a set $B(w)$ of agents who have adopted.*
$\theta \in [0,1]$	*is a uniform adoption threshold.*

To reason about the impact of knowledge on diffusion in network situations, we want to impose limiting assumptions regarding the agents' uncertainty. It is for example natural to assume that agents know who their direct neighbors are, though cases exist where it is natural that agents know more about the network. Agents may know who the neighbors of neighbors are, or maybe the whole network is even common knowledge. Likewise, the uncertainty about agents' behavior might be subject to various constraints: agents may know the behavior of their neighbors, of their neighbors' neighbors, of everybody, etc.

One way to impose restrictions on uncertainty is by giving agents an ego-centric "sphere of sight", corresponding to how far they can "see" in the network, assuming that if they can see further, they can see closer. We will say that an agent has *sight n* when she can "see" *at least n* agents away, i.e., when she knows *at least both* the network structure and the behavior of all agents within n distance. To provide a formal definition, we first fix what is meant by "n distance":

Definition 12 (n-reachable, n-distant). *Let* $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ *and let* $n \in \mathbb{N}$. *Define* $N^n : \mathcal{W} \to \mathcal{A} \to \mathcal{P}(\mathcal{A})$ *as follows, for any* $w \in \mathcal{W}$ *and any* $a, b, c \in \mathcal{A}$*:*

- $N^0(w)(a) = \{a\}$
- $N^{n+1}(w)(a) = N^n(w)(a) \cup \{b \in \mathcal{A} : \exists c \in N^n(w)(a) \text{ and } b \in N(w)(c)\}$

If $b \in N^n(w)(a)$*, then* b *belongs to the set of agents that* a *has within her sight at world* w*. Morever, if* $b \in N^n(w)(a)$ *we say that* b *is* n-reachable *from* a *in* w*.*

Definition 13 (Sight n Model[5]). *An ETM* $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ *of sight* n *is an epistemic threshold model such that, for* $n \in \mathbb{N}$ *and for any* $a, b \in \mathcal{A}$ *and* $w, v \in \mathcal{W}$ *:*

- *If* $w \sim_a v$ *and* $b \in N^{n-1}(w)(a)$*, then* $N(w)(b) = N(v)(b)$ *(agents know the network at least up to distance* n*)*
- *If* $w \sim_a v$ *and* $b \in N^n(w)(a)$*, then* $b \in B(w)$ *iff* $b \in B(v)$ *(agents know the behavior of others at least up to distance* n*).*

3.2 Knowledge-Dependent Diffusion

To remedy the problem of agents acting on information they may not possess, we introduce a revised adoption policy. It captures the intuitive idea that an agent should only be influenced by what he knows about other agents around him. This amounts to a knowledge-dependent adoption policy: agents adopt whenever *they know that* enough of their neighbors have adopted already. We call this update policy *informed update*:

Definition 14 (Informed Update). *Let* $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ *be an ETM with sight* n*. The informed adoption update of* $\mathcal{M}$ *produces ETM* $\mathcal{M}^i = (\mathcal{W}, \mathcal{A}, N, B^i, \theta, \{\sim^i_a\}_{a\in\mathcal{A}})$ *such that, for any* $a, b \in \mathcal{A}$ *and any* $w, v \in \mathcal{W}$*:*

- $B^i(w) = B(w) \cup \{a \in \mathcal{A} : \forall v \sim_a w \frac{|N(v)(a) \cap B(v)|}{|N(v)(a)|} \geq \theta\}$ *and*
- $w \sim^i_a v$ *iff i)* $w \sim_a v$ *and ii) if* $b \in N^n(w)(a)$*, then* $b \in B^i(w)$ *iff* $b \in B^i(v)$*.*

The first condition tells us that the new set of adopters at world w includes the previous set of adopters $B(w)$ (hence agents do not give up their previously adopted behavior) and it includes also all agents who, as far as they know, are certain of the fact that enough influential neighbors (given by θ) have adopted already. The second condition ensures that the informed update of an ETM with sight n is again an ETM with sight n, i.e., agents' sight is not diminished by updates.

Updating de Dicto and Updating de Re. The above informed update policy is defined using *de dicto* knowledge of others' behavior: if an agent knows that enough others will adopt, so should she, ignoring that she might not know exactly *who* will adopt.

A *de re* update is definable by setting $B^i(w) = B(w) \cup \{a \in \mathcal{A} : \frac{b \in \mathcal{A} : \forall v \sim_a w, |N(v)(a) \cap B(v)|}{|N(v)(a)|} \geq \theta\}$. While both rules are interesting, in the remainder of this paper we opt for the *de dicto* version as it expresses in a stronger sense that agents can fully utilize all their information while staying in the spirit of threshold models.

[5] We lump two notions of sight under one heading. A more general definition would be of sight (n, m), where n specifies the sight of network structure, while m specifies sight of behavior.

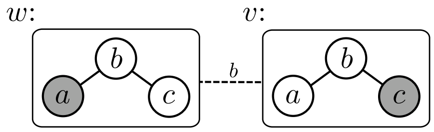

Fig. 3. Adoption *de re* vs. adoption *de dicto*. We illustrate an ETM with threshold $\theta = 1/2$ and two possible worlds. Should b adopt or not? He knows *de dicto* that enough neighbors have adopted, but he does not know so *de re*; he knows that at least half of his neighbors have adopted, but he doesn't know *which* half.

Learning the Distribution. When performing informed updates, agents may learn about the initial distribution of behavior in the network even outside their range of sight, as it may be possible to exclude possibilities based on the development of the dynamics. The learning occurs due to the restriction of the indistinguishability relation, as build into the definition of informed update. Figure 4 provides an example.

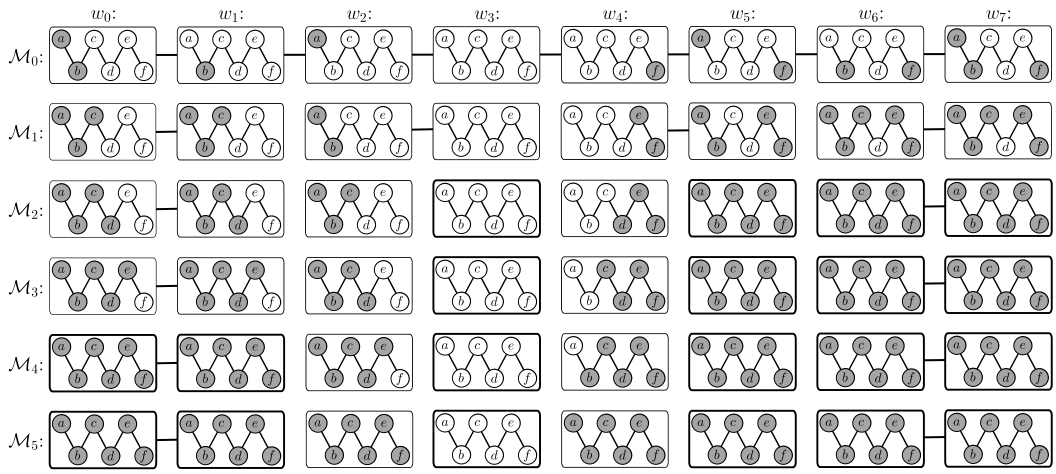

Fig. 4. The learning process for agent d (bottom center) under blind adoption, in an ETM with threshold $\theta \leq \frac{1}{2}$ and sight 1. With sight 1, the ETM contains the 8 depicted possible worlds/states. The last states to reach fixed points at time 5 are states w_2 and w_4 from the left. Epistemic relations are drawn only for d to simplify representation. Note the development of the indistinguishability relation from $\mathcal{M}_0$ to $\mathcal{M}_5$: as the updated $\sim'_d$ is a restriction of $\sim_d$ to states where both c and e's behaviors are identical, d learns about the initial distribution. Learning may or may not be complete: compare the development of states w_1 and w_2.

Implicit Information and Redundant Knowledge. Under some epistemic conditions, the epistemic and non-epistemic diffusion policies are equivalent. If each agent always knows *at least* who her neighbors are and how they are behaving, then the two policies give rise to the same diffusion dynamics, in the following sense: the diffusion dynamics resulting from the informed update on an ETM reduces to the diffusion dynamics under the initial (non-epistemic) update applied to each possible world of the ETM. This is the content of Proposition 4 below.

Proposition 4 relates two important insights. The first is that standard threshold models make the *implicit epistemic assumption* that agents know their neighborhood and its behavior. The second is that *knowledge about more distant agents is redundant* as it will not affect behavior.

To prove the result, we first define how to generate a (non-epistemic) threshold model from a possible state of an epistemic threshold model:

Definition 15 (State-Generated Threshold Model (SGM)). *Let $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a \in \mathscr{A}})$ be an ETM and let $w \in \mathscr{W}$ and $a \in \mathscr{A}$. The state-generated threshold model $\mathscr{M}(w) = (\mathscr{A}, N_{\mathscr{M}(w)}, B_{\mathscr{M}(w)}, \theta)$ is given by:*

$$N_{\mathscr{M}(w)}(a) = N(w)(a), \quad \text{and}$$
$$a \in B_{\mathscr{M}(w)} \Leftrightarrow a \in B(w).$$

Proposition 4. *Let $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a \in \mathscr{A}})$ be an ETM and $w \in \mathscr{W}$. Let $\mathscr{M}^i$ and $\mathscr{M}(w)$ be respectively the informed update and state-generated models of $\mathscr{M}$. Let $\mathscr{M}^i(w)$ be the state-generated model of $\mathscr{M}^i$ and let $\mathscr{M}(w)'$ be the non-epistemic threshold update of $\mathscr{M}(w)$. Then*

$$\text{if } \mathscr{M} \text{ has sight } n \geq 1, \text{ then}$$
$$\mathscr{M}^i(w) = \mathscr{M}(w)'.$$

Proof. As neither the non-epistemic threshold update nor the informed update changes the set of agents, the network or the threshold, it need only be shown that $B^i(w) = B(w)'$ where $B^i(w)$ is the behavior set of $\mathscr{M}^i(w)$ and $B(w)'$ is the behavior set of $\mathscr{M}(w)'$.

Assume $a \in B(w)$. Then it follows that $a \in B(w)^i$ within $\mathscr{M}^i$, by monotonicity of the informed update. Hence we also obtain $a \in B_{\mathscr{M}^i(w)}$ in $\mathscr{M}^i(w)$ by Definition 15 of SGMs. From $a \in B(w)$ it also follows that $a \in B_{\mathscr{M}(w)}$ by defintion of SGMs. By monotonicity of the non-epistemic threshold update, we have $a \in B'_{\mathscr{M}(w)}$ in $\mathscr{M}(w)'$.

Assume that $a \notin B(w)$. Then $a \notin B_{\mathscr{M}(w)}$ by definition 15 of SGMs. By definition, $a \in B(w)^i$ iff $\forall v \sim_a w : \frac{|N(v)(a)| \cup B(v)}{|N(v)(a)|} \geq \theta$. As $\mathscr{M}$ has sight $n \geq 1$, $\forall v \sim_a w$ $N(v)(a) = N(w)(a)$ and $b \in N(w)(a)$ implies $b \in B(w) \Leftrightarrow b \in B(v)$. Hence $\frac{|N(w)(a)| \cup B(w)}{|N(w)(a)|} \geq \theta$. As $N(w)(a) = N_{\mathscr{M}(w)}(a)$ and $B(w) = B_{\mathscr{M}(w)}$, it follows that $\frac{|N_{\mathscr{M}(w)}(a)| \cup B_{\mathscr{M}(w)}}{|N_{\mathscr{M}(w)}(a)|} \geq \theta$ iff $a \in B_{\mathscr{M}(w)}$. □

Proposition 4 provides a precise, but partial, interpretation of the dynamics of non-epistemic threshold models as a process of information-dependent behavior diffusion. As witnessed by its proof, only the immediate neighborhood of agents matters for the adoption behavior in a threshold model. A next step is to investigate how this changes when agents are equipped with predictive abilities; see Section 4.

The interpretation is only partial in that we *do not* obtain a full characterization of the standard threshold dynamics (see Definition 3) by requiring sight $n \geq 1$. Sight $n < 1$ *does not* imply that there will always be a *difference making* neighbor about which some agent has uncertainty. If a has uncertainty about some neighbor b's behavior but is certain that a large enough proportion of neighbors have adopted, then the model will have sight strictly less than 1 while it is still developing according to the standard threshold dynamics.

Situations in which neighbors lack knowledge of some direct neighbors' behavior are interesting in that they may cause the diffusion process to *slow down* compared to the standard update policy:

Proposition 5. *There exists an ETM $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a \in \mathscr{A}})$ with sight $n < 1$ such that*

$$B_{\mathscr{M}^i(w)} \subset B_{\mathscr{M}(w)'},$$

where $\mathscr{M}^i$ and $\mathscr{M}(w)$ are respectively the informed update and state-generated models of $\mathscr{M}$, and $\mathscr{M}^i(w)$ is the state-generated model of $\mathscr{M}^i$ and $\mathscr{M}(w)'$ is the non-epistemic update (Def. 3) of $\mathscr{M}(w)$.

Proof. By construction of a specific model: let $\mathscr{M} = ((\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a \in \mathscr{A}})$ with $\mathscr{W} = \{w, v\}$, $w \sim_a v$, $N(w)(a) = N(v)(a)$ but $\frac{|N(w)(a) \cap B(w)|}{|N(w)(a)|} \geq \theta > \frac{|N(v)(a) \cap B(v)|}{|N(v)(a)|}$. Then $a \notin B_{\mathscr{M}^i(w)}$, but $a \in B_{\mathscr{M}(w)'}$.

Figure 5 illustrates this "slower" diffusion process.

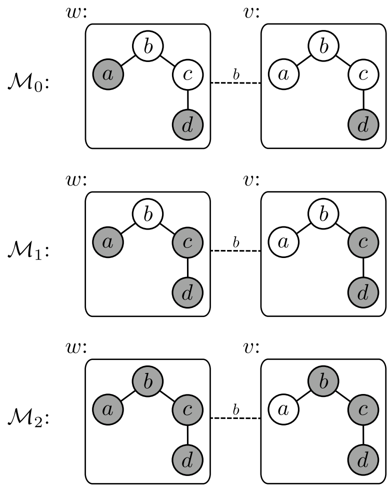

Fig. 5. A diffusion process “slowed down" by the uncertainty of agent b, with threshold $\theta = \frac{1}{2}$. Consider the situation in world w: agent a has adopted, but agent b does not know it. Therefore, agent b will not adopt immediately. The diffusion according to the informed update policy in state w will only stabilize after applying the informed update rule *twice*. Note that under the non-epistemic threshold update, or if agent b knew whether a has adopted, the situation depicted in w would stabilize after only one step (i.e. the non-epistemic threshold update of $\mathscr{M}_0(w)$ gives us directly $\mathscr{M}_2(w)$).

3.3 Knowledge and Cascades

In Section 2.3, we have shown how our language can capture complete cascades and the existence of clusters able to block diffusion, as captured by the *Cluster Theorem*: a cascade will be complete if and only if the network does not contain a cluster of non-yet-adopters of density greater than $1-\theta$.

Given proposition 4 above, the cluster theorem still holds for any epistemic threshold model with sight at least 1. Moreover, the existence of a relevant cluster will still block a cascade under the informed update policy, independently of how much agents know. However in general, considering any epistemic threshold model with any sight, the cluster theorem cannot be maintained as it was stated. What we observe is that the left to right direction of the cluster theorem still holds for epistemic threshold models with sight less than 1: indeed, if a complete cascade occurs, then the network does not contain a cluster of density greater than $1-\theta$. However, the converse of does not hold in these models with sight less than 1. We briefly explain this point in more detail. Given proposition 5 above, we know that the diffusion process, via the informed update rule, in an ETM with sight < 1 might be “slower" than the process based on the non-epistemic threshold update policy. Indeed, the lack of knowledge may for instance block a cascade, despite the absence of a cluster-obstacle. Figure 6 illustrates this difference.

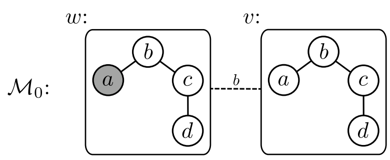

Fig. 6. A diffusion process “blocked" by the uncertainty of agent b, with $\theta = \frac{1}{2}$. Consider the situation in world w: agent a has adopted, but agent b does not know it. Therefore, agent b will not adopt (under the informed adoption rule). Note that under the non-epistemic threshold update, or if agent b knew that a has adopted, the situation depicted in state w would evolve into a complete cascade.

3.4 The Epistemic Logic of Threshold-Limited Influence

To reflect the epistemic dimension in a formal syntax, the language $\mathscr{L}$ is extended by adding the standard K_a modalities reading "agent a knows that", for each agent $a \in \mathscr{A}$.

Definition 16 (Languages $\mathscr{L}_{K[]}$ and $\mathscr{L}_K$). *Let the set of atomic propositions be given by $\{N_{ab} : a, b \in \mathscr{A}\} \cup \{\beta_a : a \in \mathscr{A}\}$ for a finite set $\mathscr{A}$. Where $a, b \in \mathscr{A}$, the formulas of $\mathscr{L}_{K[]}$ are given by*

$$\varphi := N_{ab} \mid \beta_a \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi \mid [adopt]\varphi$$

The formulas of $\mathscr{L}_K$ are those of $\mathscr{L}_{K[]}$ that do not involve the $[adopt]$ modality.

As standard, we can use the given language to define the other Boolean operators for disjunction and implication and introduce $< adopt >$ as the dual of $[adopt]$.

Definition 17 (Semantics for $\mathscr{L}_{K[]}$ with Informed Update). *Formulas $\varphi, \psi \in \mathscr{L}_{K[]}$ are interpreted over an ETM $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a\in\mathscr{A}})$ with sight n, $w, v \in \mathscr{W}$:*

$\mathscr{M}, w \models \beta_a$	iff	$a \in B(w)$
$\mathscr{M}, w \models N_{ab}$	iff	$b \in N(w)(a)$
$\mathscr{M}, w \models \neg\varphi$	iff	$\mathscr{M}, w \not\models \varphi$
$\mathscr{M}, w \models \varphi \wedge \psi$	iff	$\mathscr{M}, w \models \varphi$ and $\mathscr{M}, w \models \psi$
$\mathscr{M}, w \models K_a\varphi$	iff	for all $v \in \mathscr{W}$ such that $v \sim_a w$, $\mathscr{M}, v \models \varphi$
$\mathscr{M}, w \models [adopt]\varphi$	iff	$\mathscr{M}', w \models \varphi$, where $\mathscr{M}'$ is the informed update of $\mathscr{M}$ as specified in Def. 14 .

Axiomatization. In the specification of the epistemic reduction axioms, the following two syntactic shorthands are used:

Abbreviation. For any $k \in \mathbb{N} \geq 1$, we introduce the abbreviation N^k_{ab} by induction,

$$N^1_{ab} := N_{ab}$$
$$N^{k+1}_{ab} := N^k_{ab} \vee \bigvee_{c\in\mathscr{A}} \left(N^k_{ac} \wedge N_{cb}\right)$$

The formula N^k_{ab} expresses that b is k-reachable from a.

Abbreviation. For $\mathscr{B} \subseteq \mathscr{A}$, we introduce the abbreviation $\mathscr{B} = N^k_a\beta^+$ as follows:

$$\left(\mathscr{B} = N^k_a\beta^+\right) := \bigwedge_{b\in\mathscr{B}} \left(N^k_{ab} \wedge [adopt]\beta_b\right) \wedge \bigwedge_{b\in\mathscr{A}\setminus\mathscr{B}} \left(N^k_{ab} \to [adopt]\neg\beta_b\right).$$

The expression $\mathscr{B} = N^k_a\beta^+$ refers to the set of agents which are 1) k-reachable from a and 2) will have adopted after the next update.

Using these shorthands, the axioms for Epistemic Threshold Models and the dynamics of Informed Update are given in Table 2.

The reduction law *Ep.Red.Ax.*β states that a has adopted β after the update just in case she had already adopted it before the update, or *she knew that* she had a large enough proportion of neighbors

Network Axioms		
$\neg N_{aa}$		Irreflexivity
$N_{ab} \leftrightarrow N_{ba}$		Symmetry
$\bigvee_{b \in \mathscr{A}} N_{ab}$		Seriality
Knowledge Axioms		
$K_a\varphi \rightarrow \varphi$	(∗)	Ax.T
$K_a\varphi \rightarrow K_aK_a\varphi$	(∗)	Ax.4
$\neg K_a\varphi \rightarrow K_a\neg K_a\varphi$	(∗)	Ax.5
Reduction Axioms		
$[adopt]N_{ab} \leftrightarrow N_{ab}$		Red.Ax.N
$[adopt]\neg\varphi \leftrightarrow \neg[adopt]\varphi$		Red.Ax.¬
$[adopt]\varphi \wedge \psi \leftrightarrow [adopt]\varphi \wedge [adopt]\psi$		Red.Ax.∧
$[adopt]\beta_a \leftrightarrow \beta_a \vee K_a(\beta_{N(a)} \geq \theta)$	(∗)	Ep.Red.Ax.β
$[adopt]K_a\varphi \leftrightarrow \bigvee_{\mathscr{B} \subseteq \mathscr{A}} \big(\mathscr{B} = N_a^n\beta^+ \wedge K_a(\mathscr{B} = N_a\beta^+ \rightarrow [adopt]\varphi)\big)$	(∗)	Ep.Red.Ax.K.sight.n
Inference Rules		
From φ and $\varphi \rightarrow \psi$, infer ψ		Modus Ponens
From φ, infer $K_a\varphi$ for any $a \in \mathscr{A}$	(∗)	Nec.K_a
From φ, infer $[adopt]\varphi$		Nec.$[adopt]$

Table 2. Axioms and rules for the Epistemic Logic of Threshold-Limited Influence for sight n. Subscripts a, b are arbitrary over $\mathscr{A}$. Entries marked (∗) are new or modified relative to Table 1.

who had already adopted it before the update. *Ep.Red.Ax.K.sight.n* captures that an agent knows that φ will be the case after the update if, and only if, *she knows that*, if those very agents who actually are going to adopt do adopt, then φ will hold after the update.

Definition 18 (Epistemic Logic of Threshold-Limited Influence). *The logic $L_{\theta n}$ is comprised of the axioms and rules of propositional logic and the axioms and rules of Table 2.*

Definition 19 ($\mathscr{C}_{\theta n}$). *Let $\theta \in [0,1]$ be given. The class of ETM $\mathscr{C}_{\theta n}$ contains all and only ETM with threshold θ and sight n.*

The logic $L_{\theta n}$ is sound and complete with respect to the corresponding class of models $\mathscr{C}_{\theta n}$:

Theorem 2. *Let $\theta \in [0,1]$ and $n \in \mathbb{N}$. For any $\varphi \in \mathscr{L}_{K[]}$,*

$$\models_{\mathscr{C}_{\theta n}} \varphi \text{ iff } \vdash_{L_{\theta n}} \varphi.$$

Proof. Soundness: Let $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a\in\mathscr{A}})$ be an epistemic threshold model with sight n. Let $a, b \in \mathscr{A}$ and $w, v \in \mathscr{W}$. Then $(\mathscr{M}, w)$ satisfies the S5 axioms as all $\sim_a$ are equivalence relations and satisfies the axioms reoccuring from Table 1 for the same reasons non-epistemic threshold models satisfy them.

To see that $(\mathscr{M}, w)$ satisfies **Ep.Red.Ax.**β, let $\mathscr{M}^i$ be the informed update of $\mathscr{M}$. Then $\mathscr{M}, w \models [adopt]\beta_a$ iff $\mathscr{M}^i, w \models \beta_a$ iff $a \in B^i(w) = B(w) \cup \left\{ b \in \mathscr{A} : \forall v \sim_b w \frac{|N(v)(b)\cap B(v)|}{|N(v)(b)|} \geq \theta \right\}$ iff $\mathscr{M}, w \models \beta_a$ or $a \in \left\{ b \in \mathscr{A} : \forall v \sim_b w \frac{|N(v)(b)\cap B(v)|}{|N(v)(b)|} \geq \theta \right\}$. Using the same syntactic decoding as in the proof of Theorem 1, we obtain that $a \in \left\{ b \in \mathscr{A} : \forall v \sim_b w \frac{|N(v)(b)\cap B(v)|}{|N(v)(b)|} \geq \theta \right\}$ iff $\mathscr{M}, w \models K_a\left(\beta_{N(a)} \geq \theta\right)$. Hence $\mathscr{M}, w \models [adopt]\beta_a$ iff $\mathscr{M}, w \models \beta_a$ or $\mathscr{M}, w \models K_a\left(\beta_{N(a)} \geq \theta\right)$.

For **Ep.Red.Ax.**K**.sight.**n, let again $\mathscr{M}^i$ be the informed update of $\mathscr{M}$. Then

$$\mathscr{M}, w \models \bigvee_{\mathscr{B}\subseteq\mathscr{A}} \left((\mathscr{B} = N_a^n\beta^+) \;\wedge\; K_a((\mathscr{B} = N_a\beta^+) \to [adopt]\varphi) \right)$$

iff

$$\exists \mathscr{B} \subseteq \mathscr{A} : \mathscr{M}, w \models (\mathscr{B} = N_a^n\beta^+) \;\wedge\; K_a((\mathscr{B} = N_a\beta^+) \to [adopt]\varphi)$$

iff

$$\exists \mathscr{B} \subseteq \mathscr{A} : \mathscr{M}, w \models \bigwedge_{b\in\mathscr{B}} \left(N_{ab}^n \wedge [adopt]\beta_b\right) \wedge \bigwedge_{b\in\mathscr{A}\setminus\mathscr{B}} \left(N_{ab}^n \to [adopt]\neg\beta_b\right) \text{ and}$$

$$\mathscr{M}, w \models K_a\left(\left(\bigwedge_{b\in\mathscr{B}} \left(N_{ab}^n \wedge [adopt]\beta_b\right) \wedge \bigwedge_{b\in\mathscr{A}\setminus\mathscr{B}} \left(N_{ab}^n \to [adopt]\neg\beta_b\right) \right) \to [adopt]\varphi \right)$$

iff

$$\exists \mathscr{B} \subseteq \mathscr{A} : \mathscr{B} = N^n(w)(a) \cap B^i \text{ and}$$

$$\text{for all } v \sim_a w, \text{ if } \mathscr{B} = N^n(v)(a) \cap B^i, \text{ then } \mathscr{M}^i, v \models \varphi \qquad (*)$$

iff

$$\exists \mathscr{B} \subseteq \mathscr{A} : \mathscr{B} = N^n(w)(a) \cap B' \text{ and}$$

$$\text{if } \mathscr{B} = N^n(w)(a) \cap B^i, \text{ then } \mathscr{M}^i, w \models K_a\varphi$$

(from $(*)$ as $\mathscr{M}$ is sight n, so $N^n(v)(a) \cap B^i = N^n(w)(a) \cap B^i$ for all $v \sim_a w$)

iff

$$\mathscr{M}^i, w \models K_a\varphi$$

(as such a $\mathscr{B}$ always exists)

iff

$$\mathscr{M}, w \models [adopt]K_a\varphi.$$

Completeness (sketch): It can be shown by induction that for all $\varphi \in \mathscr{L}_{K[]}$, there exists a $\varphi' \in \mathscr{L}_K$ such that $\vdash_{L_{n\theta}} \varphi \leftrightarrow \varphi'$. Completeness then follows from the standard proof of completeness for S5 over Kripke models with equivalence relations and the straightforward insight that the network axioms characterize the imposed network conditions. □

4 Prediction Update

In defining our informed update rule based on epistemic threshold models, we ensure that agents do not act on information they do not possess. Such agents are however still limited, in that they do not take *all* their available information into account. This section investigates effects of agents that are allowed to reason about more than only the *present* behavior of the network. In particular, we focus on providing agents with *predictive power*.

Consider the ETM illustrated in Fig. 7, with a given dynamics that runs according to a blind or informed adoption policy.

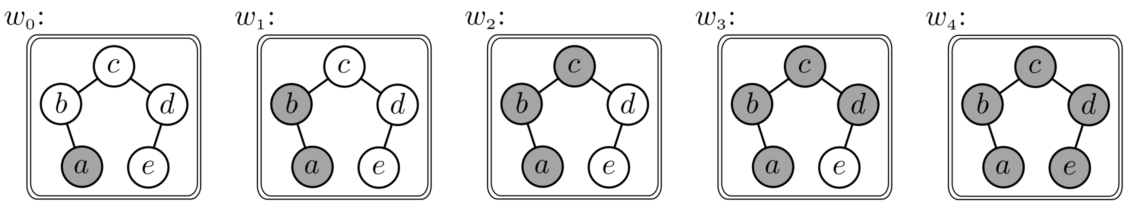

Fig. 7. An ETM with no uncertainty about the actual state w, developing according to informed update. B is marked by gray, and a threshold $\theta = 1/2$ is assumed. At time 0 (w_0), only a has adopted. According to informed adoption, b adopts at time 1. At time 2, c also adopts the behavior, etc.

If one assumes that agents (nodes) are not merely blindly influenced by their neighbors, but rather are rational agents seeking to coordinate, the dynamics in Fig. 7 seems to misfire. In particular, as the network and behavior distribution are known to c (and if the new behavior is considered the most valuable), the choice of c *not* to adopt during the first update is irrational. As c knows that a has adopted, he knows that b will adopt during the next update round. Hence c also knows that he will be better off in round 1 if he, too, has chosen to adopt. To represent this "predictive rationality" we define a new, predictive, update mechanism.

Prediction Update as the Least Fixed Point. In defining "prediction update", we make use of the notion of a *least fixed point*. Even when agents' attempt to use all their available information, each will at some point reach a conclusion about her next action. When the last agent does so, the prediction reaches a fixed point.

This fixed point may be approximated using a chain of lower level predictions. The intuitive idea of the approximation may be illustrated using Fig. 7:

> Assume agent a considers himself smart by predicting that he knows his only neighbor b is going to adopt B in round 2, if b follows the informed update policy. Then a may act preemptively, by also adopting B in round 2, rather than in round 3 as the informed update prescribes.[6] In this case, a may be thought of as a *level 1 predictor*: he assumes no-one else makes predictions, that the others are of level 0. However, a may come to think that b is as smart as he is, i.e., that

[6] If a acted according to the informed update policy, he must first see b adopt before he is influenced by b's choice

also b is a level 1 predictor. Assuming this, a now foresees that b will not adopt in round 2, but already in round 1; based on this prediction about b's predictions, a may now also adopt in round 1. In this case, a is a *level 2 predictor*, etc.

If this reasoning is pushed to its fixed point, it will "catch up with itself": in the fixed point, every agent will be a level ω predictor, predicting under the assumption that all others are the same. This is the trick we use to ensure that agents draw the most powerful conclusion available.

Common Knowledge of Predictive Rationality and Complete Information Use. Prediction update incorporates two epistemic assumptions. One is that it is common knowledge that all agents act in accordance with the prediction update policy. This assumption means that agents may not only predict the systems behavior as if everybody else was acting in accordance with informed update. Rather, agents foresee the behavior of other predictors.

Moreover, it is common knowledge that predictors predict as far into the future as possible, given their information. This means that predictors attempt to use all their available information about the network structure, the current behavior spread and information available to others when determining their next action.

Prediction Update Preliminaries. Before we define the prediction update, a few preliminaries are required.

Definition 20 (Functions Γ_g). *Let $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a\in\mathscr{A}})$ be a finite*[7] *ETM and let the set of all functions from $\mathscr{W}$ to $\mathscr{P}(\mathscr{A})$ be denoted by $\mathscr{P}(\mathscr{A})^{\mathscr{W}} = \{f \mid f : \mathscr{W} \to \mathscr{P}(\mathscr{A})\}$.*

For each $g \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$ let the function $\Gamma_g : \mathscr{P}(\mathscr{A})^{\mathscr{W}} \longrightarrow \mathscr{P}(\mathscr{A})^{\mathscr{W}}$ be given by $\forall w \in \mathscr{W}, \forall f \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$

$$\Gamma_g(f)(w) = g(w) \cup \left\{ a \in \mathscr{A} : \forall v \sim_a w, \frac{|N(v)(a) \cap f(v)|}{|N(v)(a)|} \geq \theta \right\}.$$

Lemma 1. *Let $\mathscr{M}$, $\mathscr{P}(\mathscr{A})^{\mathscr{W}}$ and Γ_g be as in Definition 20. Let $\preceq$ be a partial order on $\mathscr{P}(\mathscr{A})^{\mathscr{W}}$ such that for any $f, g \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$, all $w \in \mathscr{W}$, $f \preceq g \Leftrightarrow f(w) \subseteq g(w)$. Then*

1) *$(\mathscr{P}(\mathscr{A})^{\mathscr{W}}, \preceq)$ is a finite, complete, lattice.*

2) *For each $g \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$, the map Γ_g is order-preserving (monotonic).*

Proof. **1)** For any finite set $\mathscr{A}$, $(\mathscr{P}(A), \subseteq)$ is a finite and hence complete lattice with the order given by the set-theoretic inclusion. If $(L, \sqsubseteq)$ is a finite lattice and $\mathscr{W}$ a finite set, then $(L^{\mathscr{W}}, \leq)$ is also a finite lattice when $L^{\mathscr{W}} = \{f \mid f : \mathscr{W} \longrightarrow L\}$ and $f \leq g$ iff $\forall w \in \mathscr{W}, f(w) \sqsubseteq g(w)$. Hence, given that $\mathscr{W}$ is a finite set, also $(\mathscr{P}(\mathscr{A})^{\mathscr{W}}, \preceq)$ is a finite lattice with the order given by definition of $\preceq$. Every lattice over a finite set is also complete.

2) Let $g, f, f' \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$, and let $f \preceq f'$. Hence $\forall w \in \mathscr{W}$, $f(w) \subseteq f'(w)$. Pick an arbitrary $u \in \mathscr{W}$. Then

$$\begin{aligned}\Gamma_g(f)(u) =& g(u) \cup \left\{ a \in \mathscr{A} : \forall v \sim_a u, \frac{|N(v)(a) \cap f(v)|}{|N(v)(a)|} \geq \theta \right\}\\ \Gamma_g(f')(u) =& g(u) \cup \left\{ a \in \mathscr{A} : \forall v \sim_a u, \frac{|N(v)(a) \cap f'(v)|}{|N(v)(a)|} \geq \theta \right\}.\end{aligned}$$

Let the second terms of the unions be denoted A and A', respectively.

For all $v \in \mathscr{W}$, as $f(v) \subseteq f'(v)$, $\frac{|N(v)(a)\cap f(v)|}{|N(v)(a)|} \geq \theta$ implies $\frac{|N(v)(a)\cap f'(v)|}{|N(v)(a)|} \geq \theta$. Hence $A \subseteq A'$, so $\Gamma_g(f)(u) \subseteq \Gamma_g(f')(u)$. As u was arbitrary, $\Gamma_g(f) \preccurlyeq \Gamma_g(f')$. Hence Γ_g is order-preserving. As g was arbitrary, this holds for all $\Gamma_g, g \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$. □

[7] In a finite ETM we assume that the set of worlds $\mathscr{W}$ is finite and the set of agents $\mathscr{A}$ is finite.

Definition 21 (Least Fixed Point). *Let $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a\in\mathscr{A}})$ be a finite ETM and let $(\mathscr{P}(\mathscr{A})^{\mathscr{W}}, \preceq)$ be as in Lemma 1. Let Γ_g be as in Definition 20.*

The least fixed point *of Γ_g, $\mathtt{lfp}(\Gamma_g)$, is the unique $x \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$ such that*

$$\Gamma_g(x) = x, \text{ and}$$
$$\forall y \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}, \text{ if } \Gamma_g(y) = y, \text{ then } x \preceq y.$$

Theorem 3 ($\mathtt{lfp}$ Existence). *Let $\mathscr{M}, (\mathscr{P}(\mathscr{A})^{\mathscr{W}}, \preceq)$ and Γ_g be as in Definition 21. Then $\mathtt{lfp}(\Gamma_g)$ exists.*

Proof. The least fixed point $\mathtt{lfp}(\Gamma_g)$ exists by the Knaster-Tarski Fixed Point Theorem (see e.g. [9, p. 50]), as $(\mathscr{P}(\mathscr{A})^{\mathscr{W}}, \preceq)$ is a complete lattice (Lemma 1) and Γ_g is order-preserving (Lemma 1). □

Defining Prediction Update. Given the previous paragraph, we may now define prediction update as follows:

Definition 22 (Prediction Update). *Let $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a\in\mathscr{A}})$ be a finite ETM of sight n and let $(\mathscr{P}(\mathscr{A})^{\mathscr{W}}, \preceq)$ be as in Lemma 1. Let $\Gamma_B : \mathscr{P}(\mathscr{A})^{\mathscr{W}} \longrightarrow \mathscr{P}(\mathscr{A})^{\mathscr{W}}$ be given as per Definition 20, i.e., the function such that $\forall w \in \mathscr{W}, \forall f \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$*

$$\Gamma_B(f)(w) = B(w) \cup \left\{ a \in \mathscr{A} : \forall v \sim_a w, \frac{|N(v)(a) \cap f(v)|}{|N(v)(a)|} \geq \theta \right\}.$$

The prediction update of $\mathscr{M}$ results in the ETM $\mathscr{M}' = (\mathscr{W}, \mathscr{A}, N, \widetilde{B}, \theta, \{\sim'_a\}_{a\in\mathscr{A}})$ where $\forall w, v \in \mathscr{W}$,

$$\widetilde{B}(w) = \mathtt{lfp}(\Gamma_B)(w), \text{ and}$$

$$w \sim'_a v \text{ iff } w \sim_a v \text{ and if } b \in N^{\leq n}(w)(a), \text{ then } b \in \widetilde{B}(w) \text{ iff } b \in \widetilde{B}(v).$$

Finding the Prediction Update Fixed Point. The definition of prediction update does not provide us with a method for finding the least fixed point. The following theorem guarentees that we can find it using a bottom-up method:

Theorem 4. *Let $\mathscr{M}, (\mathscr{P}(\mathscr{A})^{\mathscr{W}}, \preceq)$ be as in Lemma 1 with bottom element $\bot$. Let Γ_B and $\widetilde{B}$ be defined as in Definition 22. Then*

$$\mathtt{lfp}(\Gamma_B) = \sup\{{\Gamma_B}^n(\bot) : n \in \mathbb{N}\}$$

Proof. This proof follows from the Knaster-Tarski Fixed Point Theorem applied to finite structures. Given that we work with a finite structure $(\mathscr{P}(\mathscr{A})^{\mathscr{W}}, \preceq)$ and that Γ_B is order-preserving, a least fixed point is reached in a constructive way in finitely many steps. The construction is similar to Proposition 3.1. of [16].

The above stated prediction update rule in definition 22 can now be used to give a new semantics to the $[adopt]$ modality in the logic language $\mathscr{L}_{K[]}$.

Definition 23 (Semantics for $\mathscr{L}_{K[]}$ with Prediction Update). *Given a finite ETM $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a\in\mathscr{A}})$ with sight n, $w \in \mathscr{W}$, and $\varphi \in \mathscr{L}_{K[]}$ truth clauses are as in Definition 17, except for $\varphi := [adopt]\psi$, $\psi \in \mathscr{L}_{K[]}$ given by*

$$\mathscr{M}, w \models [adopt]\varphi \text{ iff } \mathscr{M}', w \models \varphi, \text{ where } \mathscr{M}' \text{ is the prediction update of } \mathscr{M}.$$

Axiomatization. We provide sound axioms that govern the least fixed point behavior of the prediction update policy, but we do not provide a complete axiom system. Finding a complete logic is the aim of planned future research. For now we introduce a fixed point axiom and a least fixed point rule of inference. Note that in this section, the $[adopt]$ modality is a fixed point operator and hence may no longer be reduced away. Contrary to the informed update process, using prediction update results in a system that is strictly more expressive than its non-dynamic counterpart.

To state the proof system, we first generalize the syntactic shorthand introduced in Definition 2.2.

Abbreviation. Given a tuple of formula's $(\varphi_b)_{b\in\mathscr{A}}$, one for each agent $a \in \mathscr{A}$, we introduce the following abbreviation:

$$K_a(\varphi_{N(a)} \geq \theta) := K_a\left(\bigvee_{\left\{\mathscr{G}\subseteq\mathscr{N}\subseteq\mathscr{A}:\frac{|\mathscr{G}|}{|\mathscr{N}|}\geq\theta\right\}}\left(\bigwedge_{b\in\mathscr{N}} N_{ab} \wedge \bigwedge_{b\notin\mathscr{N}} \neg N_{ab} \wedge \bigwedge_{b\in\mathscr{G}} \varphi_b\right)\right).$$

Here $K_a(\varphi_{N(a)} \geq \theta)$ denotes that a knows that larger than a θ fraction of her neighbors has the property φ (where for instance φ_b can stand for $N_{ab} \wedge \beta_b$). In particular, $K_a([adopt]\beta_{N(a)} \geq \theta)$ expresses that a knows that at least a θ fraction of her neighbors will have adopted β after the application of the prediction update rule.

Definition 24 (Prediction Logic). *The logic $L^{predict}_{\theta n}$ is comprised of the axioms and rules of propositional logic and the axioms and rules of Table 2 with the only change that the axiom $Ep.Red.Ax.\beta$ is replaced by the Fixed Point Axiom in Table 3 and we extend the set of rules of the logic with the least fixed point inference rule in Table 3.*

Fixed Point Laws	
$[adopt]\beta_a \leftrightarrow \beta_a \vee K_a([adopt]\beta_{N(a)} \geq \theta)$	Fixed Point Axiom
$\dfrac{\vdash \{\varphi_a \leftrightarrow \beta_a \vee K_a(\varphi_{N(a)} \geq \theta)\}_{a\in\mathscr{A}}}{\vdash \{\varphi_a \rightarrow [adopt]\beta_a\}_{a\in\mathscr{A}}}$	Least Fixed Point Inference Rule

Table 3. Fixed point laws of the prediction logic $L^{predict}_{\theta n}$.

The Fixed Point axiom of Table 3 is almost identical to Ep.Red.Ax.β of Table 2, except for the inclusion of the $[adopt]$ modality on the right-hand side. This states that a will adopt after the prediction update iff she has already adopted, or if she knows that enough of her neighbors will have adopted *after having applied the same predictive reasoning she uses*.

The Least Fixed Point Inference rule reflects the fact that prediction update is defined as a least fixed point operator.

We do not provide a complete logic for the prediction update policy. It is our conjecture that the axioms and rules in definition 24 will be necessary to obtain completeness. The listed axioms and rules are sound with respect to epistemic threshold models using the prediction update rule as our semantics for the $[adopt]$ modality. For the axioms and rules not in Table 3, the proof of Theorem 2 carries over. The axiom and rule governing the fixed point behavior is shown to be sound in the following proposition.

Proposition 6. *The axiom and derivation rule of Table 3 are sound with respect to epistemic threshold models with sight n, using the prediction update as our semantics for the* $[adopt]$ *modality.*

Proof. Let $\mathscr{M}$ be a arbitrary finite ETM with sight n, domain $\mathscr{W} \ni w$ and $a, b \in \mathscr{A}$.

Fixed Point Axiom. $\mathscr{M}, w \models [adopt]\beta_a$ iff $\mathscr{M}', w \models \beta_a$ iff $a \in \widetilde{B} = B \cup \left\{ b \in \mathscr{A} : \forall v \sim_b w, \frac{|N(v)(b) \cap \widetilde{B}|}{|N(v)(b)|} \geq \theta \right\}$ iff $\mathscr{M}, w \models \beta_a$ or $\forall v \sim_a w, \frac{|N(v)(a) \cap \widetilde{B}|}{|N(v)(a)|} \geq \theta$. The right disjunct obtains iff

$$\begin{aligned}\forall v \sim_a w, \exists \mathscr{G}, \mathscr{N} \subseteq \mathscr{A} : \ & \mathscr{G} \subseteq \mathscr{N} \text{ and } \tfrac{|\mathscr{G}|}{|\mathscr{N}|} \geq \theta \text{ and} \\ & \mathscr{G} \subseteq \widetilde{B} \text{ and } \mathscr{N} = N(v)(a)\end{aligned}$$

iff

$$\begin{aligned}\forall v \sim_a w, \exists \mathscr{G}, \mathscr{N} \subseteq \mathscr{A} : \ & \mathscr{G} \subseteq \mathscr{N} \text{ and } \tfrac{|\mathscr{G}|}{|\mathscr{N}|} \geq \theta \text{ and} \\ & \forall b \in \mathscr{G}, \mathscr{M}', v \models \beta_b \text{ and } \forall b \in \mathscr{N}, \mathscr{M}', v \models N_{ab}\end{aligned}$$

iff

$$\forall v \sim_a w, \mathscr{M}', v \models \bigvee_{\left\{\mathscr{G} \subseteq \mathscr{N} \subseteq \mathscr{A} : \frac{|\mathscr{G}|}{|\mathscr{N}|} \geq \theta\right\}} \left(\bigwedge_{b \in \mathscr{N}} N_{ab} \wedge \bigwedge_{b \notin \mathscr{N}} \neg N_{ab} \wedge \bigwedge_{b \in \mathscr{G}} \beta_b \right)$$

iff

$$\forall v \sim_a w, \mathscr{M}, v \models \bigvee_{\left\{\mathscr{G} \subseteq \mathscr{N} \subseteq \mathscr{A} : \frac{|\mathscr{G}|}{|\mathscr{N}|} \geq \theta\right\}} \left(\bigwedge_{b \in \mathscr{N}} N_{ab} \wedge \bigwedge_{b \notin \mathscr{N}} \neg N_{ab} \wedge \bigwedge_{b \in \mathscr{G}} [adopt]\beta_b \right)$$

iff

$$\mathscr{M}, w \models K_a \left(\bigvee_{\left\{\mathscr{G} \subseteq \mathscr{N} \subseteq \mathscr{A} : \frac{|\mathscr{G}|}{|\mathscr{N}|} \geq \theta\right\}} \left(\bigwedge_{b \in \mathscr{N}} N_{ab} \wedge \bigwedge_{b \notin \mathscr{N}} \neg N_{ab} \wedge \bigwedge_{b \in \mathscr{G}} [adopt]\beta_b \right) \right)$$

iff

$$\mathscr{M}, w \models K_a([adopt]\beta_{N(a)} \geq \theta)$$

Hence we conclude $\mathscr{M}, w \models [adopt]\beta_a$ iff $\mathscr{M}, w \models \beta_a \vee K_a([adopt]\beta_{N(a)} \geq \theta)$.

Least Fixed Point Inference Rule. Let an arbitrary finite ETM $\mathscr{M}$ with sight n and domain $\mathscr{W}$ be given. Where $\{\varphi_a\}_{a \in \mathscr{A}}$ is a set of sentences from $\mathscr{L}_{K[]}$, let $\overline{\varphi} \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$ with $\overline{\varphi}(w) = \{a \in \mathscr{A} : \mathscr{M}, w \models \varphi_a\}$. Moreover, let $\Gamma_{\overline{\varphi}} : \mathscr{P}(\mathscr{A})^{\mathscr{W}} \longrightarrow \mathscr{P}(\mathscr{A})^{\mathscr{W}}$, given by

$$\begin{aligned}&\Gamma_{\overline{\varphi}}(f) = h \text{ such that} \\ &\forall w \in \mathscr{W}, h(w) = \overline{\varphi}(w) \cup \left\{ a \in \mathscr{A} : \forall v \sim_a w, \frac{|N(v)(a) \cap f(v)|}{|N(v)(a)|} \geq \theta \right\}.\end{aligned}$$

As shown in Lemma 1, each such $\Gamma_{\overline{\varphi}}$ is order-preserving.

Let $\overline{\beta} \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$ be determined by $\{\beta_a\}_{a \in \mathscr{A}}$ and $\overline{[]\beta} \in \mathscr{P}(\mathscr{A})^{\mathscr{W}}$ by $\{[adopt]\beta_a\}_{a \in \mathscr{A}}$. Let $\Gamma_{\overline{\beta}}$ be given by the above construction.

Given the prediction semantics of $[adopt]$ and the fact that $\widetilde{B} = \mathtt{lfp}(\Gamma_B) = \sup\{\Gamma_B{}^n(\bot) : n \in \mathbb{N}\}$ (Theorem 4), we may conclude that

$$\overline{[]\beta} = \Gamma_{\overline{\beta}}(\overline{[]\beta}) \tag{4}$$

is the least fixed point of $\Gamma_{\overline{\beta}}$.

Assume for some $\{\varphi_a\}_{a\in\mathscr{A}}$ that $\vdash \left\{\varphi_a \leftrightarrow \beta_a \vee K_a(\varphi_{N(a)} \geq \theta)\right\}_{a\in\mathscr{A}}$. This implies

$$\vdash \bigwedge_{a\in\mathscr{A}} (\varphi_a \leftrightarrow \beta_a \vee K_a(\varphi_{N(a)} \geq \theta)). \tag{5}$$

From $\{\varphi_a\}_{a\in\mathscr{A}}$ and $\{\beta_a \vee K_a(\varphi_{N(a)} \geq \theta)\}_{a\in\mathscr{A}}$ we may define functions $\overline{\varphi}$ and $\overline{\beta K}$, as specified above. Now notice that $\overline{\beta K} = \Gamma_{\overline{\beta}}(\overline{\varphi})$. Hence, for (5) to be satisfied, we have that

$$\overline{\varphi} = \Gamma_{\overline{\beta}}(\overline{\varphi}).$$

Given that (4) is the least fixed point of $\Gamma_{\overline{\beta}}$, we have that $\overline{\varphi} = \Gamma_{\overline{\beta}}(\overline{\varphi}) \Rightarrow \overline{[\,]\beta} \preceq \overline{\varphi}$, so

$$\begin{array}{ll} \forall w \forall a : a \in \overline{[\,]\beta}(w) \Rightarrow a \in \overline{\varphi}(w) & \text{so} \\ \forall w \forall a : w \models [adopt]\beta_a \Rightarrow w \models \varphi_a & \text{so} \\ \forall w \forall a : w \models [adopt]\beta_a \rightarrow \varphi_a & \end{array}$$

□

4.1 Sight and Prediction Power

Relationship between predictive power and non-epistemic update. Similarly, as we compared the informed update policy with the non-epistemic threshold model update in section 3.2, it is also natural to investigate the relationship between the 'prediction update', 'informed update' and the 'non-epistemic threshold model update' (Definition 3). Indeed, given that the prediction update policy foresees the non-epistemic deterministic development of the actual state under uncertainty, such a comparison would be rather natural. Besides comparing the cascading behavior and speed of convergence, (as illustrated in figure 4.1), other results that we expect in this investigation relate to posing conditions and finding a lower and upper bound of how far agents can predict into the future. We leave the technical details of this investigation for future work.

Bounded Rationality. Stating that prediction update is the fixed point of the informed update, as we have done in this section, corresponds to assuming that agents have unbounded rationality (maximal anticipation power given the information available). A *bounded* rationality version of the prediction update dynamics could be defined, in which agents can only anticipate a fixed finite number of steps ahead. A natural way of doing this would be by defining an update that updates to some finite level n of the prediction chain. The dynamics of bounded rationality would only differ from the unbounded dynamics for a low enough n. We leave the full exploration of technical details of the prediction update involving such boundedly rational agents for future work.

5 Alternative Adoption Policies

In the previous sections, we have presented three diffusion policies: one depending solely on whether the agents' neighbors have adopted (the "threshold model update" from Def. 3); one depending on *knowledge* of this fact (the "informed update" of Def. 14), and one depending on the *anticipation* of this fact (the "prediction update" of Def. 22). This section questions some in-built assumptions of these policies and discusses possible alternatives.

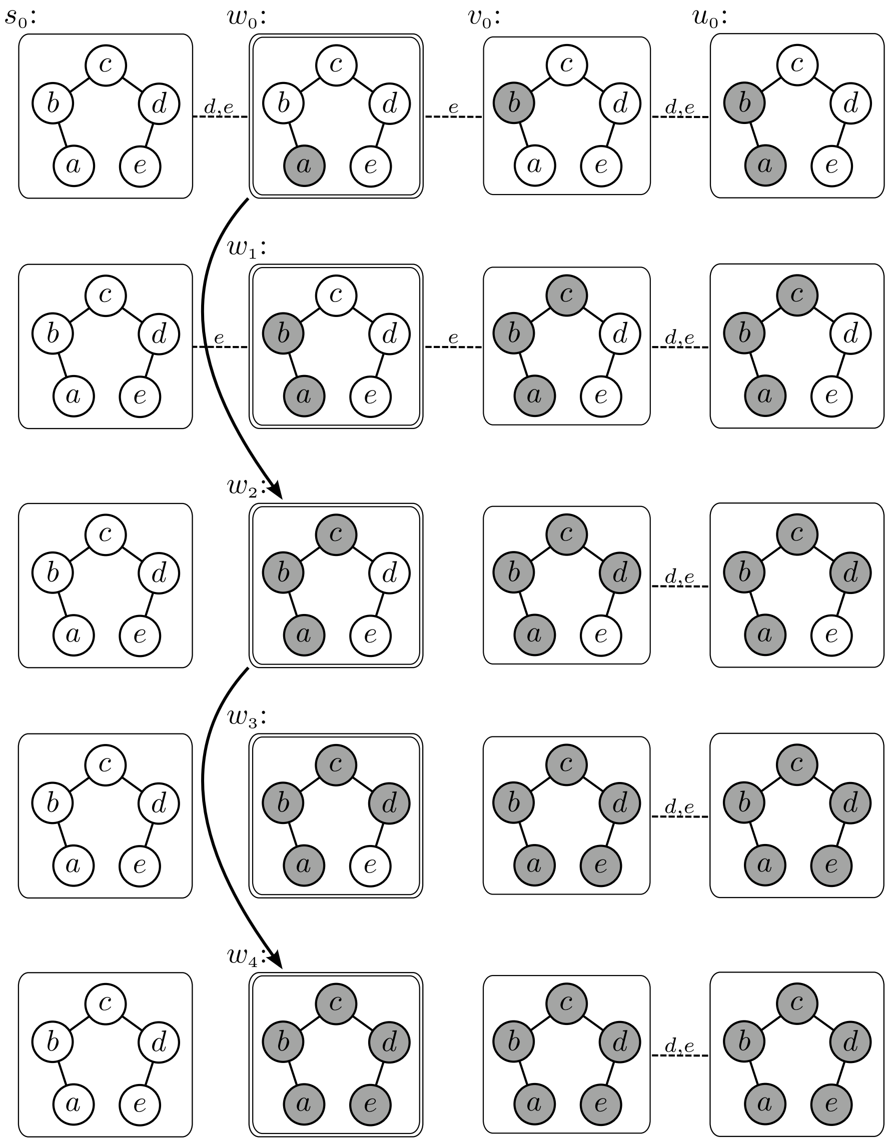

Fig. 8. We use the prediction update to regulate the dynamics of this sight 2, finite ETM with actual state w, $\theta = 1/2$ (reflexive and transitive uncertainty relations are omitted in the illustration). Agents a, b, c are endowed with additional information: they are fully informed about the actual state. The development of the states is given according to blind/informed adoption; states w_0–w_4 are from Fig. 7. The thick arrow indicates the evolution of the actual world under the specified prediction dynamics. With informed update, w reaches a fixed point after 4 updates; with prediciton update, it reaches the same fixed point after only 2 steps. Due to uncertainty, the prediction update does not jump to the fixed point of the non-epistemic update in 1 step: as d does not know whether a has adopted at time 0, she does not know that c will adopt under the prediction update. Hence, she will refrain herself from adopting until w_3. Similar considerations goes for e.

Enlarging the Sphere of Influence. The adoption policies hitherto presented rely on the idea that an agent will adopt if (she knows that) enough of her *direct neighbors* (will) have adopted.

For certain applications, decisions are made that are based not only on actions of direct neighbors, but on the population at large. A case in point is the decision of whether to *support a revolution*: the relevant factor is then whether a big enough fraction of the total population supports the revolution, not whether enough of one's direct neighbors do so.

Generally, such policies may be obtained by enlarging the "sphere of influence" of agents beyond their direct neighbors. A natural choice in the epistemic setting is to fit the "sphere of influence" to agents' "sphere of sight" (in models of sight n). The influence principles would then become: the agent adopts if (he knows that) enough of his *n-distant* neighbors (will) have adopted.

In the revolution case, agents might be influenced into adopting only if (they know that) enough agents *within the whole network* (will) have adopted. A suitable "globalized" version of the prediction update from Def. 22 may be defined as follows:

Definition 25 (Global Prediction Update). *Let $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ be a sight n finite model, and let $(F, \leq)$ be as in Def. 22.*

The global prediction update *of $\mathcal{M}$ results in the model $\mathcal{M}' = (\mathcal{W}, \mathcal{A}, N, \widetilde{B}, \theta, \{\sim'_a\}_{a\in\mathcal{A}})$ where:*

- *$\widetilde{B}$ is such that:*
  - $\forall w \in \mathcal{W}, \widetilde{B}(w) = B(w) \cup \{a \in \mathcal{A} : \forall v \sim_a w, \frac{|\mathcal{A}\cap\widetilde{B}(v)|}{|\mathcal{A}|} \geq \theta\}$
  - $\forall f \in F$, *if* $\forall w \in \mathcal{W}, f(w) = B(w) \cup \{a \in \mathcal{A} : \forall v \sim_a w, \frac{|\mathcal{A}\cap f(v)|}{|\mathcal{A}|} \geq \theta\}$, *then* $\widetilde{B} \leq f$.

  and
- *$w \sim'_a v$ iff i) $w \sim_a v$ and ii) if $b \in N^{\leq n}(w)(a)$, then $b \in \widetilde{B}(w)$ iff $b \in \widetilde{B}(v)$.*

Taking Chances. Prediction update has been defined to allow agents to take all their available information into account in their decision making. In acting upon it, agents act *conservatively*, as the information-dependent adoption policies defined require *absolute certainty*: agents will adopt only when they *know* that enough of the others (will) have adopted.

An alternative to such conservative behavior is a risky one, where agents adopt whenever they *consider it possible* that enough people (will) have adopted. In the revolution example, this means that agents would join the revolution whenever they see a chance that enough of their neighbors (or of the general population) would join.

Such chance taking behavior is captured as follows:

Definition 26 (Risky Prediction Update). *Let $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ be a sight n finite model, and let $(F, \leq)$ be as in Def. 22.*

The risky prediction update *of $\mathcal{M}$ results in the model $\mathcal{M}' = (\mathcal{W}, \mathcal{A}, N, \widetilde{B}, \theta, \{\sim'_a\}_{a\in\mathcal{A}})$ where:*

- *$\widetilde{B}$ is such that:*
  - $\forall w \in \mathcal{W}, \widetilde{B}(w) = B(w) \cup \{a \in \mathcal{A} : \exists v \sim_a w, \frac{|N(v)(a)\cap\widetilde{B}(v)|}{|N(v)(a)|} \geq \theta\}$
  - $\forall f \in F$, *if* $\forall w \in \mathcal{W}, f(w) = B(w) \cup \{a \in \mathcal{A} : \exists v \sim_a w, \frac{|N(v)(a)\cap f(v)|}{|N(v)(a)|} \geq \theta\}$, *then* $\widetilde{B} \leq f$.

  and
- *$w \sim'_a v$ iff i) $w \sim_a v$ and ii) if $b \in N^{\leq n}(w)(a)$, then $b \in \widetilde{B}(w)$ iff $b \in \widetilde{B}(v)$.*

To suitably capture e.g. a population of "risky revolutionaries", the risky prediction update should be suitably "globalized" by replacing $N(v)(a)$ with $\mathcal{A}$ everywhere in the definition.

Betting that just any uneliminated possibility is in fact the case is very risky behavior. A natural way to weaken the epistemic requirement of absolute certainty while still allowing for uncertainty to exist is to augment our framework with *beliefs*. Modeling beliefs using the *plausibility orders* of [4], a middle ground between conservative and risky prediction update could be defined. The natural definition would make agents adopt if enough neighbors (are predicted to) have adopted *in each of the worlds the agent considers most plausible*, i.e, if the agent believes enough neighbors (are predicted to) have adopted.

Trendsetters *vs.* Followers. An assumption build into threshold models in general is that agents are *followers*: even when they anticipate others' behavior with the prediction update, they only "anticipate their future following of others". Agents are thus *reacting* to others' behavior, even when they are reacting fast.

An interesting alternative would be to utilize agents' information to make them proactive instead; to have *trendsetters* instead of followers. Adding a few trendsetters to a network might induce behavior change towards B even when no-one has adopted initially.

A simple trendsetting adoption policy would state that an agent should adopt whenever she knows that *if* she were to adopt, then enough of her neighbors *will adopt afterwards*. Such an adoption policy involves both counterfactual and temporal reasoning, which complicates a predictive version. A non-predictive version may be defined as follows:

Definition 27 ((a,w)-Counterfactual Behavior). *Let $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a\in\mathscr{A}})$ be an ETM with $w \in \mathscr{W}$. Let the (a,w)-counterfactual behavior of $\mathscr{M}$ be*

$$B_{C(a,w)}(v) = \begin{cases} B(v) \cup \{a\} & \text{if } v \sim_a w \\ B(v) & \text{else} \end{cases}$$

Definition 28 (Trendsetter Update). *Let $\mathscr{M} = (\mathscr{W}, \mathscr{A}, N, B, \theta, \{\sim_a\}_{a\in\mathscr{A}})$ be an ETM and let $\{\mathscr{F}, \mathscr{T}\}$ be a partition of $\mathscr{A}$ into sets of followers and trendsetters.*

The trendsetter update *of $\mathscr{M}$ is the ETM $\mathscr{M}' = (\mathscr{W}, \mathscr{A}, N, B', \theta, \{\sim_a\}_{a\in\mathscr{A}})$ with B' given by $\forall w \in \mathscr{W}$*

$$B'(w) = B \cup \left\{ a \in \mathscr{F} : \forall v \sim_a w, \frac{|N(v)(a) \cap B|}{|N(v)(a)|} \geq \theta \right\} \cup \left\{ a \in \mathscr{T} : \forall v \sim_a w, \frac{|N(v)(a) \cap B_{C(a,v)}(v)'|}{|N(v)(a)|} \geq \theta \right\}$$

where $B_{C(a,v)}(v)'$ is the (a,v)-counterfactual behavior set of $\mathscr{M}$ after informed update.

The trendsetter update may of course also be define in global and risky versions.

6 Conclusions and Further Research

The paper has focused on two intertwined objectives. On the one hand, we have developped logical frameworks for the diffusion dynamics of the behavior of agents in social networks, and on the other hand we have developped models for the diffusion dynamics under uncertainty. We gradually have focussed our attention on agents which increasing cognitive abilities. At strart our threshold models did only focus on the adopting behavior of agents while in the following sections we have equipped agents with epistemic power and also predictive epistemic powers. In the following paragraphs, we summarize our findings.

Threshold Models. The static setting of threshold models may be described sufficiently using a propositional logic with proposition symbols that are indexed by agents. On finite networks, threshold ratios may be encoded together with other important structural notions, such as clusters of particular density. As the dynamics of threshold model update is deterministic and state dependent, these may be described using a dynamic modality reducible to the static language. The dynamic modality therefore does not add any expressive power. We have shown that the logic for threshold-limited influence is sound and complete, and as the static fragment is stated in simple propositional logic, one sees that this logic is also decidable.

Epistemic Threshold Models. Given the propositional logical representation of networks, the epistemic extension of the logic for threshold-limited influence works as expected. As both the diffusion and learning mechanism in the informed update are deterministic and state dependent, the dynamic process that is induced by the dynamic operation can be captured by a reducible dynamic modality. As such, this modality does not add any expressivity to the language. We have shown the epistemic logic of threshold-limted influence to be both sound and complete. Again we can conclude that this logic is decidable.

In epistemic threshold models, if agents' behavior is dictated by that of their direct neighbors, then knowledge of more distant agents is redundant. To act as under the standard threshold model dynamics, knowledge of neighbors' behavior is however required. If this information is not available, the diffusion speed decreases. In the limit case where no information is available, the diffusion process stops. Taken together, the most economical epistemic interpretation of standard threshold models is that their dynamics embodies an *implicit* epistemic assumption that exactly the network structure and behavior of agents in distance 1 is known.

Epistemic Threshold Models with Prediction Update. Prediction update allows agents to fully utilize their information in deciding if and when to adopt a spreading behavior. Describing the dynamics of prediction update requires a dynamic fixed point operator, which is atypical of dynamic epistemic logic. Here we have shown that formulas including this operator are not reducible to the static language. The dynamic operator which is studied in the context of our 'prediction update', thus strictly adds expressive power. The learning mechanisms of prediction update and informed update are identical, but given the fixed point construction involved in the former, obtaining a complete logic is a complex task and is left for future research. We have stated a fixed point axiom and least fixed point inference rule which are shown to be sound.

Future work. In future research we plan to work on a full comparative analysis of the different update processes that we have outlined in this paper. While convergence can be obtained for all different dynamic processes, among the ones we studied, the prediction dynamics will be the fastest in its convergence. In the limit case, where the network and behavior distribution is common knowledge, the prediction update jumps in one step to the fixed point of the standard threshold model update.

The logical treatment of threshold models and their epistemic extension undertaken also yields several more options for further development. Beyond the open question about a complete logic for prediction update, we see the three main directions for further research as the following: A) The logical apparatus and the epistemic extension of the possible generalizations of threshold models discussed is Subsection 2.4 are yet to be developed. B) The alternative diffusion processes introduced in Section 5 are to be further explored, both on the logical and on the set theoretic level. Their logics may be developed, and their dynamics may be investigated with respect to limit behavior and speed of possible stabilization. C)The epistemic and predictive treatment of non-increasing behaviors is yet to be investigated. Allowing agents to freely *unadopt* the possibly spreading behavior radically changes the limit behavior of systems by introducing the possibility of cyclic dynamics. Understanding the epistemics of such oscillating limit behavior requires tools going beyond the fixed point oriented mathematics of the current work.

References

1. K. R. Apt and E. Markakis. Diffusion in Social Networks with Competing Products. In G. Persiano, editor, *SAGT 2011*, LNCS 6982, pages 212–223. Springer, 2011. 1, 2.1
2. R. J. Aumann. Agreeing to Disagree. *Annals of Statistics*, 4:1236–1239, 1976. 3
3. A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. In *Proceedings of TARK'98 (Seventh Conference on Theoretical Aspects of Rationality and Knowledge)*, pages 43–56. Morgan Kaufmann Publishers, 1998. 2.2, 2.2
4. A. Baltag and S. Smets. A Qualitative Theory of Dynamic Interactive Belief Revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, Texts in Logic and Games, Vol. 3, pages 9–58. Amsterdam University Press, 2008. 5
5. A. Baltag, H. van Ditmarsch, and L. Moss. Epistemic Logic and Information Update. In P. Adriaans and J. van Benthem, editors, *Handbook on the Philosophy of Information*. Elsevier, 2008. 2.2, 2.2
6. J. v. Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011. 2.2, 2.2
7. Z. Christoff and J. U. Hansen. A two-tiered formalization of social influence. In D. Grossi, O. Roy, and H. Huang, editors, *Logic, Rationality, and Interaction*, volume 8196 of *Lecture Notes in Computer Science*, pages 68–81. Springer Berlin Heidelberg, 2013. 1
8. Z. Christoff and J. U. Hansen. A logic for diffusion in social networks. *Journal of Applied Logic*, 13(1):48–77, 2015. 1
9. B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, second edition, 2002. 4
10. D. Easley and J. Kleinberg. *Networks, Crowds, and Markets*. Cambridge University Press, 2010. 1, 2.1, 2.1, 2.3, 2.3, 2.3
11. S. Goyal. Interaction Structure and Social Change. *Journal of Institutional and Theoretical Economics*, 152(3):472–494, 1996. 1
12. M. Granovetter. Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978. 1
13. P. G. Hansen, V. F. Hendricks, and R. K. Rendsvig. Infostorms. *Metaphilosophy*, 44(3):301–326, 2013. 1
14. J. Hintikka. *Knowledge and Belief*. Cornell University Press, 1962. 3
15. D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *KDD 03: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM, New York, NY, USA, 2003. 1, 2.1
16. P. Kolaitis. Infinitary Logic in Finite Model Theory. In M. D. Chiara, K. Doets, D. Mundici, and J. van Benthem, editors, *Logic and Scientific Methods*, pages 9–58. Springer, 1997. 4
17. F.-H. Li, C.-T. Li, and M.-K. Shan. Labeled Influence Maximization in Social Networks for Target Marketing. *2011 IEEE Third Int'l Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third Int'l Conference on Social Computing*, pages 560–563, Oct. 2011. 1
18. C. Long and R. C.-W. Wong. Minimizing Seed Set for Viral Marketing Paper. *2011 IEEE 11th International Conference on Data Mining*, pages 427–436, Dec. 2011. 1
19. S. Morris. Contagion. *Review of Economic Studies*, 67:57–78, 2000. 1, 2.1, 2.3, 2.3
20. R. K. Rendsvig. Diffusion, Influence and Best-Response Dynamics in Networks: An Action Model Approach. In R. d. Haan, editor, *Proceedings of the ESSLLI 2014 Student Session*, pages 63–75, Tübingen, 2014. 1
21. J. Ruan and M. Thielscher. A logic for knowledge flow in social networks. In D. Wang and M. Reynolds, editors, *AI 2011: Advances in Artificial Intelligence*, volume 7106 of *Lecture Notes in Computer Science*, pages 511–520. Springer Berlin Heidelberg, 2011. 1

22. T. C. Schelling. *Micromotives and Macrobehavior*. Norton, New York, 1978. 1

23. J. Seligman, P. Girard, and F. Liu. Logical dynamics of belief change in the community, 2013. 1

24. J. Seligman, F. Liu, and P. Girard. Logic in the community. In M. Banerjee and A. Seth, editors, *Logic and Its Applications*, volume 6521 of *Lecture Notes in Computer Science*, pages 178–188. Springer Berlin Heidelberg, 2011. 1

25. P. Shakarian, S. Eyre, and D. Paulo. A Scalable Heuristic for Viral Marketing Under the Tipping Model. 2013. 1

26. P. Singh, S. Sreenivason, B. K. Szymanski, and G. Korniss. Threshold-limited spreading in social networks with multiple initiators. *Scientific Reports*, 3(2330), 2013. 1

27. T. W. Valente. Social network thresholds in the diffusion of innovations. *Social Networks*, 18(1):69–89, Jan. 1996. 1

28. H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, 2008. 2.2, 2.2, 2.2

29. L. Zhen and J. Seligman. A logical model of the dynamics of peer pressure. *Electronic Notes in Theoretical Computer Science*, 278(0):275 – 288, 2011. 1