5 Coalgebraic modalities via relation lifting

In this chapter we take an approach to coalgebraic logic which is completely uniform in the type functor T. We introduce a coalgebraic modality ∇ of which the 'arity' is the finitary version T_ω of the functor itself. That is, the set L of formulas will be closed under the following clause:

$$\text{if } \alpha \in TX \text{ for some finite set } X \text{ of formulas, then } \nabla\alpha \text{ is a formula.}$$

whereas the semantics of ∇ will be defined by *lifting* the satisfaction relation $\Vdash$ between states and formulas to the relation $\overline{T}(\Vdash)$.

In the special case where T is the powerset functor P, the nabla operator ∇ is known under the name of the *cover modality*; we discuss this case in some detail before moving on to the more general case.

Convention 5.1 Throughout this chapter we will assume that T is a smooth and standard set functor; that is, T preserves both weak pullbacks and inclusions. The first restriction is to ensure optimal behaviour of the relation lifting $\overline{T}$, while the second one is mainly for convenience. In the Facts B.35 and B.37 we list a number of properties of the operation $\overline{T}$ (all of which will be used throughout this chapter).

Convention 5.2 Throughout this chapter we will discuss T-models over some arbitrary but fixed set Q of proposition letters, each of which will feature as an atomic formula in our modal languages. As announced in Remark 1.20 we will apply all kinds of notions to these models by thinking of them as coalgebras for the functor $T_{\mathsf{Q}} := K_{P\mathsf{Q}} \times T$.

5.1 The cover modality

As we will see now, there is an interesting coalgebraic alternative for the standard formulation of basic modal logic in terms of boxes and diamonds. This alternative set-up is based on a connective ∇, sometimes referred to as the *cover* modality, which turns a *(finite) set* α of formulas into a formula $\nabla\alpha$.

Definition 5.3 Formulas of the language $\mathrm{ML}_\nabla(\mathsf{Q})$ are given by the following recursive definition:

$$a \ ::= \ p \mid \bot \mid \neg a \mid a_0 \vee a_1 \mid \nabla\alpha$$

where $p \in \mathsf{Q}$, and α denotes a finite set of formulas. $\lhd$

Observe that formulas will be denoted by lower case letters $a, b, \ldots$

For the semantics of the cover modality, observe that we may think of the forcing or satisfaction relation $\Vdash$ simply as a binary relation between states and formulas. This relation can thus be lifted to a relation $\overline{P}(\Vdash)$ between *sets* of formulas and *sets* of states.

Definition 5.4 The semantics of this modality in a Kripke model $\mathbb{S} = (S, R, V)$ is given by

$$\mathbb{S}, s \Vdash \nabla\alpha \text{ iff } (R(s), \alpha) \in \overline{P}(\Vdash),$$

where $\overline{P}(\Vdash)$ denotes the Egli-Milner relation lifting of the relation $\Vdash$. $\lhd$

In words: $\nabla\alpha$ holds at s iff every successor of s satisfies some formula in α, and every formula in α holds at some successor of s. The modality ∇ is sometimes called the *cover modality*: it holds at a state s if the set $\{[\![a]\!]^{\mathbb{S}} \mid a \in \alpha\}$ covers the collection $R(s)$ of successors of s, in the sense that $R(s) \subseteq \bigcup\{[\![a]\!]^{\mathbb{S}} \mid a \in \alpha\}$, while at the same time $R(s) \cap [\![a]\!]^{\mathbb{S}} \neq \varnothing$, for every $a \in \alpha$.

Remark 5.5 It is not so hard to see that the cover modality can be defined in the standard modal language:

$$\nabla\alpha \equiv \Box \bigvee \alpha \wedge \bigwedge \Diamond\alpha, \tag{27}$$

where $\Diamond\alpha$ denotes the set $\{\Diamond a \mid a \in \alpha\}$. Things start to get interesting once we realize that both the ordinary diamond $\Diamond$ *and* the ordinary box $\Box$ can be expressed in terms of the cover modality (and the disjunction):

$$\begin{array}{lcl} \Diamond a & \equiv & \nabla\{a, \top\}, \\ \Box a & \equiv & \nabla\varnothing \vee \nabla\{a\}. \end{array} \tag{28}$$

Here, as always, we use the convention that $\bigvee \varnothing = \bot$ and $\bigwedge \varnothing = \top$.

Given that ∇ and $\{\Diamond, \Box\}$ are mutually expressible, we arrive at the following proposition. Here we say that two languages are *effectively equi-expressive* if there are effectively definable truth-preserving translations from one language to the other, and vice versa. Recall that ML is the language of standard modal logic.

Proposition 5.6 *The languages* ML *and* ML_∇ *are effectively equi-expressive.*

A remarkable observation about the cover modality is that we can do far better than this: based on the following *modal distributive law*, we can almost completely eliminate the Boolean connective of conjunction from the language ML_∇.

Proposition 5.7 *Let α and α' be two sets of formulas. Then the following two formulas are equivalent:*

$$\nabla\alpha \wedge \nabla\alpha' \equiv \bigvee_{Z \in \alpha \bowtie \alpha'} \nabla\{a \wedge a' \mid (a, a') \in Z\}, \tag{29}$$

where $\alpha \bowtie \alpha'$ is the set of all binary relations $Z \subseteq \alpha \times \alpha'$ such that $(\alpha, \alpha') \in \overline{P}(Z)$.

Proof. For the direction from left to right, suppose that $\mathbb{S}, s \Vdash \nabla\alpha \wedge \nabla\alpha'$. Let $Z \subseteq \alpha \times \alpha'$ consist of those pairs (a, a') such that the conjunction $a \wedge a'$ is true at some successor t of s. It is then straightforward to verify that Z is full on α and α', that is: $(\alpha, \alpha') \in \overline{P}(Z)$, and that $\mathbb{S}, s \Vdash \nabla\{a \wedge a' \mid (a, a') \in Z\}$.

The converse direction is a fairly direct consequence of the definitions. QED

As a corollary of Proposition 5.7 we can restrict the use of conjunction in modal logic to that of a *special conjunction* connective $\bullet$ which may only be applied to a propositional formula and a ∇-formula.

Definition 5.8 We first define the set $\mathrm{CL}(\mathsf{Q})$ of *literal conjunctions* by the following grammar:

$$\pi \ ::= \ p \mid \neg p \mid \bot \mid \top \mid \pi \wedge \pi,$$

and then let the following grammar define the set $\mathrm{DML}_\nabla(\mathsf{Q})$ of *disjunctive modal formulas* in Q:

$$a \ ::= \ p \mid \neg p \mid \bot \mid \top \mid a_0 \vee a_1 \mid \pi \bullet \nabla\alpha.$$

Here $p \in \mathsf{Q}$, $\pi \in \mathrm{CL}(\mathsf{Q})$ and $\alpha \in P_\omega \mathrm{DML}_\nabla(\mathsf{Q})$. $\lhd$

As mentioned, the bullet connective is semantically equivalent to conjunction:

$$\mathbb{S}, s \Vdash \pi \bullet \nabla\alpha \text{ iff } \mathbb{S}, s \Vdash \pi \text{ and } \mathbb{S}, s \Vdash \nabla\alpha.$$

Note however, that this conjunction is special in the sense that it combines 'local' information about s itself with information about the unfolding of s. The point is that this is the *only* form of conjunction that the language DML_∇ allows.

Theorem 5.9 *The languages* ML *and* DML_∇ *are effectively equi-expressive.*

Proof. We will show how to rewrite a formula $a \in \mathrm{ML}$ into an equivalent formula in DML_∇. Start by rewriting a into *negation normal form*:

$$a \ ::= \ p \mid \neg p \mid \bot \mid \top \mid a_0 \vee a_1 \mid a_0 \wedge a_1 \mid \Diamond a \mid \Box a,$$

then by (28) we can find an equivalent formula a' in the language given by

$$a \ ::= \ p \mid \neg p \mid \bot \mid \top \mid a_0 \vee a_1 \mid a_0 \wedge a_1 \mid \nabla\alpha.$$

Finally, the modal distributive law (29) allows us to push down nabla's to the propositional level, and so using the propositional distributive law $(a \wedge (b_0 \vee b_1) \equiv (a \wedge b_0) \vee (a \wedge b_1))$, we can rewrite a' into an equivalent disjunctive modal formula. QED

Theorem 5.9 can be used to prove various interesting results about modal logic, such as the finite model property, or the decidability of the satisfiability problem — in *linear* time, once the formula is in disjunctive normal form. Rather than proving these corollaries here, we will prove these results in the far more general setting of the *coalgebraic* cover modality.

5.2 Moss' coalgebraic cover modality

We will now generalise the cover modality from the case where $T = P$ to the setting where T is an arbitrary smooth and standard functor. We are eager to keep our language *finitary*, in the sense that formulas will be finitary objects, with for instance finitely many subformulas. For this reason we will work with the *finitary version* of the functor T.

Recall from Convention 5.1 that T preserves inclusions; because of this we may define its *finitary* version $T_\omega : \mathsf{Set} \to \mathsf{Set}$ by putting

$$\begin{array}{lcl} T_\omega(S) & := & \bigcup\{TX \mid X \subseteq_\omega S\}, \\ T_\omega(f : S \to S') & := & (Tf)\restriction_{T_\omega S} \ . \end{array}$$

It is easy to verify that T_ω also preserves inclusions; given the definition of T_ω on functions, we may write Tf instead of $T_\omega f$ without causing confusion. Note, however, that T_ω being finitary does not necessarily mean that $T_\omega S$ is *finite*, even if S is so; a simple counterexample is given by the bag functor.

Definition 5.10 Formulas of the language $\mathrm{ML}_T(\mathsf{Q})$ are given by the following recursive definition:

$$a \ ::= \ p \mid \bot \mid \neg a \mid a_0 \vee a_1 \mid \nabla\alpha$$

where $p \in \mathsf{Q}$, and $\alpha \in T_\omega(\mathrm{ML}_T(\mathsf{Q}))$. We will often write ML_T instead of $\mathrm{ML}_T(\mathsf{Q})$ if the set Q of proposition letters is either understood or irrelevant. $\lhd$

The semantics of ML_T is defined as follows. Recall that a *T-model* is a triple (S, σ, V), where (S, σ) is a T-coalgebra and $V : \mathsf{Q} \to PS$ is a valuation.

Definition 5.11 Let $\mathbb{S} = (S, \sigma, V)$ be a T-model. Then by induction on the complexity of ML_T-formulas we define the satisfaction relation $\Vdash$:

$$\begin{array}{lll}
\mathbb{S}, s \Vdash p & \text{iff} & s \in V(p) \\
\mathbb{S}, s \Vdash \bot & : & \text{never} \\
\mathbb{S}, s \Vdash \neg a & \text{iff} & \mathbb{S}, s \nVdash a \\
\mathbb{S}, s \Vdash a_0 \vee a_1 & \text{iff} & \mathbb{S}, s \Vdash a_0 \text{ or } \mathbb{S}, s \Vdash a_1 \\
\mathbb{S}, s \Vdash \nabla\alpha & \text{iff} & (\sigma(s), \alpha) \in \overline{T}(\Vdash).
\end{array}$$

A formula $a \in \mathrm{ML}_T$ is *satisfiable* iff a is satisfiable in some state of some T-model, and *valid* if its negation is not satisfiable.

Furthermore, we say that two pointed T-models are *ML_T-equivalent* or *(modally) equivalent* if they satisfy the same ML_T-formulas, notation: $\mathbb{S}, s \equiv_{T,\mathsf{Q}} \mathbb{S}', s'$. $\lhd$

Remark 5.12 Before we consider the instantiations of this logic for some set functor T, we argue that the semantics of ML_T is *well defined*. The reader might have some worries about the inductive clause for the ∇ modality, since the definition refers to the lifting of the *full* satisfaction relation.

The point is that because of our assumptions on T, its associated relation lifting $\overline{T}$ commutes with restrictions, cf. Fact B.37. This means that

$$(\sigma(s), \alpha) \in \overline{T}(\Vdash) \text{ iff } (\sigma(s), \alpha) \in \overline{T}(\Vdash\restriction_{S \times X}), \tag{30}$$

where X is any finite set of formulas such that $\alpha \in T_\omega X$. Thus, in order to determine whether $\nabla\alpha$ holds at s or not, we only have to know the interpretation of the formulas used in the justification that $\nabla\alpha$ is a formula. Below we shall see that in fact there is a unique set $\mathit{Base}(\alpha)$ which is the smallest (finite) set X such that $\alpha \in TX$. In other words, we may replace the 'quasi-inductive' clause for ∇ in Definition 5.11 with the following, properly inductive one:

$$\mathbb{S}, s \Vdash \nabla\alpha \text{ iff } (\sigma(s), \alpha) \in \overline{T}(\Vdash\restriction_{S \times \mathit{Base}(\alpha)}). \tag{31}$$

Example 5.13 In this example we look at the interpretation of the coalgebraic cover modality instantiated for various coalgebra types T. That is, let $\mathbb{S} = (S, \sigma, V)$ be a T-model and consider an element $\alpha \in T_\omega(\mathrm{ML}_T)$; below we will explain what it means for the formula $\nabla\alpha$ to hold at s.

(a) In case $T = K_C$ is a constant functor, $\alpha \in T(\mathrm{ML}_T) = C$ is just a colour $\alpha \in C$. In this case we find $\mathbb{S}, s \Vdash \nabla\alpha$ iff $\sigma(s) = \alpha$.

(b) If $T = Id$ is the identity functor, $\alpha \in T(\mathrm{ML}_T) = \mathrm{ML}_T$ is just a formula. We obtain the *next-time operator* of linear temporal logic: $\mathbb{S}, s \Vdash \nabla\alpha$ iff $\mathbb{S}, \sigma(s) \Vdash \alpha$.

(c) For the binary tree functor $T = Id \times Id$, the semantics of nabla is as follows: given $\alpha = (a_0, a_1) \in \mathrm{ML}_T \times \mathrm{ML}_T = T(\mathrm{ML}_T)$, we have $\mathbb{S}, s \Vdash \nabla(a_0, a_1)$ iff $\mathbb{S}, t_0 \Vdash a_0$ and $\mathbb{S}, t_1 \Vdash a_1$, where t_0 and t_1 are the 'left' and 'right' successor of s, respectively, given by $\sigma(s) = (t_0, t_1)$.

(d) For the automata functor $T = 2 \times Id^C$, an element $\alpha \in T(\mathrm{ML}_T)$ is of the form $(i, \overline{a})$ with $i \in \{0, 1\}$ and $\overline{a} = (a_c)_{c \in C}$, with each $a_c \in \mathrm{ML}_T$. With $\sigma = \langle \chi, \tau \rangle$ we have $\mathbb{S}, s \Vdash \nabla(i, \overline{a})$ iff $\chi(s) = i$ and $\mathbb{S}, \tau(s)(c) \Vdash a_c$.

(e) Where $T = P$ is the power set functor, it is easy to verify that ∇_P is the cover modality discussed in the previous section.

Example 5.14 In this example we look at the ∇-logic of the distribution functor D; the nabla modality for the bag functor is defined similarly.

First of all, observe that defined as in Example B.6, D does not preserve inclusions. We can remedy this by taking a variant D' of D which takes a set S to the collection $D'S$ of *partial maps* from S to $(0, 1]$. We will not pursue this road further, but we may use it to observe that an element $\alpha \in D_\omega(\mathrm{ML}_D)$ can be represented as a finite set $\{(a_1, p_n), \ldots, (a_n, p_n)\}$ such that $p_i > 0$ for all i, and $\Sigma_i p_i = 1$.

For the definition of relation lifting $\overline{D}$, consider a relation $R \subseteq X_0 \times X_1$. We claim that the relation $\overline{D}R \subseteq DX_0 \times DX_1$ consists of those pairs (μ_0, μ_1) for which there is a distribution $\rho : R \to [0, 1]$ satisfying the 'magic square conditions':

$$\begin{array}{rl} \text{for all } x_0 \in X_0. & \mu_0(x_0) = \sum_{y_1 \in X_1} \rho(x_0, y_1), \\ \text{and for all } x_1 \in X_1. & \mu_1(x_1) = \sum_{y_0 \in X_0} \rho(y_0, x_1). \end{array}$$

For a concrete example, consider the table below:

	a	b	c	d	e	
1	□	.17	□	.08	.08	.33
2	.10	.03	.12	□	.08	.33
3	□	.00	.18	.12	.04	.34.
	.10	.20	.30	.20	.20	

Here we have $X_0 = \{a, b, c, d, e\}$, $X_1 = \{1, 2, 3\}$ and $R \subseteq X_0 \times X_1$ consists of those pairs of which the corresponding entry in the table is marked by a number (that is, not by □). The distributions μ_0 and μ_1 are given, respectively, in the bottom row and rightmost column of the table. The entries in the table represent a distribution ρ that meets the magic square conditions.

Now let $\mathbb{S} = (S, \sigma, V)$ be a D-model, and consider a formula $\nabla\alpha$ with $\alpha = \{(a_1, p_n), \ldots, (a_n, p_n)\}$. Then we find that

$$\begin{aligned}\mathbb{S}, s \Vdash \nabla\alpha \text{ iff } & \text{ there is a relation } R \subseteq \Vdash \text{ and a map } \rho : R \to [0,1] \text{ such that} \\ & \text{for all } i: \quad \sum_{\{t \mid \mathbb{S}, t \Vdash a_i\}} \rho(t, a_i) = p_i, \\ & \text{and for all } t \in S: \quad \sum_{\{i \mid \mathbb{S}, t \Vdash a_i\}} \rho(t, a_i) = \sigma(s)(t).\end{aligned}$$

5.3 Basic properties of ∇

In this section we prove some of the basic properties of the coalgebraic cover modality. We start with showing that ML_T is a finitary logic indeed, i.e., that every formula has only finitely many subformulas. The key property of finitary functors that will make this possible is that for every $\alpha \in T_\omega A$ there is a *smallest* subset $A' \subseteq A$ that *supports* α, i.e., is such that $\alpha \in T_\omega A'$.

Definition 5.15 Given a finitary functor T and an element $\alpha \in TX$, we define

$$\mathit{Base}^T_X(\alpha) := \bigcap \{Y \subseteq_\omega X \mid \alpha \in TY\}.$$

We write Base^T rather than Base^{T_ω}, and in fact omit the superscript whenever possible. $\lhd$

Example 5.16 The following examples are easy to check: $\mathit{Base}^{Id}_X : X \to P_\omega X$ is the singleton map, $\mathit{Base}^P_X : P_\omega X \to P_\omega X$ is the identity map on $P_\omega X$, $\mathit{Base}^{Id^2}_X : X \times X \to P_\omega X$ maps the pair (x_1, x_2) to the set $\{x_1, x_2\}$, and Base^D maps a finitary distribution μ to its support $\{s \in S \mid \mu(s) > 0\}$.

Proposition 5.17 *Let $T : \mathsf{Set} \to \mathsf{Set}$ be smooth and standard. Then the following hold:*
(1) for any $\alpha \in T_\omega X$, $\mathit{Base}^T_X(\alpha)$ is the smallest set supporting α;
(2) Base^T provides a natural transformation $\mathit{Base} : T_\omega \dot{\to} P_\omega$.

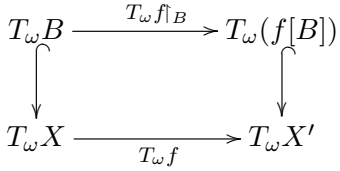
Recall that Base^T being a natural transformation means that the following diagram commutes, for any map $f : X \to Y$:

$$\begin{array}{ccc} X & T_\omega X & \xrightarrow{\mathit{Base}^T_X} P_\omega X \\ {\scriptstyle f}\downarrow & {\scriptstyle Tf}\downarrow & \downarrow{\scriptstyle P_\omega f} \\ X' & T_\omega X' & \xrightarrow[\mathit{Base}^T_{X'}]{} P_\omega X' \end{array} \tag{32}$$

Proof. Part (1) of the Proposition is an immediate consequence of Proposition B.32.

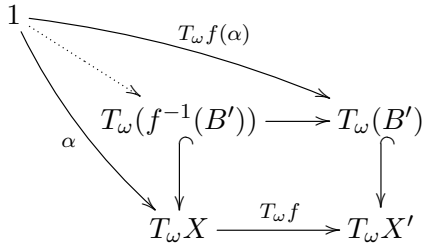
For the second part, consider a map $f : X \to X'$. We have to show $P_\omega f \circ \mathit{Base}_X = \mathit{Base}_{X'} \circ T_\omega f$. Fix $\alpha \in T_\omega X$ and write $B = \mathit{Base}_X(\alpha)$ and $B' = \mathit{Base}_{X'}(T_\omega f(\alpha))$. We need

to prove $B' = f[B]$. For the inclusion $\subseteq$, first note that the following diagram must commute (recall that hooked arrows denote inclusions):



From this it follows that $T_\omega f(\alpha)$ actually belongs to $T_\omega(f[B])$. In other words, $f[B]$ supports $T_\omega f(\alpha)$ and, as B' is by definition the smallest such set, it follows that $B' \subseteq f[B]$. •

For the opposite inclusion $\supseteq$, we consider the diagram below. Here 1 denotes some arbitrary singleton set, and the arrow α denotes any function from 1 to $T_\omega X$ mapping the unique element of 1 to α.



But since T_ω preserves weak pullbacks, the dotted arrow exists and shows that $\alpha \in T_\omega(f^{-1}(B'))$. By minimality of the base, it follows $B \subseteq f^{-1}(B')$, that is, $B' \supseteq f[B]$. QED

By Fact 5.17(1) we may find for any formula $\nabla\alpha$ a *smallest* (and finite) collection X of formulas such that $\alpha \in T_\omega X$, namely, the set $X = \mathit{Base}(\alpha)$. This means that we can define a natural notion of *subformula*.

Definition 5.18 We define the set $\mathtt{Sfor}(a)$ of *subformulas* of a formula $a \in \mathrm{ML}_T$ by the following induction:

$$\begin{array}{lcll} \mathtt{Sfor}(a) & := & \{a\} & \text{if } a \in \{p, \bot\} \\ \mathtt{Sfor}(\neg a) & := & \{\neg a\} \cup \mathtt{Sfor}(a) & \\ \mathtt{Sfor}(a_0 \vee a_1) & := & \{a_0 \vee a_1\} \cup \mathtt{Sfor}(a_0) \cup \mathtt{Sfor}(a_1) & \\ \mathtt{Sfor}(\nabla\alpha) & := & \{\nabla\alpha\} \cup \bigcup\{\mathtt{Sfor}(a) \mid a \in \mathit{Base}(\alpha)\} & \end{array}$$

The elements of $\mathit{Base}(\alpha)$ will be called the *immediate* subformulas of $\nabla\alpha$. ◁

The next properties that we consider are invariance and expressivity. In line with Convention 5.2, when we say that the language ML_T is invariant, we mean that, for each set Q of proposition letters, the language $\mathrm{ML}_T(\mathsf{Q})$ is invariant for the class of T_Q-coalgebras (T-models over Q).

Theorem 5.19 *For any smooth and standard functor T, the language ML_T is invariant: Given any two pointed T-models $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ we have*

$$(\mathbb{S}, s) \simeq_{T_\mathsf{Q}} (\mathbb{S}', s') \text{ implies } (\mathbb{S}, s) \equiv_{T,\mathsf{Q}} (\mathbb{S}', s'). \tag{33}$$

Proof. Given that the smoothness of T implies that of T_{Q}, it suffices to prove that *bisimilarity* implies modal equivalence. Assuming that $Z : \mathbb{S} \leftrightarrows \mathbb{S}'$, we will prove by induction on the complexity of ML_T-formulas that, for every ML_T-formula a:

$$\mathbb{S}, s \Vdash a \text{ iff } \mathbb{S}', s' \Vdash a, \tag{34}$$

for every pair of states $(s, s') \in Z$. Clearly this suffices to prove the proposition.

Skipping the routine parts of the proof (i.e., the base step and boolean cases of the inductive step), we focus on the case where $a = \nabla\alpha$. We only prove the direction from right to left of (34).

So, assume that $(s, s') \in Z$ and $\mathbb{S}', s' \Vdash \nabla\alpha$, and let $\Vdash \subseteq S \times \mathrm{ML}_T$ and $\Vdash' \subseteq S' \times \mathrm{ML}_T$ denote the satisfaction relations on $\mathbb{S}$ and $\mathbb{S}'$, respectively. It follows from $(s, s') \in Z$ that $(\sigma(s), \sigma'(s')) \in \overline{T}Z$, and from $\mathbb{S}', s' \Vdash \nabla\alpha$ that $(\sigma'(s'), \alpha) \in \overline{T}(\Vdash')$; but from the latter fact, together with the observation that $\alpha \in TBase(\alpha)$, we may derive that $(\sigma'(s'), \alpha) \in \overline{T}(\Vdash'{\restriction}_{S \times Base(\alpha)})$ (cf. Fact B.37). Putting these observations together with the fact that $\overline{T}$ preserves relation composition, we find that

$$(\sigma(s), \alpha) \in \overline{T}\big(Z \,;\, \Vdash'{\restriction}_{S \times Base(\alpha)}\big).$$

But by the inductive hypothesis we obtain that $Z \,;\, \Vdash'{\restriction}_{S' \times Base(\alpha)} \subseteq \Vdash$. so that it follows by the monotonicity of relation lifting that $(\sigma(s), \alpha) \in \overline{T}(\Vdash)$. From this it is immediate by the semantics of ∇ that $\mathbb{S}, s \Vdash \nabla\alpha$, as required. QED

As could be expected, the converse of this proposition only holds if we restrict attention to *image-finite* coalgebras.

Definition 5.20 A T-coalgebra $\mathbb{S} = (S, \sigma)$ is *image-finite* if $\sigma(s) \in T_\omega S$, for all $s \in S$. $\lhd$

Theorem 5.21 *For any smooth and standard functor T, the language ML_T is expressive on the class of image-finite T-models: Given any two pointed T_ω-models $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ we have*

$$(\mathbb{S}, s) \equiv_{T,\mathsf{Q}} (\mathbb{S}', s') \text{ implies } (\mathbb{S}, s) \simeq_{T_{\mathsf{Q}}} (\mathbb{S}', s'). \tag{35}$$

Proof. It suffices to show that the relation of modal equivalence is itself a *bisimulation*, when restricted to the class of image-finite coalgebra models.

Fix two T-models $\mathbb{S} = (S, \sigma, V)$ and $\mathbb{S}' = (S', \sigma', V')$, and let $\equiv \subseteq S \times S'$ denote the relation of modal equivalence between $\mathbb{S}$ and $\mathbb{S}'$. (That is, we avoid notational clutter and write $\equiv$ instead of $\equiv_{T,\mathsf{Q}}$.) Leaving the straightforward propositional part as an exercise for the reader, We will use Theorem **??** in order to prove that $\equiv$ is a bisimulation, and suppose for contradiction that $s \equiv s'$. Our goal will be to show that

$$(\sigma s, \sigma' s') \in \overline{T}(\equiv). \tag{36}$$

First of all, it follows by image-finiteness that we may define the (finite) sets $B := Base(\sigma s)$ and $B' := Base(\sigma' s')$. Since $\overline{T}$ commutes with restrictions it suffices to prove that

$$(\sigma s, \sigma' s') \in \overline{T}(\equiv {\restriction}_{B \times B'}). \tag{37}$$

Furthermore, since B and B' are finite, for every $t \in B$ we may find a formula c_t such that, for all $t' \in B'$,

$$t \equiv t' \text{ iff } \mathbb{S}', t' \Vdash c_t,$$

as the reader can easily verify. But with $H := \{c_t \mid t \in B\}$, we may think of c as a surjection $c : B \twoheadrightarrow H$ satisfying $\mathsf{Gr}(c) \subseteq \Vdash$ and

$$\equiv \restriction_{B \times B'} = \mathsf{Gr}(c) \,;\, (\Vdash_{B' \times H})^{\smile}.$$

From this we may conclude, using various properties of relation lifting, that

$$\overline{T}(\equiv \restriction_{B \times B'}) = \mathsf{Gr}(Tc) \,;\, (\overline{T}\Vdash_{B' \times H})^{\smile}. \tag{38}$$

Now let

$$\gamma := (Tc)(\sigma s),$$

then by definition we have $(\sigma s, \gamma) \in \mathsf{Gr}(Tc)$. Furthermore, from $\mathsf{Gr}(c) \subseteq \Vdash$ we derive $\mathsf{Gr}(Tc) \subseteq \overline{T}\Vdash$, so that we find $(\sigma s, \gamma) \in \overline{T}\Vdash$. This implies that the formula $\nabla\gamma$ is *true* at s, and since $s \equiv s'$, it is also true at s'. By the semantics of the ∇ modality this means that $(\sigma s', \gamma) \in \overline{T}\Vdash$. Using various properties of relation lifting we find that $(\gamma, \sigma s') \in (\overline{T}\Vdash_{B' \times H})^{\smile}$. We may then conclude that $(\sigma s, \sigma s') \in \mathsf{Gr}(Tc) \,;\, (\overline{T}\Vdash_{B' \times H})^{\smile}$, which by (38) suffices to prove (37), and thus (36). In other words, we are done. QED

The last basic property that we mention is that of *satisfiability reduction*.

Proposition 5.22 *Let $\nabla\alpha$ be a formula in ML_T. Then $\nabla\alpha$ is satisfiable iff every $a \in \mathit{Base}(\alpha)$ is satisfiable.*

Proof. For the direction from right to left, assume that every $a \in \mathit{Base}(\alpha)$ is satisfiable. That is, assume that for every $a \in \mathit{Base}(\alpha)$ there is a pointed model $(\mathbb{S}_a, s_a)$, with $\mathbb{S}_a = (S_a, \sigma_a, V_a)$ and such that $\mathbb{S}, s_a \Vdash a$, for each a. We define a new model $\mathbb{S} = (S, \sigma, V)$, where $S := \{r\} \uplus \biguplus\{S_a \mid a \in \mathit{Base}(\alpha)\}$. For the valuation V we simply define $V(p) := \bigcup_a V_a(p)$, while on an element $s \in S_a$ the coalgebra map $\sigma : S \to TS$ is defined by putting $\sigma(s) := \sigma_a(s)$. For the definition of the unfolding $\sigma(r)$ of the 'root' r, consider the map $f : \mathit{Base}(\alpha) \to S$ given by $a \mapsto s_a$, and simply put $\sigma(r) := (Tf)(\alpha)$. It is then immediate by the definition of f that $(\mathsf{Gr}f)^{\smile} \subseteq \Vdash$, so that we find, using various properties of relation lifting (cf. Fact B.35 and B.37):

$$(\sigma r, \alpha) \in (\mathsf{Gr}(Tf))^{\smile} = \overline{T}((\mathsf{Gr}f)^{\smile}) \subseteq \overline{T}(\Vdash),$$

from which it follows that $\mathbb{S}, r \Vdash \nabla\alpha$ indeed.

For the opposite direction, we need the following little fact about the functor T:

$$\text{for any } f : A \to B \text{ and any } \alpha \in TA \text{ we have } (Tf)(\alpha) \in T(f[A]). \tag{39}$$

To see why (39) holds, factorize f as the unique composition $f = \iota \circ f'$ of a surjection $f' : A \twoheadrightarrow f[A]$ and an inclusion $\iota : f[A] \hookrightarrow B$. From this factorization it follows that $Tf = (T\iota) \circ (Tf')$, where $f' : TA \twoheadrightarrow Tf[A]$ is surjective since every set functor preserves surjections,

and $T\iota : Tf[A] \hookrightarrow TB$ is an inclusion by assumption on T. From these observations (39) is easy to derive.

Now assume that $\nabla\alpha$ is satisfiable, then there is a T-model $\mathbb{S} = (S, \sigma, V)$ such that $\mathbb{S}, s \Vdash \nabla\alpha$ for some state $s \in S$. Then by definition of the semantics of ∇ we have that $(\sigma s, \alpha) \in \overline{T}(\Vdash)$, and so by definition of $\overline{T}$ there is an object $\rho \in T(\Vdash)$ such that $T\pi_0(\rho) = \sigma(s)$ and $T\pi_1(\rho) = \alpha$, where $\pi_0 : \Vdash \to S$ and $\pi_1 : \Vdash \to \mathrm{ML}_T$ are the projection functions on the relation $\Vdash$. But then it follows by (39) that $\alpha \in T(\pi_1[\Vdash]) = T(\mathsf{Ran}(\Vdash))$, so that $Base(\alpha) \subseteq \mathsf{Ran}(\Vdash)$. In other words, for every $a \in Base(\alpha)$ there is an $s \in S$ where a holds. In particular, this means that every $a \in Base(\alpha)$ is satisfiable. QED

5.4 Coalgebraic modal distributive laws

In this section we will formulate three *coalgebraic modal distributive laws (CMDLs)* describing the interaction between the coalgebraic modality ∇ on the one hand, and the boolean operations on the other. For a concise formulation of these principles it will be convenient to slightly rearrange the coalgebraic modal language, working with the finitary versions $\bigwedge$ and $\bigvee$ of the binary connectives for conjunction and disjunction. That is, in this section we will be working with the following variant of the language.

Definition 5.23 The language L_T is given by the following grammar:

$$a \ ::= \ p \mid \neg a \mid \bigwedge A \mid \bigvee A \mid \nabla\alpha$$

where $p \in \mathsf{Q}$, $A \in P_\omega L_T$ and $\alpha \in T_\omega L_T$. $\lhd$

The semantics of $\bigwedge$ and $\bigvee$ is as expected:

$$\begin{array}{lll} \mathbb{S}, s \Vdash \bigwedge A & \text{iff} & \mathbb{S}, s \Vdash a \text{ for all } a \in A \\ \mathbb{S}, s \Vdash \bigvee A & \text{iff} & \mathbb{S}, s \Vdash a \text{ for some } a \in A \end{array}$$

In particular, this means that instead of taking $\top$ and $\bot$ to be primitive symbols, we may consider them as abbreviations:

$$\begin{array}{rcl} \top & := & \bigwedge \varnothing \\ \bot & := & \bigvee \varnothing. \end{array}$$

A key aspect of the formulation of the CMDLs is the observation that we may think of the connectives $\bigwedge, \bigvee$ and $\neg$ as *maps* of the respective types $\bigwedge, \bigvee : P_\omega L_T \to L_T$ and $\neg : L_T \to L_T$. In particular, this perspective allows us to apply the functor T to these connectives, obtaining maps $T\bigwedge, T\bigvee : T_\omega P_\omega L_T \to T_\omega L_T$, and $T\neg : T_\omega L_T \to T_\omega L_T$. Thus, for any object $\Phi \in T_\omega P_\omega L_T$ we find $(T\bigvee)\Phi \in T_\omega L_T$, which means that $\nabla(T\bigvee)\Phi$ is a well-formed formula.

Convention 5.24 Since we will be dealing here with formulas and similar objects in various, closely related sets, incuding $\mathsf{Q}, L_T, T_\omega L_T, P_\omega L_T, P_\omega T_\omega L_T$ and $T_\omega P_\omega L_T$, it will be convenient to use some kind of *naming convention*, see Table 1.

In order to formulate the modal distributive laws we need some auxiliary definitions.

Set	Elements
Q	$p, q, \ldots$
L_T	$a, b, \ldots$
$T_\omega L_T$	$\alpha, \beta, \ldots$
$P_\omega L_T$	$A, B, \ldots$
$P_\omega T_\omega L_T$	$\Gamma, \Delta, \ldots$
$T_\omega P_\omega L_T$	$\Phi, \Psi, \ldots$

Table 1: Naming convention

Definition 5.25 Given a smooth and standard set functor T, we define, for every set X, a function $\lambda^T_X : TPX \to PTX$ by putting

$$\lambda^T_X(\Phi) := \{\alpha \in TX \mid (\alpha, \Phi) \in \overline{T}(\in_X)\}, \tag{40}$$

where $\in_X$ denotes the membership relation $\in$, restricted to $X \times PX$. Elements of $\lambda^T_X(\Phi)$ will be referred to as *lifted members* of Φ. The family $\lambda^T = \{\lambda^T_X\}_{X \in \mathsf{Set}}$ will be called the *T-transformation*.

A set $\Phi \in TPX$ is a *redistribution* of a set $\Gamma \in PTX$ if $\Gamma \subseteq \lambda^T_X(\Phi)$, that is, every element of Γ is a lifted member of Φ. In case $\Gamma \in P_\omega T_\omega X$, we call a redistribution Φ *slim* if $\Phi \in T_\omega P_\omega(\bigcup_{\gamma \in \Gamma} Base(\gamma))$. The set of slim redistributions of Γ is denoted as $SRD(\Gamma)$. $\lhd$

Remark 5.26 Properties of $\overline{T}$ are intimately related to those of λ^T. In order to express the connection, we need to introduce the concept of a (categorical) distributive law.

Let T be a covariant set functor. A *distributive law* of T over a (co- or contravariant) set functor M is a natural transformation $\theta : TM \dot{\to} MT$; that is, the following diagram commutes, for every map $f : X \to Y$:

$$\begin{array}{ccc} X & TMX & \xrightarrow{\theta_X} MTX \\ {\scriptstyle f}\downarrow & {\scriptstyle TMf}\downarrow & \downarrow{\scriptstyle MTf} \\ Y & TMY & \xrightarrow{\theta_Y} MTY \end{array}$$

(In case M is a covariant functor the downward arrows have to be reversed.)

Now observe that we can in fact *define* the family of functions $\lambda^T = \{\lambda^T_X\}_{X \in \mathsf{Set}}$ as in (40) for any set function T. If T preserves weak pullbacks, $\lambda^T = \{\lambda^T_X\}_{X \in \mathsf{Set}}$ is a distributive law of T over the power set functor P.

Definition 5.27 Let T be a smooth and standard set functor which restricts to finite sets. Consider the following *coalgebraic modal distributive laws*:

($DL_\bigvee$) $\nabla(T\bigvee)(\Phi) \equiv \bigvee \{\nabla\alpha \mid (\alpha, \Phi) \in \overline{T}(\in_{L_T})\}$

($DL_\bigwedge$) $\bigwedge \{\nabla\gamma \mid \gamma \in \Gamma\} \equiv \bigvee \{\nabla(T\bigwedge)(\Phi) \mid \Phi \in SRD(\Gamma)\}$

($DL_\neg$) $\neg\nabla\alpha \equiv \bigvee \{\nabla T(\bigwedge \circ P\neg)\Psi \mid \Psi \in T_\omega P_\omega Base(\alpha)$ and $(\alpha, \Psi) \notin \overline{T}(\notin)\}$ $\lhd$

Note that the restriction to finite sets is needed in order to keep the disjunctions on the right hand side of the equivalences finite

Proposition 5.28 *Let T be a smooth and standard set functor which restricts to finite sets. All three coalgebraic modal distributive laws are valid.*

Proof. In order to understand the validity of these laws, fix some T-model $\mathbb{S} = (S, \sigma, V)$.

We first consider $(DL_{\bigvee})$, proving the direction from left to right. First observe that for any $A \subseteq_\omega L_T$ we have $\mathbb{S}, s \Vdash \bigvee A$ iff $\mathbb{S}, s \Vdash a$, for some $a \in A$. Putting it differently, the relations $\Vdash ; \in$ and $\Vdash ; \bigvee^\smile$ coincide[9]. From this it follows that

$$\overline{T}(\Vdash ; \in) \;=\; \overline{T}(\Vdash ; \textstyle\bigvee^\smile). \tag{41}$$

Now fix some object $\Phi \in T_\omega P_\omega L$, and suppose that s is a state in $\mathbb{S}$ such that $s \Vdash \nabla(T\bigvee)\Phi$. By the truth definition, the pair $(\sigma(s), (T\bigvee)(\Phi))$ belongs to the relation $\overline{T}(\Vdash)$, and so $(\sigma(s), \Phi)$ belongs to $(\overline{T}\Vdash);(T\bigvee)^\smile = \overline{T}(\Vdash;\bigvee^\smile)$. But then by (41), we find $(\sigma(s), \Phi) \in \overline{T}(\Vdash;\in) = \overline{T}\Vdash;\overline{T}\in$. In other words, there is some object β such that $(\sigma(s), \beta) \in \overline{T}(\Vdash)$ and $(\beta, \Phi) \in \overline{T}(\in)$. Clearly then $s \Vdash \nabla\beta$, and so we have $s \Vdash \bigvee\{\nabla\beta \mid \beta\, \overline{T}\in\, \Phi\}$, as required.

For the validity of $(DL_{\bigwedge})$, we also confine attention to the direction from left to right. Assume that $\mathbb{S}, s \Vdash \nabla\gamma$ for all $\gamma \in \Gamma$. We need to come up with some slim redistribution Φ of Γ such that $\mathbb{S}, s \Vdash \nabla(T\bigwedge)\Phi$. For this purpose we associate, with any state $t \in S$, the finite set

$$A_t := \{a \in \bigcup_{\gamma\in\Gamma} \mathit{Base}(\gamma) \mid \mathbb{S}, t \Vdash a\}.$$

Taking A to be the map $A : S \to P_\omega L_T$, we may define $\Phi := (TA)(\sigma(s)) \in T_\omega P_\omega L_T$.

First we show that $\mathbb{S}, s \Vdash \nabla(T\bigwedge)\Phi$. Observe that by definition of the map $A : S \to P_\omega L_T$, the function $\bigwedge \circ A : S \to L_T$ is such that

$$\mathsf{Gr}(\textstyle\bigwedge \circ A) \subseteq \Vdash.$$

From this we obtain

$$\mathsf{Gr}((T\textstyle\bigwedge) \circ (TA)) \subseteq \overline{T}(\Vdash)$$

by the properties of the operation $\overline{T}$. But that means that for every element $\tau \in TS$, we have that $\big(\tau, ((T\bigwedge) \circ (TA))(\tau)\big) \in \overline{T}\Vdash$. In particular, we find that $(\sigma s, (T\bigwedge)\Phi) = (\sigma s, (T\bigwedge)(TA)(\sigma(s)) \in \overline{T}\Vdash$, showing that $\mathbb{S}, s \Vdash \nabla(T\bigwedge)\Phi$ as required.

It is left to prove that Φ is a slim redistribution of Γ. Observe that by definition of the map A, we have that

$$\mathsf{Gr}(A) ; \in^\smile = \Vdash\restriction_{S\times B},$$

where $B := \bigcup_{\gamma\in\Gamma} \mathit{Base}(\gamma)$. From this it follows by the properties of relation lifting that

$$\mathsf{Gr}(TA) ; (\overline{T}\in)^\smile = \overline{T}(\Vdash)\restriction_{TS\times TB}.$$

[9] Here we write $\bigvee$ instead of $\mathsf{Gr}(\bigvee)$

But then for each $\gamma \in \Gamma$ we may derive from the fact that $(\sigma s, \gamma) \in \overline{T}(\Vdash) \restriction_{TS \times TB}$ that there is some object Ψ such that $(\sigma s, \Psi) \in \mathsf{Gr}(TA)$ and $(\Psi, \gamma) \in (\overline{T}\in)^{\smile}$. It then easily follows that $\Psi = (TA)(\sigma s) = \Phi$ and so $(\gamma, \Phi) = (\gamma, \Psi) \in \overline{T}(\in)$. In other words, each $\gamma \in \Gamma$ is a lifted member of Φ, and so Φ is a redistribution of Γ; but then by its definition it is slim.

Finally, the validity of $(DL_\neg)$ is left as an exercise to the reader. QED

5.5 Coalgebraic Logic

We will now see that the coalgebraic modal distributive laws that we proved in the previous section are in fact quite strong principles, with important applications.

We start with the coalgebraic generalisation of the disjunctive normal form result on the cover modality, Theorem 5.9.

Definition 5.29 We let the following grammar:

$$a \ ::= \ p \mid \neg p \mid \bot \mid \top \mid a_0 \vee a_1 \mid \pi \bullet \nabla \alpha.$$

define the set $\mathrm{DML}_T(\mathsf{Q})$ of *disjunctive T-modal formulas* in Q. Here we have $p \in \mathsf{Q}$ and π is a finite conjunction of literals over Q. $\lhd$

The proof of the following theorem is completely analogous to that of Theorem 5.9.

Theorem 5.30 *Let T be a smooth and standard set functor which restricts to finite sets. The languages ML_T and DML_T are effectively equi-expressive.*

For the following result recall that a modal logic $(L, \Vdash)$ has the *finite model property* if every satisfiable L-formula is satisfiable in a finite model.

Theorem 5.31 *Let T be a smooth and standard set functor which restricts to finite sets. Then ML_T has the finite model property.*

Proof. By Theorem 5.30 it suffices to prove the finite model property for *disjunctive* formulas. We leave it as an exercise for the reader to establish this result — this goes by a straightforward proof by induction on the complexity of DML_T-formulas, of which the inductive case for the ∇ modality uses the observation underlying the proof of our satisfiability reduction result, Proposition 5.22. QED

Remark 5.32 Theorem 5.30 can also be used to obtain *decidability* results for logics ML_T. For instance, it can be proved that the satisfiability problem for the language $\mathrm{ML}_P = \mathrm{ML}_\nabla$ can be solved in *linear time*. However, since these results depend on the functor, or more specifically: on the representation of formulas of the form $\nabla\alpha$, we refrain from going into detail here.

Remark 5.33 As another corollary of Theorem 5.30 we can show that for any smooth and standard set functor T which restricts to finite sets, the logic ML has *uniform interpolation*, a strong version of Craig's interpolation property. $\bullet$

Finally, we briefly mention a sound and complete derivation system for the set of valid ML-formulas.

Definition 5.34 Let T be a smooth and standard set functor which restricts to finite sets. For the *derivation system* **M**, we start with fixing an arbitrary sound and complete set of axioms and rules for classical propositional logic; we extend this with the following derivation rule:

$$\frac{\{a \to b \mid (a,b) \in Z\}}{\nabla\alpha \to \nabla\beta}\ (\alpha,\beta) \in \overline{T}Z,$$

together with the one-sided versions of the coalgebraic modal distributive laws:

$(A_{\bigvee})$ $\nabla(T\bigvee)(\Phi) \to \bigvee\{\nabla\alpha \mid (\alpha,\Phi) \in \overline{T}(\in_X)\}$

$(A_{\bigwedge})$ $\bigwedge\{\nabla\gamma \mid \gamma \in \Gamma\} \to \bigvee\{\nabla(T\bigwedge)(\Phi) \mid \Phi \in SRD(\Gamma)\}$

$(A_{\neg})$ $\neg\nabla\alpha \to \bigvee\{\nabla T(\bigwedge \circ P\neg)\Psi \mid \Psi \in T_\omega P_\omega Base(\alpha) \text{ and } (\alpha,\Psi) \notin \overline{T}(\notin)\}$ ◁

▶ `Formulate completeness`

6 Coalgebraic modalities via predicate liftings

In this chapter we take an approach to coalgebraic modal logic where the modalities are in 1-1 correspondence with so-called *predicate liftings* for the functor T. That is, with each set Λ of such predicate liftings we will associate a modal formalism ML_Λ for T-coalgebras. As a result this set-up is not completely uniform in the coalgebra type T, but it has some advantages over the approach based on relation lifting. First of all, the language of ML_Λ is completely standard, with a syntax that adds to the language of propositional logic an n-ary modality $\heartsuit_\lambda$ for each n-ary predicate lifting $\lambda \in \Lambda$. Second, there is no reason to restrict attention to functors that are smooth (preserve weak pullbacks). And finally, predicate liftings provide a uniform framework to many well-known variants of standard modal logic (including monotone and probabilistic modal logic, which were already mentioned in section 1.5).

Before we introduce the approach in full generality, we outline our convention on the use of proposition letters, and we briefly discuss a few other concrete variants of standard modal logic that are covered by the approach.

Convention 6.1 As in the previous chapter we will consider T-models over some fixed set Q of proposition letters, and as announced in Remark 1.20 we will apply all kinds of notions to these models by thinking of them as coalgebras for the functor $T_\mathsf{Q} := K_{P\mathsf{Q}} \times T$.

6.1 Variants of modal logic

Example 6.2 (1) The *next-time operator* $\bigcirc$ of *linear time logic* is perhaps the most simple example. For its definition, consider models of the form (ω, V), where $V : \mathsf{Q} \to P(\omega)$ is a valuation on the set ω of natural numbers; the modality $\bigcirc$ is interpreted as follows:

$$\omega, V, n \Vdash \bigcirc\varphi \text{ iff } \omega, V, n+1 \Vdash \varphi.$$

Clearly the semantics of this operator can be generalised to arbitrary T-models for the identity functor $T = Id$.

(2) Similarly, on the binary tree 2^ω we can interpet two modalities $\bigcirc_0$ and $\bigcirc_1$, with the following interpretation:

$$2^\omega, V, u \Vdash \bigcirc_i\varphi \text{ iff } 2^\omega, V, u \cdot i \Vdash \varphi,$$

where $i \in \{0,1\}$ and $V : \mathsf{Q} \to P(2^\omega)$ is a valuation on the set of finite words over the alphabet $2 = \{0,1\}$.

The semantics of these operators can be generalised to arbitrary models for the binary tree functor $T = Id \times Id$.

(3) Graded modal logic is a version of modal logic that allows statements about the *number* of successors that satisfy a certain formula. Formally, interpreted in Kripke models, the modality $\Diamond_{\geq k}$ has the following semantics:

$$\mathbb{S}, n \Vdash \Diamond_{\geq k}\varphi \text{ iff } s \text{ has at least } k \text{ } \varphi\text{-successors},$$

where a φ-successor of s is a state $t \in R(s)$ where φ holds. If we restrict attention to image-finite Kripke models, it also makes sense to introduce the following 'majority modality' M:

$$\mathbb{S}, n \Vdash M(\varphi, \psi) \text{ iff } s \text{ has more } \varphi\text{-successors than } \psi\text{-successors.}$$

Note that these modalities are not bisimulation invariant if we consider Kripke frames as coalgebras for the powerset functor. However, as we will see below, we may also see Kripke frames as coalgebras for the *bag* functor B (see the appendix for its definition), and for that functor both modalities will turn out to be invariant.

As we will see in this section, the common semantic pattern in many of these formalisms can be captured rather nicely in a coalgebraic framework by the notion of a *predicate lifting*.

6.2 Modalities via predicate liftings

To introduce the notion of a predicate lifting, we consider the example of probabilistic modal logic. In Example 1.8 we defined the semantics of the modality $\Diamond_q$ (with q a rational number in $[0, 1]$) in a model $\mathbb{S} = (S, \sigma, V)$ for the distribution functor D, as follows:

$$\mathbb{S}, s \Vdash \Diamond_q \varphi \text{ iff } \sum_{u \in [\![\varphi]\!]^{\mathbb{S}}} \sigma(s)(u) > q, \tag{42}$$

where we recall that $[\![\varphi]\!]^{\mathbb{S}}$ denotes the *extension* of φ in $\mathbb{S}$, that is, the set $[\![\varphi]\!]^{\mathbb{S}} = \{t \in S \mid \mathbb{S}, t \Vdash \varphi\}$ of states in $\mathbb{S}$ where φ is true. The way that we will be thinking of this definition now is as

$$\mathbb{S}, s \Vdash \Diamond_q \varphi \text{ iff } \sigma(s) \in \big\{\mu \in D(S) \mid \sum_{u \in [\![\varphi]\!]^{\mathbb{S}}} \mu(u) > q\big\}, \tag{43}$$

or, in fact, as

$$\mathbb{S}, s \Vdash \Diamond_q \varphi \text{ iff } \sigma(s) \in \theta^q_S([\![\varphi]\!]^{\mathbb{S}}), \tag{44}$$

where $\theta^q_S : PS \to PDS$ is defined by

$$\theta^q_S : U \mapsto \{\mu \in D(S) \mid \sum_{u \in U} \mu(u) > q\}.$$

In other words, we may think of the semantics of the modality $\Diamond_q$ as being indexed by a family θ^q of maps $\theta^q_S : PS \to PDS$, where each θ^q_S *lifts* a predicate on S (i.e., a subset of S) to a predicate on DS.

Now in principle we may associate a modality with each such family θ. However, as we will see below, it will make a lot of sense to impose the following uniformity condition on the family of maps: We will require that, for each map $f : S' \to S$, the following diagram commutes:

$$\begin{array}{ccc} S' & PS' & \xrightarrow{\theta_{S'}} PDS' \\ {\scriptstyle f}\downarrow & \uparrow{\scriptstyle \check{P}f} & \uparrow{\scriptstyle \check{P}Df} \\ S & PS & \xrightarrow[\theta_S]{} PDS \end{array}$$

That is, we will require a 'proper' predicate lifting for the distribution functor to be a *natural transformation* $\theta : \breve{P} \dot{\to} \breve{P}D$, where $\breve{P}$ is the *contravariant* powerset functor. In general, for an arbitrary set functor T we introduce the concept of a *predicate lifting* of some arbitrary but fixed finite arity, as follows.

Definition 6.3 A *predicate lifting* is a natural transformation of the form $\lambda : \breve{P}^n \dot{\to} \breve{P}T$, for some number $n \in \omega$ which we shall refer to as the *arity* of λ, notation: $n = \mathsf{ar}\lambda$. $\lhd$

Recall that the *naturality condition* on predicate liftings means that the following diagram commutes, for every function $f : S' \to S$:

$$\begin{array}{ccc} S' & (PS')^n \xrightarrow{\lambda_{S'}} & PTS' \\ {\scriptstyle f}\downarrow & {\scriptstyle (\breve{P}f)^n}\uparrow & \uparrow{\scriptstyle \breve{P}Tf} \\ S & (PS)^n \xrightarrow[\lambda_S]{} & PTS \end{array} \tag{45}$$

The idea now is that with each predicate lifting λ we may associate a modality $\heartsuit_\lambda$, of the same arity as λ, and that the semantics of $\heartsuit_\lambda$ in a coalgebra (S, σ) is defined in terms of the predicate lifting λ itself (and the coalgebra map σ). The logics that we about to introduce are thus parametrised by a collection of predicate liftings; such collections we shall refer to as *modal signatures*.

Definition 6.4 A *modal signature* for a set functor T is nothing but a set of predicate liftings for T. Given such a collection Λ, and a set Q of proposition letters, the formulas of the modal logic $\mathrm{ML}_\Lambda(\mathsf{Q})$ are given by the following grammar:

$$\varphi \ ::= \ p \mid \bot \mid \neg\varphi \mid \varphi_0 \vee \varphi_1 \mid \heartsuit_\lambda(\varphi_0, \ldots, \varphi_{n-1}),$$

where $p \in \mathsf{Q}$, and λ is an n-ary predicate lifting for T. $\lhd$

We will also use the connectives $\wedge$, $\to$ and $\leftrightarrow$ as the standard abbreviations.

The semantics of the languages ML_Λ is defined in a uniform way, with the modality $\heartsuit_\lambda$ being interpreted 'by λ itself'.

Definition 6.5 Let $\mathbb{S} = (S, \sigma, V)$ be a T-model for some set functor T. We define the satisfaction relation $\Vdash_{\mathbb{S}} \subseteq S \times \mathrm{ML}_\Lambda(\mathsf{Q})$ by induction on the complexity of $\mathrm{ML}_\Lambda(\mathsf{Q})$-formulas. With all other clauses of this definition being standard, we only mention the clause for the coalgebraic modalities:

$$\mathbb{S}, s \Vdash \heartsuit_\lambda(\varphi_0, \ldots, \varphi_{n-1}) \text{ iff } \sigma(s) \in \lambda([\![\varphi_0]\!]^{\mathbb{S}}, \ldots, [\![\varphi_{n-1}]\!]^{\mathbb{S}}). \tag{46}$$

The notions of satisfiability, validity, and ML_Λ-equivalence are all defined in the obvious way • (cf. Definition 5.11 for the ∇-logic). We say that two pointed T-models are Λ, Q-*equivalent* if they satisfy the same $\mathrm{ML}_T(\mathsf{Q})$-formulas, notation: $\mathbb{S}, s \equiv_{\Lambda,\mathsf{Q}} \mathbb{S}', s'$. $\lhd$

Remark 6.6 A succinct way of defining the semantics of the modality $\heartsuit_\lambda$ (46) is as follows:

$$[\![\heartsuit_\lambda(\varphi_0,\ldots,\varphi_{n-1})]\!]^{\mathbb{S}} := (\breve{P}\sigma)\big(\lambda([\![\varphi_0]\!]^{\mathbb{S}},\ldots,[\![\varphi_{n-1}]\!]^{\mathbb{S}})\big). \tag{47}$$

Example 6.7 (1) The box and diamond modalities of standard modal logic can be seen as the coalgebraic modalities associated with the unary predicate liftings $\lambda^\Box, \lambda^\Diamond : \breve{P} \dot{\to} \breve{P}P$ given by

$$\begin{array}{ll} \lambda^\Box_S : & U \mapsto \{X \in PS \mid X \subseteq U\}, \\ \lambda^\Diamond_S : & U \mapsto \{X \in PS \mid X \cap U \neq \varnothing\}. \end{array}$$

We quickly verify that $\lambda^\Box$ satisfies the naturality condition (45). In order to show that $\lambda^\Box_{S'} \circ \breve{P}f = (\breve{P}Pf) \circ \lambda^\Box_S$, it suffices to show that the following identies hold, for every $U \in PS$:

$$\begin{array}{rcll} \lambda^\Box_{S'}(\breve{P}f(U)) & = & \{X' \in PS' \mid X' \subseteq \breve{P}f(U)\} & (\text{definition } \lambda^\Box_{S'}) \\ & = & \{X' \in PS' \mid fx' \in U, \text{ all } x' \in X'\} & (\text{obvious}) \\ & = & \{X' \in PS' \mid Pf(X') \subseteq U\} & (\text{obvious}) \\ & = & \{X' \in PS' \mid Pf(X') \in \lambda^\Box_S(U)\} & (\text{definition } \lambda^\Box_{S'}) \\ & = & (\breve{P}Pf)(\lambda^\Box_S(U)). & (\text{definition } \breve{P}Pf) \end{array}$$

(2) In the case of monotone modal logic, the box and diamond modalities are induced by the following predicate liftings $\mu^\Box, \mu^\Diamond : \breve{P} \dot{\to} \breve{P}M$:

$$\begin{array}{ll} \mu^\Box_S : & U \mapsto \{\sigma \in MS \mid U \in \sigma\}, \\ \mu^\Diamond_S : & U \mapsto \{\sigma \in MS \mid (S \setminus U) \notin \sigma\}. \end{array}$$

(3) The next-time operator $\bigcirc$ of linear temporal logic is obtained from the *identity lifting* $\lambda^\bigcirc : \breve{P} \to \breve{P}$:

$$\lambda^\bigcirc_S : U \mapsto U.$$

Example 6.8 Similarly to the case of probabilistic modal logic, for each $k \in \mathbb{N}$ we can define a predicate lifting $\lambda^{\geq k} : \breve{P} \to \breve{P}B$ for the bag functor B:

$$\lambda^{\geq k}_S : U \mapsto \{\mu : S \to \mathbb{N}^\infty \mid \sum_{u \in U} \mu(u) \geq k\}.$$

Now consider a Kripke model $\mathbb{S} = (S, \sigma, V)$, with $\sigma : S \to PS$. We may think of σ as a coalgebra map $\sigma^\circ : S \to BS$ for the functor B, by putting

$$\sigma^\circ(s)(t) := \begin{cases} 1 & \text{if } t \in \sigma(s), \\ 0 & \text{if } t \notin \sigma(s). \end{cases}$$

It is then straightforward to verify that for any Kripke model $\mathbb{S}$ we have

$$\mathbb{S}, n \Vdash \heartsuit_{\lambda^{\geq k}}\varphi \text{ iff } s \text{ has at least } k \text{ } \varphi\text{-successors},$$

so that we can indeed think of graded modal logic as a coalgebraic logic.

Similarly, the 'majority modality' M of Example 6.2(3) can be seen as the coalgebraic modality that is induced by the *binary* predicate lifting $\lambda^M : \breve{P}^2 \dot{\to} \breve{P}B_\omega$ given by

$$\lambda^M_S : (U_0, U_1) \mapsto \{\mu : S \to \mathbb{N} \mid \sum_{u \in U_0} \mu(u) > \sum_{u \in U_1} \mu(u)\}.$$

Remark 6.9 Nullary predicate liftings exist. To unravel their meaning, note that we may think of any set of the form $(PS)^0$ as a *singleton* (more precisely, as the singleton consisting of the unique map $!_S : 0 \to PS$, where $0 = \varnothing$ is the empty set). Hence, we may identify a map $\lambda_S : (\breve{P}S)^0 \to \breve{P}TS$ with a *distinguished element* $\lambda_S(!_S)$ of the set $\breve{P}TS$, i.e., a subset of TS, and the naturality condition states that

$$(PTf)(\lambda_{S'}(!_{S'})) = \lambda_S(!_S),$$

for any map $f : S \to S'$.

Now suppose that λ is such a nullary predicate lifting, then the nullary modality $\heartsuit_\lambda$ associated with λ can be seen as a modal *constant*:

$$\mathbb{S}, \sigma, s \Vdash \heartsuit_\lambda \text{ iff } \sigma(s) \in \lambda_S(!_S).$$

Below we give two natural examples of this phenomenon.

Example 6.10 A natural example of a nullary modality is the constant $\surd$ that is sometimes used to indicate that a state in a finite deterministic automaton is accepting. Recall that these devices are coalgebras for the functor $2 \times Id^C$, and consider the nullary predicate lifting $\lambda^{\surd} : \breve{P}^0 \dot{\to} \breve{P}(2 \times Id^C)$ given by

$$\lambda^{\surd}_S(!_{S'}) := \{(i, f) \in 2 \times S^C \mid i = 1\}.$$

We obtain, for any state s in a deterministic automaton $\mathbb{S} = (S, \sigma)$, that s is accepting iff $\sigma(s) \in \lambda^{\surd}_S(!_S)$, so that we may think of the predicate lifting $\lambda^{\surd}$ as inducing the modality $\surd$.

Example 6.11 Let T be a functor, and Q a set of proposition letters. Recall from Remark 1.20 that we may see a T-model (S, σ, V) over Q as a coalgebra (S, σ_V) for the functor $T_{\mathsf{Q}} = K_{P\mathsf{Q}} \times T$, where $\sigma_V : S \to P\mathsf{Q} \times TS$ is defined by putting

$$\sigma_V(s) := (V^\flat(s), \sigma(s)).$$

Now fix a proposition letter $q \in \mathsf{Q}$, and consider the following *nullary* predicate lifting $\lambda^q : \breve{P}^0 \dot{\to} \breve{P}T_{\mathsf{Q}}$ for this functor:

$$\lambda^q_S(!_S) := \{(c, \tau) \in P\mathsf{Q} \times TS \mid q \in c\}.$$

Furthermore, observe that the *modality* associated with this predicate lifting is also nullary, that is, a constant; its semantics in a T_{Q}-coalgebra (X, ξ) is given by

$$\begin{array}{lll} X, \xi, x \Vdash \heartsuit_{\lambda^q} & \text{iff} & \xi x \in \lambda^q_X(!_X). \\ & \text{iff} & q \in \pi_0(\xi x). \end{array}$$

In particular, this means that if (X, ξ) is of the form (S, σ_V) for some T-model (S, σ, V), we obtain that

$$(S, \sigma, V), s \Vdash q \text{ iff } (S, \sigma_V), s \Vdash \heartsuit_{\lambda^q}. \tag{48}$$

Based on this equivalence, we may think of proposition letters as modalities associated with nullary predicate liftings.

Remark 6.12 Given a set functor T and a set Q of proposition letters, we can now make the connection explicit between modal languages for T-models over Q on the one hand, and for T_{Q}-coalgebras on the other.

Based on Example 6.11, we see that there is a 1-1 connection between proposition letters in Q and nullary predicate liftings for T_{Q} that disect the '$P\mathsf{Q}$-part' c of an arbitrary object $(c, \tau) \in T_{\mathsf{Q}}S$.

To finish the picture, we now associate with an arbitary n-ary predicate lifting $\lambda : P^n \dot{\to} PT$ for T, an n-ary predicate lifting $\lambda' : P^n \dot{\to} PT_{\mathsf{Q}}$ for T_{Q} as follows:

$$\lambda'_S(X_0, \ldots, X_{n-1}) := \big\{(c, \tau) \in T_{\mathsf{Q}} \mid \tau \in \lambda_S(X_0, \ldots, X_{n-1})\big\}.$$

Then, given a modal signature Λ for T and a set Q of proposition letters, we define the signature $\Lambda \oplus \mathsf{Q}$ for the functor T_{Q} by putting

$$\Lambda \oplus \mathsf{Q} := \{\lambda' \mid \lambda \in \Lambda\} \cup \{\lambda^q \mid q \in \mathsf{Q}\},$$

and we leave it for the reader to verify that with this definition, we can see the language $\mathrm{ML}_\Lambda(\mathsf{Q})$ (for T-models over Q) and $\mathrm{ML}_{\Lambda\oplus\mathsf{Q}}(\varnothing)$ (for T_{Q}-coalgebras) as notational variants of one another. In the sequel we will use this observation and use the language $\mathrm{ML}_\Lambda(\mathsf{Q})$ for T_{Q}-coalgebras; in particular, we will always simply write q instead of $\heartsuit_{\lambda^q}$.

6.3 Basic properties of ML_Λ

In this subsection we make some first observations about the modal logic of predicate liftings. First we show how the naturality condition (45) implies invariance under behavioural equivalence.

Theorem 6.13 *Let Λ be a modal signature for the set functor T. Then the language ML_Λ is invariant: Given any two pointed T-models $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ we have*

$$(\mathbb{S}, s) \simeq_{T_{\mathsf{Q}}} (\mathbb{S}', s') \text{ implies } (\mathbb{S}, s) \equiv_{\Lambda,\mathsf{Q}} (\mathbb{S}', s'). \tag{49}$$

Proof. Given the definition of behavioural equivalence, it suffices to prove that for any T-model morphism $f : \mathbb{S} \to \mathbb{S}'$, and any state $s \in S$, we have $(\mathbb{S}, s) \equiv_{\Lambda,\mathsf{Q}} (\mathbb{S}', fs)$. So fix $f : \mathbb{S} \to \mathbb{S}'$; we will show that every formula $\varphi \in \mathrm{ML}_\Lambda(\mathsf{Q})$ satisfies the following:

$$\mathbb{S}, s \Vdash \varphi \text{ iff } \mathbb{S}', fs \Vdash \varphi, \text{ for all states } s \in S,$$

or equivalently,

$$[\![\varphi]\!]^{\mathbb{S}} = (\breve{P}f)[\![\varphi]\!]^{\mathbb{S}'}. \tag{50}$$

We will prove (50) by a straightforward formula induction. Leaving the routine cases as an exercise to the reader, we focus on the case where $\varphi = \heartsuit_\lambda(\psi_0, ..., \psi_{n-1})$. Our proof makes use of the fact that the diagram below commutes. This should be obvious for the left rectangle,

witnessing the naturality of λ; observe that the right rectangle is obtained by applying the (contravariant!) functor $\breve{P}$ to the diagram indicating that $f : S \to S'$ is a coalgebra morphism.

$$\begin{array}{ccccc} (PS)^n & \xrightarrow{\lambda_S} & PTS & \xrightarrow{\breve{P}\sigma} & PS \\ {\scriptstyle (\breve{P}f)^n}\uparrow & & \uparrow{\scriptstyle \breve{P}Tf} & & \uparrow{\scriptstyle \breve{P}f} \\ (PS')^n & \xrightarrow[\lambda_{S'}]{} & PTS' & \xrightarrow{\breve{P}\sigma'} & PS' \end{array} \tag{51}$$

Now consider the following calculation:

$$\begin{array}{rll} [\![\heartsuit_\lambda(\ldots,\psi_i,\ldots)]\!]^{\mathbb{S}} & = (\breve{P}\sigma)\lambda_S(\ldots,[\![\psi_i]\!]^{\mathbb{S}},\ldots) & \text{(semantics } \heartsuit_\lambda \text{ in } \mathbb{S}\text{, as in (47))} \\ & = (\breve{P}\sigma)\lambda_S(\ldots,(\breve{P}f)[\![\psi_i]\!]^{\mathbb{S}'},\ldots) & \text{(induction hypothesis)} \\ & = (\breve{P}\sigma)(\breve{P}Tf)\lambda_{S'}(\ldots,[\![\psi_i]\!]^{\mathbb{S}'},\ldots) & \text{(naturality of } \lambda\text{, (51) left)} \\ & = (\breve{P}f)(\breve{P}\sigma)\lambda_{S'}(\ldots,[\![\psi_i]\!]^{\mathbb{S}'},\ldots) & (f \text{ is a morphism, (51) right)} \\ & = (\breve{P}f)[\![\heartsuit_\lambda(\ldots,\psi_i,\ldots)]\!]^{\mathbb{S}'} & \text{(semantics } \heartsuit_\lambda \text{ in } \mathbb{S}'\text{, as in (47))} \end{array}$$

showing that (50) holds for $\varphi = \heartsuit_\lambda(\psi_0, ..., \psi_{n-1})$ indeed. QED

Concerning the property of expressiveness, we may obtain a general result if we put some constraints on the signature Λ.

Definition 6.14 Let Λ be a modal signature for a set functor T. We say that Λ is *separating* for T, if for all sets S and all pairs of distinct objects $\sigma_0, \sigma_1 \in TS$ there is a $\lambda \in \Lambda$ and a tuple $(A_0, \ldots, A_{n-1})$ such that *exactly* one of the two objects σ_i belongs to the set $\lambda_S(A_0, \ldots, A_{n-1})$. $\lhd$

Example 6.15 (1) The box relation lifting $\lambda^\Box$ is separating on its own, for the powerset functor P. To see this, consider two subsets $X, Y \in PS$. If X and Y are distinct, suppose without loss of generality that $Y \not\subseteq X$, so that $Y \notin \lambda^\Box(X) = \{U \in PS \mid U \subseteq X\}$.

(2) The predicate liftings associated with the graded modalities are jointly separating. To see this, consider two bags $\beta_0, \beta_1 : S \to \mathbb{N}^\infty$ over some set S. If β_0 and β_1 are distinct, then we must have $\beta_0(s) \neq \beta_1(s)$, for some $s \in S$. Without loss of generality we may assume that $\beta_0(s) < \beta_1(s)$, so that in particular, $\beta_0(s)$ belongs to $\mathbb{N}$, say, $\beta_0(s) = m$. Recall that $\lambda_S^{m+1}(\{s\}) = \{\beta \in BS \mid \beta(s) \geq m+1\}$, so that we find $\beta_0 \notin \lambda_S^{m+1}(\{s\})$ but $\beta_1 \in \lambda_S^{m+1}(\{s\})$.

Theorem 6.16 *Let Λ be a separating modal signature for the set functor T. Then the language ML_Λ is expressive on the class of image-finite T-coalgebras: Given any two pointed T_ω-models $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ we have*

$$(\mathbb{S}, s) \equiv_{\Lambda,\mathsf{Q}} (\mathbb{S}', s') \text{ implies } (\mathbb{S}, s) \simeq_{T_\mathsf{Q}} (\mathbb{S}', s'). \tag{52}$$

Remark 6.17 In preparation for the proof we isolate two observations on the use of the contravariant power set functor to represent some basic set-theoretic operations.

First of all, consider two sets X and S such that $X \subseteq S$, and let $\iota : X \hookrightarrow S$ be the associated inclusion map. We leave it for the reader to check that the map $\breve{P}\iota : PS \to PX$ corresponds to the operation of taking intersections with the set X:

$$X \cap U = (\breve{P}\iota)U.$$

Observe as well that, in case T is standard, it follows that, for any subset $U \subseteq TS$:

$$TX \cap U = (\breve{P}T\iota)U.$$

Second, assume that $\equiv$ is some equivalence relation on the set S, and that $\overline{S}$ is the induced partition consisting of all $\equiv$-cells. Then for any collection A of $\equiv$-cells it holds that

$$\bigcup A = (\breve{P}q)A,$$

where $q : S \to \overline{S}$ is the quotient map associated with $\equiv$. The proof of this equation is immediate by the observation that $s \in \bigcup A$ iff $q(s) \in A$.

Proof of Theorem 6.16. For notational simplicity we will confine ourselves to a setting where all predicate liftings in Λ are unary, leaving the (routine) generalisation to an arbitrary signature as an exercise. It will also be convenient to assume that T preserves inclusions. Furthermore, we will only treat the special case where the coalgebras $\mathbb{S}$ and $\mathbb{S}'$ coincide (i.e., $\mathbb{S} = \mathbb{S}'$); the general case, where the two coalgebras are distinct, can easily be reduced to this by considering their disjoint union.

So let Λ be a separating set of unary predicate liftings for some set functor T, and let $\mathbb{S} = (S, \sigma, V)$ be an image-finite T-model, that is, $\sigma : S \to T_\omega S$. To avoid notational clutter we will simply write $\equiv$ for the equivalence relation $\equiv_{\Lambda,\mathbf{Q}}$.

Our aim is to prove that $s_0 \equiv s_1$ implies $s_0 \simeq s_1$, for all $s_0, s_1 \in \mathbb{S}$. Clearly then it suffices to show that the relation $\equiv$ is contained in the kernel of some coalgebra morphism. We will show that in fact we may define coalgebra structure $\overline{\sigma} : \overline{S} \to T\overline{S}$ on the set $\overline{S}$ of $\equiv$-cells in such a way that the quotient map $q : S \to \overline{S}$ becomes a coalgebra morphism:

$$\begin{array}{ccc} S & \xrightarrow{q} & \overline{S} \\ {\scriptstyle\sigma}\downarrow & & \vdots\,{\scriptstyle\overline{\sigma}} \\ TS & \xrightarrow[Tq]{} & T\overline{S} \end{array}$$

Now suppose that we can show that

$$q(s_0) = q(s_1) \text{ implies } (Tq)(\sigma s_0) = (Tq)(\sigma s_1), \tag{53}$$

then putting

$$\overline{\sigma}(\overline{s}) := (Tq)(\sigma s)$$

would give a correctly defined map, for which the quotient map q is trivially a coalgebra morphism.

That is, we have reduced our problem to finding a proof for (53), and to this aim we reason by contraposition: Assuming that

$$(Tq)(\sigma s_0) \neq (Tq)(\sigma s_1), \tag{54}$$

we will show that $q(s_0) \neq q(s_1)$ (that is, $s_0 \not\equiv_\Lambda s_1$). It follows by separation from (54) that there is some $\lambda \in \Lambda$ and some $A \subseteq \overline{S}$ such that (without loss of generality) we have

$$(Tq)(\sigma s_0) \in \lambda_{\overline{S}}(A), \text{ but } (Tq)(\sigma s_1) \notin \lambda_{\overline{S}}(A). \tag{55}$$

Our purpose is now is find a formula witnessing this, in the sense that this formula holds at s_0 but not at s_1: this would show indeed that $s_0 \not\equiv s_1$, and so $q(s_0) \neq q(s_1)$. Note that $A \subseteq \overline{S}$ simply means that A is a collection of equivalence classes.

Now if we would have $\bigcup A = [\![\varphi]\!]$ we would be done immediately since in this case we could prove that $\mathbb{S}, s_0 \Vdash \heartsuit_\lambda \varphi$ while $\mathbb{S}, s_1 \nVdash \heartsuit_\lambda \varphi$. To see this, note that we have $s_i \Vdash \heartsuit_\lambda \varphi$ iff $\sigma(s_i) \in \lambda_S([\![\varphi]\!])$ and that $\bigcup A = [\![\varphi]\!]$ implies $\lambda_S([\![\varphi]\!]) = \lambda_S(\bigcup A) = \lambda_S(\breve{P}q(A)) = (\breve{P}Tq)(\lambda_{\overline{S}}(A))$, where the last identity is by naturality of λ. It is thus immediate by (55) that $\sigma s_0 \in \lambda_S([\![\varphi]\!])$ but $\sigma s_1 \notin \lambda_S([\![\varphi]\!])$, as required.

In the general case we need to work harder. But since $\mathbb{S}$ is a T_ω-coalgebra, and T preserves inclusions, there is a *finite* subset $X \subseteq S$ such that both σs_0 and σs_1 belong to the set TX. We leave it for the reader to verify that there is a formula $\varphi \in \mathrm{ML}_\Lambda$ that characterizes, *within* X, the union $\bigcup A$, in the sense that

$$\text{for all } x \in X: \quad \mathbb{S}, x \Vdash \varphi \text{ iff } x \in \bigcup A,$$

or equivalently, since $x \in \bigcup A$ is another way of saying that $q(x) \in A$:

$$X \cap [\![\varphi]\!] \;=\; X \cap (\breve{P}q)A. \tag{56}$$

We now claim that, for this formula φ, we have

$$\mathbb{S}, s_0 \Vdash \heartsuit_\lambda \varphi \text{ but } \mathbb{S}, s_1 \nVdash \heartsuit_\lambda \varphi. \tag{57}$$

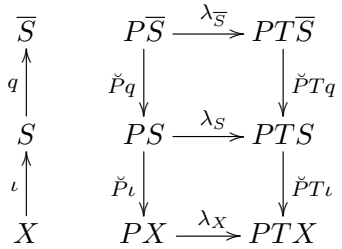
To prove this, we first observe that by the semantics of $\heartsuit_\lambda$ we have that $\mathbb{S}, s_i \Vdash \heartsuit_\lambda \varphi$ iff $\sigma(s_i) \in \lambda_S([\![\varphi]\!])$, while it follows from (55) that $\sigma(s_0) \in (\breve{P}Tq)\lambda_{\overline{S}}(A)$ but $\sigma(s_1) \notin (\breve{P}Tq)\lambda_{\overline{S}}(A)$. Hence, because both $\sigma(s_0)$ and $\sigma(s_1)$ belong to TX, it suffices to show that

$$TX \cap (\breve{P}Tq)\lambda_{\overline{S}}(A) \;=\; TX \cap \lambda_S[\![\varphi]\!]. \tag{58}$$

We will establish this by chasing the diagram below, where we use a trick to interpret the intersection with X in (56) and with TX in (58) using the inclusion map $\iota : X \hookrightarrow S$. That is, we observe that $\breve{P}\iota : PS \to PX$ is given by $U \mapsto X \cap U$; as a consequence another way of formulating (56) is:

$$(\breve{P}\iota)([\![\varphi]\!]) = (\breve{P}\iota)(\breve{P}q)(A). \tag{59}$$

Similarly, since T preserves inclusions we have that $T\iota : TX \to TS$ is the inclusion map witnessing that $TX \subseteq TS$, and so $\breve{P}T\iota : PTS \to PTX$ is given by $\Sigma \mapsto (TX) \cap \Sigma$.



$$
\begin{array}{lll}
TX \cap (\breve{P}Tq)\lambda_{\overline{S}}(A) & & \\
\quad = (\breve{P}T\iota)(\breve{P}Tq)\lambda_{\overline{S}}(A) & \text{(just discussed)} \\
\quad = \lambda_X(\breve{P}\iota)(\breve{P}q)(A) & \text{(naturality of } \lambda) \\
\quad = \lambda_X(\breve{P}\iota)([\![\varphi]\!]) & \text{(see (59))} \\
\quad = (\breve{P}T\iota)\lambda_S([\![\varphi]\!]) & \text{(naturality of } \lambda) \\
\quad = TX \cap \lambda_S[\![\varphi]\!] & \text{(just discussed)}
\end{array}
$$

This proves (58), and therefore (57). That is, we have shown that $s_0 \not\equiv s_1$, on the assumption that $(Tq)(\sigma s_0) \neq (Tq)(\sigma s_1)$. This means that (53) holds, and as we argued already, this suffices to prove the Theorem. QED

•

6.4 Finite model property

In this subsection we will show that the coalgebraic modal logic ML_Λ has the (strong) finite model property. That is, we will show that any satisfiable ML_Λ-formula φ can in fact be satisfied in a *finite* coalgebra of which the size (number of states) is exponentially bounded by the size of φ. We will prove this result by adapting the method of filtration, which is well known in the theory of standard modal logic, to the more general coalgebraic setting.

First we need some preliminary definitions.

Definition 6.18 The collection $\mathtt{Sfor}(\varphi)$ of *subformulas* of a ML_Λ-formula φ is defined in the standard way. The *size* $|\varphi|$ of a φ is defined as its number of subformulas: $|\varphi| := |\mathtt{Sfor}(\varphi)|$.

A set of formulas Σ is called *subformula-closed* if it is closed under taking subformulas, that is, if $\mathtt{Sfor}(\varphi) \subseteq \Sigma$ for all $\varphi \in \Sigma$. ◁

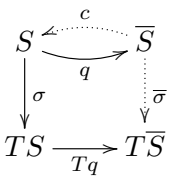
The idea behind the method of filtration is fairly simple: given a subformula-closed set Σ and a T-model $\mathbb{S} = (S, \sigma, V)$, define on S a suitable equivalence relation $\equiv_\Sigma$ of finite index, build a new T-model $\overline{\mathbb{S}}$ on the finite set of $\equiv_\Sigma$-cells, and show that any state $s \in S$ satisfies the same formulas in $\mathbb{S}$ as its cell $\overline{s}$ does in the filtrated model $\overline{\mathbb{S}}$.

Definition 6.19 Let Σ be a finite, subformula closed set of formulas in $\mathrm{ML}_\Lambda(\mathsf{Q})$, and let $\mathbb{S} = (S, \sigma, V)$ be a T-model. We define $\equiv_\Sigma \subseteq S \times S$ as the equivalence relation given by

$$s \equiv_\Sigma t \text{ iff for all } \varphi \in \Sigma : \mathbb{S}, s \Vdash \varphi \iff \mathbb{S}, t \Vdash \varphi,$$

and denote the $\equiv_\Sigma$-cell of a state s as $\overline{s}$. We also let $\overline{S} := \{\overline{s} \mid s \in S\}$ denote the set of cells, and let $q : s \mapsto \overline{s}$ denote the quotient map $q : S \to \overline{S}$. ◁

In order to define a coalgebra map $\overline{\sigma} : \overline{S} \to T\overline{S}$, we may pick any *choice function* $c : \overline{S} \to S$, and define $\overline{\sigma} := Tq \circ \sigma \circ c$, cf. the diagram below:



Here we call $c : \overline{S} \to S$ a choice function if c picks an element from each $\equiv_\Sigma$-cell; in other words, we require that $q \circ c = \mathsf{id}_{\overline{S}}$ and $\ker(c \circ q) \subseteq \equiv_\Sigma$.

Note that while the 'outer' rectangle of the above diagram commutes by definition, the 'inner' one need not commute: it will generally not be possible to define a coalgebra map on $\overline{S}$ for which the quotient map q is a coalgebra morphism. Fortunately, for our purposes we don't need such a coalgebra morphism.

Definition 6.20 Let Σ be a finite, subformula closed set of formulas in $\mathrm{ML}_\Lambda(\mathsf{Q})$, and let $\mathbb{S} = (S, \sigma, V)$ be a T-model. A Σ-*filtration* of $\mathbb{S}$ is any T-model $\overline{\mathbb{S}} = (\overline{S}, \overline{\sigma}, \overline{V})$ such that
(1) $\overline{S} = S/{\equiv_\Sigma}$ is the class of $\equiv_\Sigma$-cells,
(2) $\overline{\sigma} = Tq \circ \sigma \circ c$ for some choice function $c : \overline{S} \to S$, and
(3) $\overline{V}(q) = \{\overline{s} \mid s \in V(q)\}$ for $q \in \Sigma \cap \mathsf{Q}$. $\lhd$

Observe that filtrations are not unique, they depend in particular on the choice of the choice function c.

We can now prove the following Filtration Lemma.

Theorem 6.21 (Filtration Lemma) *Let Λ be a modal signature for a set functor T, and let $\Sigma \subseteq \mathrm{ML}_\Lambda$ be a finite subformula-closed set of formulas. Furthermore, let $\mathbb{S} = (S, \sigma, V)$ be a T-model, and let $\overline{\mathbb{S}} = (\overline{S}, \overline{\sigma}, \overline{V})$ be a Σ-filtration of $\mathbb{S}$. Then for all formulas $\varphi \in \Sigma$ we have*

$$\mathbb{S}, s \Vdash \varphi \text{ iff } \overline{\mathbb{S}}, \overline{s} \Vdash \varphi, \text{ for all states } s \in S. \tag{60}$$

Proof. First, note that the statement (60) is equivalent to

$$[\![\varphi]\!]^{\mathbb{S}} = (\breve{P}q)[\![\varphi]\!]^{\overline{\mathbb{S}}}. \tag{61}$$

This means that we may use (61) as the inductive hypothesis in our inductive proof.

Second, since we have $s \equiv c(\overline{s})$ for all $s \in S$ the statement (60) is also equivalent to

$$[\![\varphi]\!]^{\overline{\mathbb{S}}} = (\breve{P}c)[\![\varphi]\!]^{\mathbb{S}}, \tag{62}$$

as long as we restrict to formulas $\varphi \in \Sigma$. Since we only want to establish (60) for formulas in Σ, it therefore suffices to prove (62). We will do so by a straightforward formula induction.

Turning to the actual inductive proof of (62), we only consider the inductive step where $\varphi = \heartsuit_\lambda(\psi_0, \ldots, \psi_{n-1})$, and for notational simplicity we confine ourselves to the case where λ is unary, i.e, $\varphi = \heartsuit_\lambda \psi$ for some formula ψ to which the inductive hypothesis applies. Now consider the following diagram:

$$\begin{array}{ccccc} P\overline{S} & \xrightarrow{\lambda_{\overline{S}}} & PT\overline{S} & \xrightarrow{\breve{P}\overline{\sigma}} & P\overline{S} \\ {\scriptstyle \breve{P}q}\big\downarrow & & {\scriptstyle \breve{P}Tq}\big\downarrow & & \big\uparrow{\scriptstyle \breve{P}c} \\ PS & \xrightarrow[\lambda_S]{} & PTS & \xrightarrow[\breve{P}\sigma]{} & PS \end{array} \tag{63}$$

$$\begin{aligned} [\![\heartsuit_\lambda\psi]\!]^{\overline{\mathbb{S}}} &= (\breve{P}\overline{\sigma})\lambda_{\overline{S}}([\![\psi]\!]^{\overline{\mathbb{S}}}) && \text{(semantics of } \heartsuit_\lambda) \\ &= \big(\breve{P}(Tq \circ \sigma \circ c)\big)\lambda_{\overline{S}}([\![\psi]\!]^{\overline{\mathbb{S}}}) && \text{(definition } \overline{\sigma}) \\ &= (\breve{P}c)(\breve{P}\sigma)(\breve{P}Tq)\lambda_{\overline{S}}([\![\psi]\!]^{\overline{\mathbb{S}}}) && \text{(functoriality)} \\ &= (\breve{P}c)(\breve{P}\sigma)\lambda_S(\breve{P}q)([\![\psi]\!]^{\overline{\mathbb{S}}}) && \text{(naturality of } \lambda) \\ &= (\breve{P}c)(\breve{P}\sigma)\lambda_S([\![\psi]\!]^{\mathbb{S}}) && \text{(induction hypothesis (61))} \\ &= (\breve{P}c)[\![\heartsuit_\lambda\psi]\!]^{\mathbb{S}} && \text{(semantics of } \heartsuit_\lambda) \end{aligned}$$

In other words, we have established (62) for $\varphi = \heartsuit_\lambda\psi$, as required. QED

Corollary 6.22 (Strong Finite Model Property) *Let φ be a formula in ML_Λ, where Λ is a modal signature for a set functor T. If φ is satisfiable in some T-model, then it is satisfiable in a finite T-model (S, σ, V) such that $|S| \leq 2^{|\varphi|}$.*

Proof. Fix an ML_Λ-formula φ, and let $\Sigma := \mathtt{Sfor}(\varphi)$. It follows by the filtration lemma that φ, if satisfiable in some pointed T-model $(\mathbb{S}, s)$, also holds at the state $\overline{s}$, in any filtration $\overline{\mathbb{S}}$ of $\mathbb{S}$. This proves the theorem, since it easily follows from the definion of the relation $\equiv_\Sigma$ that $|\overline{S}| \leq |P\Sigma| = 2^{|\varphi|}$. QED

6.5 Predicate liftings as coalgebra type changers

This short section presents a slightly different perspective in which predicate liftings provide natural ways to transform T-coalgebras to neighbourhood frames. We first consider a simplified version, in which T-coalgebras are transformed to Kripke frames.

Example 6.23 Define a *natural relation* for T to be a natural transformation $\mu : T \dot{\to} P$.

Given such a natural relation $\mu : T \dot{\to} P$, we can transform a T-coalgebra $\mathbb{S} = (S, \sigma)$ into a Kripke frame $\mathbb{S}^\mu := (S, \mu_S \circ \sigma)$. By naturality of μ, any T-homomorphism $f : \mathbb{S} \to \mathbb{S}'$ is also a bounded morphism $f : \mathbb{S}^\mu \to (\mathbb{S}')^\mu$. To check this, one may easily verify that $Pf \circ (\mu_S \circ \sigma) = (\mu_{S'} \circ \sigma') \circ f$ by chasing the diagram below:

$$\begin{array}{ccccc} S & \xrightarrow{\sigma} & TS & \xrightarrow{\mu_S} & PS \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle Tf} & & \downarrow{\scriptstyle Pf} \\ S' & \xrightarrow{\sigma'} & TS' & \xrightarrow{\mu_{S'}} & PS' \end{array}$$

Connecting this to logic, with any natural relation μ we may associate a modality $\langle\mu\rangle$ for T-coalgebras, with the following interpretation:

$$\mathbb{S}, s \Vdash \langle\mu\rangle\varphi \text{ iff } [\![\varphi]\!]^{\mathbb{S}} \cap \mu_S\sigma(s) \neq \varnothing.$$

As an example, recall from Fact 5.17(2) that for smooth and standard functors T, we have a natural transformation $Base^T : T_\omega \dot{\to} P_\omega$. Hence, we may take the *Base* operation as a way to transform T_ω-coalgebras into P_ω-coalgebras, that is, image-finite Kripke frames. Naturality ensures that every morphism between T-coalgebras is also a bounded morphism between the underlying Kripke frames.

As we will see now, predicate liftings can be seen as generalisations of this phenomenon, where we move from Kripke frames to the more general setting of neighbourhood frames. For this general setting we introduce the *transpose* of a predicate lifting, cf. Definition A.4. This notion is based on the correspondence between maps $A \to PB$ and maps $B \to PA$ — we have seen this correspondence already in the coalgebraic presentation of a valuation $V : \mathsf{Q} \to PS$ as a colouring $V^\flat : S \to P\mathsf{Q}$.

Definition 6.24 Given a map $\alpha : A \to PB$ we define its *transposed map* $\alpha^\flat : B \to PA$ by putting $\alpha^\flat(b) := \{a \in A \mid b \in \alpha(a)\}$.

Extending this definition, given an n-ary predicate lifting for the set functor T, we define its *transpose* $\lambda^\flat$ as the set-indexed family of maps

$$\lambda^\flat_S : TS \to P(P^n(S))$$

given by $\lambda^\flat_S(\sigma) := (\lambda_S)^\flat(\sigma)$. ◁

By the naturality of predicate liftings we obtain the following proposition, which shows that predicate liftings indeed generalise the natural relations of Example 6.23.

Proposition 6.25 *If $\lambda : P^n \dot{\to} PT$ then $\lambda^\flat$ is a natural transformation*

$$\lambda^\flat : T \dot{\to} \breve{P} \circ \breve{P}^n.$$

It follows from this proposition that any predicate lifting λ induces a transformation of T-coalgebras to n-ary neighbourhood frames. We confine attention to the unary case.

Definition 6.26 Let $\lambda : \breve{P} \dot{\to} \breve{P}T$ be a unary predicate lifting for the set functor T. Given a T-coalgebra $\mathbb{S} = (S, \sigma)$, we let $\mathbb{S}^\lambda$ denote the neighbourhood frame $\mathbb{S}^\lambda := (S, \lambda^\flat \circ \sigma)$; given a function $f : S \to S'$ we define $f^\lambda := f$. $\lhd$

The following proposition is easy to verify.

Proposition 6.27 *Let $\lambda : \breve{P} \dot{\to} \breve{P}T$ be a unary predicate lifting for the set functor T. Then the construction $(\cdot)^\lambda$ is a functor from the category of T-coalgebras to the category of neighbourhood frames (N-coalgebras).*

The following proposition, which is easy to verify, provides a slightly different perspective on the concept of separation.

Proposition 6.28 *Let Λ be a modal signature for a set functor T. Then Λ is separating iff, for every set S, the collection $(\lambda^\flat_S)_{\lambda \in \Lambda}$ of transposed functions is jointly injective (i.e., for any pair of distinct objects $\tau_0, \tau_1 \in TS$ there is a $\lambda \in \Lambda$ such that $\lambda^\flat_S(\tau_0) \neq \lambda^\flat_S(\tau_1)$).*

6.6 Predicate liftings and the Yoneda lemma

In the final subsection of this chapter we take a slightly different perspective on the modalities that are given by predicate liftings, seeing them as describing certain *admissible patterns*. This perspective will also reveal *how many* predicate liftings of each different arity there are.

The key observation here is that there is a natural bijection between PS (subsets of S) and 2^S (characteristic functions on S).

Definition 6.29 Given a subset $X \subseteq S$, we define the *characteristic function of X* as the map $\chi^S_X : S \to 2$ given by

$$\chi^S_X(s) := \begin{cases} 1 & \text{if } s \in X \\ 0 & \text{if } s \notin X. \end{cases}$$

In case S is understood, we may drop the superscript 'S'. Conversely, given a map $\chi : S \to 2$, we shall call $\chi^{-1}(1) \in PS$ the *subset determined by χ*. $\lhd$

Remark 6.30 This correspondence reaches far enough for us to think of the contravariant power set functor as the functor 2^- that associates with a set S the collection 2^S of functions from S to $2 = \{0, 1\}$, and with an arrow $f : S' \to S$ the 'precomposition map' that assigns to an arbitrary characteristic function $\chi : S \to 2$ the function $\chi \circ f : S' \to 2$.

From this point of view, a unary predicate lifting is a way of transforming arrows $S \to 2$ into arrows $TS \to 2$. In particular, any arrow $\gamma : T2 \to 2$ (that is, any characteristic function χ_Γ corresponding to a subset $\Gamma \subseteq T2$), induces a unary predicate lifting: Given an arrow $\chi : S \to 2$, simply consider the arrow $\gamma \circ T\chi$, as in the diagram below:

$$\begin{array}{ccccc} S & \xrightarrow{\chi} & 2 & & \\ & & & \overset{\gamma\circ T\chi}{\nearrow} & \\ TS & \xrightarrow{T\chi} & T2 & \xrightarrow{\gamma} & 2 \end{array}$$

Formulated in terms of subsets rather than characteristic functions, we arrive at the following definition.

Definition 6.31 Given an object $\Gamma \subseteq T2$, let $\widehat{\Gamma}$ be the following set-indexed family of operations. For a set S, we define

$$\widehat{\Gamma}_S : PS \to PTS,$$

by putting, for any $X \subseteq S$:

$$\widehat{\Gamma}_S(X) := \{\sigma \in TS \mid (T\chi^S_X)(\sigma) \in \Gamma\},$$

where $\chi^S_X : S \to 2$ is the characteristics map associated with X. Where $\widehat{\Gamma}$ is a predicate lifting, we shall denote its associated modality as $\heartsuit_\Gamma$ rather than as $\heartsuit_{\widehat{\Gamma}}$. $\lhd$

Remark 6.32 Taking a glance at the *modalities* that are induced by subsets of the set $T2$, we consider the interpretation of the formula $\heartsuit_\Gamma\varphi$ in the T-model $\mathbb{S} = (S, \sigma, V)$. If we represent the set $[\![\varphi]\!]$ with its characteristic function $\chi^S_{[\![\varphi]\!]}$, applying the functor T to this arrow we obtain

$$TS \xrightarrow{T\chi^S_{[\![\varphi]\!]}} T2\ .$$

Now think of the elements of $T2$ as 'T-patterns', then the above arrow associates a T-pattern with each object $\tau \in TS$. We can say that the formula $\heartsuit_\Gamma\varphi$ holds at s if the pattern $\big(T\chi^S_{[\![\varphi]\!]}\big)(\sigma(s))$ associated with $\sigma(s) \in TS$ is *admissible*, i.e., belongs to the set Γ, or, equivalently, that $\big(\chi^{T2}_\Gamma \circ T\chi^S_{[\![\varphi]\!]}\big)(\sigma(s)) = 1$.

Example 6.33 As an example, consider the binary tree functor $Id \times Id$. The set $T2$ consists of four patterns: $T2 = \{(0,0),(0,1),(1,0),(1,1)\}$. As an example, take the set $\Gamma = \{(0,0),(1,1)\}$; it is not hard to see that it induces the predicate lifting $\widehat{\Gamma} : \breve{P} \dot{\to} \breve{P}\circ(Id\times Id)$ given by

$$\widehat{\Gamma}_S(X) := X \times X \cup (S \setminus X) \times (S \setminus X).$$

The associated modality $\heartsuit_\Gamma$ has the following semantics:

$$\mathbb{S}, s \Vdash \heartsuit_\Gamma\varphi \text{ iff the two successors of } s \text{ either both satisfy or both falsify } \varphi.$$

Proposition 6.34 *For any $\Gamma \subseteq T2$, the collection $\widehat{\Gamma}$ of maps constitutes a predicate lifting $\widehat{\Gamma} : \breve{P} \dot{\to} \breve{P}T$.*

Proof. Given a map $f : S' \to S$, we need to check that the following diagram commutes:

$$\begin{array}{ccc} S' & PS' & \xrightarrow{\widehat{\Gamma}_{S'}} PTS' \\ {\scriptstyle f}\downarrow & {\scriptstyle \breve{P}f}\uparrow & \uparrow{\scriptstyle \breve{P}Tf} \\ S & PS & \xrightarrow[\widehat{\Gamma}_S]{} PTS \end{array}$$

This follows by the following chain of identities, for an arbitrary subset $X \subseteq S$:

$$\begin{array}{rcll} (\widehat{\Gamma}_{S'} \circ \breve{P}f)(X) & = & \widehat{\Gamma}_{S'}(\breve{P}f(X)) & \text{(obvious)} \\ & = & \{\sigma' \in TS' \mid (T\chi^{S'}_{\breve{P}f(X)})(\sigma') \in \Gamma\} & \text{(definition } \widehat{\Gamma}_{S'}\text{)} \\ & = & \{\sigma' \in TS' \mid (T(\chi^{S}_{X} \circ f))(\sigma') \in \Gamma\} & (*) \\ & = & \{\sigma' \in TS' \mid (T\chi^{S}_{X})((Tf)(\sigma')) \in \Gamma\} & \text{(functoriality of } T\text{)} \\ & = & \{\sigma' \in TS' \mid (Tf)(\sigma') \in \widehat{\Gamma}_S(X)\} & \text{(definition } \widehat{\Gamma}_S\text{)} \\ & = & (\breve{P}Tf)(\widehat{\Gamma}_S(X)) & \text{(definition } \breve{P}Tf\text{)} \\ & = & (\breve{P}Tf \circ \widehat{\Gamma}_S)(X) & \text{(obvious)} \end{array}$$

Here the identity (*) is immediate by the observation that

$$\chi^{S}_{X} \circ f = \chi^{S'}_{\breve{P}f(X)}$$

as is revealed by a straightforward verification, for an arbitrary element $s' \in S'$: $\chi^{S}_{X} \circ f(s') = 1$ iff $f(s') \in X$ iff $s' \in f^{-1}(X) = (\breve{P}f)(X)$ iff $\chi^{S'}_{\breve{P}f(X)}(s') = 1$. QED

Interestingly, we may prove that *all* unary predicate liftings are of this form, and this result generalises to predicate liftings of arbitrary arity. This is the main content of the following theorem, which is in fact the instantiation of the well-known *Yoneda Lemma* to the setting of predicate liftings.

Theorem 6.35 *For any set functor T there is a natural bijection between the set of n-ary predicate liftings for T and the power set of $T(Pn)$.*

The key observations underlying the proof of this theorem are the natural correspondences

$$(PS)^n \sim (2^S)^n \sim (2^n)^S \sim (Pn)^S$$

between n-tuples of subsets of S, n-tuples of characteristic functions, maps from S to 2^n, and maps from S to Pn.

Proof. We confine the proof of this result to providing an n-ary predicate lifting for each subset of TPn, and vice versa.

So let $\Gamma \subseteq T(Pn)$ be a set of admissible n-ary T-patterns. The associated predicate lifting $\widehat{\Gamma} : P \dot{\to} PT$ is obtained by a straightforward generalisation of Definition 6.31. For the

details, take an arbitrary tuple $X = (X_0, \ldots, X_{n-1}) \in (PS)^n$ of subsets of S, and represent this tuple as the map $\chi_X := (\chi_0, \ldots, \chi_{n-1})$ of associated characteristic functions, then we have $\chi_X : S \to 2^n$, and so $T\chi_X : TS \to T(2^n)$. Now put

$$\widehat{\Gamma}_S(X_0, \ldots, X_{n-1}) := \{\tau \in TS \mid (T\chi_X)(\tau) \in \Gamma\}.$$

We leave it for the reader to verify that this collection of maps indeed provides a natural transformation $\widehat{\Gamma} : \breve{P}^n \to \breve{P}T$.

For the opposite direction, let $\lambda : \breve{P}^n \dot{\to} \breve{P}T$ be an n-ary predicate lifting. Our aim is to find a subset $\Theta_\lambda \subseteq T(2^n)$ such that $\lambda = \widehat{\Theta_\lambda}$.

The easiest way to proceed is by thinking of λ_S as a way to transform arrows $S \to 2^n$ into arrows $TS \to 2$, that is, $\lambda_S : (2^n)^S \to 2^{TS}$. To find the object Θ_λ, we take a *special* set S, viz., the set 2^n itself, and, as input for λ, a *special* arrow, viz., the *identity* arrow $\mathsf{id}_{2^n} : 2^n \to 2^n$. Then we let $\Theta_\lambda \subseteq T(2^n)$ be the subset of $T(2^n)$ that is determined by the image of this arrow id_{2^n} under λ.

If we think of λ as a set-indexed family of maps $\lambda_S : (PS)^n \to PTS$, we may define

$$\Theta_\lambda := \lambda(\mathcal{U}),$$

where $\mathcal{U}$ is the distinguished element of the set $P^n(Pn)$ corresponding to the identity map id_{2^n}, that is, $\mathcal{U} = (\mathcal{U}_0, \ldots, \mathcal{U}_{n-1})$ with

$$\mathcal{U}_i := \{U \subseteq n \mid i \in U\}.$$

As mentioned, we leave it for the reader to verify that the maps just defined are each other's inverse, i.e., that $\lambda = \widehat{\Theta_\lambda}$ for all predicate liftings λ, and that $\Gamma = \Theta_{\widehat{\Gamma}}$ for all $\Gamma \in PTn$, some $n \in \omega$. QED