Dynamic Logics for Threshold Models and their Epistemic Extension

Zoé Christoff* and Rasmus K. Rendsvig**

Accepted for ELISIEM 2014

Abstract

We take a logical approach to threshold models*, used to study the diffusion of e.g. new technologies or behaviours in social networks. In short, threshold models consist of a network graph of agents connected by a social relationship and a threshold to adopt a possibly cascading behaviour. Agents adopt new behaviour when the proportion of their neighbours who have already adopted it meets the threshold. Under this adoption policy, threshold models develop dynamically with 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 neighbours' behaviours allows agents to anticipate changes in behaviour of their closer neighbours. It is shown that this epistemic prediction dynamics is equivalent to the non-epistemic threshold model dynamics if and only if agents know exactly their neighbours' behaviour. We further show results regarding fixed points and convergence speed, and provide a partial set of reduction laws, venues for further research, and graphical representations of the dynamics.*

1 Introduction

An individual's actions or opinions may be influenced by the actions of people around her [7]. 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, the notion of "threshold influence" as presented in e.g. [5]. Threshold influence relies on a simple imitation or conformity pressure effect: agents adopt a behaviour/product/like/fashion whenever some given threshold of the agents they are related to in their social network, their *neighbours*, have adopted it already. In this sense, investigating diffusion is investigating how agents are locally influenced and how they tend to become more similar to their neighbours. The so-called *threshold models*, first introduced by [6,13], are used precisely to represent the dynamics of diffusion under threshold influence. This type of models has received much attention in recent literature [5,8,11,16].

The paper has two main goals. The first one is to design a logic to represent the traditional view of threshold influence and to reason about diffusion phenomena in social networks. This is the topic of section 2. After recalling standard threshold models in Subsection 2.1, a dynamic logic for modeling threshold influence within social networks is introduced in Subsection 2.2. While conceptually in line with [4,12,14,15,22] in using logic to model social influence effects within networks structures, this framework differs by avoiding the use of static modalities or hybrid logic tools. In this sense, the logic introduced is "minimal": propositional logic is used to specify both network structure and agent behaviour, and a single dynamic modality is used to represent "threshold influence update". Moreover, while [4,12,14,15,22] focus on the limit thresholds of 100% (all neighbours) and non-0% (at least one neighbour), here any (uniform) adoption threshold can be used, as is standard with respect to threshold models. Subsections 2.3 and 2.4 exemplify how the logic allows to reason about clusters, cascading effects and the minimal seed problem.

The second goal of the paper is to extend threshold models and their dynamic logic with an epistemic dimension. This is done in Section 3. Subsection 3.1 introduces epistemic threshold models and their update procedure. This corresponds to a conceptual jump from a minimal modeling of influence as "blind adoption" to a more so-

* University of Amsterdam, zoe.christoff@gmail.com
** Lund University, rendsvig@gmail.com

phisticated model of it as "informed adoption". In particular, the epistemic dimension allows agents to predict the future development of the dynamics. Subsection 3.2 presents a series of results regarding such "prediction updates". In Subsection 3.3, an epistemic extension of the minimal language is presented, together with a partial set of reduction laws in the standard way of dynamic epistemic logic [2,19,20]. Section 4 concludes with a discussion of venues for future research.

2 Threshold Models and their Dynamic Logic

This section aims is to design a logic to capture the dynamic of threshold models. Subsection 2.1 first reminds the reader of the standard definition of threshold models, makes explicit some assumptions and fixes some notation.

2.1 Threshold Models for Social Influence

A social network may be seen as a graph, where the nodes represent agents and the edges represent a binary social relationship among them. This paper restricts itself to finite, undirected connected[1] graphs without self-loops, i.e., considers only symmetric, irreflexive social relationships, like being neighbours or friends.

Definition 1 (Network). *A network is a pair $(\mathscr{A}, N)$ where $\mathscr{A}$ is a finite set of* agents *and $N : \mathscr{A} \to \mathscr{P}(\mathscr{A})$ assigns a set $N(a)$ to each $a \in \mathscr{A}$, such that:*
– $a \notin N(a)$ (Irreflexivity),
– $b \in N(a)$ if and only if $a \in N(b)$ (Symmetry),
– for any $a, b \in \mathscr{A}$, there is a k such that $b \in N^k(a)$ with $N^k(a) = N^{k-1}(N(a))$ (Connectedness).

A threshold model consists of a network together with a unique behaviour B (or fashion, or product, or "like-able item") distributed over $\mathscr{A}$ and a fixed uniform adoption threshold θ. A threshold model represents the current spread of B throughout the network, while containing the adoption threshold which prescribes how the current state will evolve.

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

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

Definition 3 (Threshold model update). *The update of threshold model $\mathscr{M}_n = (\mathscr{A}, N, B_n, \theta)$ is the threshold model $\mathscr{M}_{n+1} = (\mathscr{A}, N, B_{n+1}, \theta)$ identical to $\mathscr{M}_n$, except possibly for B_{n+1}, given by*

$$B_{n+1} = B_n \cup \{i \in \mathscr{A} : \frac{|N(i) \cap B_n|}{|N(i)|} \geq \theta\}.$$

Fact. For any threshold model $\mathscr{M}$, the successive updates of $\mathscr{M}$ reaches a fixed point in finite time. I.e. for some $n \in \mathbb{N}$, $B_n = B_{n+1}$. This follows as $\mathscr{A}$ is finite and B is inflating in n.[2]

Fig. 1 illustrates how a behaviour spreads step by step up to the fixed point.

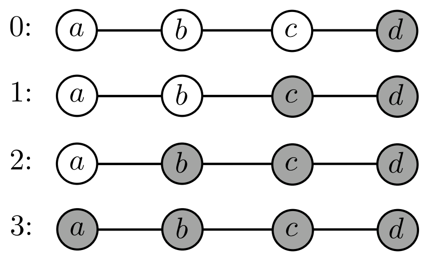

Fig. 1. Diffusion of the gray behaviour with threshold $\theta \leq \frac{1}{2}$.

Model Interpretation. Threshold models and their dynamics may be interpreted in two ways. One interpretation assumes that agents are mere automata, their behaviour forced upon them by their environment. This interpretation suits models to be used e.g. in epidemiology: viral infection "just happens". Alternatively, agents may be interpreted as being rational and aiming towards coordinating with their neighbours. In fact, the above update rule corresponds to the best response dynamics of an associated coordination game [11], with the assumption of a 'seed' set of players that always, possibly irrationally, play B [5].

[1] There exists a path between any two agents. It is not assumed that graphs are *complete*, i.e. that all agents are directly connected.

[2] A set is inflating in n if $B_n \subseteq B_{n+1}$ for all n.

Numerous variations of threshold models exist in the literature, including infinite networks [11], non-inflating behaviour adoption [11], agent-specific threshold [8], weighted links [8] and multiple behaviours [1].

2.2 Minimal Threshold Influence Logic

This section introduces a minimal logic to model the standard notion of threshold influence. To describe the *situation* of a social network at a given moment, the static language must describe two things: who is related to whom and who is displaying behaviour B. Both features are encoded using propositional variables. To describe the *change* of situation of a social network, the language includes a dynamic modality, representing how agents adopt the behaviour of their neighbours, whenever the given threshold is reached, i.e., whenever enough neighbours present it.

Definition 4 (Minimal threshold influence language $\mathscr{L}_{TI}$). *Let $\mathscr{A}$ be finite and let atoms be given by $\Phi = \{N_a b : a, b \in \mathscr{A}\} \cup \{B_a : a \in \mathscr{A}\}$. The language $\mathscr{L}_{TI}$ is then given by:*

$$\varphi := N_a b \mid B_a \mid \neg\varphi \mid \varphi \wedge \varphi \mid [adopt]\varphi$$

Truth is defined in a straightforward way.

Definition 5 (Truth clauses for $\mathscr{L}_{TI}$). *Given a model $\mathscr{M} = (\mathscr{A}, N, B, \theta)$, $N_a b, B_a \in \varphi$, and $\varphi, \psi \in \mathscr{L}_{TI}$:*
$\mathscr{M} \vDash N_a b$ *iff* $b \in N(a)$
$\mathscr{M} \vDash B_a$ *iff* $a \in B$
$\mathscr{M} \vDash \neg\varphi$ *iff* $\mathscr{M} \nvDash \varphi$
$\mathscr{M} \vDash \varphi \wedge \psi$ *iff* $\mathscr{M} \vDash \varphi$ *and* $\mathscr{M} \vDash \psi$
$\mathscr{M} \vDash [adopt]\varphi$ *iff* $\mathscr{M}' \vDash \varphi$*, where* $\mathscr{M}'$ *is the update of* $\mathscr{M}$ *cf. Def. 3.*

From this, the following proposition is straightforwardly obtained.

Proposition 1. *For any model $\mathscr{M}$, there exists an n for which*

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

for any $\varphi \in \mathscr{L}_{TI}$.[3]

Proof. The threshold dynamics introduced include only inflating behaviours, and are therefore guaranteed to reach a fixed point. I.e., for some n, $\mathscr{M}_n = \mathscr{M}_{n+1}$. Hence $\mathscr{M}_n$ and $\mathscr{M}_{n+1}$ are guaranteed to satisfy the same formulas, whereby $\mathscr{M} \vDash [adopt]^n\varphi \longleftrightarrow [adopt]^{n+1}\varphi$.

Axiomatization. As each threshold model in effect is a propositional logic model, the static logic is captured by the axioms of propositional logic, with the addition of axioms capturing the requirements on network structures:

Definition 6 (Network axioms).
$\neg N_a a$ *(Irreflexivity)*
$N_a b \longleftrightarrow N_b a$ *(Symmetry)*
$N_a b \bigvee_{G \subseteq \mathscr{A}} (\bigvee_{c \in G} N_a c \wedge \bigvee_{c' \in G} N_c c' \wedge \bigvee_{c'' \in G} N_{c'} c'' \wedge \cdots \wedge \bigvee_{c^* \in G} N_{c^*} b)$ *(Connectedness)*

As the $[adopt]$ modality does only affects the extension of behaviour B, the axiomatization of the dynamic part of the logic is straightforward:

Definition 7 (Reduction axioms).
$[adopt]N_a b \longleftrightarrow N_a b$ *(Red.Ax.N)*
$[adopt]\neg\varphi \longleftrightarrow \neg[adopt]\varphi$ *(Red.Ax.¬)*
$[adopt]\varphi \wedge \psi \longleftrightarrow [adopt]\varphi \wedge [adopt]\psi$ *(Red.Ax.∧)*
$[adopt]B_a \longleftrightarrow B_a \vee (\bigvee_{k \in \mathscr{K}} \bigwedge_{b \in k} B_b)$,
with $\mathscr{K} = \{k \subseteq \{b \in \mathscr{A} : N_a b\} : |k| \geq t \cdot \left|\{b \in \mathscr{A} : N_a b\}\right|\}$ *(Red.Ax.B)*

The reduction law (Red.Ax.B) utilizes the fact that threshold model updates are deterministic, as this entails that whatever change will occur following an update is pre-encoded before the update. The axiom states that a will have adopted B after the update just in case she either had adopted it before the update, or if she currently has a large enough proportion of neighbours that have already adopted it.

Definition 8 (L_{TI}). *The minimal threshold influence logic L_{TI} is comprised of all instantiations of propositional tautologies, the network axioms of Def. 6, the reduction axioms of Def. 7 and the derivation rule of Modus Ponens (MP).*

The logic L_{TI} is sound and complete with respect to the class of threshold models. As the proof utilizes only well-known techniques, only a sketch is presented.

Theorem 1. *For every $\varphi \in \mathscr{L}_{TI}$,*

$$\vDash \varphi \text{ iff } \vdash \varphi$$

[3] The formula is model relative. To obtain a similar validity over the class of threshold models, the language could be extended with the Kleene star.

Sketch of proof. For soundness, it is enough to note that soundness of the reduction axioms follows trivially from the given semantics. For completeness, proceed as for public announcement logic (see e.g. [20, Ch. 7]): Define a translation t allowing a reduction of all dynamic formulas to static ones, mimicking the form of the reduction axioms. Then define a complexity measure c forcing $t(\varphi)$ to be reducible to $t(\varphi')$ with $c(\varphi') < c(\varphi)$. To obtain this, it is enough to tweak the complexity of a dynamic formula to be a "big enough" function of the size of set of agents $\mathscr{A}$. Defining c of Boolean formulas as in [20, Ch. 7], and $\bigvee$ and $\bigwedge$ by grouping formulas to the left, the complexity of a dynamic formula needs to be bigger than $c(B_a \vee \bigvee_{k\in\mathscr{K}} \bigwedge_{b\in k} B_b) = 3 \cdot |\mathscr{P}(\mathscr{A}\backslash a)| + |\mathscr{A}\backslash a|$. It is therefore sufficient to set $c([adopt]\varphi) := (3 \cdot |\mathscr{P}(\mathscr{A}\backslash a)| + |\mathscr{A}|) \cdot c(\varphi)$.
Using this fact, one can show by induction on c that each φ is provably equivalent to $t(\varphi)$, and therefore, by soundness, also semantically equivalent to $t(\varphi)$. Finally, assuming that a formula φ is valid, it follows that its static translation $t(\varphi)$ is valid too. By completeness of propositional logic, $t(\varphi)$ is a theorem and since $t(\varphi) \longleftrightarrow \varphi$ is a theorem, φ is a theorem too. □

2.3 Clusters and Cascades

An agent adopting B may be influenced by some of her neighbours to adopt at the next moment, which in turn may causes further agents to follow suit. Such chain reaction is termed a *cascade* in the literature (see e.g. [5, Ch. 19]). As threshold model updates 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 threshold influence logic.

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 group members support one another in disagreeing with the opinion. Tightly connected components of a network might block the diffusion of a behaviour when it stems from outside this component. In other words, dense components of a network may prevent complete cascades and the denser a group, the better it resists change from the outside. The required precise notion of a "dense" group is that of a *d-cohesive set* [11], also referred to as a a *cluster of density d* [5]. A cluster of density d is a set of agents such that for each agent in the set, the proportion of her neighbours which are also in the group is at least d. Formally:

Definition 9 (Cluster of density d). *Given a network $(\mathscr{A}, N)$ a cluster of density d is any group $C \subseteq \mathscr{A}$ such that for all $i \in C$,*

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

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

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

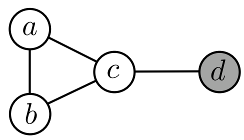

Fig. 2. A social network with a cluster of density $\frac{2}{3}$.

The minimal language can be used to express the existence of a cluster. Notice that if C is a cluster, then for each a in C, there is a big enough subgroup of C which are a's neighbours. This is expressible in $\mathscr{L}_{TI}$.

Proposition 2. *The group C is a cluster of density d in $(\mathscr{A}, N)$ iff $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ satisfies*

$$\bigwedge_{i\in C} \bigvee_{G'\subseteq C} \bigwedge_{j\in G'} N_i j \tag{1}$$

for $G' \subseteq \mathscr{A}$ such that $\frac{|G'|}{|\{j : \mathscr{M} \models N_i j\}|} \geq d$.
The proof may be found the appendix.

Given Proposition 2, it is easy to see that the following characterizes the existence of a cluster of density d in which all members have the property described by a formula $\varphi \in \mathscr{L}_{TI}$.

$$\exists C_d(\varphi) := \bigvee_{G\subseteq\mathscr{A}} \bigwedge_{i\in G} \bigvee_{G'\subseteq G} \bigwedge_{j\in G'} (N_i j \wedge \varphi)$$

for G' such that $\frac{|G'|}{|\{j:\mathscr{M}\models N_i j\}|} \geq d$.

***Example: Clusters, cont.*.** The model illustrated in Fig. 1 contains a cluster $C = \{a, b, c\}$ of density $\frac{2}{3}$, and moreover, no agents in C have adopted behaviour B. Hence, the model should satisfy

$$\exists C_{\frac{2}{3}}(\neg B_i) := \bigvee_{G\subseteq\mathscr{A}} \bigwedge_{i\in G} \bigvee_{G'\subseteq G} \bigwedge_{j\in G'} (N_i j \wedge \neg B_i). \quad (2)$$

To verify this, assume C is the group required to satisfy the outmost disjunction. Then there is a G' such that $\frac{|G'|}{|\{j:N_i j\}|} \geq \frac{2}{3}$ for which $\mathscr{M}$ satisfies

$$\bigwedge_{i\in C} \bigvee_{G'\subseteq C} \bigwedge_{j\in G'} (N_i j \wedge \neg B_i). \quad (3)$$

To see this, regard first agent c, for whom it should be satisfied that

$$\bigvee_{G'\subseteq C} \bigwedge_{j\in G'} (N_c j \wedge \neg B_c)$$

As $|\{j : \mathscr{M} \models N_c j\}| = 3$, we must identify a group $G' \subseteq C$ with $|G'| \geq 2$ such that for all $j \in G'$, $\mathscr{M} \models N_c j$. Such a G' exists, being $\{a, b\}$. Finally, indeed $\mathscr{M} \models \neg B_c$, and hence the conjunct for c is satisfied. Similar reasoning shows that the conjuncts for a and b also hold. This gives us (3), and hence (2).

The following theorem [11],[5, Ch.19.3] characterizes the possibility of a complete adoption cascade in a network:

> Given threshold θ, if the set of agents that have adopted B is A, then all agents will eventually adopt B *if, and only if*, there does not exist a cluster of density greater than $1 - \theta$ in $\mathscr{A}\backslash A$.

As the existence of clusters can be expressed in $\mathscr{L}_{TI}$, this theorem can be given the following succinct form:

Theorem 2. *Let $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ and $n \geq k$ for k such that $\mathscr{M}_k = \mathscr{M}_{k+1}$,*

$$\mathscr{M} \models [adopt]^n \bigwedge_{i\in\mathscr{A}} B_i \leftrightarrow \neg\exists C_{1-\theta}(\neg B_i).$$

2.4 The Minimal Seed Problem

In recent literature on diffusion processes, the so-called *minimal seed problem* [3] has received much attention [9,10,18,17,21]. Stated loosely, the problem requires finding a smallest possible group that when converted to the new behaviour will eventually ensure a complete cascade. It is known that the problem is NP hard [10], and the typical approach to the problem is to define greedy and effective algorithms to approximate an answer within a reasonable bound of error.

Given the completeness of the introduced logic, the minimal seed problem can be solved without error, though the procedure may be computationally costly. Given a specification of a network by a propositional formula φ, it can be checked for which smallest $C \subseteq \mathscr{A}$ it holds that

$$\varphi \vdash \bigwedge_{i\in C}(L_i\beta) \rightarrow \langle adopt\rangle^n \bigwedge_{j\in\mathscr{A}}(L_i\beta),$$

where n is the longest non-cyclic path in the network specified by φ, and hence an upper bound on the number of updates required before reaching the fixed point.

3 Epistemic Threshold Models and Their Dynamic Logic

The previous section defined a minimal logic for modeling the evolution of threshold models. In such models, it is assumed that agents only *react* to their environment: they are always influenced by their direct neighbours and only by them. They are not considering options, they are not anticipating behavioural changes, they are never influenced by anybody further away than who they are in direct contact with. We will refer to this type of influence and adoption as “blind”.

In what follows, the standard threshold models and the corresponding logic will be augmented to produce a more refined adoption policy. This section investigates what happens when agents may know more than only the present behaviour of their direct neighbours. Providing agents with more information about the current state of diffusion in turn allows agents to anticipate change of behaviour in others.

3.1 Threshold Models with Uncertainty

Consider the dynamics illustrated in Fig. 2, run according to blind adoption. If one assumes that nodes are not

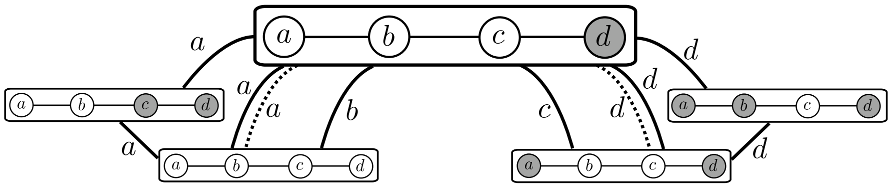

Fig. 3. An epistemic threshold model with 5 states and epistemic relations drawn for sight 1 (bold lines), sight 2 (dotted lines) and sight $k \geq 3$ (disconnected). Reflexive loops are omitted. With sight 1, agent a does not know the behaviour of c, but with sight 2, she does.

merely blindly influenced by their neighbours, but rather are rational agents seeking to coordinate, the dynamics seem to misfire. In particular, if network and behaviour distribution are known and the new behaviour is considered the most valuable, the choice of b *not* to change during the first update is irrational. As b knows that d has adopted, he knows that c will adopt during the next update, whereby b should also know that he will be better off in round 2 if he, too, has chosen to adopt. One way to represent this "predictive rationality" is to add an epistemic dimension to threshold models and define a new, predictive, update mechanism.

In non-epistemic threshold models, agents set, social network, current distribution of behaviours and threshold may be seen as being common knowledge.[4] While this is a non-neglectable simplifying assumption, it is justified by the fact that in the associated dynamics, only behaviours of an agent's direct neighbours affect her behaviour. Hence, further knowledge of the network structure itself is irrelevant.

Theoretically, agents could be uncertain about any combination of the parameters that specify a threshold model. However, this paper restricts uncertainty to a unique parameter, namely the distribution of a given behaviour among agents more or less closely related within the network structure. In other words, the radius within which agents know the behaviour of others, their "line of sight", may vary.

Uncertainty is modeled in by introducing equivalence relation to partition a set of threshold models (cf Def. 2) into epistemic alternatives. $\mathbb{M}$ denotes an epistemic threshold model with typical state $\mathscr{M}$, a threshold model.

Definition 10 (Epistemic threshold model w. sight k). *Let an epistemic threshold model with sight k be a tuple $\mathbb{M} = (dom\mathbb{M}, \{\sim_i\}_{i\in\mathscr{A}})$ where*
– $dom\mathbb{M}$ is a set of threshold models such that for all $\mathscr{M}, \mathscr{M}' \in dom\mathbb{M}$, if $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ and $\mathscr{M}' = (\mathscr{A}', N', B', \theta')$, then $\mathscr{A} = \mathscr{A}'$, $N = N'$ and $\theta = \theta'$.
– For each $i \in \mathscr{A}$, $\sim_i$ is an equivalence relation on $dom\mathbb{M}$ such that $\mathscr{M} \sim_i \mathscr{M}'$ iff $\forall j \in N^k(i) \cup \{i\} : j \in B_{\mathscr{M}} \Leftrightarrow j \in B'_{\mathscr{M}'}$, where $N^k(i)$ is the set of k-reachable neighbours of i.

The constraint on the domain of epistemic threshold models prevents uncertainty about the network structure or the adoption threshold, while the constraint on the indistinguishability relations implies that agents always know the behaviour of themselves and all agents within distance k in the network structure.[5] Fig. 3 illustrates an epistemic threshold model.

Informed dynamics. Notice that it is no problem to apply the model update procedure defined for threshold models to the epistemic extension. One only needs to 1) update the set B in each state according to the defined rule, and 2) update the indistinguishability relations in order to satisfy the definition of epistemic network model. Such an update does not utilize the potential knowledge of agents.

Definition 11 (Blind adoption update). *Let $\mathbb{M}_n$ be an ep. threshold model. The blind adoption update of $\mathbb{M}_n$ produces $\mathbb{M}_{n+1}$, identical to $\mathbb{M}_n$ in all respects except that*
– For $\mathscr{M} = (\mathscr{A}, N, B_n, \theta) \in \mathbb{M}_n$, $B_{n+1} := B_n \cup \{i \in \mathscr{A} : \frac{|N(i)\cap B_n|}{|N(i)|}\}$
– All relations $\sim_i$ are restricted to satisfy the requirement for epistemic threshold models.

[4] Alternatively, one may interpret standard threshold models as letting each agent know exactly the behaviour of their neighbours, see below.

[5] As each threshold model indirectly includes a valuation for the propositional variables, each $\mathbb{M}$ can be recast as a standard S5 Kripke model with an associated global threshold.

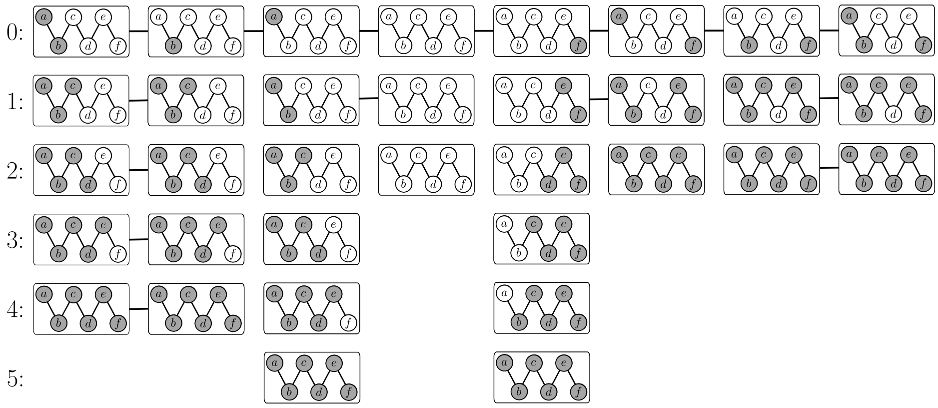

Fig. 4. The learning process for agent d (bottom center) under blind adoption with threshold $\theta \leq \frac{1}{2}$. With sight radius 1, the epistemic threshold model contains at most the 8 depicted states. The last states to reach fixed points at time 5 are states 3 and 5 from the left. Epistemic relations are drawn only for d to simplify representation. Note the development of the indistinguishability relation: as the updated $\sim'_d$ is a restriction of $\sim_d$ to states where both c and e's behaviours are identical, d learns about the distribution. Learning may or may not be complete: compare the development of states 2 and 3.

As blind adoption does not utilize the introduced uncertainty, all states in an epistemic threshold model will develop exactly as they would have as isolated threshold models. Learning in a blind adoption scenario is illustrated in Fig. 4.

For agents to utilize their knowledge, we introduce a refined update procedure, by which agents use the information available about their surrounding agents' behaviour. In particular, this additional information allows agents to partially anticipate the system's development. The definition recursively defines higher orders of prediction, taking as base blind adoption. Hence, a 1st level predictor assumes all others are blind adopters, a 2nd level predictor assumes all others are 1st level, etc. Different prediction levels produce different conjectures regarding the next state, as is exemplified below.

Definition 12 (k-level prediction update). *Let $\mathbb{M}_n$ be given and let B_n be from $\mathscr{M} \in dom\mathbb{M}$. The k-level prediction update of $\mathbb{M}_n$ produces $\mathbb{M}_{n+1}$, identical to $\mathbb{M}_n$ in all respects except that*

– The k-level prediction update of B_n is given by

$$B^k_{n+1} := B_n \cup \{a \in \mathscr{A} : \frac{|N(a) \cap K_a B^{k-1}_{n+1}|}{|N(a)|} \geq \theta\}$$

where $K_a B^{k-1}_{n+1}$ is the set of agents such that a knows that if these agents updated in accordance with $k-1$ level prediction update, then they will adopt in the next round. The set is given by

$$K_a B^{k-1}_{n+1} := \{j \in \mathscr{A} : \forall \mathscr{M}' \sim_a \mathscr{M}, j \in B^{k-1}_{n+1}\}$$

with B^0_{n+1} the behaviour set obtained if blind adopt update is applied to $\mathscr{M}$.

– All relations $\sim_i$ are restricted to satisfy the requirement for epistemic threshold models.

Fig. 5 gives an example of k-level prediction dynamics.

Two important features about prediction update are stated in the following proposition. Point 1 captures that agents possess knowledge of a deterministic dynamics: if they are able to make predictions, then these predictions are guaranteed correct. Point 2 captures that agents may still be surprised by future developments: if agents in $N^k(a)$ choose to adopt due to influence from agents in $\mathscr{A} \backslash N^k(a)$, a will not possess the required information to predict their behavioural change.

Proposition 3. *Let $\mathbb{M}$ be a k-sight epistemic threshold model with actual world $\mathscr{M} = (\mathscr{A}, N, B, \theta)$. Then*

1. Predictions are correct: $K_a B^m_n \subseteq B^m_n$ for all $a \in \mathscr{A}$, all $m, n \in \mathbb{Z}^+$.

2. Predictions are not necessarily locally complete: possibly, $B^m_n \cap N^k(a) \not\subseteq K_a B^m_n$.

Proof. **1.:** If $j \in K_a B^m_n$, then $j \in B'^m_n$ for all $\mathscr{M}' \sim_a \mathscr{M}$. Hence $j \in B^m_n$. **2.:** For an example, see Fig. 4 above.

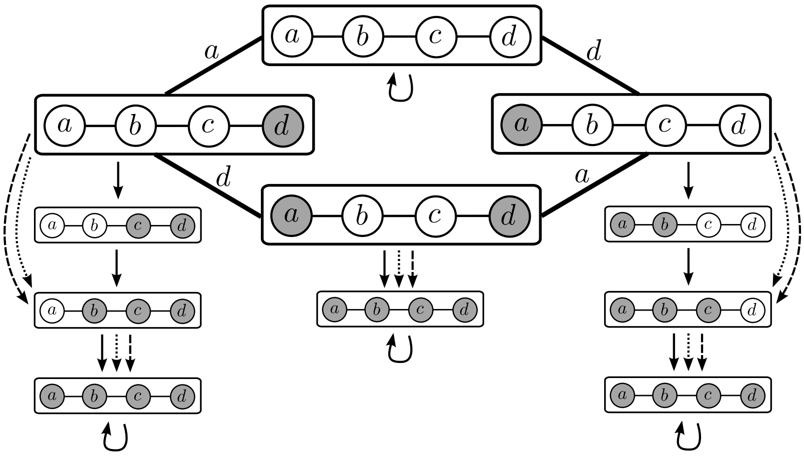

Fig. 5. Prediction dynamics for prediction levels $k \in \{0, 1, 2\}$, with threshold $\theta \leq \frac{1}{2}$ and sight 2. Full arrows show transitions for level 0, dotted arrows level 1 and dashed arrows level 2. Loops indicate fixed points. Bold lines represent indistinguishability. All agents learn the distribution after one update. The same fixed point is reached regardless of prediction level, but is reached faster. Prediction levels 1 and 2 are step-wise identical, cf. Thm. 4.

3.2 Results: Properties of Prediction Update

One immediate worry regarding the combination of extended sight and prediction dynamics may be whether these in fact correspond to blind adoption in a more connected network. However, as the following proposition shows, this worry is unwarranted.

Proposition 4. *Prediction dynamics with sight k is not equivalent with blind adoption dynamics with all k-reachable neighbours set to also be direct neighbours.*

Proof. The development of the left-most state of Fig. 5 serves as an example. For $k = 2$: if a was direct neighbours with c, then c would not be sufficiently influenced by d to adopt. □

The following Lemma is used in the proofs of the below theorems. It shows two important features of prediction dynamics. First, increasing agents' predictive power either leaves the dynamics unchanged or reduces the number of updates required before meeting the fixed point. This was illustrated in Fig. 5. Secondly, it shows that an increase in prediction power only makes agents more informed.

Lemma 1. *For any n, the following hold if $m \geq k$:*
1. Increased prediction does not slow dynamics: $B^k_n \subseteq B^m_n$.
2. Knowledge is inflating in prediction level: $K_b B^k_n \subseteq K_b B^m_n$.

The proof of Lemma 1 and the following propositions may be found in the appendix.

Our first main result pertaining to prediction dynamics show that providing agents with additional information about the deterministic dynamics does not change the outcome. A corollary to this theorem is that the cluster theorem also applies to the epistemic dynamics.

Theorem 3. *All prediction dynamics are fixed point equivalent to blind adoption dynamics. Specifically, for all k, if $B^k_n = B^k_{n+1}$ and $B^0_m = B^0_{m+1}$, then $B^k_n = B^0_m$.*

Our second main result relates the range of sight with the capabilities acquired by higher levels of prediction. It shows that dynamics run using levels of prediction greater or equal to the agents' line of sight will all be equivalent. Not only will they reach the same fixed point, but will also be step-wise equivalent. In other word, extra anticipation power will be superfluous. See Fig. 5 for an example.

Theorem 4. *Prediction is limited by sight. If $\sim_i$ is defined using $N^k(i)$, then if $m \geq k-1$, then $B^m_n = B^{k-1}_n$, for all n.*

As a corollary to Theorem 4, we obtain a relation between between blind dynamics and prediction dynamics. Specifically, prediction dynamics are step-wise equivalent to blind dynamics just in case the line of sight is set to 1.

3.3 Epistemic Threshold Influence Logic

As in the previous section, after introducing the epistemic version of our threshold models, a logic to reason about how those models change.

The minimal language $\mathscr{L}_{TI}$ is extended by adding the standard K_a modality reading "agent a knows that" and replacing the dynamic modality $[adopt]$ by dynamic modalities $[adopt^k]$, one for each prediction level k. As a corollary to Theorem 4 , we see that the relevance of prediction levels is limited by the smallest k such that all agents are k-reachable from each other. This k will always be less than $|\mathscr{A}|$. This means that given $\mathscr{A}$, the language can be restricted to $[adopt^k]$-operators for $k < |\mathscr{A}|$.

Definition 13 (Epistemic threshold influence language $\mathscr{L}_{KTI}$). *Let $\mathscr{A}$ be a finite a set of agents and $k < |\mathscr{A}|$. Let the set of atomic propositions be given by $\varphi = \{N_a b : a, b \in \mathscr{A}\} \cup \{B_a : a \in \mathscr{A}\}$. The epistemic threshold influence language ($\mathscr{L}_{KTI}$) is defined as follows:*

$$\varphi := N_a b \mid B_a \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi \mid [adopt^k]\varphi$$

Definition 14 (Truth clauses for $\mathscr{L}_{KTI}$). *Given an epistemic threshold model $\mathbb{M} = (dom\mathbb{M}, \sim_i)_{i\in\mathscr{A}}$, a threshold model $\mathscr{M} = (\mathscr{A}, N, B, \theta) \in dom\mathbb{M}$, and formulas $N_a b, B_a \in \varphi$, and $\varphi, \psi \in L_{KTI}$:*
$\mathbb{M}, \mathscr{M} \models N_a b$ *iff* $b \in N(a)$
$\mathbb{M}, \mathscr{M} \models B_a$ *iff* $a \in B$
$\mathbb{M}, \mathscr{M} \models K_a\varphi$ *iff for all* $\mathscr{M}' \sim_a \mathscr{M}$, $\mathbb{M}, \mathscr{M}' \models \varphi$
$\mathbb{M}, \mathscr{M} \models [adopt^k]\varphi$ *iff* $\mathbb{M}', \mathscr{M}' \models \varphi$, *where* $\mathbb{M}'$ *is the k-level prediction update of* $\mathbb{M}$ *with* $\mathscr{M}'$ *identical to* $\mathscr{M}$ *except with updated* B, *cf. Def. 12.*

Towards an Axiomatization. Where the reduction laws for the non-epistemic logic were simple, the interplay between the $[adopt^k]$ and K_i modalities are far more subtle.

Before reaching a fixed point, any prediction update performs two actions simultaneously. Firstly, it induces a factual change in a subset of the propositional variables by changing the truth values of some B_i from false to true. Secondly, it enforces an epistemic update, whereby agents learn the new behaviour of agents within their range of sight. The combination of these two roles with the in-built asymmetry that B_i may only turn from false to true makes constructing a complete set of reduction laws more complicated than in the case of blind adoption. The task is as of yet unfinished, but below we present reduction laws for the essential cases.

Claim 1 (Reductive validities). *The following formulas are valid over the class of epistemic threshold models.*
1. $[adopt^k]B_a \leftrightarrow B_a \vee (\neg B_a \rightarrow K_a[adopt^{k-1}](\bigvee_{k\in\mathscr{K}} \bigwedge_{b\in k} B_b))$,
$\mathscr{K} = \{k \subseteq \{b \in \mathscr{A} : N_a b\} : |k| \geq t \cdot \big|\{b \in \mathscr{A} : N_a b\}\big|\}$.
2. $[adopt^0]K_a B_b \leftrightarrow K_a B_a \vee (N_a b \wedge \neg B_b \wedge [adopt^0]B_b) \vee (\bigvee_{j:N_a^k j} N_j a \wedge \neg B_j \wedge K_a([adopt^0]B_j \rightarrow B_b))$
3. $[adopt^0]K_a\neg B_b \leftrightarrow (N_a b \wedge \neg B_b \wedge [adopt^0]\neg B_b) \vee (\bigvee_{j:N_a^k j} N_j a \wedge \neg B_j \wedge K_a([adopt^0]\neg B_j \rightarrow \neg B_b))$

The first formula reflects the recursive nature of prediction dynamics, in the sense that level k prediction is reducible to knowledge of the behaviour of level $k-1$ predictors. Using this law, formulas containing $[adopt^k]$ modalities may be reduced to ones containing only $[adopt^{k-1}]$ modalities. As $[adopt^0]$ is independent of further prediction levels, the reduction axioms from L_{TI} may be used to reduce it to a propositional formula.

The last two formulas, which allow pushing the $[adopt^0]$ modality into the scope of K_a operators, illustrate the problem of obtaining general laws. Such laws must be sensitive to negations of propositional variables under the scope of knowledge operators, as knowledge of $\neg B_b$ may be lost, whereas knowledge of B_b is permanent. We trust that a complete set of reduction laws is obtainable, but that its construction will require a vast amount of clauses to manage syntactical decomposition.

4 Further Research

A built-in dynamic asymmetry. One particular feature of threshold models is the strong assumption that once agents adopt behaviour B, they can not unadopt it. This essential asymmetry comes attached with both nice dynamic properties and with formal difficulties. On the one hand, since the set of agents adopting the behaviour is inflating, the dynamics are guaranteed to stabilize.[6] Consequently, the set of formulas needed to describe the dynamics of a given threshold models is finite. On the other hand, the lack of unadoption creates an asymmetry between B_i and $\neg B_i$, as the truth value of B_i can only change from false to true, while the one of $\neg B_i$ can only

[6] Allowing unadoption might lead to looping situations which arise e.g. in the settings of [15,4] in which agents may find themselves in a situation where they switch forever between two options.

change from true to false This is reflected e.g. in the difference between the reduction rules for $[adopt^k]K_aB_b$ and $[adopt^k]K_a\neg B_b$, and results in some difficulty for getting a generalized reduction rule for negation of arbitrary φ. While the search for a complete axiomatization of L_{KTI} is left for future research, this difficulty is an interesting reflection of the built-in asymmetry of threshold dynamics. The difficulty could be avoided by adopting a symmetrical dynamics, e.g. by changing the update of behaviours from B to B':

$$B_{n+1} = B_n \cup \{i : \tfrac{|N(i)\cap B|}{|N(i)|}\} \qquad B'_{n+1} = \{i : \tfrac{|N(i)\cap B|}{|N(i)|}\}.$$

Agent uniformity assumptions. As mentioned in Section 2, there exists many variations of threshold models, including weighted edges and agent-specific thresholds. As such, the present paper only deals with a special case of threshold models, and falls neatly between two existing approaches taken in the logic literature to date. On the one hand, it is more general threshold-wise than the works [15,14,22,4], since it allows for any uniform proportional threshold, but is threshold-wise a restriction of the logic proposed by [12],where thresholds are agent-relative. The task of logically representing weighted edges remains untackled, and is an obvious candidate for future research.

Another uniformity assumption which was made concerns agents' knowledge. Not only was it assumed that the network structure and the threshold are common knowledge, it was also assumed that the minimal knowledge that agents have about their k-distant network neighbour is uniform, i.e., all agents know (at least) the behaviour of all agents which are at distance k or less within the network structure. This assumption creates another asymmetry: the uniformity concerns only the minimum amount of knowledge. Some agents might know more, and hence the total of what they know is still agent-relative. Again, this asymmetry results in some difficulty in axiomatization and on could think about the benefit of adopting one of the two following policies instead: either making the levels of knowledge fully uniform among agents, forcing them to know all and only the behaviour of all agents at distance at most k, or on the contrary, dropping the uniformity assumption on the minimal level of knowledge, allowing some agents to know the behaviour of only some of the k-reachable agents. These investigations are also left for further research.

Followers or trendsetters? So far, agents are only reacting to their environment, they are pure "followers", they do not make any reasoned choice. Because of this, the notion of goal-oriented action, necessary to talk about whether agents are rational or not, seems to be missing in the representation given here. This is troubling since the initial threshold models are naturally representing some coordination game. What disappeared in this logical treatment is the notion of payoff, and therefore the notion of "good" or "bad" outcome for the agents. But another type of change could be defined in a way entirely similar to what has been done in this paper, coming closer to goal-orientedness: adopt something when you know that if you do, your neighbour will adopt it in the next round. In this way, agents are not followers any more but "trendsetters" On the other hand, even the "follower" type of agents can be seen as goal-oriented agents, assuming that a certain level of conformity with their neighbours is their goal.

References

1. K. R. Apt and E. Markakis. Diffusion in Social Networks with Competing Products. In G. Persiano, editor, *SAGT 2011*, LNCS 6982, pages 212–223. Springer, 2011.
2. A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. Technical report, CWI, 1999.
3. S. P. Borgatti. Identifying sets of key players in a social network. *Computational and Mathematical Organization Theory*, 12(1):21–34, Apr. 2006.
4. Z. Christoff and J. U. Hansen. A two-tiered formalization of social influence. In D. Grossi, O. Roy, and H. Huang, editors, *Logic, Rationality, and Interaction*, volume 8196 of *Lecture Notes in Computer Science*, pages 68–81. Springer Berlin Heidelberg, 2013.
5. D. Easley and J. Kleinberg. *Networks, Crowds, and Markets*. Cambridge University Press, 2010.
6. M. Granovetter. Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978.
7. P. G. Hansen, V. F. Hendricks, and R. K. Rendsvig. Infostorms. *Metaphilosophy*, 44(3):301–326, 2013.
8. D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *KDD 03:*

Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, pages 137–146. ACM, New York, NY, USA, 2003.

9. F.-H. Li, C.-T. Li, and M.-K. Shan. Labeled Influence Maximization in Social Networks for Target Marketing. *2011 IEEE Third Int'l Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third Int'l Conference on Social Computing*, pages 560–563, Oct. 2011.
10. C. Long and R. C.-W. Wong. Minimizing Seed Set for Viral Marketing Paper. *2011 IEEE 11th International Conference on Data Mining*, pages 427–436, Dec. 2011.
11. S. Morris. Contagion. *Review of Economic Studies*, 67:57–78, 2000.
12. J. Ruan and M. Thielscher. A logic for knowledge flow in social networks. In D. Wang and M. Reynolds, editors, *AI 2011: Advances in Artificial Intelligence*, volume 7106 of *Lecture Notes in Computer Science*, pages 511–520. Springer Berlin Heidelberg, 2011.
13. T. C. Schelling. *Micromotives and Macrobehavior*. Norton, New York, 1978.
14. J. Seligman, P. Girard, and F. Liu. Logical dynamics of belief change in the community, 2013.
15. J. Seligman, F. Liu, and P. Girard. Logic in the community. In M. Banerjee and A. Seth, editors, *Logic and Its Applications*, volume 6521 of *Lecture Notes in Computer Science*, pages 178–188. Springer Berlin Heidelberg, 2011.
16. P. Shakarian, S. Eyre, and D. Paulo. A Scalable Heuristic for Viral Marketing Under the Tipping Model. 2013.
17. P. Shakarian, L. Gerdes, and H. Lei. Circle-Based Tipping Cascades in Social Networks. In *WSDM 2014 Workshop on Diffusion Networks and Cascade Analytics*, 2014.
18. P. Shakarian and D. Paulo. Large social networks can be targeted for viral marketing with small seed sets. Technical report, 2013.
19. J. van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011.
20. H. van Ditmarsch, B. Kooi, and W. van der Hoek. *Dynamic Epistemic Logic*. Springer, 2006.
21. P. Zhang, W. Chen, X. Sun, Y. Wang, and J. Zhang. Minimizing Seed Set Selection with Probabilistic Coverage Guarantee in a Social Network. *arXiv:1402.5516v1*, 2014.
22. L. Zhen and J. Seligman. A logical model of the dynamics of peer pressure. *Electronic Notes in Theoretical Computer Science*, 278(0):275 – 288, 2011.

Appendix: Proofs

As a notational shorthand, set $|C/b| := \frac{|N(b)\cap C|}{|N(b)|}$ for any set C and agent b.

Proposition 2. *The group C is a cluster of density d in $(\mathscr{A}, N)$ iff $\mathscr{M} = (\mathscr{A}, N, B, \theta)$ satisfies*

$$\bigwedge_{i\in C}\bigvee_{G'\subseteq C}\bigwedge_{j\in G'} N_i j \tag{4}$$

for $G' \subseteq \mathscr{A}$ such that $\frac{|G'|}{|\{j:\mathscr{M}\models N_i j\}|} \geq d$.

Proof. Left to right: Assume C is a cluster of density d in $(\mathscr{A}, N)$. Then by definition, for all $i \in C$, $\frac{|N(i)\cap C|}{|N(i)|} \geq d$. As $\mathscr{M}$ is based on $(\mathscr{A}, N)$, $\{j : \mathscr{M} \models N_i j\} = N(i)$. Hence for $G' = N(i) \cap C$, $\frac{|G'|}{|j:\mathscr{M}\models N_i j|} \geq d$ and $\mathscr{M} \models \bigwedge_{j\in G'} N_i j$. So $\mathscr{M}$ satisfies (4).

Right to left: Assume that $\mathscr{M}$ satisfies (4) for some $C \subseteq \mathscr{A}$ and some $d \in [0, 1]$. Then for each $i \in C$, there is a $G' \subseteq C$ with $\frac{|G'|}{|\{j:\mathscr{M}\models N_i j\}|} \geq d$, such that for all $j \in G'$, $\mathscr{M} \models N_i j$. As $\{j : \mathscr{M} \models N_i j\} = N(i)$, it follows that $\frac{|G'|}{|N(i)|} = \frac{|G'|}{|\{j:\mathscr{M}\models N_i j\}|} \geq d$. Moreover, $G' \subseteq N(i)$. As also $G' \subseteq C$, it follows that $G' \subseteq N(i) \cap C$. Hence $\frac{|N_{\mathscr{M}}(i)\cap C|}{|N_{\mathscr{M}}(i)|} \geq \frac{|G'|}{|N_{\mathscr{M}}(i)|} \geq d$. As i was arbitrary from C, C is indeed a cluster of density d in $(\mathscr{A}, N)$.

Lemma 1. *For any n, the following properties hold if $m \geq k$:*
1. Increased prediction does not slow down dynamics: $B^k_n \subseteq B^m_n$.
2. Knowledge is inflating in prediction level: $K_b B^k_n \subseteq K_b B^m_n$

Proof. Let $\mathbb{M}$ be an epistemic threshold model with actual state $\mathscr{M} = (\mathscr{A}, N, B, \theta)$.

1.: We show by induction that if $b \in B^k_n$, then $b \in B^{k+1}_n$. **Base:** Assume $b \in B^0_n$. Then $|B^0_{n-1}/b| \geq \theta$. As $N(b) \cap B^0_{n-1} \subseteq K_a B^0_n$, it follows that $|K_b B^0_n/b| \geq \theta$. Hence $b \in B^1_n$. **Step:** Assume as induction hypothesis that for all $l < k$, $B^l_n \subseteq B^{l+1}_n$. Assume that $b \in B^k_n$. Then $|K_b B^{k-1}_n/b| \geq \theta$. By the induction hypothesis, $B^{k-1}_n \subseteq B^k_n$, so $K_b B^{k-1}_n \subseteq K_b B^k_n$. Hence $|K_b B^k_n/b| \geq \theta$, so $b \in B^{k+1}_n$.

2.: Assume $j \in K_b B^k_n$ in $\mathscr{M}$. Then by definition, for all $\mathscr{M}' \sim_b \mathscr{M}$, $j \in B'^k_n \cap N(b)$. So $j \in B'^k_n$. By 1. of this lemma, $B'^k_n \subseteq B'^m_n$. As this holds for all $\mathscr{M}' \sim_b \mathscr{M}$, by definition $j \in K_b B^m_n$ in $\mathscr{M}$.

Theorem 3. *All prediction dynamics are fixed point equivalent to blind adoption dynamics. Specifically, for all k, if $B^k_n = B^k_{n+1}$ and $B^0_m = B^0_{m+1}$, then $B^k_n = B^0_m$.*

Proof. Let $\mathbb{M}$ be an epistemic threshold model with actual world $\mathscr{M} = (\mathscr{A}, N, B, \theta)$. The inclusion $B^0_m \subseteq B^k_n$ follows from Lemma 1, point 1. To show that $B^k_n \subseteq B^0_m$, we decompose B^k_n and show that if $b \in B^k_n$, then $b \in B^0_l$ for some l. As B^0 is inflating, this shows that b will be in the fixed point, B^0_m.

Assume that $b \in B^k_n$ with n the fixed point of B^k at $\mathscr{M}$. We distinguish two cases for why b would belong to B^k_n:

1. b is a new member of B^k, i.e. $b \notin B^k_{n-1}$.
2. b is an old member of B^k, i.e. $b \in B^k_{n-1}$.

Case 1: We argue that since b entered B_n^k, he would also have entered B_l^0, for some l. Assume $b \in B_n^k$. Then $|K_b B_n^{k-1}/b| \geq \theta$. There are two sub-cases: either $k-1=0$ or $k-1>0$.

$k-1=0$: If $k-1=0$, then $B_n^{k-1}=B_n^0$, so $K_b B_n^{k-1} \subseteq B_n^0$. Hence $|B_n^0/b| \geq |K_b B_n^{k-1}/b| \geq \theta$. Hence $b \in B_n^0$.

$k-1>0$: We argue that this case effectively reduces to the former. If $k-1>0$, then $|K_b B_n^{k-1}/b| \geq \theta$, for which reason $K_b B_n^{k-1}$ must be non-empty. Let $c \in K_b B_n^{k-1}$. Then $c \in B_n^{k-1}$. Hence $|K_c B_n^{k-2}/c| \geq \theta$. In this case, the two sub-cases reappear: either $k-2=0$ or $k-2>0$. If the former, then $c \in B_n^0$ by the argument for $k-1=0$. As c was arbitrary among the neighbours of b that caused b's membership in B_n^k, it follows that $b \in B_{n+1}^0$. If $k-2>0$, the present argument may be reapplied, reducing the case to $k-3$. Continuing this reduction will lead to $k-k=0$. Then apply the argument for $k-1=0$ and conclude that $b \in B_l^0$ for some l.

Case 2: There are again two sub-cases: either $n-1=0$ or $n-1>0$. for the latter, find the least l such that $b \in B_l^k$, but $b \notin B_{l+1}^k$ and apply the argument for case 1. For the former, if $n-1=0$, then b must have been a seed, i.e. $b \in B_0^0$.

Theorem 4. *Prediction is limited by sight. If $\sim_i$ is defined using $N^k(i)$, then if $m \geq k-1$, then $B_n^m = B_n^{k-1}$, for all n.*

Proof. The case for $B_n^{k-1} \subseteq B_n^m$ follows from Lemma 1, point 1. For the inclusion $B_n^m \subseteq B_n^{k-1}$, then this follows if we can show that $K_b B_n^{k-1} = K_b B_n^m$ for arbitrary b.[7] From Lemma 1, point 2, we have that $K_b B_n^{k-1} \subseteq K_b B_n^m$. We now argue that given $\sim_b$ is limited by $N^k(b)$, then $K_b B_n^{k-1}$ is the largest set of agents b may know anything about.

Assume j is reachable from b in at least k steps, so that j is one the boarder of $N^k(b)$. Then j is in $K_b B_n^{k-1}$ iff

1. j is already in B_{n-1}, or
2. j has so few neighbours outside $N^k(b)$ that if $N^k(b) \backslash N(j) \subseteq B_{n-1}$, then $j \in B_n$. I.e., j would enter B_n by blind adoption alone.

Each point clearly implies that $j \in K_b B_n^{k-1}$. To see that no other options exist, notice that it makes no difference what assumption b makes about j's prediction level: If j would not adopt by blind adoption, then b has insufficient information about $N^k(j)$ to determine j's future actions. Hence $j \notin K_b B_n^{k-1}$.

A similar argument can be made for j', reachable from b in at least $k-2$ steps: it makes no difference what assumptions b makes about j''s prediction level, as long as b assumes j' is of level 1 or higher. For j'' reachable in at least $k-3$ steps, assumptions above level 2 makes no difference, etc., until we reach the conclusion that no assumption above level $k-1$ makes a difference in predicting the behaviour of agents in $N^k(b)$. Hence, if $j \in K_b B_n^m$, then this is caused by information available to b already at prediction level $k-1$. Hence $j \in K_b B_n^{k-1}$.

[7] Because if $K_b B_n^{k-1} = K_b B_n^m$, then $\frac{|N(b) \cap K_b B_n^{k-1}|}{|N(b)|} = \frac{|N(b) \cap K_b B_n^m|}{|N(b)|}$. Hence if the latter is above θ, so is the former. Hence membership of B_n^m implies membership of B_n^{k-1}.