Alexandru Baltag
Zoé Christoff
Rasmus K. Rendsvig
Sonja Smets

Dynamic Epistemic Logics of Diffusion and Prediction in Social Networks

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 diffusion 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. [15, 37]. Threshold-limited influence relies 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 [19, 38]. So-called *threshold models*, having gained early, wide-spread attention through

Accepted for publication in *Studia Logica*, 2018.
Special Issue: Logic and the Foundations of Game and Decision Theory (LOFT 12, 2016)
Edited by **Giacomo Bonanno, Wiebe van der Hoek and Andrés Perea**

[18, 33, 34], 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 [1, 15, 17, 20, 23, 25, 26, 36], also within the logic community [8, 9, 12, 10, 11, 24, 28, 29, 32, 35, 39].

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 2.3. While conceptually in line with [10, 11, 24, 28, 32, 35, 39] 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 [10, 11, 24, 35, 39] 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 2.4 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 the behavior. This is a conceptual jump from the initial minimal modeling of influence from Section 2 to a more sophisticated (information dependent) diffusion policy: Agents change from adopting the behavior whenever sufficiently many neighbors have done so to adopting the behavior only if they *know* that sufficiently many neighbors have done so. We then relate these two diffusion policies by showing under which epistemic conditions their diffusion dynamics is stepwise 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.

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 2.1 (Network). A network is a pair $(\mathcal{A}, N)$ where $\mathcal{A}$ is a non-empty finite set of *agents* and the function $N : \mathcal{A} \to \mathcal{P}(\mathcal{A})$ assigns a set $N(a)$ to each $a \in \mathcal{A}$, such that

*While the case of networks without symmetry is also interesting, for instance to model influence on "followers", most of the existing literature on threshold models mentioned in the introduction concerns the symmetric and irreflexive case only. This is why these restrictions are imposed here.

- $a \notin N(a)$ (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 a network together with the extension of a unique behavior (or opinion, fashion, or product) distributed over the agents, and a fixed uniform adoption threshold. A threshold model thus represents the current spread of the behavior throughout the network, while containing the adoption threshold which prescribes how this spread will evolve.

Remark 2.2. Throughout the text, we identify the behavior with its extension, i.e., with a designated subset B of $\mathcal{A}$ of agents that have adopted the behavior. Moreover, the verb "adopt" is used with "the behavior" as implicit object: When writing "Agent a has adopted", we imply that a has adopted the unique behavior in question.

Definition 2.3 (Threshold Model). A threshold model is a tuple $\mathcal{M} = (\mathcal{A}, N, B, \theta)$ where $(\mathcal{A}, N)$ is a network, $B \subseteq \mathcal{A}$ 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. Therefore, 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 2.4 (Threshold Model Update). The update of threshold model $\mathcal{M} = (\mathcal{A}, N, B, \theta)$ is the threshold model $\mathcal{M}' = (\mathcal{A}, N, B', \theta)$, where B' is given by

$$B' = B \cup \left\{ a \in \mathcal{A} : \frac{|N(a) \cap B|}{|N(a)|} \geq \theta \right\}. \tag{1}$$

This definition captures the idea that the new set of agents B' who adopted the behavior (in the new updated model $\mathcal{M}'$) does include the set of agents B who had already adopted the behavior before and it includes those agents who have enough neighbors (given by the number θ) that have adopted already. This definition is set in line with the standard approach on adoption rules in the literature [15].

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

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

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

Proposition 2.1. Let $\mathcal{S}_{\mathcal{M}}$ be a diffusion sequence. For some $n \in \mathbb{N} < |\mathcal{A}|$, we reach a fixed point $\mathcal{M}_n = \mathcal{M}_{n+1}$ in the sequence $\mathcal{S}_{\mathcal{M}}$.

PROOF. The fact that there is a $n \in \mathbb{N}$ such that $\mathcal{M}_n = \mathcal{M}_{n+1}$ follows immediately from the fact that $\mathcal{A}$ is finite and $B_n \subseteq B_{n+1}$ for all $n \in \mathbb{N}$. The fact that $n < |\mathcal{A}|$ 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 $\left|B_{|\mathcal{A}|-1}\right| = |\mathcal{A}|$. ⊣

2.2. 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, where agents are undeliberatively infected through viral contagion. Under this interpretation, the update procedure corresponds to that of a deterministic Susceptible-Infected (SI) model. It is closely related to a deterministic Susceptible-Infected-Susceptible (SIS) model, which also allows *unadoption* of the behavior in question. An SIS model diffusion policy given by Eq. 2, where the right term in the union captures a conservative tie-breaking rule:

$$B' = \left\{a : \frac{|N(a) \cap B|}{|N(a)|} > \theta\right\} \cup \left\{a : \frac{|N(a) \cap B|}{|N(a)|} = \theta \text{ and } a \in B\right\}. \quad (2)$$

Since Eq. 2 does not cause B to inflate, this alternative rule allows the possibility of *loops* in behavior, i.e. where $B = B'' \neq B'$. Thereby repeated updates according to Eq. 2 do *not* necessarily reach a fixpoint.

Dynamics as Induced by Game Play. Alternatively, agents may be interpreted as rational beings aiming towards coordination with their neighbors. In fact, Equations 1 and 2 correspond to the best response dynamics of agents playing an instance of a coordination game

	B	¬B
B	x, x	$0, 0$
¬B	$0, 0$	y, y

with each of their neighbors at each timestep, under the constraint that at each timestep, each agent may pick only one strategy to play simultaneously in all instances. The utility of a play round for an agent a is the sum of utilities of the individual coordination games played by a in that round. With B the set of agents currently playing action B, B is thus a best response for agent a iff

$$x \cdot \frac{|N(a) \cap B|}{|N(a)|} \geq y \cdot \frac{|N(a) \cap \neg B|}{|N(a)|} \Leftrightarrow \frac{|N(a) \cap B|}{|N(a)|} \geq \frac{y}{x+y}.$$

Defining θ as $\frac{y}{x+y}$, we specifically obtain that Eq. 2 captures such plays' best response dynamics with conservative tie-breaking [26]: B' as given by Eq. 2 is exactly the set of agents for whom B is a best response. Hence the diffusion dynamics arising from updating a network using best response analysis is step-wise equivalent with those given by Eq. 2. Moreover, for any $\theta \in [0, 1]$, there exists coordination game payoffs that yield best response dynamics equivalent to those of Eq. 2 instantiated with the given θ.

Equation 1, which we stick to as the foundation for the main diffusion policies of this paper, captures the same game-theoretic dynamics, but with two variations: First, the tie-breaking rule in Eq. 1 favors adopting the behavior over not doing so, in contrast with Eq. 2's conservative tie-breaking.† Second, with discriminating tie-breaking, but also the added assumption that the initial agents playing B will never stop doing so, either do to irrationality or because they sufficiently mutually support one another in that choice, cf. the Cluster Theorem of Sec. 2.4. See [15, 26] for game-theoretic details and [28] for action model-based logical treatments.

This paper focuses on the dynamics given by Eq. 1, for which we find the game-theoretic interpretation natural as a basis for rationality considerations.

Model Variations. There are several variations to threshold models as given by Def. 2.3 that one may wish to consider, given the application in mind. Numerous such exist in the literature, including infinite networks [26], networks with non-inflating behavior adoption [26], agent-specific thresholds [20], weighted links [20] and multiple behaviors [1]. For simplicity, we stick to

†Note that with a finite set of agents, the adoption threshold/coordination game payoffs could always be chosen as to eliminate any possibility that tie-breaking need be used.

the finite threshold models defined. This also holds throughout the epistemic extensions.

For the results presented, of these variations, only assuming the network infinite would give rise to revisions, cf. the comments concerning finiteness and reaching fixed points when applying predictive update polices in Sec. 4. In particular, both weighted links and agent-specific thresholds may be incorporated from the game-theoretic underpinnings of Eq. 1 by setting agents to play game with *non-symmetric* payoffs, possibly *varying* across neighbors. This will be relevant when modeling diffusion in networks where the relation is non-symmetric. Then a lower – or zero – weighting may be chosen for given interactions, thus obtaining non-symmetry. Though details should be revised, this variation would not cause significant difficulties for the presented.

2.3. 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 2.6 (Languages $\mathcal{L}_{[]}$ and $\mathcal{L}$). Let $\mathcal{A}$ be a finite set and let atoms be given by $\Phi = \{N_{ab} : a, b \in \mathcal{A}\} \cup \{\beta_a : a \in \mathcal{A}\}$. The language $\mathcal{L}_{[]}$ is then given by:

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

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

Disjunction and material implication are defined in the standard way. $\mathcal{L}_{[]}$ 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

[‡]Informationally, both features could be broken down using more fine-grained modelling tools, using e.g. hybrid [39] or term-modal logical [16, 27] approaches. For simplicity, we refrain from doing so.

atomic formulas. The $[adopt]$ modality is interpreted as is standard in dynamic epistemic logic[§] [4, 3, 6, 14]: Intuitively, we evaluate $[adopt]\varphi$ as true in a given model if and only if φ is true in the model after a given change occurs. Here, this change is that all agents simultaneously update their behavior according to the threshold update of Definition 2.4.

Definition 2.7 (Truth Clauses for $\mathcal{L}_{[]}$). Given a model $\mathcal{M} = (\mathcal{A}, N, B, \theta)$, $N_{ab}, \beta_a \in \Phi$, and $\varphi, \psi \in \mathcal{L}_{[]}$:

$\mathcal{M} \vDash \beta_a$	iff	$a \in B$
$\mathcal{M} \vDash N_{ab}$	iff	$b \in N(a)$
$\mathcal{M} \vDash \neg\varphi$	iff	$\mathcal{M} \nvDash \varphi$
$\mathcal{M} \vDash \varphi \wedge \psi$	iff	$\mathcal{M} \vDash \varphi$ and $\mathcal{M} \vDash \psi$
$\mathcal{M} \vDash [adopt]\varphi$	iff	$\mathcal{M}' \vDash \varphi$, where $\mathcal{M}'$ is the updated threshold model (Definition 2.4).

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:

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

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. 2.1) 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:

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

Proposition 2.2. Let $\mathcal{M} = (\mathcal{A}, N, B, \theta)$ be a threshold model. There exists $n \in \mathbb{N} < |\mathcal{A}|$ such that, for any $\varphi \in \mathcal{L}_{[]}$:

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

PROOF. As noted in the proof of Prop. 2.1, in the diffusion sequence $\mathcal{S}_{\mathcal{M}}$, for some $n \in \mathbb{N} < |\mathcal{A}|$, $\mathcal{M}_n = \mathcal{M}_{n+1}$. Hence $\mathcal{M}_n$ and $\mathcal{M}_{n+1}$ are guaranteed to satisfy the same formulas, whereby $\mathcal{M} \models [adopt]^n \varphi \leftrightarrow [adopt]^{n+1} \varphi$. ⊣

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 [4, 3, 6, 14].

DEFINITION 2.8 (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 2.9 ($\mathcal{C}_\theta$). Let the threshold $\theta \in [0, 1]$ be given. The class of threshold models $\mathcal{C}_\theta$ contains all and only models with the same threshold 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 $\mathcal{C}_\theta$: ¶

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

Network Axioms	
$\neg N_a a$	Irreflexivity
$N_{ab} \leftrightarrow N_{ba}$	Symmetry
$\bigvee_{b \in \mathcal{A}} N_{ab}$	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 $[adopt]\varphi$	$Nec_{[adopt]}$
From φ and $\psi \leftrightarrow \chi$, infer $\varphi[\psi/\chi]$	Repl. of Equiv.

Table 1. Hilbert-style proof system L_θ. In Replacement of Equivalents, $\varphi[\psi/\chi]$ is the formula resulting from replacing, in φ, every occurence of the subformula ψ with the formula χ.

Theorem 2.1 (Completeness). Let $\theta \in [0, 1]$. For any $\varphi \in \mathcal{L}$,

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

Proof. *Soundness:* Let $\mathcal{M} = (\mathcal{A}, N, B, \theta)$ be an arbitrary threshold model with $a, b \in \mathcal{A}$. Then $\mathcal{M}$ satisfies Irreflexivity (Symmetry/seriality) directly by the semantics and the assumption of irreflexivity (symmetry/seriality) of the network. $\mathcal{M} \models [adopt]N_{ab} \leftrightarrow N_{ab}$ 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 $\mathcal{M}$ satisfies Red.Ax.β, let $\mathcal{M}'$ be the adoption update of $\mathcal{M}$. Then $\mathcal{M} \models [adopt]\beta_a$ iff $\mathcal{M}' \models \beta_a$ iff $a \in B' = B \cup \left\{b \in \mathcal{A} : \frac{N(b) \cap B}{N(b)} \geq \theta\right\}$ iff $\mathcal{M} \models \beta_a$ or $a \in \left\{b \in \mathcal{A} : \frac{N(b) \cap B}{N(b)} \geq \theta\right\}$. A syntactic decoding following Definition 2.3 of the large, right-hand disjunct of Red.Ax.β (called $\beta_{N(a) \geq \theta}$) shows that it is satisfied iff $a \in \left\{b \in \mathcal{A} : \frac{N(b) \cap B}{N(b)} \geq \theta\right\}$: The outer disjunction requires/ensures the existence of two sets of agents, $\mathcal{G}$ and $\mathcal{N}$, such that $\mathcal{G} \subseteq \mathcal{N}$ and $\frac{|\mathcal{G}|}{|\mathcal{N}|} \geq \theta$.

The inner conjunction in Definition 2.3 is satisfied iff $\mathcal{N} = N(a)$ and $\mathcal{G} \subseteq B$. Hence φ is satisfied iff $\exists \mathcal{G} \subseteq N(a) \cap B : \frac{|\mathcal{G}|}{|N(a)|} \geq \theta$ iff $\frac{|N(a) \cap B|}{|N(a)|} \geq \theta$ iff $a \in \left\{ b \in \mathcal{A} : \frac{N(b) \cap B}{N(b)} \geq \theta \right\}$. Hence $\mathcal{M} \models [adopt]\beta_a$ iff $\mathcal{M} \models \beta_a$ or $\mathcal{M} \models \beta_{N(a) \geq \theta}$.

Completeness: The proof goes via translation of the dynamic language into the static part of the language, in the usual way (see e.g. [14, Ch. 7]). $\dashv$

2.4. 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. [15, 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 2.1:

Definition 2.10. The sentence abbreviated by '*cascade*' expresses that all agents will have adopted eventually:

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

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* [26], also referred to as a *cluster of density d* [15]. 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.

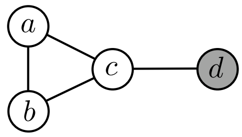

Figure 1. A social network with a cluster of density $\frac{2}{3}$.

Definition 2.11 (Cluster of density d). Given a network $(\mathcal{A}, N)$, a cluster of density d is any group $C \subseteq \mathcal{A}$ 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 $\mathcal{A}$, and that each singleton $\{a\} \subseteq \mathcal{A}$ is a cluster of density 0 (by irreflexivity).

Example: Clusters. Let $\mathcal{M}$ be the model illustrated in Figure 1, 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.

The language $\mathcal{L}$ 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 2.3. The group C is a cluster of density d in $(\mathcal{A}, N)$ iff $\mathcal{M} = (\mathcal{A}, N, B, \theta)$ satisfies

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

Proof. Left to right: Let $\mathcal{M} = (\mathcal{A}, N, B, \theta)$ and assume C is a cluster of density d in $(\mathcal{A}, N)$. Then by definition, for all $a \in C$, $\frac{|N(a) \cap C|}{|N(a)|} \geq d$. As $\mathcal{M}$ is based on $(\mathcal{A}, N)$, $\{b : \mathcal{M} \models N_{ab}\} = N(a)$ for all $a \in \mathcal{A}$. Let a be given and pick $\mathcal{N} = N(a)$ and $\mathcal{G} = N(a) \cap C$. Then $\frac{|\mathcal{G}|}{|\mathcal{N}|} \geq d$. Given the choice of $\mathcal{N}$, $\mathcal{M} \models \bigwedge_{b \in \mathcal{N}} N_{ab} \wedge \bigwedge_{b \notin \mathcal{N}} \neg N_{ab}$. So $\mathcal{M}$ satisfies (3).

Right to left: Assume that $\mathcal{M}$ satisfies (3) for some $C \subseteq \mathcal{A}$ and some $d \in [0, 1]$. Then for each $a \in C$, there is are sets $\mathcal{G}$ and $\mathcal{N}$ with $\mathcal{G} \subseteq \mathcal{N}$ and $\frac{|\mathcal{G} \cap C|}{|\mathcal{N}|} \geq d$, such that $\mathcal{N} = \{\mathcal{M} \models N_{ab}\} = N(a)$. Hence $\frac{|\mathcal{G} \cap C|}{|N(a)|} = \frac{|\mathcal{G} \cap C|}{|\mathcal{N}|} \geq d$. As $\mathcal{G} \cap C \subseteq \mathcal{N} = N(a)$, $\frac{|N(a) \cap C|}{|N(a)|} \geq d$. As a was arbitrary from C, C is indeed a cluster of density d in $(\mathcal{A}, N)$. $\dashv$

Given Proposition 2.3, 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 $\exists C_{\geq d}\neg\beta$):

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

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, continued. The model illustrated in Figure 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 $\exists C_{\frac{2}{3}}\neg\beta$:

$$\bigvee_{C\subseteq\mathcal{A}} \bigwedge_{a\in C} \bigvee_{\left\{\mathcal{G}\subseteq\mathcal{N}\subseteq\mathcal{A}:\frac{|\mathcal{G}\cap C|}{|\mathcal{N}|}\geq \frac{2}{3}\right\}} \left(\bigwedge_{b\in\mathcal{N}} N_{ab} \wedge \bigwedge_{b\notin\mathcal{N}} \neg N_{ab} \wedge \bigwedge_{b\in\mathcal{G}} \neg\beta_b \right). \tag{4}$$

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

$$\bigwedge_{b\in\mathcal{N}} N_{ab} \wedge \bigwedge_{b\notin\mathcal{N}} \neg N_{ab} \wedge \bigwedge_{b\in\mathcal{G}} \neg\beta_b. \tag{5}$$

To see that $\mathcal{M}$ satisfies (5), regard first agent c, for whom the appropriate $\mathcal{N}$ is $\{a, b, d\}$. As $|\mathcal{N}| = 3$, we must identify a group $\mathcal{G} \subseteq C$ with $|\mathcal{G}| \geq 2$ such that for all $b \in \mathcal{G}$, $\mathcal{M} \models N_{cb}$. Such a $\mathcal{G}$ exists, being $\{a, b\}$. Finally, indeed $\mathcal{M} \models \neg\beta_a \wedge \neg\beta_b$, and hence the conjunct for c is satisfied. Similar reasoning shows that the conjuncts for a and b also hold. This gives us (4).

The Cluster Theorem. The following theorem from [26], [15, Ch.19.3] characterizes the possibility of a complete adoption cascade in a network:

> Given a threshold model $\mathcal{M}$ with threshold $\theta \neq 0$ and a set $B \subset \mathcal{A}$ 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 $\mathcal{A}\backslash B$.

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

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

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

2.5. 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. 2.3) 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 $\theta : \mathcal{A} \to [0,1]$ in the definition of threshold models (Def. 2.3); 2) replace θ by $\theta(a)$ in the definition of the update (Def. 2.4) 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. 2.3) to let $\mathcal{B}$ be a finite *set* of behaviors ($\mathcal{B} = \{B_1, B_2, ...B_n\}$) and define $\theta : \mathcal{A} \times \mathcal{B} \to [0,1]$; 2) Relativize the definition of the update to each behavior B_i; 3) extend our set of atomic propositions: $\Phi = \{N_{ab} : a, b \in \mathcal{A}\} \cup \{\beta_{ia} : a \in \mathcal{A}, i \in 1, ...n\}$; 4) relativize the semantic clause in the obvious way: $\mathcal{M} \vDash \beta_{ia}$ iff $a \in B_i$, and replace the reduction axiom *Red.Ax.*β by *Red.Ax.*β_i accordingly. The "signature" of the resulting logic will then be given by $[\theta, \mathcal{A}, \mathcal{B}]$. 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

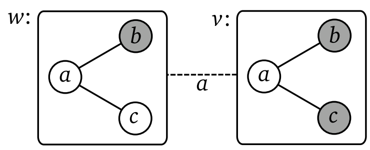

Figure 2. A situation of uncertainty. Agent a cannot tell whether world w or world v is the actual one, as indicated by the dashed line (when representing indistinguishability relations we omit reflexive and transitive links). Hence, a does not know whether c has adopted or not. Assume that the threshold 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!

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 Figure 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 knowledge 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 Figure 2. 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 3.1 (Epistemic Threshold Model (ETM)). An epistemic threshold model (ETM) is a tuple $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ where

- $\mathcal{W}$ is a finite, non-empty set of possible worlds (or states),
- $\mathcal{A}$ is a finite non-empty set of agents,
- $\sim_a \subseteq \mathcal{W} \times \mathcal{W}$ is an equivalence relation, for each agent $a \in \mathcal{A}$,
- $N : \mathcal{W} \to (\mathcal{A} \to \mathcal{P}(\mathcal{A}))$ assigns a neighborhood $N(w)(a)$ to each $a \in \mathcal{A}$ in each $w \in \mathcal{W}$, 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)
- $B : \mathcal{W} \to \mathcal{P}(\mathcal{A})$ assigns to each $w \in \mathcal{W}$ 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 distance n. To provide a formal definition, we first fix what is meant by "distance n":

[‖]Uncertainty concerning the adoption threshold could also be considered, as one reviewer points out. Mathematically, this would be straightforward, and if our remaining assumption below are kept in effect, nothing hinders this extension. We omit it in the name of simplicity, cf. Sec. 2.2

DEFINITION 3.2 (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 n sight at world w. Morever, if $b \in N^n(w)(a)$ we say that b is *n-reachable* from a in w.

DEFINITION 3.3 (Sight n Model**). 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 $b \in N^n(w)(a)$, then $b \in B(w)$ iff $b \in B(v)$ (agents know the behavior of others at least up to distance n).
- 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)

In other words: in an ETM of sight n, *the structure of the network and the others' behavior are known at least up to distance* n, and this is common knowledge. Note, though, that the n sight is common knowledge does not imply that all agents have equal sight: Some may see further.

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 she knows about other agents around her. 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 3.4 (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 (n sight) informed adoption update of $\mathcal{M}$ produces results in an ETM $\mathcal{M}^i = (\mathcal{W}, \mathcal{A}, N, B^i, \theta, \{\sim^i_a\}_{a\in\mathcal{A}})$ where for all $a \in \mathcal{A}$ and all $w, w' \in \mathcal{W}$, we put:

**Note that 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.

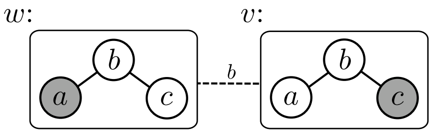

Figure 3. Adoption *de re* vs. adoption *de dicto*. We illustrate an ETM with threshold $\theta = 1/2$ and two possible worlds. Should b adopt or not? He knows *de dicto* that enough neighbors have adopted, but he does not know so *de re*; he knows that at least half of his neighbors have adopted, but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

- $B^i(w) = B(w) \cup \left\{a \in \mathcal{A} : \forall v \sim_a w \frac{|N(v)(a) \cap B(v)|}{|N(v)(a)|} \geq \theta\right\}$ and
- $w \sim_a^i w'$ iff i) $w \sim_a w'$ and ii) $\forall b \in N^n(w)(a)\colon b \in B^i(w) \Leftrightarrow b \in B^i(w')$.

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 of their own 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 can still see the (new) behavior of n-distant neighbors after the update.

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. For an illustration, see Figure 3.

A *de re* update is definable by setting

$$B^i(w) = B(w) \cup \left\{a \in \mathcal{A} : \frac{b \in \mathcal{A} : \forall v \sim_a w, |N(v)(a) \cap B(v)|}{|N(v)(a)|} \geq \theta\right\}.$$

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.

Learning the Distribution. When performing informed updates, agents may learn about the initial distribution of behavior in the network. Figure 4 provides an example. The learning occurs as agents' information cells may be restricted when other agents change their behavior. If such sufficient such restrictions occur, an agent may be left with a singleton information cell that allows only for a unique initial state. In this case, the agent will have learned

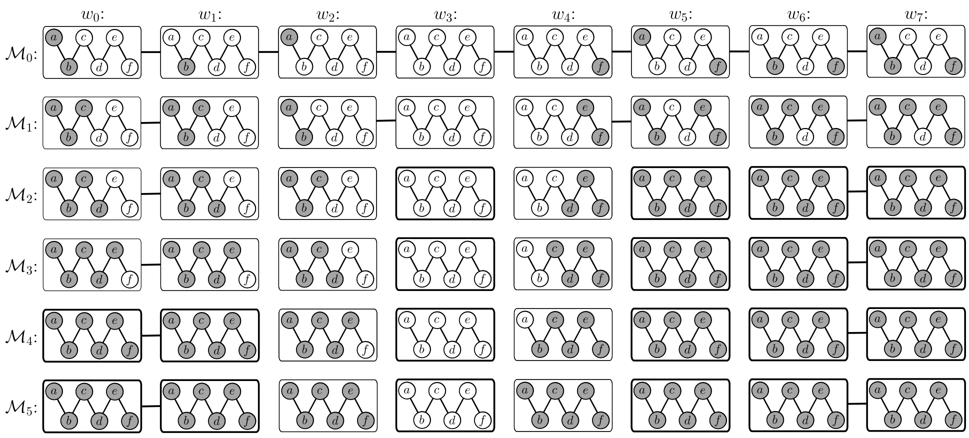

Figure 4. The learning process for agent d (bottom center) under informed adoption, in an ETM with threshold $\theta \leq \frac{1}{2}$ and sight 1. With sight 1, the ETM contains the 8 depicted possible worlds/states. The last states to reach fixed points at time 5 are states w_2 and w_4. Epistemic relations 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.

the initial distribution. This occurs in the initial state w_2 of Figure 4, but not in w_1. The information conveyed through perceiving the dynamics of the informed update policy may thus teach agents of the network at distances greater than their initial sight.

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 3.1 below.

Proposition 3.1 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 3.5 (State-Generated Threshold Model (SGM)). Let an ETM $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ be given and let $w \in \mathcal{W}$ and $a \in \mathcal{A}$. The state-generated threshold model $\mathcal{M}(w) = (\mathcal{A}, N_{\mathcal{M}(w)}, B_{\mathcal{M}(w)}, \theta)$ is given by:

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

Proposition 3.1. Let $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ be an ETM and $w \in \mathcal{W}$. Let $\mathcal{M}^i$ and $\mathcal{M}(w)$ be respectively the informed update and state-generated models of $\mathcal{M}$. Let $\mathcal{M}^i(w)$ be the state-generated model of $\mathcal{M}^i$ and let $\mathcal{M}(w)'$ be the non-epistemic threshold update of $\mathcal{M}(w)$. Then

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

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 $\mathcal{M}^i(w)$ and $B(w)'$ is the behavior set of $\mathcal{M}(w)'$.

Assume $a \in B(w)$. Then it follows that $a \in B(w)^i$ within $\mathcal{M}^i$, by monotonicity of the informed update. Hence we also obtain $a \in B_{\mathcal{M}^i(w)}$ in $\mathcal{M}^i(w)$ by Definition 3.5 of SGMs. From $a \in B(w)$ it also follows that $a \in B_{\mathcal{M}(w)}$ by defintion of SGMs. By monotonicity of the non-epistemic threshold update, we have $a \in B'_{\mathcal{M}(w)}$ in $\mathcal{M}(w)'$.

Assume that $a \notin B(w)$. Then $a \notin B_{\mathcal{M}(w)}$ by definition 3.5 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 $\mathcal{M}$ 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 $N(w)(a) = N_{\mathcal{M}(w)}(a)$ and $B(w) = B_{\mathcal{M}(w)}$, it follows that $\frac{|N_{\mathcal{M}(w)}(a)| \cup B_{\mathcal{M}(w)}}{|N_{\mathcal{M}(w)}(a)|} \geq \theta$ iff $a \in B_{\mathcal{M}(w)}$. ⊣

Proposition 3.1 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 partial, since the restriction to the case of sight $n \geq 1$ does not fully characterize the standard threshold dynamics as given in Def. 2.4. In the case of *no sight* ($n = 0$), the agent may have uncertainty about some neighbor b's behavior, and might not even know exactly who are all her neighbors; but she might still know that a large enough proportion of these neighbors have adopted B: in which case she will still update 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 3.2. There exists an ETM $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ with sight $n < 1$ such that

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

where $\mathcal{M}^i$ and $\mathcal{M}(w)$ are respectively the informed update and state-generated models of $\mathcal{M}$, and $\mathcal{M}^i(w)$ is the state-generated model of $\mathcal{M}^i$ and $\mathcal{M}(w)'$ is the non-epistemic update (Def. 2.4) of $\mathcal{M}(w)$.

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

Figure 5 illustrates this "slower" diffusion process.

3.3. Knowledge and Cascades

In Section 2.4, 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 3.1 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 does not hold in these models with sight less than 1. We briefly explain this point in more detail. Given proposition 3.2 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

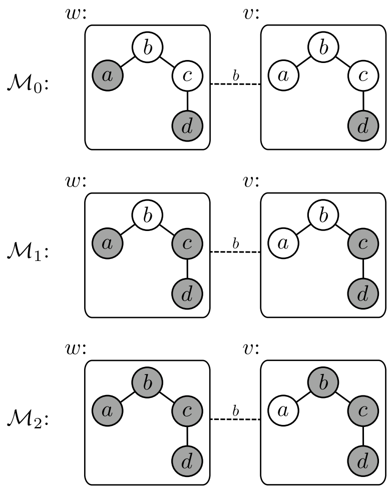

Figure 5. A diffusion process "slowed down" by the uncertainty of agent b, with threshold $\theta = \frac{1}{2}$. Consider the situation in world w: agent a has adopted, but agent b does not know it. Therefore, agent b will not adopt immediately. The diffusion according to the informed update policy 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 $\mathcal{M}_0(w)$ gives us directly $\mathcal{M}_2(w)$).

cascade, despite the absence of a cluster-obstacle. Figure 6 illustrates this difference.

3.4. The Epistemic Logic of Threshold-Limited Influence

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

DEFINITION 3.6 (Languages $\mathcal{L}_{K[]}$ and $\mathcal{L}_K$). Let the set of atomic propositions be given by $\{N_{ab} : a, b \in \mathcal{A}\} \cup \{\beta_a : a \in \mathcal{A}\}$ for a finite set $\mathcal{A}$. Where $a, b \in \mathcal{A}$, the formulas of $\mathcal{L}_{K[]}$ 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 the "static" fragment $\mathcal{L}_K$ are those of $\mathcal{L}_{K[]}$ 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 $\langle adopt\rangle$ as the dual of $[adopt]$.

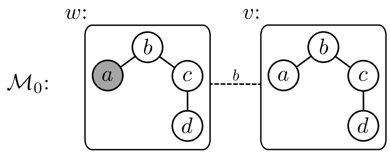

Figure 6. A diffusion process "blocked" by the uncertainty of agent b, with $\theta = \frac{1}{2}$. Consider the situation in world w: agent a has adopted, but agent b does not know it. Therefore, agent b will not adopt (under the informed adoption rule). Note that under the non-epistemic threshold update,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

DEFINITION 3.7 (Semantics for $\mathcal{L}_{K[]}$ with Informed Update). Formulas φ, ψ from $\mathcal{L}_{K[]}$ are interpreted over an ETM $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a \in \mathcal{A}})$ with sight n, and $w, v \in \mathcal{W}$:

$\mathcal{M}, w \models \beta_a$	iff	$a \in B(w)$
$\mathcal{M}, w \models N_{ab}$	iff	$b \in N(w)(a)$
$\mathcal{M}, w \models \neg\varphi$	iff	$\mathcal{M}, w \not\models \varphi$
$\mathcal{M}, w \models \varphi \wedge \psi$	iff	$\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
$\mathcal{M}, w \models K_a\varphi$	iff	for all $v \in \mathcal{W}$ such that $v \sim_a w$, $\mathcal{M}, v \models \varphi$
$\mathcal{M}, w \models [adopt]\varphi$	iff	$\mathcal{M}', w \models \varphi$, where $\mathcal{M}'$ is the informed update of $\mathcal{M}$ as specified in Def. 3.4 .

3.4.1. 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, with the formula N^k_{ab} expressing that b is k-reachable from a.

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

Abbreviation. For $\mathcal{B} \subseteq \mathcal{A}$, we introduce the abbreviation $\mathcal{B} = N^k_a\beta^+$ refering to the set of agents which are 1) k-reachable from a and 2) will have adopted after the next update:

$$\left(\mathcal{B} = N_a^k \beta^+\right) := \bigwedge_{b\in\mathcal{B}} \left(N_{ab}^k \wedge [adopt]\beta_b\right) \wedge \bigwedge_{b\in\mathcal{A}\setminus\mathcal{B}} \left(N_{ab}^k \to [adopt]\neg\beta_b\right).$$

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 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 3.8 (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 3.9 (Class: $\mathcal{C}_{\theta n}$). For $\theta \in [0,1]$ and $n \in \mathbb{N}$, the class $\mathcal{C}_{\theta n}$ consists of all ETM's with threshold θ and sight n.

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

Theorem 3.1 (Soundness, Completeness, Expressivity and Decidability). Let $\theta \in [0,1]$ and $n \in \mathbb{N}$. For any $\varphi \in \mathcal{L}_{K[]}$,

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

The language $\mathcal{L}_{K[]}$, endowed with the informed update semantics, has the same expressivity as its static counterpart $\mathcal{L}_K$. Moreover, $L_{\theta n}$ is decidable.

Proof. *Soundness:* Let $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ be an epistemic threshold model with sight n. Let $a, b \in \mathcal{A}$ and $w, v \in \mathcal{W}$. Then $(\mathcal{M}, w)$ 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 $(\mathcal{M}, w)$ satisfies **Ep.Red.Ax.**β, let $\mathcal{M}^i$ be the informed update of $\mathcal{M}$. Then $\mathcal{M}, w \models [adopt]\beta_a$ iff $\mathcal{M}^i, w \models \beta_a$ iff $a \in B^i(w) = B(w) \cup \left\{b \in \mathcal{A} : \forall v \sim_b w \frac{|N(v)(b)\cap B(v)|}{|N(v)(b)|} \geq \theta\right\}$ iff $\mathcal{M}, w \models \beta_a$ or $a \in \left\{b\in\mathcal{A}:\forall v\sim_b w \frac{|N(v)(b)\cap B(v)|}{|N(v)(b)|} \geq\theta\right\}$. Using the same syntactic decoding as in the proof of Theorem 2.1, we obtain that $a \in \left\{b\in\mathcal{A}:\forall v\sim_b w \frac{|N(v)(b)\cap B(v)|}{|N(v)(b)|} \geq\theta\right\}$ iff

Network Axioms		
$\neg N_{aa}$		Irreflexivity
$N_{ab} \leftrightarrow N_{ba}$		Symmetry
$\bigvee_{b \in \mathcal{A}} N_{ab}$		Seriality
Knowledge Axioms		
$K_a\varphi \to \varphi$	$(*)$	Ax.T
$K_a\varphi \to K_aK_a\varphi$	$(*)$	Ax.4
$\neg K_a\varphi \to K_a\neg K_a\varphi$	$(*)$	Ax.5
Knowledge-Network Axioms		
$(N^n_{ab} \wedge \beta_b) \to K_a\beta_b$	$(*)$	Known Behavior
$(N^{n-1}_{ab} \wedge N_{bc}) \to K_aN_{bc}$	$(*)$	Known Neighbors
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.$\wedge$
$[adopt]\beta_a \leftrightarrow \beta_a \vee K_a(\beta_{N(a)} \geq \theta)$	$(*)$	Ep.Red.Ax.β
$[adopt]K_a\varphi \leftrightarrow \bigvee_{\mathcal{B} \subseteq \mathcal{A}} (\mathcal{B} = N^n_a\beta^+ \wedge K_a (\mathcal{B} = N_a\beta^+ \to [adopt]\varphi))$	$(*)$	Ep.Red.Ax.K.sight.n
Inference Rules		
From φ and $\varphi \to \psi$, infer ψ		Modus Ponens
From φ, infer $K_a\varphi$ for any $a \in \mathcal{A}$	$(*)$	Nec.K_a
From φ, infer $[adopt]\varphi$		Nec.$[adopt]$
From φ and $\psi \leftrightarrow \chi$, infer $\varphi[\psi/\chi]$		Repl. of Equiv.

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

$\mathcal{M}, w \models K_a\left(\beta_{N(a)} \geq \theta\right)$. Hence $\mathcal{M}, w \models [adopt]\beta_a$ iff $\mathcal{M}, w \models \beta_a$ or $\mathcal{M}, w \models K_a\left(\beta_{N(a)} \geq \theta\right)$.

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

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

iff

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

iff

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

$$\mathcal{M}, w \models$$

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

iff

$\exists\mathcal{B} \subseteq \mathcal{A} : \mathcal{B} = N^n(w)(a) \cap B^i$ and
for all $v \sim_a w$, if $\mathcal{B} = N^n(v)(a) \cap B^i$, then $\mathcal{M}^i, v \models \varphi$
$(*)$

iff

$\exists\mathcal{B} \subseteq \mathcal{A} : \mathcal{B} = N^n(w)(a) \cap B'$ and
if $\mathcal{B} = N^n(w)(a) \cap B^i$, then $\mathcal{M}^i, w \models K_a\varphi$
(from $(*)$ as $\mathcal{M}$ 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

$\mathcal{M}^i, w \models K_a\varphi$
(as such a $\mathcal{B}$ always exists)

iff

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

Completeness (sketch): It can be shown by induction that for all $\varphi \in \mathcal{L}_{K[]}$, there exists a $\varphi' \in \mathcal{L}_K$ 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.

Expressivity and Decidability (sketch): May be shown by the reduction axioms. ⊣

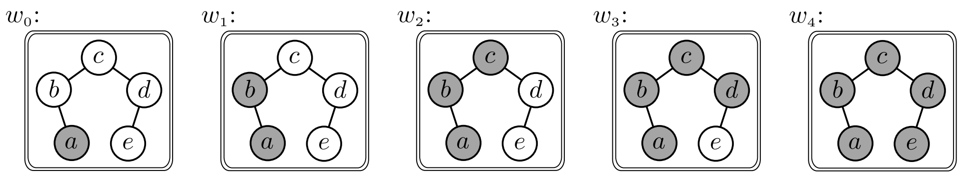

Figure 7. An ETM with no uncertainty about the actual state w, developing according to informed update. B is marked by gray, and a threshold $\theta = 1/2$ is assumed. At time 0 (w_0), only a has adopted. According to informed adoption, b adopts at time 1. At time 2, c also adopts the behavior, etc.

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 Figure 7, with a given dynamics that runs according to a non-epistemic or informed adoption policy.
If one assumes that agents (nodes) are not merely nomothetically influenced by their neighbors, but rather are rational agents seeking to coordinate behavior with their neighbors [26], the dynamics in Figure 7 seems to miss the target. 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 successfully coordinate with more neighbors and thus 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*. This is necessary because of the circular character of prediction update: an agent adopts based on the predicted behavior of her neighbors, but that behavior is in its turn based on their predictions about the first agent's behavior (among others), etc.

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

> Consider agent d's reasoning about the behavior of agent c, with d assuming that c acts in accordance with the (non-predictive) informed update policy. Then d may predict that c will adopt during the second round of updates. Hence, as d seeks to coordinate with c and e and has an adoption threshold of $\theta = 1/2$, d may act *preemptively*: To maximize the number of rounds in which she has adopted if a θ fraction of neighbors have adopted, d may update already in round 2, together with c.
>
> In this case, d may be thought of as a *level 1 predictor*: She assumes no-one else makes predictions, that the others are of level 0. However, d may come to think that c is as smart as she is, i.e., that also c is a level 1 predictor. Assuming this, d now foresees that c will not wait till round 2 to adopt, but will instead adopt B already in round 1; based on this prediction about c's predictions, d may now also adopt in round 1. In this case, d 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 of 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 do 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.

The second assumption is that it is common knowledge that predictors use all their available information (about the network structure, the current behavior spread and information available to others) as far into the future as possible when determining their next action.

Prediction Update Preliminaries. Before we define the prediction update, a few preliminaries are required. An example of prediction update is given in Figure 8 which follows the definitions.

Definition 4.1 (The Lattice of Behaviors and the Informed-Update Map). For a given ETM $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ let $\mathcal{P}(\mathcal{A})^{\mathcal{W}}$ be the set of all possible "behaviors", i.e. all functions $f : \mathcal{W} \to \mathcal{P}(\mathcal{A})$. We can convert this set into a lattice, by defining a partial order $\preceq$ on $\mathcal{P}(\mathcal{A})^{\mathcal{W}}$, given by:

$$f \preceq g \;\Leftrightarrow\; f(w) \subseteq g(w).$$

The *informed-update map* is a function

$$\Gamma_B : \mathcal{P}(\mathcal{A})^{\mathcal{W}} \longrightarrow \mathcal{P}(\mathcal{A})^{\mathcal{W}},$$

mapping any behavior $f \in \mathcal{P}(\mathcal{A})^{\mathcal{W}}$ to some behavior $\Gamma_B(f)$, given by, for all $w \in \mathcal{W}$,

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

Lemma 4.1. Let $\mathcal{M}$, $\mathcal{P}(\mathcal{A})^{\mathcal{W}}$, $\preceq$ and Γ_B be as in Definition 4.1. Then
1) $(\mathcal{P}(\mathcal{A})^{\mathcal{W}}, \preceq)$ is a finite, and hence complete, lattice.
2) Informed update Γ_B is an order-preserving (monotonic) map.

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

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

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

For all $v \in \mathcal{W}$, 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 $\Gamma_B(f)(u) \subseteq \Gamma_B(f')(u)$. As u was arbitrary, $\Gamma_B(f) \preccurlyeq \Gamma_B(f')$. Hence Γ_B is order-preserving. ⊣

Definition 4.2 (Least Fixed Point). Let $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a \in \mathcal{A}})$ be an ETM, and $\mathcal{P}(\mathcal{A})^{\mathcal{W}}, \preceq$ and Γ_B be as in Definition 4.1. The *least fixed point* of Γ_B, denoted by $\mathtt{lfp}(\Gamma_B)$, is the unique behavior $x \in \mathcal{P}(\mathcal{A})^{\mathcal{W}}$ such that

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

Theorem 4.1 (`lfp` Existence, Uniqueness and Approximation). Let be given $\mathcal{M}$, $(\mathcal{P}(\mathcal{A})^{\mathcal{W}}, \preceq)$ and Γ_B as in Definition 4.1. Then $\texttt{lfp}(\Gamma_B)$ exists. Moreover, this least fixed point is unique, and it can be reached by finite iterations of the informed-update map starting on the bottom element of the lattice. More precisely: if we put

$$\Gamma_B^0 = \bot, \text{ where } \bot(w) = \emptyset \text{ for all } w \in \mathcal{W},$$
$$\Gamma_B^{n+1} = \Gamma_B(\Gamma_B^n), \text{ for all } n \geq 1,$$

then there exists some $N \in \mathbb{N}$, such that the sequence stabilizes at stage N, and we have: $\texttt{lfp}(\Gamma_B)(w) = \Gamma_B^N(w) = \Gamma_B^{N+1}(w)$.

Proof. By Lemma 4.1, Γ_B is a monotonic map on the complete lattice $(\mathcal{P}(\mathcal{A})^{\mathcal{W}}, \preceq)$. Hence, the least fixed point $\texttt{lfp}(\Gamma_B)$ exists by the Knaster-Tarski Fixed Point Theorem (see e.g. [13, p. 50]). Moreover, since our lattice is finite, the proof of that theorem shows in fact that $\texttt{lfp}(\Gamma_B)$ is reached at some finite iteration Γ_B^N. ⊣

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

Definition 4.3 (Prediction Update). Let $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ be an ETM of sight n, and let $(\mathcal{P}(\mathcal{A})^{\mathcal{W}}, \preceq)$ be as in Lemma 4.1. Let $\Gamma_B : \mathcal{P}(\mathcal{A})^{\mathcal{W}} \longrightarrow \mathcal{P}(\mathcal{A})^{\mathcal{W}}$ be given as in Definition 4.1.

The *prediction update* of $\mathcal{M}$ is the ETM $\mathcal{M}^p = (\mathcal{W}, \mathcal{A}, N, B^p, \theta, \{\sim_a^p\}_{a\in\mathcal{A}})$ where $\forall w, w' \in \mathcal{W}$,

$$B^p(w) = \texttt{lfp}(\Gamma_B)(w), \text{ and}$$
$$w \sim_a^p w' \text{ iff } \text{i) } w \sim_a w', \text{ and}$$
$$\text{ii) } \forall b \in N^{\leq n}(w)(a) : b \in B^p(w) \Leftrightarrow b \in B^p(w')$$

Theorem 4.1 is important, since it ensures first, that our prediction update is well-defined, and second that, when engaged in prediction update agents do *not* run the risk of falling into infinite chains of reasoning about each other (which presumably would take an infinite time): they can compute the resulting prediction (and update) in finitely many steps.

Example, Sanity Check and Proof of Concept. The "irrational" behavior illustrated in Figure 7 is solved by prediction update. The dynamics are illustrated in Figure 8. Notice that now c adopts B *as soon as she knows* B *is preferred.*

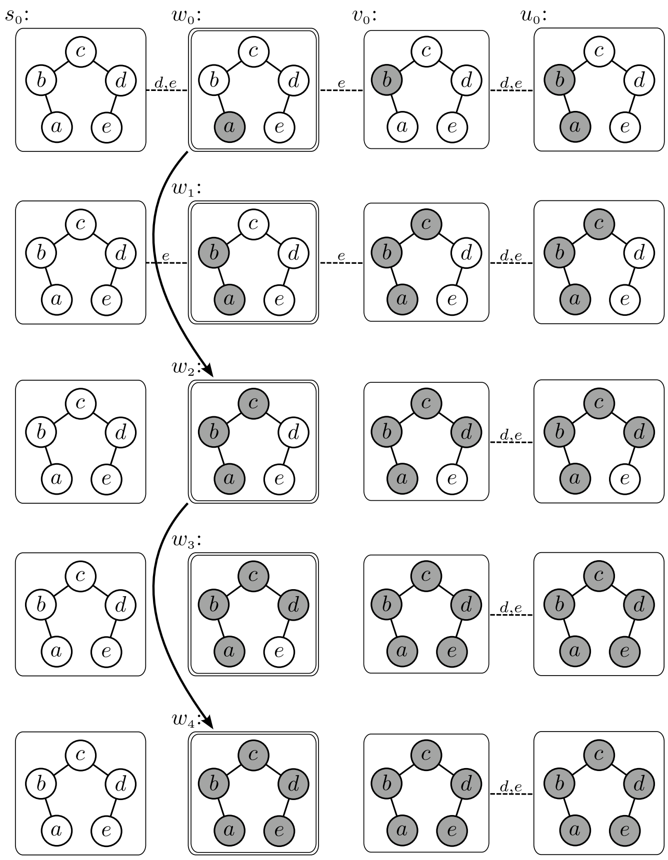

Figure 8. The prediction update of a sight 2 ETM with actual state w, $\theta = 1/2$. Agents a, b, c know the actual state; d, e are uncertain. The development of the five stages is given according to informed adoption; states w_0–w_4 are from Figure 7. The arrow shows the prediction update of the actual world. With informed update, w reaches a fixed point after 4 updates; with prediction update, it reaches the same fixed point after only 2 steps. Due to uncertainty, the prediction update does not jump to the fixed point in 1 step: as d does not know whether a has adopted at time 0, she does not know that c will adopt under prediction update. Hence, she will refrain herself from adopting until w_3. Similar considerations goes for e.

Bounded Rationality. Prediction assumes that agents have unbounded rationality (maximal predictive and reasoning power given the available information). A *bounded* rationality version of prediction update could be defined, in which agents can only compute a fixed finite number n of steps of the prediction chain. A natural way of doing this would be by defining an update that uses Γ_B^n instead of $\mathtt{lfp}(\Gamma_B)$. When n is low enough, the dynamics of bounded-rationality update would differ from the dynamics of unbounded prediction update. We leave the exploration of bounded-rationality updates for future work.

Iterated Dynamics, Fixed Point, Cascades, Speed of Convergence. When any of our adoption updates is iterated, a long-term dynamics is produced, in the form of an infinite sequence of models $\mathcal{M}, \mathcal{M}^{(1)}, \mathcal{M}^{(2)}, \ldots$. Since all the update rules considered in this paper are inflationary, a fixed point is always eventually reached: a stage N such that $\mathcal{M}^{(N)} = \mathcal{M}^{(N+1)}$. The *extent of the cascade* produced by each update type on an initial model $\mathcal{M}$ is given by the behavior $B^{(N)}$ in the fixed point $\mathcal{M}^{(N)}$, which comprises the set of all agents who will eventually adopt B (in a given world). A *full cascade* is produced if all agents will eventually adopt B, i.e. $B^N(w) = \mathcal{A}$. It is easy to see that *prediction update accelerates the cascading behavior in comparison to informed updated*: the fixed point of the adoption process is typically reached earlier if the agents use prediction update than if they use informed update. A full analysis of the relationship between the three types of update is left for future work. But a concrete example in this sense is given below.

4.1. On the Fixed-Point Logic of Prediction Update

The above stated prediction update rule in Definition 4.3 can now be used to give a new semantics to the [*adopt*] modality in the logic language $\mathcal{L}_{K[]}$.

Definition 4.4 (Prediction Update Semantics). Given $\theta \in [0,1]$, $n \in \mathbb{N}$ and any ETM $\mathcal{M} \in \mathcal{C}_{\theta n}$, the satisfaction relation for the *prediction update semantics* can be defined using the same truth clauses as in Def. 3.7, except for the formulas of the form $[adopt]\varphi$, for which we put:

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

Axiomatization. We present an axiomatic system that is *sound* for the logic of prediction update, although *completeness remains an open question.*

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

Table 3. Fixed point laws of prediction update logic $L^{predict}_{\theta n}$. The fixed point axiom takes the place of the informed update reduction axiom and the least fixed point inference rule is added.

Note that in this section, the [*adopt*] modality is a fixed point operator and hence may no longer be reduced away. In contrast to the informed update logic, the prediction update logic appears to be strictly more expressive than its static counterpart.

To state the proof system, we first generalize the syntactic shorthand introduced in Definition 2.3.

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

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

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 4.5 (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.*β 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.

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.

The proposition below establishes soundness of $L_{\theta n}^{predict}$. As for completeness, the proof would not go through the standard methods used in the previous sections. We therefore leave it for future research. However, we have reasons to make the following

Conjecture: The system $L_{\theta n}^{predict}$ is a complete axiomatization of predictive update logic over the class $\mathcal{C}_{\theta n}$.

Proposition 4.1. 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 $\mathcal{M}$ be a arbitrary finite ETM with sight n, domain $\mathcal{W}$ containing state w and $a, b \in \mathcal{A}$.

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

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

iff

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

iff

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

iff

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

iff

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

iff

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

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

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

$$\Gamma_{\overline{\varphi}}(f) = h \text{ such that}$$

$$\forall w \in \mathcal{W}, h(w) = \overline{\varphi}(w) \cup \left\{ a \in \mathcal{A} : \forall v \sim_a w, \frac{|N(v)(a) \cap f(v)|}{|N(v)(a)|} \geq \theta \right\}.$$

The same reasoning used in the proof of Lemma 4.1 shows that each such $\Gamma_{\overline{\varphi}}$ is order-preserving.

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

Given the prediction semantics of $[adopt]$ and the fact (cf. Theorem 4.1) that $B^p = \mathtt{lfp}(\Gamma_B) = \sup\{{\Gamma_B}^n(\bot) : n \in \mathbb{N}\}$, we may conclude that

$$\overline{[]\beta} = \Gamma_{\overline{\beta}}(\overline{[]\beta}) \tag{6}$$

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

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

$$\vdash \bigwedge_{a\in\mathcal{A}} (\varphi_a \leftrightarrow \beta_a \vee K_a(\varphi_{N(a)} \geq \theta)). \tag{7}$$

From $\{\varphi_a\}_{a\in\mathcal{A}}$ and $\{\beta_a \vee K_a(\varphi_{N(a)} \geq \theta)\}_{a\in\mathcal{A}}$ 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 (7) to be satisfied, we have that

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

Given that (6) is the least fixed point of $\Gamma_{\overline{\beta}}$, we have that $\overline{\varphi} = \Gamma_{\overline{\beta}}(\overline{\varphi})$ implies $\overline{[]\beta} \preceq \overline{\varphi}$, so

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

$\dashv$

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. 2.4); one depending on *knowledge* of this fact (the "informed update" of Def. 3.4), and one depending on the *anticipation* of this fact (the "prediction update" of Def. 4.3). This section questions some in-built assumptions of these policies and discusses possible alternatives.

5.1. 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. 4.3 may be defined as follows:

Definition 5.1 (Global Prediction Update). Let be given a finite sight n model, $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a \in \mathcal{A}})$, and let $(F, \leq)$ be as in Definition 4.3. The *global prediction update* of $\mathcal{M}$ is then $\mathcal{M}' = (\mathcal{W}, \mathcal{A}, N, \widetilde{B}, \theta, \{\sim'_a\}_{a \in \mathcal{A}})$ 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 $b \in \widetilde{B}(v)$.

5.2. 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 5.2 (Risky Prediction Update). Let a finite sight n model $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a \in \mathcal{A}})$ and $(F, \leq)$ be as in Definition 4.3. 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 [5], 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.

5.3. 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 5.3 ((a, w)-Counterfactual Behavior). Let be given an ETM $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ with $a \in \mathcal{A}, w \in \mathcal{W}$. Then the (a, w)-*counterfactual* behavior of $\mathcal{M}$ is

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

Definition 5.4 (Trendsetter Update). Let $\mathcal{M} = (\mathcal{W}, \mathcal{A}, N, B, \theta, \{\sim_a\}_{a\in\mathcal{A}})$ be an ETM and let $\{\mathcal{F}, \mathcal{T}\}$ be a partition of $\mathcal{A}$ into sets of followers and trendsetters. The *trendsetter update* of $\mathcal{M}$ results in the model $\mathcal{M}' = (\mathcal{W}, \mathcal{A}, N, B', \theta, \{\sim_a\}_{a\in\mathcal{A}})$ with B' given by $\forall w \in \mathcal{W}$

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

where $B_{C(a,v)}(v)'$ is the (a, v)-*counterfactual behavior set of* $\mathcal{M}$ *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 developed models for the diffusion dynamics under uncertainty, based on two natural epistemic variants of the standard threshold adoption rule: the informed update, and the prediction update. On the other hand, we presented logical frameworks for reasoning about diffusion dynamics. We proved soundness and completeness for the logic of informed update, and proposed a sound system for the logic of prediction update. The problem of completeness for the later logic is an open question. 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-limited 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 better coordinate with their neighbors in adopting a spreading behavior, by using their information about the others' future behavior. As a result, prediction-update agents increase a network's speed of convergence. In the extreme case when the network and behavior distribution are common knowledge, the prediction update jumps in one step to the fixed point of the standard threshold model update. But in general, even describing the one-step dynamics of prediction update requires a dynamic fixed point operator, which is atypical of dynamic epistemic logic. As a consequence, the logic of prediction update does not have full reduction axioms: the dynamic modality seems to genuinely add expressivity in this case. This poses technical challenges to obtaining a completeness proof.

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. We plan to tackle in another paper the open problem about completeness of the logic of prediction update. Besides this question, there are five other main directions for further research:

A) develop the logical apparatus and the epistemic extension of the possible generalizations of threshold models discussed is Subsection 2.5; B) explore the alternative diffusion processes introduced in Section 5, both on the logical, set theoretic and game theoretic levels. Their logics may be developed, and their dynamics may be investigated with respect to limit behavior and speed of possible stabilization; (C) explore the dynamics induced by boundedly-rational versions of predictive update; (D) explore the game theoretic perspectives of game play on networks under uncertainty and in particular the game structure underpinning the intuitive rationality of prediction update; (E) investigate the epistemic and predictive versions of the non-inflationary adoption rules, such as the policy given by regular coordination game play on networks [26]. Such rules, that allow agents to *unadopt* an already adopted behavior, can lead to very different limit behavior, e.g. to a cyclic dynamics. Understanding the epistemic aspects of such oscillating behavior will require logical tools going beyond the fixed point theory used in this paper.[††]

Acknowledgements. We would like to thank the anonymous reviewers for their insightful comments. Zoé Christoff acknowledges support for this research by the Engineering and Physical Sciences Research Council (EPSRC) under grant EP/M015815/1 "Foundations of Opinion Formation in Autonomous Systems", as well as from the Deutsche Forschungsgemeinschaft (DFG) and Grantová agentura České republiky (GAČR) joint project RO 4548/6–1 "From Shared Evidence to Group Attitudes". The contribution of Rasmus K. Rendsvig to the research reported in this article was funded by the Swedish Research Council through the framework project 'Knowledge in a Digital World' (Erik J. Olsson, PI). The Center for Information and Bubble Studies is sponsored by the Carlsberg Foundation.

References

[1] Apt, Krzysztof R., and Evangelos Markakis, 'Diffusion in Social Networks with Competing Products', in G. Persiano, (ed.), *SAGT 2011*, LNCS 6982, Springer, 2011, pp. 212–223.

[2] Aumann, Robert J., 'Agreeing to Disagree', *Annals of Statistics*, 4 (1976), 1236–1239.

[3] Baltag, Alexandru, Hans van Ditmarsch, and Lawrence Moss, 'Epistemic Logic and Information Update', in Pieter Adriaans, and Johan van Benthem, (eds.), *Handbook on the Philosophy of Information*, Elsevier, 2008.

[4] Baltag, Alexandru, Lawrence Moss, and Slawomir 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)*, Morgan Kaufmann Publishers, 1998, pp. 43–56.

[5] Baltag, Alexandru, and Sonja Smets, 'A Qualitative Theory of Dynamic Interactive Belief Revision', in Giacomo Bonanno, Wiebe van der Hoek, and Michael Wooldridge, (eds.), *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, Texts in Logic and Games, Vol. 3, Amsterdam University Press, 2008, pp. 9–58.

[††] Some very recent work goes towards this direction: [7] investigates the logic of oscillations, [30] defines model transformers for DEL dynamical systems with limit cycles while [21, 22, 31] and investigates their convergence and recurrence properties, and [12] proposes an epistemic logic for unadoptable behavior.

[6] van Benthem, Johan, *Logical Dynamics of Information and Interaction*, Cambridge University Press, 2011.

[7] van Benthem, Johan, 'Oscillations, Logic, and Dynamical Systems', in Sujata Ghosh, and Jakub Szymanik, (eds.), *The Facts Matter. Essays on Logic and Cognition in Honour of Rineke Verbrugge*, College Publications, 2015, pp. 9–22.

[8] Christoff, Zoé, *Dynamic Logics of Networks: Information Flow and the Spread of Opinion*, Ph.D. thesis, Institute for logic, Language and Computation, University of Amsterdam, Amsterdam, The Netherlands, 2016. ILLC Dissertation Series DS-2016-02.

[9] Christoff, Zoé, and Davide Grossi, 'Stability in binary opinion diffusion', in Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, (eds.), *Logic, Rationality, and Interaction: 6th International Workshop, LORI 2017, Sapporo, Japan, September 11-14, 2017, Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2017, pp. 166–180.

[10] Christoff, Zoé, and Jens Ulrik Hansen, 'A two-tiered formalization of social influence', in Davide Grossi, Olivier Roy, and Huaxin Huang, (eds.), *Logic, Rationality, and Interaction*, vol. 8196 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2013, pp. 68–81.

[11] Christoff, Zoé, and Jens Ulrik Hansen, 'A logic for diffusion in social networks', *Journal of Applied Logic*, 13 (2015), 1, 48–77.

[12] Christoff, Zoé, Jens Ulrik Hansen, and Carlo Proietti, 'Reflecting on social influence in networks', *Journal of Logic, Language and Information*, (2016). To appear.

[13] Davey, B. A., and H. A. Priestley, *Introduction to Lattices and Order*, second edn., Cambridge University Press, 2002.

[14] van Ditmarsch, Hans, Wiebe van der Hoek, and Barteld Kooi, *Dynamic Epistemic Logic*, Springer, 2008.

[15] Easley, David, and Jon Kleinberg, *Networks, Crowds, and Markets*, Cambridge University Press, 2010.

[16] Fitting, Melvin, Lars Thalmann, and Andrei Voronkov, 'Term-Modal Logics', *Studia Logica*, 69 (2001), 133–169.

[17] Goyal, Sanjeev, 'Interaction Structure and Social Change', *Journal of Institutional and Theoretical Economics*, 152 (1996), 3, 472–494.

[18] Granovetter, Mark, 'Threshold Models of Collective Behavior', *American Journal of Sociology*, 83 (1978), 6, 1420–1443.

[19] Hansen, Pelle G., Vincent F. Hendricks, and Rasmus K. Rendsvig, 'Infostorms', *Metaphilosophy*, 44 (2013), 3, 301–326.

[20] Kempe, David, Jon Kleinberg, and Éva 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*, ACM, New York, NY, USA, 2003, pp. 137–146.

[21] Klein, Dominik, and Rasmus K. Rendsvig, 'Convergence, Continuity and Recurrence in Dynamic Epistemic Logic', in Alexandru Baltag, and Jeremy Seligman, (eds.), *Logic, Rationality, and Interaction (LORI 2017, Sapporo)*, vol. 10455 of *LNCS*, Springer, 2017.

[22] Klein, Dominik, and Rasmus K. Rendsvig, 'Metrics for Formal Structures, with an Application to Kripke Models and their Dynamics', *arXiv:1704.00977*, (2017).

[23] Li, Fa-Hsien, Cheng-Te Li, and Man-Kwan 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*, (2011), 560–563.

[24] Liu, Fenrong, Jeremy Seligman, and Patrick Girard, 'Logical dynamics of belief change in the community', *Synthese*, 191 (2014), 11, 2403–2431.

[25] Long, Cheng, and Raymond Chi-Wing Wong, 'Minimizing Seed Set for Viral Marketing Paper', *2011 IEEE 11th International Conference on Data Mining*, (2011), 427–436.

[26] Morris, Stephen, 'Contagion', *Review of Economic Studies*, 67 (2000), 57–78.

[27] Rendsvig, Rasmus K., 'Epistemic Term-Modal Logic', in Marija Slavkovik, (ed.), *Proceedings of the 15th Student Session of The European Summer School in Logic, Language and Information*, 2010, pp. 37–46.

[28] Rendsvig, Rasmus K., 'Diffusion, Influence and Best-Response Dynamics in Networks: An Action Model Approach', in Ronald de Haan, (ed.), *Proceedings of the ESSLLI 2014 Student Session*, Tübingen, 2014, pp. 63–75.

[29] Rendsvig, Rasmus K., 'Pluralistic ignorance in the bystander effect: Informational dynamics of unresponsive witnesses in situations calling for intervention', *Synthese*, 191 (2014), 11, 2471–2498.

[30] Rendsvig, Rasmus K., 'Model Transformers for Dynamical Systems of Dynamic Epistemic Logic', in Wiebe van der Hoek, Wesley H. Holliday, and Wen-fang Wang, (eds.), *Logic, Rationality, and Interaction (LORI 2015, Taipei)*, LNCS, Springer, 2015, pp. 316–327.

[31] Rendsvig, Rasmus K., *Logical Dynamics and Dynamical Systems*, Ph.D. thesis, Theoretical Philosophy, Lund University, Lund, Sweden, 2018. Lund University Publications.

[32] Ruan, Ji, and Michael Thielscher, 'A logic for knowledge flow in social networks', in Dianhui Wang, and Mark Reynolds, (eds.), *AI 2011: Advances in Artificial Intelligence*, vol. 7106 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2011, pp. 511–520.

[33] Schelling, T. C., 'Models of segregation', *The American Economic Review*, 59 (1969), 2, 488–493.

[34] Schelling, Thomas C., *Micromotives and Macrobehavior*, Norton, New York, 1978.

[35] Seligman, Jeremy, Fenrong Liu, and Patrick Girard, 'Logic in the community', in Mohua Banerjee, and Anil Seth, (eds.), *Logic and Its Applications*, vol. 6521 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2011, pp. 178–188.

[36] Shakarian, Paulo, Sean Eyre, and Damon Paulo, 'A Scalable Heuristic for Viral Marketing Under the Tipping Model', *arXiv:1309.2963v1*, (2013).

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

[38] Valente, Thomas W., 'Social network thresholds in the diffusion of innovations', *Social Networks*, 18 (1996), 1, 69–89.

[39] Zhen, Liang, and Jeremy Seligman, 'A logical model of the dynamics of peer pressure', *Electronic Notes in Theoretical Computer Science*, 278 (2011), 0, 275 – 288.

Alexandru Baltag
Institute for Logic, Language and Computation
University of Amsterdam
Amsterdam, The Netherlands
A.Baltag@uva.nl

Zoé Christoff
Department of Philosophy
University of Bayreuth
Bayreuth, Germany
Zoe.Christoff@uni-bayreuth.de

Rasmus K. Rendsvig
Center for Information and Bubble Studies
University of Copenhagen
Copenhagen, Denmark
R.K.Rendsvig@hum.ku.dk

Sonja Smets
Institute for Logic, Language and Computation
University of Amsterdam
Amsterdam, The Netherlands
S.J.L.Smets@uva.nl