Ultrafilter Unions: an exercise in modal definability

Yde Venema
Institute for Logic, Language and Computation
University of Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam
yde@wins.uva.nl

Abstract

We introduce the notion of the ultrafilter union of a family of pointed modal models over an ultrafilter of the index set and we argue that this operation is the modal analogue of taking ultraproducts in the model theory of first-order logic. We use this notion to give a structural characterization of the modally definable classes of pointed models.

1 Introduction

For many years modal model theory mostly consisted of modal frame theory; for instance, the notion of modal definability was almost exclusively studied on the level of frames. One important result in this area is a theorem by R. Goldblatt and S.K. Thomason (1974) stating that an elementary class of frames is modally definable if and only if it has certain closure properties: it should be closed under taking bounded morphic images, disjoint unions and generated subframes, while its complement must be closed under taking ultrafilter extensions. The first three concepts are best understood as frame derivatives of bisimulation, while the notion of an ultrafilter extension is often conceived as a rather esoteric trace of the duality theory between frames and Boolean algebras with operators. Later on we will see that taking the ultrafilter extension of a *model* is in fact a very natural operation — this is of course known, but it does not seem to be *well* known.

Definability results concerning classes of modal *models* are more recent. In his dissertation, M. de Rijke (1993) proves a number of results that are inspired by his 'equation' stating that bisimulations are to modal logic what partial isomorphisms are to first order logic. One of the results that de Rijke proves is a theorem concerning classes of pointed models (a pointed model is a modal model together with a designated point): he shows that such classes are modally definable if and only if they are closed under taking bisimilar pointed models and ultraproducts, while their complement is closed under ultrapowers. The key technical result that de Rijke uses here is his Bisimulation Theorem stating that two models are modally equivalent if and only if they have bisimilar ultrapowers.

What I want to do here is push this ‘modalizing’ of model definability results one step further, adding the ‘equation’ that ultrafilter extensions of models are to modal logic what ultrapowers are to first order logic. The idea is that taking the ultrafilter extension of a model is a way to modally saturate it, in the same way that the ultrapower over a free ultrafilter saturates a first order model. I will also introduce a new structural operation, viz., that of taking the *ultrafilter union* of a family of (pointed) models over a given ultrafilter of the index set; this operation is in some sense the modal analogue of forming ultraproducts. Using these operations, one can prove the following result which is a clear analogue of the well-known (corollary of the) Keisler-Shelah Theorem in first-order model theory characterizing the elementary and the basic elementary classes of models.

Theorem 1 *Let* K *be a class of pointed models.*

1. K *is definable by a set of modal formulas if and only if it is closed under taking ultrafilter unions and bisimilar models, while its complement is closed under taking ultrafilter extensions.*

2. K *is definable by a single modal formula if and only if both it and its complement are closed under taking ultrafilter unions and bisimilar models.*

One can state analogous results for ordinary modal models (i.e., models without a designated point), but for reasons of space limitations I will not do so here.

2 Basics

We briefly review the background material needed to understand Theorem 1.

Modal semantics We work in the basic modal language having one diamond $\Diamond$. Given a fixed set of proposition letters, the set of (modal) formulas is given by the usual inductive definition stating that (i) proposition letters are formulas, and (ii) whenever φ and ψ are formulas, then so are $\neg\varphi$, $\varphi \wedge \psi$ and $\Diamond\psi$. We will use the standard abbreviations; in particular, we write $\Box\varphi$ for $\neg\Diamond\neg\varphi$.

A (modal) *model* for such a language is a triple $\mathcal{M} = (W, R, V)$ such that W is some set; R is a binary relation on W; and V is a *valuation*, that is, a function mapping proposition letters to subsets of W. Given a model $\mathcal{M} = (W, R, V)$, we inductively define the notion of *truth* or *satisfaction* of a formula at a point of the model:

$$\begin{array}{rll}
\mathcal{M}, s \Vdash p & \text{if} & s \in V(p), \\
\mathcal{M}, s \Vdash \neg\varphi & \text{if} & \mathcal{M}, s \nVdash \varphi, \\
\mathcal{M}, s \Vdash \varphi \wedge \psi & \text{if} & \mathcal{M}, s \Vdash \varphi \text{ and } \mathcal{M}, s \Vdash \psi, \\
\mathcal{M}, s \Vdash \Diamond\varphi & \text{if} & \mathcal{M}, t \Vdash \varphi \text{ for some } t \text{ with } Rst.
\end{array}$$

We denote the set of points where a formula φ is true by $V(\varphi)$. Given two points s and s' in models $\mathcal{M}$ and $\mathcal{M}'$, respectively, we say that s and s' are *modally equivalent*, notation: $\mathcal{M}, s \stackrel{\diamond}{\sim} \mathcal{M}', s'$, if for all formulas φ, $\mathcal{M}, s \Vdash \varphi$ iff $\mathcal{M}', s' \Vdash \varphi$.

In some applications, such as process theory, pointed models rather than plain models are studied. A *pointed model* is simply a model $\mathcal{M}$ together with a designated point r in $\mathcal{M}$; formally, such a structure is denoted as $(\mathcal{M}, r)$. A formula φ is said to *hold of* a pointed model $\mathcal{M}$ if $\mathcal{M}, r \Vdash \varphi$; a set of formulas holds of a pointed model if each formula of the set holds of it. Two pointed models are called *modally equivalent* if $\mathcal{M}, s \stackrel{\diamond}{\sim} \mathcal{M}', s'$. A set Δ of modal formulas is said to *define* a class of pointed models K if any pointed model belongs to K if and only if Δ holds of it. In case Δ is a singleton $\{\delta\}$ we will say that δ rather than $\{\delta\}$ defines K. A class of pointed models is *modally definable* if there is some set of formulas defining it.

Bisimulations We first define the notion of a bisimulation between two models. Let $\mathcal{M} = (W, R, V)$ and $\mathcal{M}' = (W', R', V')$ be two models. A non-empty relation $Z \subseteq W \times W'$ is called a *bisimulation between $\mathcal{M}$ and $\mathcal{M}'$*, notation: $Z : \mathcal{M} \leftrightarroweq \mathcal{M}'$, if the following conditions are satisfied:

(prop) if sZs' then s and s' satisfy the same proposition letters,

(forth) if sZs' and Rst then there is a t' in $\mathcal{M}'$ such that tZt' and $R's't'$,

(back) if sZs' and $R's't'$ then there is a t in $\mathcal{M}$ such that tZt' and Rst.

Truth of modal formulas is invariant under bisimulations, that is:

Proposition 2 *Let Z be a bisimulation between the models $\mathcal{M}$ and $\mathcal{M}'$. Then $Z \subseteq \stackrel{\diamond}{\sim}$.*

We say that two pointed models $(\mathcal{M}, s)$ and $(\mathcal{M}', s')$ are bisimilar if there is a bisimulation linking the respective designated points, notation: $(\mathcal{M}, s) \leftrightarroweq (\mathcal{M}', s')$. A class K of pointed models is *closed under bisimilar models* if every pointed model that is bisimilar to a model in K belongs to K itself. It is immediate from Proposition 2 that modally definable classes are closed under bisimilar models.

Disjoint unions and generated submodels For ordinary models we can define the notion of *the disjoint union* of a family of models. Let $\{\mathcal{M}_i \mid i \in I\}$ be a collection of models, say $\mathcal{M}_i = (W_i, R_i, V_i)$. Assume that the universes of these models are pairwise disjoint — if this is not the case, then we proceed with some canonically defined isomorphic copies that *do* have this property. The *disjoint union* $\biguplus_{i \in I} \mathcal{M}_i$ of this collection is defined as the model (W, R, V) with $W = \bigcup_i W_i$, $R = \bigcup_i R_i$, and V being defined by $V(p) = \bigcup_i V_i(p)$.

It is easy to see that the identity relation between one of the models $\mathcal{M}_i$ and the disjoint union $\biguplus_{i \in I} \mathcal{M}_i$ is in fact a bisimulation. From this the following proposition is immediate.

Proposition 3 *Let $\{\mathcal{M}_i \mid i \in I\}$ be a collection of models. Then for each $i \in I$ and each point s_i in $\mathcal{M}_i$ we have that $(\mathcal{M}_i, s) \overset{\diamond}{\sim} (\biguplus_{i \in I} \mathcal{M}_i, s_i)$.*

For pointed models the notion of disjoint union seems to be inappropriate since there is no natural way to single out a designated point.

A model $\mathcal{M}' = (W', R', V')$ is a *submodel* of $\mathcal{M} = (W, R, V)$ if $W' \subseteq W$, while R' and V' are the restrictions of R and V, respectively, to W'. It is a *generated submodel* if W' satisfies in addition that $R[w'] \subseteq W'$ for all $w' \in W'$, that is, W' is closed under taking R-successors. It is easy to see that the identity between a generated submodel and the big model is in fact a bisimulation. From this the following is immediate.

Proposition 4 *Let $\mathcal{M}'$ be a generated submodel of the model $\mathcal{M}$, and let s' be a point in $\mathcal{M}'$. Then $\mathcal{M}', s' \overset{\diamond}{\sim} \mathcal{M}, s'$.*

Ultrafilter extensions The *ultrafilter extension* of a modal model can be seen as a kind of modal completion or saturation of it. In order to define this notion, we assume familiarity with the concept of an *ultrafilter* over a set.

Given a model $\mathcal{M} = (W, R, V)$, its ultrafilter extension $\mathcal{M}^*$ will be defined as the model (Uf_W, R^*, V^*). In this definition, Uf_W is the set of ultrafilters over W, V^* is given by $V^*(p) = \{u \in Uf_W \mid V(p) \in u\}$, and R^* is defined by putting R^*uv iff $m_R(X) \in u$ for all $X \in v$. Here the operation m_R on the power set of W is given as $m_R(X) = \{s \in W \mid Rst \text{ for some } t \in W\}$.

The ultrafilter extension of a pointed model $(\mathcal{M}, s)$ is the pointed model $(\mathcal{M}^*, \pi_s)$, where π_s denotes the *principal* ultrafilter generated by s. Since we have Rst iff $R^*\pi_s\pi_t$, we can indeed consider $\mathcal{M}^*$ as an extension of (an isomorphic copy of) $\mathcal{M}$. As a corollary of the following property, modally definable classes of pointed models are invariant under taking ultrafilter extensions.

Proposition 5 *Let $\mathcal{M}$ be a model. Then*

1. *for any modal formula φ and any ultrafilter u, $V(\varphi) \in u$ iff $\mathcal{M}^*, u \Vdash \varphi$;*
2. *hence, $\mathcal{M}, s \overset{\diamond}{\sim} \mathcal{M}^*, \pi_s$ for any state s in $\mathcal{M}$.*

3 Ultrafilter unions

In order to give a structural characterization of the modally definable classes of pointed models we need an operation that amalgamates a family of pointed modal models, in the same way that ultraproducts take together a collection of first-order structures.

For a precise definition, we need some auxiliary notions. Fix an index set I and a family $\{(\mathcal{M}_i, s_i) \mid i \in I\}$ of pointed models. Write $\mathcal{M}_i = (W_i, R_i, V_i)$ and let $\mathcal{M} = (W, R, V)$ be the disjoint union of these models; in order not to clutter up notation we assume once more that the universes of the models are mutually disjoint.

Also fix an ultrafilter F over I. A subset A of the disjoint union $W = \biguplus_i W_i$ is called *F-compatible* if the set of indices i for which $A \cap W_i$ is non-empty belongs to

F. An ultrafilter u over the set W is called *F-compatible* if each of its elements is an F-compatible subset of W. As an example of an F-compatible ultrafilter, let $\bar{s} = (s_i)_{i \in I}$ be the sequence of designated points, and consider the set

$$\pi_F(\bar{s}) \;=\; \{A \subseteq W \mid \{i \in I \mid s_i \in A\} \in F\}.$$

The ultrafilter union of $\{(\mathcal{M}_i, s_i) \mid i \in I\}$ over F is defined as the pair

$$(\biguplus_F \mathcal{M}_i, \pi_F(\bar{s})).$$

Here $\biguplus_F \mathcal{M}_i$, the ultrafilter union of the models without designated points, is defined as the structure (W_F, R_F, V_F), where W_F denotes the set of F-compatible ultrafilters over W and R_F and V_F denote the restriction of R^* and V^* to W_F respectively.

Note that if F is the principal ultrafilter over I generated by the index i, then the ultrafilter union of the family $\{(\mathcal{M}_i, s_i) \mid i \in I\}$ over F is isomorphic to the ultrafilter extension of $\mathcal{M}_i$. It is also fairly obvious that any ultrafilter union of a collection of models is a submodel of the ultrafilter extension of the disjoint union of the collection. In fact, it is a *generated* submodel, a fact of which we will make good use.

Proposition 6 *Let $\{\mathcal{M}_i \mid i \in I\}$ be a family of models, and let F be an ultrafilter over I. Then $\biguplus_F \mathcal{M}_i$ is a generated submodel of $(\biguplus_{i \in I} \mathcal{M}_i)^*$.*

PROOF. Obviously, $\biguplus_F \mathcal{M}_i$ is a submodel of $(\biguplus \mathcal{M}_i)^*$. In order to prove that it is a generated submodel, suppose that u is an F-compatible ultrafilter and that R^*uv. We have to show that v is F-compatible as well, so take an arbitrary element A of v and assume for contradiction that A is not F-compatible. Write $A_i = A \cap W_i$ (where W_i is the universe of $\mathcal{M}_i$. F-incompatibility of A means that the set $\{i \in I \mid A_i \neq \varnothing\}$ does not belong to F. Since F is an ultrafilter this implies that $\{i \in I \mid A_i = \varnothing\} \in F$. From the fact that for any relation R we have $m_R(\varnothing) = \varnothing$ and the fact that ultrafilters are upwards closed we may conclude that $\{i \in I \mid m_{R_i}(A_i) = \varnothing\} \in F$, so $\{i \in I \mid m_{R_i}(A_i) \neq \varnothing\} \notin F$.

However, it is easy to see that $m_R(A) \cap W_i = m_{R_i}(A_i)$, so we have proved that $m_R(A)$ is not F-compatible. But $m_R(A)$ belongs to u since R^*uv and $A \in v$. Thus we have proved that u is not an F-compatible ultrafilter which gives the desired contradiction. QED

The following proposition can be seen as the modal analogue of Łos' Lemma.

Proposition 7 *Let $\{\mathcal{M}_i \mid i \in I\}$ be a family of models, and let F be an ultrafilter over I. Furthermore, let φ be a modal formula and let $\{s_i \mid i \in I\}$ be a collection of states such that $\{i \in I \mid \mathcal{M}_i, s_i \Vdash \varphi\}$ belongs to F. Then $\biguplus_F \mathcal{M}_i, \pi_F(\bar{s}) \Vdash \varphi$.*

PROOF. It follows from Proposition 3 that the set $\{i \in \mid \biguplus_i \mathcal{M}_i, s_i \Vdash \varphi\}$ belongs to F, so (with U denoting the valuation on the disjoint union) we can write this as $\{i \in I \mid s_i \in U(\varphi)\} \in F$. From the definition of $\pi_F(\bar{s})$ it then follows that $U(\varphi)$ belongs to $\pi_F(\bar{s})$.

Now Proposition 5 yields that φ *holds* at $\pi_F(\bar{s})$ in the ultrafilter extension $(\biguplus_{i\in I} \mathcal{M}_i)^*$ of the disjoint union. But the ultrafilter *union* over F is a generated submodel of this model by Proposition 6, so since $\pi_F(\bar{s})$ is a point of this generated submodel, we may use Proposition 4 to conclude that indeed $\biguplus_F \mathcal{M}_i, \pi_F(\bar{s}) \Vdash \varphi$. QED

4 M-saturation

It is well-known that the converse of Proposition 2 does not hold in general: points may be modally equivalent without being bisimilar. A class of models K has the *Hennessy-Milner property* if any bisimulation between models in K. This notion was introduced by R. Goldblatt (1995) (for single models). Many natural classes of models have this property, for instance, the class of image finite models in which every point has a finite number of successors. A more general sufficient condition involves the notion of *m-saturation*, which was introduced by K. Fine (1975) under the name 'modally saturated$_2$'.

Let Σ be a set of modal formulas, $\mathcal{M} = (W, R, V)$ a model, and A a subset of W. We say that Σ is *satisfiable in* A if there is a point s in A where Σ is satisfied, and *finitely satisfiable in* A if every finite subset of Σ is satisfiable in A. We call $\mathcal{M}$ *m-saturated* if the following holds for every state s in the model. Suppose that a set of formulas Σ is finitely satisfiable in the collection $R[s]$ of successors of s; then we require that Σ is also satisfiable in $R[s]$. A pointed model $(\mathcal{M}, s)$ is m-saturated if $\mathcal{M}$ is m-saturated.

Proposition 8 *Let $\mathcal{M}$ and $\mathcal{M}'$ be two m-saturated models, and s and s' two points in $\mathcal{M}$ and $\mathcal{M}'$, respectively. Then $s \leftrightarroweq s'$ if and only if $s \overset{\diamond}{\sim} s'$.*

The reason why we call the ultrafilter extension $\mathcal{M}^*$ the *modal saturation* of $\mathcal{M}$ is that ultrafilter extensions are m-saturated. This result can be generalized to ultrafilter unions.

Proposition 9 *Let $\{(\mathcal{M}_i, s_i) \mid i \in I\}$ be a family of models, and let F be an ultrafilter over I. Then the ultrafilter union $(\biguplus_F \mathcal{M}_i, \pi_F(\bar{s}))$ is m-saturated. In particular, the ultrafilter extension $\mathcal{M}^*$ of any model $\mathcal{M}$ is m-saturated.*

Proof. The fact that ultrafilter extensions are m-saturated, is due to Goldblatt (1995). The easiest way to prove that ultrafilter unions are m-saturated is then by observing that any generated submodel of an m-saturated model is itself m-saturated, and using Proposition 6. QED

5 Proof of Theorem 1

In this section we will prove Theorem 1. For both parts, the easy direction, that is, the one from left to right, can easily be derived from the Propositions 2, 5 and 7.

For the other direction of part 1, assume that K has all the closure properties specified in the statement of the Theorem. We will prove that the modal theory Θ_{K} of K, consisting of the formulas that hold of each pointed model in K, in fact defines K. That is, we will show that for any pointed model $(\mathcal{M}, s)$:

(1) $$(\mathcal{M}, s) \Vdash \Theta_{\mathsf{K}} \text{ iff } (\mathcal{M}, s) \text{ belongs to } \mathsf{K}.$$

Obviously, the right to left direction of this equation follows trivially from the definition of Θ_{K}. The remainder of the proof is devoted to establishing the other direction of (1). Hence, suppose that $(\mathcal{M}, s)$ is a pointed model such that $\mathcal{M}, s \Vdash \Theta_{\mathsf{K}}$; write $\mathcal{M} = (M, R, V)$.

Let Σ be the set of formulas true in $\mathcal{M}$ at s:

$$\Sigma = \{\sigma \mid \mathcal{M}, s \Vdash \sigma\}.$$

We will show that Σ is satisfiable in some m-saturated model in K. First we observe that for every $\sigma \in \Sigma$ there is some pointed model $\mathcal{N}_\sigma$ in K of which σ holds. This is in fact rather easy to see, for if it were not the case for some $\sigma_0 \in \Sigma$, then $\neg\sigma_0$ would be true at the designated point of each model in K and hence would belong to Θ_{K}; but then by our assumption on $(\mathcal{M}, s)$ this would mean that $\mathcal{M}, s \Vdash \neg\sigma_0$, contradicting the fact that $\sigma_0 \in \Sigma$.

Assume that Σ is enumerated as $\{\sigma_i \mid i \in \omega\}$, and define ψ_n as the formula $\sigma_0 \wedge \ldots \wedge \sigma_n$, for any $n \in \omega$. Note that every ψ_i belongs to Σ; hence, there is a family of pointed models $\{(\mathcal{M}_i, t_i) \mid i \in \omega\}$ in K such that each ψ_i holds at t_i in $\mathcal{M}_i$. Now let F be some non-principal ultrafilter over ω; define

$$\mathcal{N} := \biguplus_F \mathcal{N}_i, \qquad t := \pi_F(\bar{t}).$$

Obviously, $(\mathcal{N}, t)$ belongs to K since it is an ultrafilter union of models in K.

We will now prove that Σ holds of $(\mathcal{N}, t)$. Fix some $n \in \omega$; we will show that $\mathcal{N}, t \Vdash \sigma_n$. It follows from the definition of the ψ formulas and our assumption that $\mathcal{N}_i, t_i \Vdash \psi_i$ for each i, that $\mathcal{N}_i, t_i \Vdash \sigma_n$ holds for co-finitely many i. Hence, the set $\{i \in I \mid \mathcal{N}_i, t_i \Vdash \sigma_n\}$ belongs to F since F is non-principal. But then it is immediate from Proposition 7 that $\mathcal{N}, t \Vdash \sigma_n$; since this applies to all n we may conclude that $\mathcal{N}, t \Vdash \Sigma$.

From $\mathcal{N}, t \Vdash \Sigma$ we may infer that $(\mathcal{M}, s) \stackrel{\diamond}{\sim} (\mathcal{N}, t)$: let φ be an arbitrary formula. If $\mathcal{M}, s \Vdash \varphi$, then $\varphi \in \Sigma$ and thus, $\mathcal{N}, t \Vdash \varphi$. If, on the other hand, $\mathcal{M}, s \not\Vdash \varphi$, then the negation of φ belongs to Σ, and thus, $\mathcal{N}, t \Vdash \neg\varphi$ whence $\mathcal{N}, t \not\Vdash \varphi$.

But now we are almost finished: since both $\mathcal{M}$ and $\mathcal{N}$ are m-saturated, it follows from Proposition 8 that the relation $\stackrel{\diamond}{\sim}$ is a bisimulation between $\mathcal{N}$ and $\mathcal{M}$ which links t to s. From the closure of K under bisimulation we may then deduce that indeed, the pointed model $(\mathcal{M}, s)$ belongs to K. This proves the first part of the theorem.

In order to prove the second part of Theorem 1, assume that K is a class of pointed models such that both K and its complement $\overline{\mathsf{K}}$ are closed under bisimulations and ultrafilter unions. We will show that K is definable by a single modal formula.

To start with, it follows from the first part of the theorem that K is definable by a set of modal formulas. Assume for contradiction that K is not definable by a single formula. Then we may assume that there is a defining set Σ for K which can be enumerated as a collection $\Sigma = \{\sigma_i \mid i \in \omega\}$ of ever stronger axioms in the sense that $\sigma_i \to \sigma_j$ is a modal validity if and only if i is strictly larger than j. From the fact that no $\sigma_i \to \sigma_{i+1}$ is a modal validity it follows that there is a family of pointed models $\{(\mathcal{N}_i, s_i) \mid i \in I\}$ such that for each i we have that $\sigma_i \wedge \neg\sigma_{i+1}$ holds at s_i in $\mathcal{N}_i$. Note that this means that each $\mathcal{N}_i$ belongs to $\overline{\mathsf{K}}$.

Now let i be an arbitrary index. From the fact that $\sigma_j \to \sigma_i$ is a modal validity for all $j > i$ we may derive that σ_i holds of co-finitely many of the pointed models. Reasoning in a similar way as before we can show that if we take a non-principal ultrafilter F over I, we will find that σ_i will hold of the ultrafilter union $(\mathcal{N}, t) = (\biguplus_F gN_i, \pi_F(\bar{s}))$. Since i was arbitrary this means that in fact $\mathcal{N}, t \Vdash \Sigma$, whence $\mathcal{N}$ belongs to K. This gives the desired contradiction to the fact that the class $\overline{\mathsf{K}}$ should be closed under taking ultrafilter unions — these results have a somewhat more involved formulation.

6 Conclusion

The construction of taking ultrafilter unions seems to be a natural way of amalgamating a family of (pointed) modal models. In this paper, I used it to give a purely 'modal' structural characterization of the classes of pointed modal models that are modally definable. I have similar results for the modally definable classes of ordinary modal models, but these have a slightly more involved formulation. In the future I hope to put this construction to further use, for instance to prove modal interpolation results.

Generalizing these characterization results to extended modal formalisms like multi-modal logic, tense logic or modal logic with polyadic operators are straightforward. It seems to be much harder to find a structural characterization of the classes of (pointed) modal models that can be defined using (sets of) formula(s) of *dynamic* logic; it would be nice if some version of the ultrafilter union construction could be of help here.

References

K. Fine. 1975. Some connections between elementary and modal logic. In S. Kanger, editor, *Proceedings of the Third Scandinavian Logic Symposium. Uppsala 1973*, Amsterdam. North-Holland.

R.I. Goldblatt and S.K. Thomason. 1974. Axiomatic classes in propositional modal logic. In J. Crossley, editor, *Algebra and Logic*, Lecture Notes in Mathematics 450, pages 163–173. Springer, Berlin.

R. Goldblatt. 1995. Saturation and the Hennessy-Milner property. In A. Ponse, M. de Rijke, and Y. Venema, editors, *Modal Logic and Process Algebra*, volume 53 of *Lecture Notes*. CSLI Publications.

M. de Rijke. 1993. *Extending Modal Logic.* Ph.D. thesis, ILLC, University of Amsterdam.