Model definability, purely modal

Yde Venema

Contents

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 [3] 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 is 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 [6], Maarten de Rijke 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 the ultrafilter extension of a model modally saturates it, in the same way that the ultrapower over a free ultrafilter saturates a first order model. To be more precise, the aim of this paper is to prove the following two results:

Theorem 1 *A class of modal models is modally definable if and only if it is closed under taking disjoint unions, surjective bisimulations and ultrafilter extensions, while its complement is closed under taking ultrafilter extensions.*

Theorem 2 *A class of pointed models is modally definable if and only if it is closed under bisimulation and ultrafilter unions, while its complement is closed under taking ultrafilter extensions.*

All these notions will be defined and explained below, and the paper is entirely self-contained.

2 Basics

We first review the background material needed to understand Theorem 1 and its proof. None of this material in this section is original, with the exception of

Theorem 3 (which is a straightforward consequence of known results).

Modal semantics

Although all of the results in the paper can be proved, with minor and obvious adaptations, for extended modal languages such as polymodal logic (with more than one diamond), Priorean tense logic (which has both a forward and a backward looking operator) or languages with polyadic operators such as arrow logic, for reasons of notational simplicity I confine myself to the basic modal language.

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 $V(\varphi)$. A formula φ is *globally true in* or *true throughout* $\mathcal{M}$, notation: $\mathcal{M} \Vdash \varphi$, if φ is true at every point of the model. Since we will be comparing the sets of formulas holding at points in different models, we need the following definition: s in $\mathcal{M}$ and s' in $\mathcal{M}'$ are *modally equivalent*, notation: $\mathcal{M}, s \overset{\diamond}{\sim} \mathcal{M}', s'$ if for all formulas φ, $\mathcal{M}, s \Vdash \varphi$ iff $\mathcal{M}', s' \Vdash \varphi$.

For sets of formulas and classes we use analogous or obvious definitions and notation. For instance, a set of formulas Σ is *true at* or *satisfied in* a point s of a model $\mathcal{M}$ if every formula in Σ is true at s. Σ is *satisfiable in a class of models* if there is some point in some model in the class where Σ is satisfied.

The central notion in this paper is that of a set of formulas defining a class of models.

Definition 1 A set Δ of modal formulas is said to *define* a class of models K if any model $\mathcal{M}$ belongs to K if and only if Δ is true throughout $\mathcal{M}$. A class of modal models is *modally definable* if there is some set of formulas defining it.

Theorem 1 is a structural characterization of the classes of models that are modally definable. The main concepts that we will be using are those of a bisimulation between two models, m-saturation, and the ultrafilter extension of a model. These notions will be discussed now.

Bisimulation

The notion of a bisimulation, which was introduced in Johan van Benthem's dissertation under the name 'p-relation', is of fundamental importance in the model theory of modal logic. It also plays a crucial role in theoretical computer science, as the basic relation of observational indistinguishability between (graph representations of) processes.

Definition 2 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.

If Z is a bisimulation between $\mathcal{M}$ and $\mathcal{M}'$, and if wZw' then we say that w and w' are bisimilar, notation: $w \leftrightarroweq w'$ (or $\mathcal{M}, s \leftrightarroweq \mathcal{M}', s'$ if we need to make the models explicit). We call a bisimulation $Z : \mathcal{M} \leftrightarroweq \mathcal{M}'$ *surjective* if every point in $\mathcal{M}'$ belongs to the range of Z.

The following proposition states that bisimilar points are modally equivalent.

Proposition 1 *Let* $\mathcal{M}$ *and* $\mathcal{M}'$ *be two modal models, and* s *and* s' *two states in* $\mathcal{M}$ *and* $\mathcal{M}'$*, respectively. Then* $s \leftrightarroweq s'$ *implies* $s \overset{\diamond}{\sim} s'$*.*

Proof. Let $\mathcal{M} = (W, R, V)$ and $\mathcal{M}' = (W', R', V')$ be two models, and assume that there is a bisimulation Z between $\mathcal{M}$ and $\mathcal{M}'$. By an easy formula induction one can prove that any modal formula φ satisfies the property that for any two points s and s' with sZs' we have that $\mathcal{M}, s \Vdash \varphi$ iff $\mathcal{M}', s \Vdash \varphi$. From this the Proposition is immediate. QED

This result can already be used to prove one small part of our main Theorem:

Proposition 2 *Modally definable classes of models are closed under surjective bisimulations.*

Proof. Suppose that the class K is defined by the set Σ, and assume that Z is a surjective bisimulation between the models $\mathcal{M}$, belonging to K, and $\mathcal{M}'$. We will show that $\mathcal{M}'$ belongs to K as well.

In order to arrive at a contradiction, suppose that this is not the case. Since Σ defines K, this means that there is some formula $\sigma \in \Sigma$ and a point s' in $\mathcal{M}'$ such that $\mathcal{M}', s' \not\Vdash \sigma$. But since Z is surjective, there must be a point s in $\mathcal{M}$ such that sZs'. By Proposition 1, this implies that $\mathcal{M}, s \not\Vdash \sigma$. But then Σ cannot be globally true in $\mathcal{M}$, which gives the desired contradiction. QED

One of the other closure property of modally definable classes, viz. the one involving disjoint unions, is also proved using bisimulations. Let us first define what disjoint unions *are*:

Definition 3 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 disjoint — if this is not the case, then we proceed with some canonically defined isomorphic copies that *do* have a non-empty intersection. which the universes are disjoint. The *disjoint union* $\biguplus_{i\in I} \mathcal{M}_i$ of this collection is defined as the model (W, R, V) with W being the union $\bigcup_i W_i$, R being the union $\bigcup_i R_i$, and V being defined by $V(p) = \bigcup_i V_i(p)$.

Proposition 3 *Modally definable classes are closed under taking disjoint unions.*

PROOF. Obvious by the observation that for any $j \in I$, the relation $\{(x, x) \mid x \in W_j\}$ is a bisimulation between $\mathcal{M}_j$ and $\biguplus_{i\in I} \mathcal{M}_i$. QED

Before we move on to the other closure properties mentioned in Theorem 1, we discuss another structural operation related to bisimulations, namely that of a generated subframe.

Definition 4 Let $\mathcal{M} = (W, R, V)$ be some model. A *submodel* of $\mathcal{M}$ is a model (W', R', V') such that W' is a subset of W, and R' and V' are the restrictions of R and V to W, respectively. That is, $R' = R \cap (W \times W)$ and $V'(p) = V(p) \cap W$ for each proposition letter p. A submodel $\mathcal{M}' = (W', R', V')$ of $\mathcal{M}$ is a *generated submodel* if W' is R-closed; that is, if a point w belongs to W' then each of its successors must belong to W' as well.

Given a point s of W, there is a smallest generated submodel of $\mathcal{M}$ containing s; this model is denoted by $\mathcal{M}_s$. It is easy to see that the universe of $\mathcal{M}_s$ consists of all points that can be reached from s via a finite R-path.

The following Proposition, which has a very simple proof, will play a (minor) role in the proof of Theorem 1.

Proposition 4

1. *For any model $\mathcal{M}$ and any point s in $\mathcal{M}$, there is a surjective bisimulation between $\mathcal{M}$ and $\mathcal{M}_s$.*
2. *For any model $\mathcal{M} = (W, R, V)$ there is a surjective bisimulation between $\bigcup_{s\in W} \mathcal{M}_s$ and $\mathcal{M}$.*
3. *Given a modally definable class K, any model $\mathcal{M}$ belongs to K if and only each of its point-generated subframes belong to K.*

M-saturation

It is well-known that the converse of Proposition 1 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 [2] (for single models) and has been studied further by M. Hollenberg [4]. 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 [1] under the name modally saturated$_2$. For a formal definition we need to fine-tune the definition of satisfiability somewhat. 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.

Definition 5 We call a model *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 $R[s]$.

Proposition 5 *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 $w \leftrightarrows w'$ if and only if $w \overset{\diamond}{\sim} w'$.*

Proof. The left to right direction of this proposition has already been proved in Proposition 1.

For the other direction, suppose that $\mathcal{M}$ and $\mathcal{M}'$ are m-saturated models. We will prove that the relation $\overset{\diamond}{\sim}$ itself is a bisimulation. It is easy to see that $\overset{\diamond}{\sim}$ satisfies the propositional clause of the definition of a bisimulation, so we concentrate on the back clause (the forth clause is of course symmetric). Suppose that s, s' and t' are points in $\mathcal{M}$, $\mathcal{M}'$ and $\mathcal{M}'$ respectively, such that $s \overset{\diamond}{\sim} s'$ and that $Rs't'$. We have to find a successor t of s such that $t \overset{\diamond}{\sim} t'$. Define Σ to be the set of formulas true at t' in $\mathcal{M}'$. We will show that Σ is satisfied at some successor of s in $\mathcal{M}$. This suffices to prove the Proposition, since any x that satisfies Σ is modally equivalent to t'.

But by m-saturatedness of $\mathcal{M}$ it suffices to show that Σ is finitely satisfiable in the set of successors of s. Hence, let Σ_0 be some finite subset of Σ; obviously, $\mathcal{M}', t' \Vdash \bigwedge \Sigma_0$. But then by $R's't'$ it follows that $\mathcal{M}', s' \Vdash \Diamond \bigwedge \Sigma_0$. Hence, from $s \overset{\diamond}{\sim} s'$ we may infer that $\mathcal{M}, s \Vdash \Diamond \bigwedge \Sigma_0$; this in its turn implies the existence of a point t_0 with Rst_0 and $\mathcal{M}, t_0 \Vdash \bigwedge \Sigma_0$. But this clearly means that Σ_0 is satisfiable in t_0, a successor of s, as required. QED

Ultrafilter extensions

Since not every model is m-saturated, and we have seen that m-saturated models have some nice properties, we are looking for a way to 'modally saturate' models. Now the basic defect of an unsaturated model is that states may lack certain successors. Hence, the basic idea underlying the modal saturation of a model will be to complete the model by adding certain points to it. To make this more precise, we need the notion of an ultrafilter.

Let W be some set; a collection u of subsets of W is an *ultrafilter* over W if it is (i) closed under taking intersections: if $A \in u$ and $B \in u$, then $A \cap B$ belongs to u as well, (ii) upwards closed: if $A \in u$ and $A \subseteq B \subseteq W$, then B belongs to u as well; and (iii) contains, for any subset A of W, either A or its complement $W \setminus A$.

Given an element s of W, it is easy to see that the collection

$$\pi_s \;:=\; \{A \subseteq W \mid s \in A\}$$

is an ultrafilter over W. We will refer to this set as the *principal ultrafilter associated with* s; ultrafilters that are not principal are called *free*.

In the sequel we will frequently need to prove the existence of an ultrafilter satisfying certain properties. We will each time do this by an appeal to the Ultrafilter Theorem yielding that any collection E of subsets of a set W can be extended to an ultrafilter $u \supseteq E$ if E has the *finite intersection property*, meaning that for any finite number of sets $A_1, \ldots, A_n \in E$ we have that $A_1 \cap \cdots \cap A_n \neq \varnothing$.

The modal saturation $\mathcal{M}^*$ of a model $\mathcal{M}$ will be called its *ultrafilter extension*. On the level of frames, this concept stems from Goldblatt & Thomason [3], while the name is due to Johan van Benthem (as far as I know). The universe of $\mathcal{M}^*$ will consist of the collection Uf_W of all ultrafilters over W — the name *extension* is appropriate since $\mathcal{M}$ itself will be (isomorphic to) a submodel of $\mathcal{M}^*$ via the correspondence $s \sim \pi_s$. In order to understand the definition of the accessibility relation R^* and the valuation V^* of $\mathcal{M}^*$ it is useful to take as a guideline our wish to prove the following equivalence (holding for every formula φ):

(1) $\qquad V(\varphi) \in u$ if and only if $\mathcal{M}^*, u \Vdash \varphi$.

For, this immediately provides the definition of V^*:

$$V^*(p) \;:=\; \{u \in \mathit{Uf}_W \mid V(p) \in u\},$$

while it also gives a useful requirement on the definition of R^*. Namely, if R^*uv holds then we need that for any formula φ, $V(\Box\varphi) \in u$ implies $V(\varphi) \in v$. (It is more convenient to switch to $\Box$-notation here.) Now consider the following operation l_R on the power set $\mathcal{P}(W)$ of W:

$$l_R(A) \;:=\; \{x \in W \mid \; y \in A \text{ for all } y \text{ such that } Rxy\},$$

that is, a point $x \in W$ belongs to $l_R(A)$ if each of its R-successors belongs to A. Clearly, the operation l_R corresponds to the truth definition of $\Box$, in the sense that $V(\Box\varphi) = l_R(V(\varphi))$. In other words, the requirement that we just mentioned can be rephrased as follows: R^*uv implies that $V(\varphi) \in v$ whenever $l_R(V(\varphi)) \in u$. This inspires the following definition for R^*:

$$R^* \;=\; \{(u, v) \in \mathit{Uf}_W \times \mathit{Uf}_W \mid \text{ for } \mathit{all}\ A \subseteq W\text{: } l_R(A) \in u \text{ only if } A \in v\}.$$

Definition 6 Let $\mathcal{M} = (W, R, V)$ be some model. Then the *ultrafilter extension* $\mathcal{M}^*$ of $\mathcal{M}$ is defined as the model $(\mathit{Uf}_W, R^*, V^*)$.

As we have seen, the concept of an ultrafilter extension is in fact tailored towards making the following proposition true.

Proposition 6 *Let $\mathcal{M} = (W, R, V)$ be some model. Then (1) holds for any ultrafilter u over W and any formula φ. Hence, for every state s of $\mathcal{M}$ we have that s (in $\mathcal{M}$) and π_s (in $\mathcal{M}^*$) are modally equivalent:*

$$\mathcal{M}, s \overset{\diamond}{\sim} \mathcal{M}^*, \pi_s.$$

Proof. The second part of the Proposition follows immediately from the first part, which is proved by formula-induction. We will only treat the case of the inductive step in which φ is of the form $\Diamond\psi$.

First assume that $\mathcal{M}^*, u \Vdash \Diamond\psi$. By definition then, there is an ultrafilter v over W such that R^*uv and $\mathcal{M}^*, v \Vdash \psi$. It follows from the inductive hypothesis that $V(\psi) \in v$; but then it can be easily proved from R^*uv and the defining properties of ultrafilters that $V(\Diamond\psi) \in u$.

The other implication requires a bit more work. Assume that $V(\Diamond\psi) \in u$; we will show that the set

$$E \;=\; \{V(\psi)\} \cup \{A \mid l_R(A) \in u\}$$

has the finite intersection property. Since the collection $\{A \mid l_R(A) \in u\}$ is closed under taking intersections, it suffices to prove that $V(\varphi) \cap A \neq \varnothing$ for any $A \subseteq W$ satisfying $l_R(A) \in u$. Take an arbitrary such A: it follows from $V(\Diamond\psi) \in u$ and $l_R(A) \in u$ that $V(\Diamond\psi) \cap l_R(A) \neq \varnothing$. Let s be some point in this intersection. From $s \in V(\Diamond\psi)$ it follows that there is some point t with Rst and $t \in V(\psi)$. But then $s \in l_R(A)$ implies that $t \in A$, whence $V(\varphi) \cap A \neq \varnothing$.

But if E has the finite intersection property, we can extend it to some ultrafilter v. Clearly $\{A \mid l_R(A) \in u\} \subseteq v$ implies that R^*uv, and from $V(\psi) \in v$ it follows by the inductive hypothesis that $\mathcal{M}^*, v \Vdash \psi$. This shows that indeed $\mathcal{M}^*, u \Vdash \Diamond\psi$. QED

Proposition 6 paves the way for establishing the remaining closure properties of modally definable classes.

Proposition 7 *Modally definable classes of models and their complements are closed under taking ultrafilter extensions.*

Proof. This proof follows immediately from the following observation, which in its turn is a consequence of the previous Proposition:

(2) for any formula φ and any model $\mathcal{M}$: $\mathcal{M} \Vdash \varphi$ iff $\mathcal{M}^* \Vdash \varphi$.

For a proof of (2), first assume that $\mathcal{M} \Vdash \varphi$; this means that $V(\varphi) = W$. But W is a member of every ultrafilter over W, so it follows from Proposition 6 that $\mathcal{M}^*, u \Vdash \varphi$ for every ultrafilter u. That is, $\mathcal{M}^* \Vdash \varphi$.

In order to prove the other direction, we reason by contraposition: assume that φ is *not* true throughout $\mathcal{M}$. That is, there is some point s in $\mathcal{M}$ where φ is false; but then it follows from Proposition 6 that φ is false at π_s in $\mathcal{M}^*$ as well, so $\mathcal{M}^* \not\Vdash \varphi$. QED

The next proposition, due to Goldblatt, states that ultrafilter extensions can indeed be seen as modal completions or saturations of models.

Proposition 8 *For any model $\mathcal{M}$, the model $\mathcal{M}^*$ is m-saturated.*

PROOF. Let $\mathcal{M} = (W, R, V)$ be a modal model. We will prove that its ultrafilter extension $\mathcal{M}^* = (Uf_W, R^*, V^*)$ is m-saturated.

Consider an ultrafilter u over W, and a set Σ that is finitely satisfiable in the set of R^*-successors of u. We will show that the set

$$Q \ = \ \{V(\sigma) \mid \sigma \in \Sigma\} \cup \{A \mid l_R(A) \in u\}.$$

has the finite intersection property. For this it suffices to prove that for an arbitrary formula σ in Σ and an arbitrary A in u we have that $V(\sigma) \cap A \neq \varnothing$, since both $\{V(\sigma) \mid \sigma \in \Sigma\}$ and $\{A \mid l_R(A) \in u\}$ are closed under taking intersections. But if σ belongs to Σ, then by definition of Σ there is an R^*-successor w of u such that σ holds at w, or, equivalently, such that $V(\sigma) \in w$. But $l_R(A) \in u$ implies that $A \in w$, by definition of R^*. Hence, w belongs to the intersection of $V(\sigma)$ and A, and therefore this intersection cannot be the empty set.

But if Q has the finite intersection property, by the Ultrafilter Theorem we can extended it to some ultrafilter v over W. For this v we have that $\mathcal{M}^*, v \Vdash \Sigma$ because $V(\sigma) \in v$ for every $\sigma \in \Sigma$, and R^*uv because $A \in v$ whenever $l_R(A) \in u$. Hence, we have satisfied Σ in some successor of u; since u and Σ were arbitrary, this means that indeed, $\mathcal{M}^*$ is m-saturated. QED

As a digression, we mention that using the above two Propositions, we can prove the following alternative to de Rijke's Bisimulation Theorem. Like that result, Theorem 3 below identifies modal equivalence as 'bisimilarity somewhere else', but now this 'somewhere else' is not situated in some ultrapowers of the models, but in the respective ultrafilter extensions.

Theorem 3 *Let $\mathcal{M}$ and $\mathcal{N}$ be two models, and s and t two points in $\mathcal{M}$ and $\mathcal{N}$, respectively. Then s and t are modally equivalent if and only if their corresponding principal ultrafilters are bisimilar:*

$$\mathcal{M}, s \overset{\diamond}{\sim} \mathcal{N}, t \textit{ if and only if } \mathcal{M}^*, \pi_s \ \leftrightarroweq \ \mathcal{N}^*, \pi_t.$$

PROOF. Let $\mathcal{M}$, $\mathcal{N}$, s and t be as in the statement of the Theorem. It follows from Proposition 6 that

$$\mathcal{M}, s \ \overset{\diamond}{\sim} \ \mathcal{N}, t \text{ iff } \mathcal{M}^*, \pi_s \ \overset{\diamond}{\sim} \ \mathcal{N}^*, \pi_t,$$

and since both $\mathcal{M}^*$ and $\mathcal{N}^*$ are m-saturated by Proposition 8, it follows from Proposition 5 that

$$\mathcal{M}^*, \pi_s \ \overset{\diamond}{\sim} \ \mathcal{N}^*, \pi_t \text{ iff } \mathcal{M}^*, \pi_s \ \leftrightarroweq \ \mathcal{N}, \pi_t.$$

From these two observations the Theorem is immediate. QED

3 Modal definability

The aim of this section is to provide a proof of Theorem 1.

PROOF OF THEOREM 1. The easy, left to right, direction of the Theorem is immediate by the Propositions 2, 3 and 7 above. For the other direction of the Theorem, 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 are true throughout each model in K, in fact defines K. That is, we will show that for any model $\mathcal{M}$:

(3) $$\mathcal{M} \models \Theta_{\mathsf{K}} \text{ iff } \mathcal{M} \text{ belongs to } \mathsf{K}.$$

Obviously, the right to left direction of this equation follows trivially from the definitions. The remainder of the proof is devoted to establishing the other direction of (3), for an increasingly wide range of models.

CASE 1 We first consider the case that $\mathcal{M} = (W, R, V)$ is an m-saturated and point-generated model in which Θ_{K} is globally true. Assume that $\mathcal{M}$ is generated from the point s; that is, $\mathcal{M} = \mathcal{M}_s$. Let Σ be the set of formulas true in $\mathcal{M}$ at s:

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

We claim that Σ is satisfiable in some m-saturated model in K.

We first show that for every $\sigma \in \Sigma$ there is some model $\mathcal{N}_\sigma$ in K in which σ is satisfied. This is in fact rather easy to see, for suppose that it were not the case for some $\sigma_0 \in \Sigma$. This would imply that $\neg\sigma_0$ is true throughout all models in K, whence $\neg\sigma_0$ would belong to the theory of K. But by our assumption on $\mathcal{M}$ this would mean that $\mathcal{M}, s \Vdash \neg\sigma_0$, contradicting the fact that $\sigma_0 \in \Sigma$.

Now define

$$\mathcal{N} = \biguplus_{\sigma \in \Sigma} \mathcal{N}_\sigma,$$

say $\mathcal{N} = (X, P, U)$. Obviously, $\mathcal{N}$ belongs to K since it is the disjoint union of models in K. We will prove that the ultrafilter extension $\mathcal{N}^*$ of $\mathcal{N}$ is in fact the model that we are looking for.

For each $\sigma \in \Sigma$, recall that $U(\sigma) = \{x \in X \mid \mathcal{N}, x \Vdash \sigma\}$ denotes the extension of σ in $\mathcal{N}$, and define the set $E \subseteq \mathcal{P}(X)$ by

$$E = \{U(\sigma) \mid \sigma \in \Sigma\}.$$

It is not difficult to prove that E does not contain the empty set and that it is closed under intersections; for the latter, it suffices to observe that $U(\sigma_1) \cap U(\sigma_2) = U(\sigma_1 \wedge \sigma_2)$. But then E has the finite intersection property and thus it follows from the Ultrafilter Theorem that E is contained in some ultrafilter u over X. It follows from Proposition 6 and the fact that $U(\sigma) \in u$ for each $\sigma \in \Sigma$, that $\mathcal{N}^*, u \Vdash \Sigma$. But by Proposition 8, $\mathcal{N}^*$ is m-saturated, and by the assumption that K is closed under taking ultrafilter extensions, $\mathcal{N}^*$ belongs to K.

Hence, we have found our m-saturated model $\mathcal{N}^*$ in K in which Σ is satisfiable, namely at the point u. It easily follows that s (in $\mathcal{M}$) and u (in $\mathcal{N}^*$) are modally equivalent; but then by the fact that both $\mathcal{M}$ and $\mathcal{N}^*$ are m-saturated, Proposition 5 guarantees the existence of a bisimulation Z between $\mathcal{N}^*$ and $\mathcal{M}$ linking u and s. Hence, if we can prove that Z is in fact surjective, we are finished, since $\mathcal{N}^*$ belongs to K and we have assumed that K is closed under surjective bisimulations. But the surjectivity of Z is a straightforward consequence of the fact that $\mathcal{M}$ is generated from s: a rather easy inductive proof on the distance of a point from s will reveal that any point in $\mathcal{M}$ belongs to the range of Z.

CASE 2 Now we consider the case that $\mathcal{M}$ is an arbitrary m-saturated model in which Θ_{K} is globally true. It follows from Proposition 4 that Θ_{K} is globally true in every point-generated submodel of $\mathcal{M}$. But also, an easy argument shows that each such model is m-saturated. Hence, the argument of the previous, special, case reveals that every point-generated submodel of $\mathcal{M}$ belongs to K. But then by Proposition 4 and the fact that K is closed under generated subframes and disjoint unions we may infer that $\mathcal{M}$ itself also belongs to K.

CASE 3 Finally, we can consider the general case in which Θ_{K} is true throughout some arbitrary model $\mathcal{M}$. It follows from Proposition 7 and the assumption that $\mathcal{M} \Vdash \Theta_{\mathsf{K}}$ that $\mathcal{M}^* \Vdash \Theta_{\mathsf{K}}$. But we have just seen that this implies the membership of $\mathcal{M}^*$ in K, since $\mathcal{M}^*$ is m-saturated. And since K reflects ultrafilter extensions, this means that $\mathcal{M}$ must belong to K as well. QED

4 Pointed models

In some applications, such as process theory, pointed models rather than ordinary models simpliciter 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 $\mathcal{M}, r \Vdash \varphi$; for sets of formulas, and classes of pointed models, the obvious analogues apply. A bisimulation between two pointed models is simply a bisimulation between the models that links the designated points. The ultrafilter extension of a pointed model consists of the ultrafilter extension of the model, together with the principal ultrafilter associated with the designated point of the original model.

In the last section of this paper we consider the question which classes of pointed models are modally definable. This problem is on the one hand simpler and on the other hand more complex than the version for ordinary models. Simpler because we do not need the surjectivity condition on bisimulations, but more complex since we cannot work with simple disjoint unions. This is because there is no obvious *definition* of the disjoint union of a collection of pointed models: what would be the designated point of such a structure? Fortunately, when searching for the appearance of disjoint unions in the proof of Theorem 1, we find only one crucial use. This is in the construction of a model in which a certain set Σ is satisfiable, from a collection of models in which each finite part of Σ is satisfiable. Since this construction is essentially a combination of a disjoint unions and ultrafilter extensions, why not combine these two operations into a new, more complex one?

Definition 7 Let $\{(\mathcal{M}_i, s_i) \mid i \in I\}$ be a set of pointed models (with disjoint universes). An ultrafilter u over the set $W = \bigcup_{i \in I} W_i$ is called *point evening* if every cofinite subset of $S = \{s_i \mid i \in I\}$ belongs to u (or, equivalently, if u is a superset of some free ultrafilter over S). Any structure of the form $((\bigcup_{i \in I} \mathcal{M}_i)^*, u)$ with u a point evening ultrafilter is called an *ultrafilter union* of $\{(\mathcal{M}_i, s_i) \mid i \in I\}$.

Observe that by taking as the designated point of the ultrafilter union an ultrafilter containing all cofinite subsets of the set of designated points of the individual models, we obtain indeed some kind of amalgamation of the pointed models. In particular, suppose that some formula σ holds of cofinitely many pointed models of the set $\{(\mathcal{M}_i, s_i) \mid i \in I\}$. Then $V(\varphi)$ belongs to u for any point-evening ultrafilter u (where V is the valuation of the disjoint union of the models), so using Proposition 6 we may infer that σ holds of any ultrafilter union of the models.

We are now ready to prove Theorem 2.

Proof of Theorem 2. We only prove the hard direction (from right to left). Suppose that K has the mentioned closure properties and define Δ as the set of formulas holding of any pointed model in K. We claim that Δ defines K.

Following the same proof strategy as before, we take an arbitrary pointed model $(\mathcal{M}, s)$ satisfying $\mathcal{M}, s \Vdash \Delta$, and prove that it belongs to K. We may assume that $\mathcal{M}$ is m-saturated — as in the proof of Theorem 1, the more general case can be reduced to this one, using the fact that the complement of K is closed under taking ultrafilter extensions.

Let $\Sigma = \{\sigma_n \mid n \in \mathbb{N}\}$ be the set of formulas true in $\mathcal{M}$ at s, and define for every n, ψ_n to be the formula $\sigma_0 \wedge \cdots \wedge \sigma_n$. It is easy to see that for every n, the formula ψ_n holds of some pointed model $(\mathcal{N}_n, t_n)$ in K. This means that every formula $\sigma \in \Sigma$ is true at *cofinitely* many of these, so by the argument given above, we conclude that Σ holds of the ultrafilter union of the models. Since K is closed under taking ultrafilter unions, this shows that there is an m-saturated model $\mathcal{N}$ and a point t such that such that $(\mathcal{N}, t)$ belongs to K and Σ is true at t in $\mathcal{N}$. From the fact that both $\mathcal{M}$ and $\mathcal{N}$ are m-saturated it then follows that $\mathcal{M}, s \leftrightarrows \mathcal{N}, t$. But then it is immediate by the closure properties of K that $(\mathcal{M}, s)$ belongs to K. QED

From the proof of Theorem 1 it is seen that indeed, ultrafilter extensions are to modal logic what ultrapowers are to first order logic, and analyzing the proof of Theorem 2 I would conclude that ultrafilter unions play the role of ultraproducts. Perhaps these operations can be put to other use in the model theory of modal logic — certainly in answering the question, when a class of models is definable by a *single* modal formula.

References

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

[2] R. Goldblatt. Saturation and the Hennessy-Milner property. In Ponse et al. [5].

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

[4] M.J. Hollenberg. Hennessy-Milner classes and process algebra. In Ponse et al. [5].

[5] A. Ponse, M. de Rijke, and Y. Venema, editors. *Modal Logic and Process Algebra*, volume 53 of *Lecture Notes*. CSLI Publications, 1995.

[6] M. de Rijke. *Extending Modal Logic*. PhD thesis, ILLC, University of Amsterdam, 1993.