Distributed Knowledge

Floris Roelofsen

Institute for Logic, Language, and Computation
University of Amsterdam, Netherlands

f.roelofsen@uva.nl

ABSTRACT. *This paper provides a complete characterization of epistemic models in which distributed knowledge complies with the principle of full communication [HOE 99, GER 99]. It also introduces an extended notion of bisimulation and corresponding model comparison games that match the expressive power of distributed knowledge operators.*

KEYWORDS: *epistemic logic, distributed knowledge, full communication, expressive power.*

1. Introduction

Distributed knowledge is a standard notion in epistemic logic [FAG 95, MEY 95]. Intuitively, a formula φ is distributed knowledge among a group of agents B iff φ follows from the knowledge of all individual agents in B put together. Semantically, φ is distributed knowledge among B iff φ is true in all worlds that *every* agent in B considers possible. This paper addresses two issues concerning distributed knowledge.

Full communication. Van der Hoek, van Linder, and Meyer [HOE 99] argued that, to be of any use at all, a notion of group knowledge should comply with what they call the *principle of full communication*: whenever φ is considered group knowledge, it should be possible for the members of the group to establish φ through communication (this will be made more precise below). Van der Hoek, van Linder, and Meyer [HOE 99] and Gerbrandy [GER 99] showed that distributed knowledge does not generally comply with the principle of full communication, but does in certain special model classes. It is not clear, however, *why* distributed knowledge does not generally comply with the principle of full communication, and why it does in these special model classes. Moreover, it is not known whether these model classes are *complete*, that is, whether they comprise all models in which distributed knowledge complies with the principle of full communication. We will provide a simple analysis of the problem and a complete characterization of the class of models in which distributed knowledge complies with the principle of full communication.

Expressive power. A standard notion of structural equivalence between epistemic models is that of *bisimilarity*. This notion (to be defined below) perfectly matches the expressive power of basic epistemic formulas (formulas without distributed knowledge operators): if two models are bisimilar, then they satisfy exactly the same basic formulas. But adding distributed knowledge to the basic language yields a more expressive language, whose formulas may be able to distinguish bisimilar models. Is there a natural extended notion of bisimulation that matches the expressive power of the language with distributed knowledge? This question is the first of a list of open problems in a recent survey by van Benthem [BEN 05]. We will define and analyze a suitable extended notion of bisimulation, corresponding model comparison games, and a closely related extended notion of modal saturation.

The paper is organized as follows. Section 2 reviews some basic notions from epistemic logic. Section 3 is concerned with the extent to which distributed knowledge complies with the principle of full communication, and section 4 introduces notions of bisimulation and saturation, as well as related model comparison games, to capture the expressive power of distributed knowledge operators. Sections 3 and 4 each conclude with a short summary and pointers to related work.

2. Epistemic Logic

The following notions are all standard in epistemic logic [FAG 95, MEY 95]. A countable set of proposition letters $\mathcal{P}$ and a finite set of agents $\mathcal{A}$ is assumed to be given throughout our general discussion and clear from the context in particular examples.

Languages. The basic epistemic language consists of all formulas that can be built from proposition letters in $\mathcal{P}$ using conjunction, negation, and a modal operator K_a for every agent $a \in \mathcal{A}$. $K_a\varphi$ stands for *agent a knows that φ is true*. The basic epistemic language is denoted by $\mathcal{L}_K$:

$$\varphi \quad ::= \quad p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a\varphi$$

One standard way to extend the basic language is to add a modal operator D_B for every group of agents $B \subseteq \mathcal{A}$. $D_B\varphi$ stands for *φ is distributed knowledge among B*. The resulting language is denoted by $\mathcal{L}_D$:

$$\varphi \quad ::= \quad p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a\varphi \mid D_B\varphi$$

Models. A model M is a triple (W, R, V), where:

– W is a non-empty set of worlds,

– $R : \mathcal{A} \rightarrow \wp(W \times W)$

– $V : W \rightarrow \wp(\mathcal{P})$

R assigns to every agent $a \in \mathcal{A}$ a so-called *accessibility relation* on W. Intuitively, $(w, v) \in R(a)$ means that in world w, agent a considers world v possible. Accessibility relations are often assumed to be equivalence relations, or to have other less restrictive properties, but for sake of generality we do not commit ourselves to any such specific assumptions here. V associates every world $w \in W$ with a subset of $\mathcal{P}$, the proposition letters that are true in w. If $M = (W, R, V)$ is a model and w is a particular world in W, then (M, w) is called a pointed model, and w is called its actual world. We will often simply refer to pointed models as models. The *information state* $[M, w]_a$ of an agent a in (M, w) is the set of worlds that a considers possible in (M, w). Similarly, the information state $[M, w]_B$ of a group of agents B in (M, w) is the set of worlds that every agent $a \in B$ considers possible in (M, w):

$$\begin{array}{lcl} [M, w]_a & = & \{v \in W \mid (w, v) \in R_a\} \\ [M, w]_B & = & \{v \in W \mid (w, v) \in R_a \text{ for all } a \in B\} \end{array}$$

Semantics. The satisfaction relation $\vDash$ between pointed models and formulas in $\mathcal{L}_K$ or $\mathcal{L}_D$ is recursively defined as follows:

$$\begin{array}{lll} M, w \vDash p & \text{iff} & p \in V(w) \\ M, w \vDash \neg\varphi & \text{iff} & M, w \nvDash \varphi \\ M, w \vDash \varphi \wedge \psi & \text{iff} & M, w \vDash \varphi \text{ and } M, w \vDash \psi \\ M, w \vDash K_a\varphi & \text{iff} & M, v \vDash \varphi \text{ for all } v \in [M, w]_a \\ M, w \vDash D_B\varphi & \text{iff} & M, v \vDash \varphi \text{ for all } v \in [M, w]_B \end{array}$$

Intuitively, the $K_a\varphi$ clause says that an agent knows φ to be true just in case φ is true in all worlds she considers possible. Similarly, the $D_B\varphi$ clause says that φ is distributed knowledge among B just in case φ is true in all worlds that *every* agent in B considers possible. A set of formulas Φ *entails* a formula φ, $\Phi \Vdash \varphi$, iff every pointed model that satisfies all formulas in Φ also satisfies φ. A set of formulas Φ is *consistent* or *satisfiable* iff there is a pointed model that satisfies all formulas in Φ. One set of formula Φ is consistent with another set of formulas Σ iff $\Phi \cup \Sigma$ is consistent. A set of formulas Φ is satisfiable in an information state iff that information state contains a world that satisfies all formulas in Φ. The *theory* of a world in a model is the set of all formulas true in that world. A world is consistent with a set of formulas Φ iff its theory is consistent with Φ.

3. Full Communication

Van der Hoek, van Linder, and Meyer [HOE 99] argue that, to be of any use at all, a notion of group knowledge should comply with what they call the *principle of full communication*: whenever φ is considered group knowledge, it should be possible for the members of the group to establish φ through communication. To make this more precise, they define the *knowledge set* of a group of agents B in a model (M, w) to be the set of all $\mathcal{L}_K$-formulas that at least one agent in B knows to be true in (M, w):

$$\mathrm{KS}_B(M, w) = \{\psi \in \mathcal{L}_K \mid M, w \vDash K_a\psi \text{ for some } a \in B\}$$

Then they take it that a formula can be established through communication by a group of agents iff that formula is entailed by the knowledge set of that group. So distributed knowledge complies with the principle of full communication iff for all $\varphi \in \mathcal{L}_K$:

$$M, w \vDash D_B\varphi \quad \Rightarrow \quad \mathrm{KS}_B(M, w) \Vdash \varphi \tag{1}$$

Van der Hoek, van Linder, and Meyer [HOE 99] as well as Gerbrandy [GER 99] show that (1) does *not* generally hold. To see this, consider the model depicted below (all accessibility relations are equivalence relations here). Take $B = \{a, b\}$. Then, p is distributed knowledge among B in w, but p is not entailed by B's knowledge set in w.

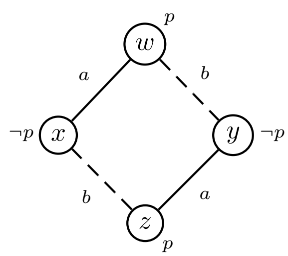

In [HOE 99] and [GER 99] various classes of models are identified in which distributed knowledge *does* comply with the principle of full communication. Van der Hoek, van Linder, and Meyer define a model M to be *distinguishing* iff for every world w in M there is an $\mathcal{L}_K$-formula that is true in w and no-where else in M. They show that distributed knowledge complies with the principle of full communication in finite, distinguishing models. Gerbrandy generalizes this result. He defines M to be *locally distinguishing* iff for every w in M, every $v \in \bigcup_{a \in \mathcal{A}} [M, w]_a$, and every $a \in \mathcal{A}$, there is an $\mathcal{L}_K$-formula φ_a such that $M, v \vDash \varphi_a$ iff $v \in [M, w]_a$. He shows that distributed knowledge complies with the principle of full communication in locally distinguishing models. Gerbrandy also defines a model M to be *full* iff for every group $B \subseteq \mathcal{A}$ and every $w \in W$, every set of $\mathcal{L}_D$-formulas that is consistent with $\mathrm{KS}_B(M, w)$ is satisfiable in $[M, w]_B$. He proves that distributed knowledge complies with the principle of full communication in full models.

It is not clear, however, *why* distributed knowledge does not generally comply with the principle of full communication, and why it does in models that are finite, (locally) distinguishing and/or full. Moreover, it is not clear whether *all* models in which distributed knowledge complies with the principle of full communication are finite, (locally) distinguishing and/or full. And if not, whether a complete characterization of such models can be given.

Section 3.1 gives a simple analysis of why distributed knowledge does not always comply with the principle of full communication. This analysis yields a rather general class of models in which distributed knowledge *does* comply with the principle. Section 3.2 relates the defining properties of this class, called *tightness* and *epistemic saturation*, with notions familiar from modal logic. In particular, epistemic saturation is shown to be equivalent with modal saturation. In section 3.3 we compare our model class with the model class defined by van der Hoek, van Linder, and Meyer. We show that every finite, distinguishing model is tight and saturated, and moreover, that there is an interesting intermediate model class consisting of generated submodels of the canonical model. In section 3.4 we show that our model class can be generalized in various ways, which finally leads to a complete characterization of models in which distributed knowledge complies with the principle of full communication. Interestingly, the characteristic property of such models turns out to be a weak version of Gerbrandy's fullness property.

3.1. *Distributed Knowledge and Full Communication*

Why does distributed knowledge not comply with the principle of communication? One answer to this question is that in general *an agent's information state may not contain every world that is consistent with her knowledge set*. Consider, for example, the model depicted above: everything a knows in w is true in y. So y is consistent with a's knowledge set in w. Still, y does not belong to a's information state in w. Similarly, x does not belong to b's information state in w, even though everything b knows in w is true in x, that is, x is consistent with b's knowledge set in w. As a consequence, the intersection of two information states may sometimes yield more information than the union of the corresponding knowledge sets. To continue the above example, a and b's knowledge sets in w are identical. Thus, taking their union does not yield any new information. However, a and b's information states in w are different, and their intersection yields a new, more informative state. As a result, p is distributed knowledge among a and b in w, even though it is not entailed by the union of their knowledge sets.

This explanation, trivial as it may seem, leads us to the characterization of a rather general class of models in which distributed knowledge *does* comply with the principle of full communication, namely, the class of models in which every agent's information state *does* contain every world that is consistent with her knowledge set.

This requirement can be split into two sub-requirements. First, every set of formulas that is consistent with an agent's knowledge set must be satisfiable in the agent's information state. We call this requirement *epistemic saturation*.

DEFINITION 1 (EPISTEMIC SATURATION). — *A model* $M = (W, R, V)$ *is* epistemically saturated *iff for all* $w \in W$ *and all* $a \in \mathcal{A}$*, every set of* $\mathcal{L}_K$*-formulas that is consistent with* $KS_a(M, w)$ *is satisfiable in* $[M, w]_a$.

Second, every world that is consistent with an agent's knowledge set must be contained in the agent's information state. We call this requirement *tightness*.

DEFINITION 2 (TIGHTNESS). — *A model* $M = (W, R, V)$ *is* tight *iff for all* $w \in W$ *and all* $a \in \mathcal{A}$*, every world in* M *that is consistent with* $KS_a(M, w)$ *is in* $[M, w]_a$.

Notice that epistemic saturation and tightness are complementary requirements: the former requires that a model contains *enough worlds*; the latter requires that a model provides for *enough access*.

We now show that distributed knowledge complies with the principle of full communication in tight and epistemically saturated models.

PROPOSITION 3. — *In tight and epistemically saturated models distributed knowledge complies with the principle of full communication.*

PROOF. — Let $M = (W, R, V)$ be tight and epistemically saturated. To prove:

$$M, w \vDash D_B\varphi \quad \Rightarrow \quad \mathrm{KS}_B(M, w) \Vdash \varphi$$

Suppose that $\mathrm{KS}_B(M, w) \nVdash \varphi$. Then $\{\neg\varphi\} \cup \mathrm{KS}_B(M, w)$ is consistent. Let $a \in B$. Then $\{\neg\varphi\} \cup \mathrm{KS}_B(M, w)$ is consistent with $\mathrm{KS}_a(M, w)$ so, by epistemic saturation of M, $[M, w]_a$ contains a world w_a satisfying all formulas in $\{\neg\varphi\} \cup \mathrm{KS}_B(M, w)$. Thus, w_a is consistent with $\mathrm{KS}_b(M, w)$ for all $b \in B$, and by tightness of M, we must have $w_a \in [M, w]_b$ for all $b \in B$. This means that $w_a \in [M, w]_B$, and thus, that $M, w \nvDash D_B\varphi$, as desired. ■

3.2. *Tightness and Saturation*

Tightness is familiar from modal logic (see, for example, [FIN 75]). Epistemic saturation seems new, but turns out to be equivalent with another very familiar notion from modal logic, namely, modal saturation (see, for example, [BLA 01]). A model $M = (W, R, V)$ is modally saturated iff for every $w \in W$, every $a \in \mathcal{A}$, and every set of $\mathcal{L}_K$-formulas Σ, if every finite subset of Σ is satisfiable in $[M, w]_a$, then Σ itself is satisfiable in $[M, w]_a$.

PROPOSITION 4. — *A model is epistemically saturated iff it is modally saturated.*

PROOF. — Let $M = (W, R, V)$ be a model, $a \in \mathcal{A}$, and $w \in W$.

($\Rightarrow$) Suppose M is epistemically saturated. Let Σ be a set of $\mathcal{L}_K$-formulas such that every finite subset of Σ is satisfiable in $[M, w]_a$. Then every finite subset of $\Sigma' = \Sigma \cup \mathrm{KS}_a(M, w)$ is also satisfiable in $[M, w]_a$. It follows, by compactness, that Σ' is satisfiable and therefore consistent. Then, by epistemic saturation of M, Σ' must be satisfiable in $[M, w]_a$. This means that Σ is satisfiable in $[M, w]_a$, and thus that M is modally saturated.

($\Leftarrow$) Suppose M is modally saturated. Let Σ be a set of $\mathcal{L}_K$-formulas that is consistent with $\mathrm{KS}_a(M, w)$. Then every finite subset of Σ is satisfiable in $[M, w]_a$. To see this, let Σ' be a finite subset of Σ and suppose that Σ' is *not* satisfiable in $[M, w]_a$. Then $\neg \bigwedge \Sigma' \in \mathrm{KS}_a(M, w)$, which contradicts the assumption that Σ is consistent with $\mathrm{KS}_a(M, w)$. So every finite subset of Σ is satisfiable in $[M, w]_a$. By modal saturation, it follows that Σ is satisfiable in $[M, w]_a$, and thus that M is epistemically saturated. ■

Henceforth, we will simply refer to modal and epistemic saturation as *saturation*.

3.3. *Tight & Saturated versus Finite & Distinguishing*

Next, we show how tightness and saturation are related to the special model properties defined by van der Hoek, van Linder, and Meyer. We prove that every finite, distinguishing model is tight and saturated, and also identify an interesting intermediate class of models.

PROPOSITION 5. — *Finite, distinguishing models are saturated and tight.*

PROOF. — Finite models are always saturated [BLA 01]. Now suppose a model M is finite and therefore saturated, but *not* tight. Then for some agent a there must be two worlds w and v in M such that everything a knows in w is true in v, but v is not contained in $[M, w]_a$. Let Γ be the set of formulas true in v. Clearly, Γ is consistent with $\text{KS}_a(M, w)$. M is saturated so there must be a world u in $[M, w]_a$ that satisfies all formulas in Γ. But this means that u and v satisfy exactly the same formulas. So M is not distinguishing. We conclude that, if M is finite and distinguishing, then it must be saturated and tight. ■

The class of saturated and distinguishing models clearly subsumes the class of finite and distinguishing models, but is itself subsumed by the class of tight and saturated models. Interestingly, a model is saturated and distinguishing if and only if it is a generated submodel of the so-called *canonical* model. To show this, let us first recall the relevant definitions, which are all standard in modal logic [BLA 01].

A set of $\mathcal{L}_K$-formulas Σ is *maximally consistent* iff it is consistent and for all $\varphi \in \mathcal{L}_K$, either $\varphi \in \Sigma$ or $\neg\varphi \in \Sigma$. Note that the theory of a pointed model is always maximally consistent. The *canonical model* M^c is a triple (W^c, R^c, V^c) where:

$$
\begin{array}{lcl}
W^c & = & \{w_\Sigma \mid \Sigma \text{ is a maximally consistent set of } \mathcal{L}_K\text{-formulas}\} \\
R^c(a) & = & \{(w_\Sigma, w_\Delta) \mid K_a\varphi \in \Sigma \text{ implies } \varphi \in \Delta\} \\
V^c(w_\Gamma) & = & \{p \in \mathcal{P} \mid p \in \Sigma\}
\end{array}
$$

It is a well-known result that for every $\varphi \in \mathcal{L}_K$, $M^c, w_\Sigma \models \varphi$ iff $\varphi \in \Sigma$.

Given a model $M = (W, R, V)$ and a set of worlds $X \subseteq W$, the model $M|X = (X, R|X, V|X)$ is called the *restriction* of M to X. $M|X$ is a *generated submodel* of M iff, whenever $w \in X$ and $v \in [M, w]_a$ for some $a \in \mathcal{A}$, then also $v \in X$. So a generated submodel is a submodel that preserves information states.

Finally, two models (M, w) and (M', w') are isomorphic, $M, w \cong M', w'$, iff there is a *bijective* bisimulation (i.e., an *isomorphism*) $\mathcal{Z}$ between M and M' such that $w\mathcal{Z}w'$.

PROPOSITION 6. — *A model is saturated and distinguishing iff it is isomorphic to a generated submodel of the canonical model.*

PROOF. — Let $M = (W, R, V)$ be a model.

($\Rightarrow$) Suppose M is saturated and distinguishing. Then, every world in M has a unique, maximally consistent theory, and therefore uniquely corresponds to a world in M^c. Let W' be the set of all worlds w' in M^c that correspond to a world w in M, and let $M' = (W', R', V')$ be the restriction of M^c to W'. We show (1) that M' is a generated submodel of M^c and (2) that M and M' are isomorphic.

For (1), suppose that $w' \in W'$ and $(w', v') \in R^c(a)$ for some $a \in \mathcal{A}$ and some $v' \in W^c$. We must show that $v' \in W'$. By definition of the canonical model v' satisfies every formula in $\text{KS}_a(M^c, w')$. So $\Gamma(M^c, v')$ is consistent with $\text{KS}_a(M^c, w')$, and therefore also with $\text{KS}_a(M, w)$. Then, by epistemic saturation of M, $\Gamma(M^c, v')$ must be satisfiable in $[M, w]_a$. But this means that $[M, w]_a$ must contain a world v with $\Gamma(M^c, v')$ as its theory. It follows that v' is in W', and thus, that M' is a generated submodel of $M_{\mathcal{A},\mathcal{P}}$.

For (2), let $\mathcal{Z}$ be the bijection that relates every world $w \in W$ with its corresponding world $w' \in W'$. Clearly, w and w' satisfy the same proposition letters. It remains to check the back and forth clauses. We do the forth clause; the back clause is completely analogous. Let $a \in \mathcal{A}$ and $v \in [M, w]_a$. Clearly, $\Gamma(M, v)$ is consistent with $\text{KS}_a(M, w)$, and therefore $\Gamma(M', v')$ is also consistent with $\text{KS}_a(M, w)$. But then, by definition of the canonical model, $v' \in [M', w']_a$, as desired.

($\Leftarrow$) Suppose M is isomorphic to a generated submodel $M' = (W', R', V')$ of M^c. It is clear that no two worlds in M can have the same theory, so M must be distinguishing. To show that M is saturated it suffices to show that M' is saturated. M^c is well-known to be saturated [BLA 01]. Moreover, generated submodels preserve information states and saturation thereof. So M' and M are saturated. ■

We are not aware of any earlier statements of this result, although Fine [FIN 75] made several observations in this direction, and we would not be surprised if others had proven similar results, be it for different purposes and presumably in terms of modal saturation rather than epistemic saturation.

3.4. *A Complete Characterization of Full Communication*

So far, we have established that finite, distinguishing models are canonical, that canonical models are saturated and tight, and that saturation and tightness are sufficient conditions for compliance with the principle of full communication. Now we could ask whether they are also *necessary*, that is, whether *all* models in which distributed knowledge complies with the principle of full communication are saturated and tight. The answer is *no*. To see this, reconsider the proof of proposition 3. The role of saturation here is to ensure that every $\mathcal{L}_K$-formula φ that is consistent with the knowledge set of a group of agents is satisfiable in the information state of every agent in the group. Saturation is sufficient here, but not necessary: it is concerned with *sets* of formulas, where it could equally well just be concerned with *single* formulas. Then, once it has been established that φ is satisfiable in the information state of every agent in the group, the role of tightness is to ensure that there is at least one world satisfying φ which is contained in the information state of *all* agents in the group. Again, tightness is sufficient to fulfil this role, but not necessary. These observations give rise to the following definition:

DEFINITION 7 (FULL COMMUNICATION MODEL). — *A model* $M = (W, R, V)$ *is a full communication model iff for all* $w \in W$ *and all* $B \subseteq \mathcal{A}$*, every* $\mathcal{L}_K$*-formula that is consistent with* $\text{KS}_B(M, w)$ *is satisfiable in* $[M, w]_B$.

Notice that full models, as defined by Gerbrandy, are full communication models: where Gerbrandy quantifies over *sets of* $\mathcal{L}_D$*-formulas*, we quantify over *single* $\mathcal{L}_K$*-formulas*. The following proposition establishes that tight and saturated models are full communication models as well.

PROPOSITION 8. — *Tight and saturated models are full communication models.*

PROOF. — Let M be tight and saturated, w in M, $B \in \mathcal{A}$, and Σ a set of $\mathcal{L}_K$-formulas that is consistent with $\text{KS}_B(M, w)$. We must show that Σ is satisfiable in $[M, w]_B$. Let $a \in B$. Then $\Sigma \cup \text{KS}_B(M, w)$ is consistent with $\text{KS}_a(M, w)$, so by saturation of M, $[M, w]_a$ must contain a world w_a that satisfies all formulas in $\Sigma \cup \text{KS}_B(M, w)$. Thus, w_a is consistent with $\text{KS}_b(M, w)$ for all $b \in B$, and by tightness of M, we must have $w_a \in [M, w]_b$ for all $b \in B$. This means that $w_a \in [M, w]_B$. So Σ is satisfiable in $[M, w]_B$ which means that M is a full communication model. ■

Now we prove that distributed knowledge complies with the principle of full communication in full communication models, and moreover, that full communication models are the *only* models in which distributed knowledge complies with the principle of full communication.

PROPOSITION 9. — *Distributed knowledge complies with the principle of full communication in* M *if and only if* M *is a full communication model.*

PROOF. — The *if* part was, implicitly, already established by Gerbrandy [GER 99]: suppose M is a full communication model, w a world in M, B a group of agents, and φ a formula in $\mathcal{L}_K$ that is *not* entailed by $\text{KS}_B(M, w)$. Then $\neg\varphi$ is consistent with $\text{KS}_B(M, w)$, and therefore, as M is a full communication model, $\neg\varphi$ is satisfiable in $[M, w]_B$. But this means that φ cannot be distributed knowledge among B in (M, w). So distributed knowledge complies with the principle of full communication in M.

For the *only if* part, suppose that M is *not* a full communication model. Then, for some w in M and some $B \subseteq \mathcal{A}$, there is a formula φ in $\mathcal{L}_K$ that is consistent with $\text{KS}_B(M, w)$ but not satisfiable in $[M, w]_B$. For this to be the case, $\neg\varphi$ must not be entailed by $\text{KS}_B(M, w)$. On the other hand, $\neg\varphi$ must hold throughout $[M, w]_B$ and therefore be distributed knowledge among B in (M, w). So distributed knowledge does *not* comply with the principle of full communication in M. ■

3.5. *Conclusion*

Van der Hoek, van Linder, and Meyer [HOE 99] as well as Gerbrandy [GER 99] observed that distributed knowledge does not always comply with the principle of full communication. They also observed that distributed knowledge *does* comply with the principle in finite, distinguishing models, in locally distinguishing models, and in full models. It was not clear, however, why these properties were sufficient and whether they were necessary for compliance with the principle of full communication.

We established the following hierarchy of models in which distributed knowledge complies with the principle of full communication. Every class in the hierarchy contains the ones below it.

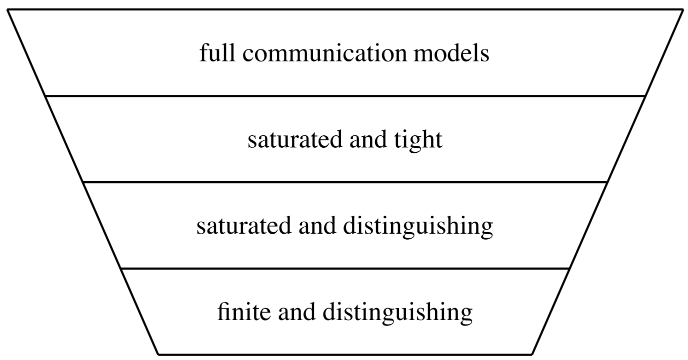

Our point of departure was the class of saturated and tight models, in which every agent's information state contains all worlds that are consistent with her knowledge set. We established that this class contains all finite and distinguishing models and all canonical (saturated and distinguishing) models. On the other hand, it is subsumed by the class of full communication models, which was shown to consist of all models in which distributed knowledge complies with the principle of full communication. Interestingly, the characteristic property of full communication models turned out to be a weak version of Gerbrandy's fullness property. One more thing to notice about the hierarchy above is that less restrictive properties are concerned with more fine-grained structural aspects of a model. Finiteness and distinguishability are concerned with the structure of a model as a whole, saturation and tightness are concerned with information states of individual agents, and full communication models impose restrictions on the information states of groups of agents. Distributed knowledge itself is defined in terms of information states of groups of agents. In the light of this observation it is not so surprising that finiteness, distinguishability, tightness, and saturation only led to a partial characterization of the class of full communication models.

As a final remark we would like to point out that the principle of full communication considered here presupposes a particular kind of communication, namely *private* communication. Other kinds of communication may be considered as well. Van Benthem [BEN 02], for example, is interested in the information that can be obtained by a group of agents through *public* communication. Parikh and Pacuit [PAC 04] are interested in the information that can be established by a group of agents relative to a restricted communication network. Chapter 4 and 5 of [ROE 05] are also relevant in this respect. In general, a dynamic view on epistemic logic [DIT 06] gives rise to many interesting notions of group knowledge, most of which remain to be explored.

We leave such explorations for another occasion, however, and now turn to the second topic of the paper.

4. Expressive Power

Bisimulation, which was introduced in a general modal setting by van Benthem [BEN 76], is a standard measure of structural equivalence between epistemic models. Intuitively, two pointed models are bisimilar iff (1) they assign the same truth values to all proposition letters and (2) they assign equivalent information states to all agents.

DEFINITION 10 (BISIMULATION). — *Let $M = (W, R, V)$ and $M' = (W', R', V')$ be two epistemic models. A non-empty relation $\mathcal{Z} \subseteq W \times W'$ is a bisimulation between M and M' iff for every $w \in W$ and $w' \in W'$ such that $w\mathcal{Z}w'$ we have:*

1) $V(w) = V'(w')$

2) For every agent $a \in \mathcal{A}$:

$$\begin{array}{llll} \textbf{forth} & \forall v \in [M, w]_a & : \exists v' \in [M', w']_a & : v\mathcal{Z}v' \\ \textbf{back} & \forall v' \in [M', w']_a & : \exists v \in [M, w]_a & : v\mathcal{Z}v' \end{array}$$

Two pointed models (M, w) and (M', w') are bisimilar, $(M, w) \simeq (M', w')$, if and only if there is a bisimulation $\mathcal{Z}$ between M and M' such that $w\mathcal{Z}w'$.

Bisimilarity matches the expressive power of the basic epistemic language: if $(M, w) \simeq (M', w')$ then (M, w) and (M', w') satisfy the same $\mathcal{L}_K$-formulas. The language with distributed knowledge operators, however, is more expressive. To see this, consider the models depicted in figure 1 and 2 (where all accessibility relations are equivalence relations).

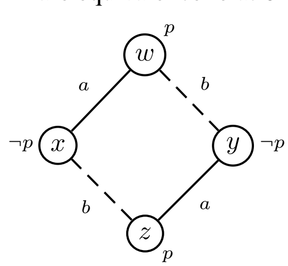

Figure 1. *Model M.*

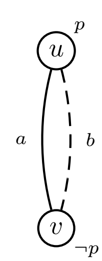

Figure 2. *Model N.*

(M, w) and (N, u) are bisimilar, but do not satisfy the same formulas in $\mathcal{L}_D$. For example, with $B = \{a, b\}$ we have:

$$\begin{array}{lcl} M, w & \models & D_B p \\ N, u & \not\models & D_B p \end{array}$$

Is there a natural extended notion of bisimulation that matches the expressive power of distributed knowledge operators? This question is the first of a list of open problems recently put together by van Benthem [BEN 05]. Section 4.1 proposes an extended notion of bisimulation, and proves its adequacy. Section 4.2 investigates related model comparison games, and section 4.3 explores a corresponding notion of saturation.

4.1. *Collective Bisimulation*

Bisimilarity requires two models to associate equivalent information states with every single agent. But distributed knowledge involves information states of *groups* of agents. Thus, in order to obtain a suitable notion of bisimulation, it seems natural to require two models to associate equivalent information states not just with every individual agent, but with every *group* of agents.

DEFINITION 11 (COLLECTIVE BISIMULATION). — *Let $M = (W, R, V)$ and $M' = (W', R', V')$ be two models. A non-empty relation $\mathcal{Z} \subseteq W \times W'$ is a collective bisimulation between M and M' iff for every $w \in W$ and $w' \in W'$ such that $w\mathcal{Z}w'$:*

1) $V(w) = V'(w')$

2) For every group *of agents $B \subseteq \mathcal{A}$:*

$$\begin{array}{llll} \textbf{forth} & \forall v \in [M, w]_B & : \exists v' \in [M', w']_B & : v\mathcal{Z}v' \\ \textbf{back} & \forall v' \in [M', w']_B & : \exists v \in [M, w]_B & : v\mathcal{Z}v' \end{array}$$

Two models (M, w) and (M', w') are collectively bisimilar, $(M, w) \simeq_c (M', w')$, iff there is a collective bisimulation $\mathcal{Z}$ between M and M' such that $w\mathcal{Z}w'$.

Collective bisimilarity generalizes ordinary bisimilarity. The latter only requires the back and forth conditions to hold for singleton groups of agents. So if two models are collectively bisimilar, then they are also bisimilar, but not vice versa: the models depicted in figure 1 and 2, for example, are bisimilar, but not collectively bisimilar.

We now show that collective bisimulation indeed matches the expressive power of $\mathcal{L}_D$ and establish some of its basic properties.

PROPOSITION 12 (ADEQUACY). — *If two pointed models are collectively bisimilar, then they satisfy exactly the same formulas in $\mathcal{L}_D$.*

PROOF. — Let $M = (W, R, V)$ and $M' = (W', R', V')$ be two models and let $w \in W$ and $w' \in W'$ be such that $(M, w) \simeq_c (M', w')$. We must prove that for every $\varphi \in \mathcal{L}_D$:

$$M, w \models \varphi \quad \text{iff} \quad M', w' \models \varphi$$

Clearly, it suffices to prove either the *if* or the *only if* part of the statement. We do the latter by induction on the complexity of φ. The base case and the induction steps for negation, conjunction and individual knowledge operators are standard [BLA 01]. Here is the induction step for distributed knowledge operators. Let φ be of the form $D_B\psi$ and suppose that $M, w \not\models D_B\psi$. Then $M, v \models \neg\psi$ for some $v \in [M, w]_B$. As

$(M, w) \simeq_c (M', w')$, there is a world $v' \in [M', w']_B$ such that $(M, v) \simeq_c (M', v')$. By the induction hypothesis, then, $M', v' \models \neg\psi$, and thus $M', w' \not\models D_B\psi$. ■

PROPOSITION 13. — *Collective bisimulation is an equivalence relation.*

PROOF. — Reflexivity and symmetry are clear. Transitivity amounts to:

$$(M, w) \simeq_c (M', w') \text{ and } (M', w') \simeq_c (M'', w'') \text{ imply } (M, w) \simeq_c (M'', w'')$$

Let $\mathcal{Z}$ be a collective bisimulation between (M, w) and (M', v'), and let $\mathcal{Z}'$ be a collective bisimulation between (M', w') and (M'', v''). Then the relation $\mathcal{Z}''$ between worlds in M and worlds in M'' given by $\{(w, w'') \mid \exists w' \in M' : w\mathcal{Z}w' \text{ and } w'\mathcal{Z}'w''\}$ is a collective bisimulation between (M, w) and (M'', v''). To see this, first observe that (M, w) and (M'', v'') should satisfy the same proposition letters, by $(M, w) \simeq_c (M', w')$ and $(M', w') \simeq_c (M'', w'')$. Now let v be a world in $[M, w]_B$ for some arbitrary group of agents B. Then, by $(M, w) \simeq_c (M', w')$, there is a world v' in $[M', w']_B$ such that $v\mathcal{Z}v'$. Moreover, by $(M', w') \simeq_c (M'', w'')$, there is a world v'' in $[M'', w'']_B$ such that $v'\mathcal{Z}'v''$. Then, by definition of $\mathcal{Z}''$, we have $v\mathcal{Z}''v''$. This establishes the forth clause. The back clause can be proven analogously. So we may conclude that $\mathcal{Z}''$ is a collective bisimulation between (M, w) and (M'', w''). ■

PROPOSITION 14. — *The set of collective bisimulations between any two models is closed under taking arbitrary (finite or infinite) unions.*

PROOF. — Let $\{\mathcal{Z}_i \mid i \in I\}$ be a non-empty set of collective bisimulations between two models M and M'. To show that $\mathcal{Z} = \bigcup_{i \in I} \mathcal{Z}_i$ is again a collective bisimulation between M and M', let w in M and w' in M' be any two worlds such that $w\mathcal{Z}w'$. Then we must have $w\mathcal{Z}_iw'$ for some $i \in I$. Now let v be a world in $[M, w]_B$ for some arbitrary group of agents B. Then, there must be a world v' in $[M', w']_B$ such that $v\mathcal{Z}_iv'$. But $v\mathcal{Z}_iv'$ means that $v\mathcal{Z}v'$ holds as well. This establishes the forth clause, and the back clause can be proven analogously. We conclude that $\mathcal{Z}$ is a collective bisimulation between M and M'. ■

COROLLARY 15. — *If there is a collective bisimulation between two models M and M', then there is always a* maximal *collective bisimulation between M and M': one that includes all other collective bisimulations between M and M'.*

PROOF. — Take the union of all collective bisimulations between M and M'. By proposition 14, this is again a collective bisimulation between M and M' and clearly, it includes all others. ■

4.2. *Model Comparison Games*

If two models are bisimilar, then they satisfy exactly the same formulas in $\mathcal{L}_K$. The opposite is only true for certain special classes of models. If two *finite* models, for example, satisfy the same formulas in $\mathcal{L}_K$, then they are bisimilar. However, *infinite* models may very well satisfy exactly the same formulas in $\mathcal{L}_K$ even though they are not bisimilar. Intuitively, this is because whether or not a model (M, w)

satisfies a formula φ only depends on worlds in M that can be reached from w in a finite number of steps along the accessibility relations in M. This number of relevant steps is bounded by the so-called *modal depth* of φ.

DEFINITION 16 (MODAL DEPTH). — *The modal depth $d(\varphi)$ of a formula $\varphi \in \mathcal{L}_K$ is defined recursively as follows:*

$$\begin{array}{lcll} d(p) & = & 0 & \textit{for all } p \in \mathcal{P} \\ d(\neg\varphi) & = & d(\varphi) & \\ d(\varphi \wedge \psi) & = & \mathit{max}(d(\varphi), d(\psi)) & \\ d(K_a\varphi) & = & d(\varphi) + 1 & \textit{for all } a \in \mathcal{A} \end{array}$$

Now, for two models (M, w) and (M', w') to satisfy the same formulas in $\mathcal{L}_K$ it is sufficient that (M, w) and (M', w') satisfy the same proposition letters and that for every *finite* path starting from w in M, we can find a corresponding path starting from w' in M' (and vice versa). Bisimulation requires something stronger, namely that (M, w) and (M', w') satisfy the same proposition letters and that for every (possibly *infinite*) path starting from w in M, we can find a corresponding path starting from w' in M' (and vice versa). This is why two infinite models that satisfy the same formulas in $\mathcal{L}_K$ are not necessarily bisimilar.

Model comparison games can be seen as finite approximations of a bisimulation. A model comparison game is played on two models, say (M, w) and (M', w'), by two players called spoiler and duplicator, and consists of a fixed number of rounds. Spoiler tries to establish that (M, w) and (M', w') satisfy different formulas; duplicator tries to show that they satisfy exactly the same formulas. The number of rounds of the game yields a bound on the length of the paths starting from w in M and from w' in M' that are available to spoiler as possible evidence for a structural difference between M and M'. If spoiler cannot win the n-round model comparison game on (M, w) and (M', w'), then these models satisfy exactly the same formulas in $\mathcal{L}_K$ up to depth n. So if duplicator has a winning strategy for all model comparison games on (M, w) and (M', w'), then (M, w) and (M', w') satisfy the same $\mathcal{L}_K$-formulas of arbitrary depth. And now the converse also holds: if (M, w) and (M', w') satisfy the same $\mathcal{L}_K$-formulas, then duplicator has a winning strategy for all model comparison games on (M, w) and (M', w'). This works out because games themselves have a finite number of rounds (just like formulas have finite modal depth), but for every n, there is a game with n rounds (just as for every n, there are formulas with modal depth n). Model comparison games characteristic for $\mathcal{L}_K$ are standard and can be found, for example, in [DIT 06]. Here, we define model comparison games characteristic for $\mathcal{L}_D$, called $\mathcal{L}_D$-games, and show that duplicator has a winning strategy in all $\mathcal{L}_D$-games on two models iff those two models satisfy exactly the same $\mathcal{L}_D$-formulas. Just as $\mathcal{L}_K$-games can be seen as finite approximations of a bisimulation, $\mathcal{L}_D$-games can be seen as finite approximations of a collective bisimulation.

DEFINITION 17 ($\mathcal{L}_D$-GAMES). — *Let $M = (W, R, V)$ and $M' = (W', R', V')$ be two models, let $w \in W$ and let $w' \in W'$. The rules of the n-round $\mathcal{L}_D$-game on (M, w) and (M', w') are as follows:*

1) If $n = 0$ then duplicator wins iff $V(w) = V'(w')$.

2) If $n \neq 0$ then spoiler can do either one of the following moves:

forth-move *Spoiler picks a group of agents $B \subseteq \mathcal{A}$ and a world $v \in [M, w]_B$. Duplicator responds by picking a world $v' \in [M', w']_B$. The rest of the game is the $(n-1)$-round $\mathcal{L}_D$-game on (M, v) and (M', v').*

back-move *Spoiler picks a group of agents $B \subseteq \mathcal{A}$ and a world $v' \in [M', w']_B$. Duplicator responds by picking a world $v \in [M, w]_B$. The rest of the game is the $(n-1)$-round $\mathcal{L}_D$-game on (M, v) and (M', v').*

3) If a player cannot make any further move, she loses the game.

PROPOSITION 18. — *Let (M, w) and (M', w') be two models, and let $\mathcal{P}$ be finite. Then (M, w) and (M', w') satisfy exactly the same $\mathcal{L}_D$-formulas iff for all $n \geq 0$, duplicator has a winning strategy for the n-round $\mathcal{L}_D$-game on (M, w) and (M', w').*

PROOF. — Define the depth $d(\varphi)$ of a formula $\varphi \in \mathcal{L}_D$ just as for formulas in $\mathcal{L}_K$ (see definition 16) with the following additional clause:

$$d(D_B\varphi) = d(\varphi) + 1 \text{ for all } B \subseteq \mathcal{A}$$

Write $(M, w) \equiv_n (M', w')$ iff (M, w) and (M', w') satisfy exactly the same formulas in $\mathcal{L}_D$ with depth n. Observe that, as $\mathcal{P}$ and $\mathcal{A}$ are both assumed to be finite, for every $n \geq 0$, there are only finitely many formulas with depth n, up to logical equivalence. Now we prove, by induction on n, that duplicator has a winning strategy for the n-round $\mathcal{L}_D$-game on (M, w) and (M', w') iff $(M, w) \equiv_n (M', w')$.

The base case ($n = 0$) follows directly from the definitions.

The induction step involves two directions, which we treat separately:

($\Rightarrow$) Suppose duplicator has a winning strategy for the $(n+1)$-round $\mathcal{L}_D$-game on (M, w) and (M', w'). Then we must show that for all $\varphi \in \mathcal{L}_D$ such that $d(\varphi) = n+1$, we have $M, w \models \varphi$ iff $M', w' \models \varphi$. We do so by induction on φ, and only treat the non-standard case in which φ is of the form $D_B\phi$, where $d(\phi) = n$. Suppose $M, w \models D_B\phi$. We show that $M', w' \models D_B\phi$ must hold as well: let v' be any world in $[M', w']_B$. Suppose spoiler chooses v' in a first back-move of the $(n+1)$-round $\mathcal{L}_D$-game on (M, w) and (M', w'). Then, by assumption, duplicator can pick a world v in $[M, w]_B$ such that she (duplicator) has a winning strategy for the remaining n-round $\mathcal{L}_D$-game on (M, v) and (M', v'). By the induction hypothesis, $(M, v) \equiv_n (M', v')$. In particular, we have that $M', v' \models \phi$, and as v' was an arbitrary world in $[M, w]_B$, we may conclude that $M', w' \models D_B\phi$.

($\Leftarrow$) Now suppose that $(M, w) \equiv_{n+1} (M', w')$. We must show that duplicator has a winning strategy in the $(n+1)$-round $\mathcal{L}_D$-game on (M, w) and (M', w'). Suppose spoiler starts with a back-move and picks a world v' in $[M', w']_B$. Now, toward a contradiction, suppose that there is no world v in $[M, w]_B$ such that $(M, v) \equiv_n (M', v')$

(which is what duplicator needs for a winning strategy). Then, for every v in $[M,w]_B$, there is a formula $\psi_v \in \mathcal{L}_D$ of depth n such that $M,v \models \psi_v$ but $M',v' \not\models \psi_v$. Define:

$$\psi_w = \bigvee_{v \in [M,w]_B} \psi_v$$

We observed that there are only finitely many formulas with depth n, up to logical equivalence. So ψ_w corresponds to a finite disjunction, again of depth n. We have $M,w \models D_B\psi_w$, but $M',w' \not\models D_B\psi_w$, even though $d(D_B\psi_w) = n+1$. This contradicts our assumption that $(M,w) \equiv_{n+1} (M',w')$. So we may conclude that duplicator has a winning strategy for the $(n+1)$-round $\mathcal{L}_D$-game on (M,w) and (M',w'). ■

4.3. *Saturation and Fullness*

Recall the notions of saturation and fullness from section 3. A model is saturated iff every set of $\mathcal{L}_K$-formulas that is consistent with the knowledge set of a single agent is satisfiable in the information state of that single agent. A model is full iff every set of $\mathcal{L}_D$-formulas that is consistent with the knowledge set of a group of agents is satisfiable in the information state of that group of agents. Notice that fullness generalizes saturation just like collective bisimulation generalizes ordinary bisimulation and distributed knowledge operators generalize individual knowledge operators: where saturation, bisimulation, and individual knowledge operators are concerned with the information states of individual agents, fullness, collective bisimulation and distributed knowledge operators are concerned with the information states of *groups* of agents.

There is an important and well-known connection between saturation, bisimilarity, and the expressive power of $\mathcal{L}_K$: saturated models are bisimilar iff they satisfy the same $\mathcal{L}_K$-formulas [BLA 01]. The following proposition establishes a similar connection between fullness, collective bisimilarity, and the expressive power of $\mathcal{L}_D$.

PROPOSITION 19. — *Full models are collectively bisimilar iff they satisfy the same $\mathcal{L}_D$-formulas.*

PROOF. — Proposition 12 already established that, in general, collectively bisimilar models satisfy the same formulas in $\mathcal{L}_D$. Here, we show that full models which satisfy the same formulas in $\mathcal{L}_D$ are collectively bisimilar. Let M and M' be full, and let w in M and w' in M' be such that (M,w) and (M',w') satisfy the same formulas in $\mathcal{L}_D$. Let $\mathcal{Z}$ be the relation between W and W' that consists of all pairs of worlds that satisfy exactly the same formulas in $\mathcal{L}_D$. Clearly, $w\mathcal{Z}w'$. We will show that $\mathcal{Z}$ is a collective bisimulation between M and M'. To do so, let B be an arbitrary group of agents and let v be any world in $[M,w]_B$. Let Γ be the set of $\mathcal{L}_D$-formulas true in v. Then Γ is consistent with $\text{KS}_B(M,w)$, and therefore also consistent with $\text{KS}_B(M',w')$. As M' is full, Γ must be satisfiable in $[M',w']_B$. But this means that there must be a world v' in $[M',w']_B$ that satisfies exactly the same $\mathcal{L}_D$-formulas as v. So we have $v\mathcal{Z}v'$, which means that $\mathcal{Z}$ is indeed a collective bisimulation between M and M'. ■

This can be taken a bit further. The following propositions establish that full models are collectively bisimilar iff they are bisimilar and that with respect to full models, $\mathcal{L}_D$ has no more expressive power than $\mathcal{L}_K$.

PROPOSITION 20. — *Full models are bisimilar iff they are collectively bisimilar.*

PROOF. — Clearly, collectively bisimilar models are always bisimilar. We must show that full, bisimilar models are always collectively bisimilar. To do so, let M and M' be full and let w in M and w' in M' be such that (M, w) and (M', w') are bisimilar. Let $\mathcal{Z}$ be the *maximal* bisimulation between (M, w) and (M', w') (the existence of which is established in [BLA 01]). We will show that $\mathcal{Z}$ is also a *collective* bisimulation between (M, w) and (M', w'). To do so, let B be an arbitrary group of agents and let v be any world in $[M, w]_B$. Let Γ be the set of $\mathcal{L}_D$-formulas true in v. Then Γ is consistent with $\text{KS}_B(M, w)$. (M, w) and (M', w') are bisimilar, so they satisfy the same formulas in $\mathcal{L}_K$. Therefore, Γ is also consistent with $\text{KS}_B(M', w')$. M' is full, so Γ must be satisfiable in $[M', w']_B$. This means that there must be a world v' in $[M', w']_B$ that satisfies exactly the same $\mathcal{L}_D$-formulas as v. In particular, v and v' satisfy exactly the same $\mathcal{L}_K$-formulas. M and M' are both saturated, so (M, v) and (M', v') must be bisimilar. But then, as $\mathcal{Z}$ was assumed to be maximal, we must have $v\mathcal{Z}v'$, and this establishes that $\mathcal{Z}$ is indeed a collective bisimulation. ■

PROPOSITION 21. — *Full models satisfy the same formulas in $\mathcal{L}_K$ iff they satisfy the same formulas in $\mathcal{L}_D$.*

PROOF. — Clearly, it is enough to prove the *only if* part. Let (M, w) and (M', w') be full. Suppose (M, w) and (M', w') satisfy the same formulas in $\mathcal{L}_K$. Since both models are saturated, they must be bisimilar [BLA 01]. Then, by proposition 20, they must be collectively bisimilar. But then it follows from proposition 12 that they must satisfy exactly the same formulas in $\mathcal{L}_D$. ■

4.4. *Conclusion*

We have proposed an extended notion of bisimulation that matches the expressive power of the basic epistemic language extended with distributed knowledge operators. We have also defined related model comparison games and established their adequacy. Finally, we showed that fullness generalizes saturation just as collective bisimulation generalizes ordinary bisimulation and established that full models are collectively bisimilar iff they are bisimilar iff they satisfy exactly the same $\mathcal{L}_K$-formulas iff they satisfy exactly the same $\mathcal{L}_D$-formulas. In particular, $\mathcal{L}_D$ and $\mathcal{L}_K$ have equal expressive power with respect to full models.

As mentioned before, the issue addressed here is the first of a list of open problems in a recent survey by van Benthem [BEN 05]. The model comparison games defined here are variations on standard games for the basic modal language as described, for instance, by van Ditmarsch, van der Hoek, and Kooi in [DIT 06]. The connection between saturation, bisimulation, and the expressive power of $\mathcal{L}_K$ is a standard result in modal logic [BLA 01].

Acknowledgements

I would like to thank Johan van Benthem, Paul Dekker, Hans van Ditmarsch, Jan van Eijck, Michael Franke, Jelle Gerbrandy, Patrick Girard, Jeroen Groenendijk, Paul Harrenstein, Andreas Herzig, Wiebe van der Hoek, Samson de Jager, Barteld Kooi, Fenrong Liu, Maricarmen Martinez, Siewert van Otterloo, Olivier Roy, Tomasz Sadzik, Josh Snyder, Yanjing Wang, and Audrey Yap for inspiration and discussion. I would also like to thank the anonymous referees for helpful feedback.

5. References

[BEN 76] VAN BENTHEM J., "Modal Correspondence Theory", PhD thesis, University of Amsterdam, 1976.

[BEN 02] VAN BENTHEM J., "One is a Lonely Number: on the Logic of Communication", *Technical Report PP-2002-27, ILLC Amsterdam*, , 2002.

[BEN 05] VAN BENTHEM J., "Open Problems in Logical Dynamics", *Technical Report PP-2005-06, ILLC Amsterdam*, , 2005.

[BLA 01] BLACKBURN P., DE RIJKE M., VENEMA Y., *Modal Logic*, Cambridge University Press, 2001.

[DIT 06] VAN DITMARSCH H., VAN DER HOEK W., KOOI B., *Dynamic Epistemic Logic*, Springer-Kluwer, 2006.

[FAG 95] FAGIN R., HALPERN J., MOSES Y., VARDI M., *Reasoning about Knowledge*, MIT Press, 1995.

[FIN 75] FINE K., "Some Connections between Elementary and Modal Logic", *Proceedings of the Third Scandinavian Logic Symposium*, 1975, p. 15–31.

[GER 99] GERBRANDY J., "Bisimulations on Planet Kripke", PhD thesis, University of Amsterdam, 1999.

[HOE 99] VAN DER HOEK W., VAN LINDER B., MEYER J.-J., "Group knowledge is not always distributed (neither is it always implicit)", *Mathematical social sciences*, vol. 38, 1999, p. 215–240.

[MEY 95] MEYER J.-J., VAN DER HOEK W., *Epistemic Logic for AI and Computer Science*, Campbridge University Press, 1995.

[PAC 04] PACUIT E., PARIKH R., "The Logic of Communication Graphs", *Declarative Agent Languages and Technologies (DALT 04)*, 2004.

[ROE 05] ROELOFSEN F., "Exploring Logical Perspectives on Distributed Information and its Dynamics", *Master Thesis, University of Amsterdam*, , 2005.