Relation Liftings in Coalgebraic Modal Logic

MSc Thesis (*Afstudeerscriptie*)

written by

Johannes Marti
(born December 19, 1986 in Basel, Switzerland)

under the supervision of **Prof Dr Yde Venema**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
September 8, 2011

Members of the Thesis Committee:
Dr Alexandru Baltag
Prof Dr Dick de Jongh
Raul Andres Leal
Prof Dr Yde Venema

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

In this thesis we study relation liftings in the context of coalgebraic modal logic. In the first part of the thesis we look for conditions on relation liftings that can be used to define a notion of bisimilarity between states in coalgebras, such that two states are bisimilar if and only if they are behaviorally equivalent. We show that this is the case for relation liftings that are lax extensions and additionally preserve diagonal relations. In the second part of the thesis we develop a coalgebraic nabla logic for an arbitrary lax extension. For this logic we prove that, under additional conditions, bisimulation quantifiers are definable in the nabla logic. This has a Uniform Interpolation Theorem as consequence.

Contents

1 Introduction

Coalgebras are functions $\xi : X \to TX$ from a set of states X to the set TX given by some endofunctor T in the category of sets. By varying the functor T one can study many different types of structures in the unified framework of coalgebras. These include numerous examples from computer science such as finite automata, infinite data structures or transition systems.

The theory of coalgebras is also important for modal logic since Kripke frames and Neighborhood frames can be represented as coalgebras. Kripke frames, that are used as the standard semantics for normal modal logics, are coalgebras for the covariant powerset functor $\mathcal{P}$. With every state x in a Kripke frames one associates the set of its successors, which is an subset of the set of states. Neighborhood frames, that are used in the semantics of classical modal logic, are coalgebras for the double contravariant powerset functor $\mathcal{N} = \breve{\mathcal{P}}\breve{\mathcal{P}}$. A neighborhood frame specifies for every state the set of its neighborhoods. In between classical modal logic and normal modal logic there is monotone modal logic. The standard semantics for monotone modal logic is given on monotone neighborhood frames that are coalgebras for the monotone neighborhood functor $\mathcal{M}$. This functor is a restriction of $\mathcal{N}$ in which one requires that the set of neighborhoods associated to a state is upwards-closed.

It is not only that the models of modal logic are coalgebras but it has also turned out that coalgebraic modal logics, which generalize standard modal logics, are an adequate tool to reason about any type of coalgebras. Researchers working on coalgebraic modal logic develop logics for coalgebras of any functor. There are two main approaches of how this is usually done. The first one uses so called predicate liftings to define a modal language, similar to standard modal logic with boxes and diamond, for any functor T. See [18] for an up-to-date example of how this works. The other approach originates from the work of Moss [13], who started the study of coalgebraic modal logic. Moss uses so called nabla modalities $\nabla\alpha$ for $\alpha \in T\mathcal{L}$, where $\mathcal{L}$ is the set of all formulas, to describe properties of coalgebras for the functor T. These nabla modalities are somewhat unusual because they incorporate the functor T into the syntactic shape of modal formulas. Despite of their somewhat peculiar syntax, modal logics using the nabla modalities have very strong normal forms. One can show under relatively weak restrictions that every formula in the logic with nablas is equivalent to an formula in which negations and conjunctions occur only on the propositional non-modal level.

In this thesis we study so called relation liftings in the context of coalgebraic modal logic. A relation lifting L for a functor T maps every relation $R : X \nrightarrow Y$ between the sets X and Y to a relation $LR : TX \nrightarrow TY$ between the sets TX and TY. A relation lifting that figures very prominently in the theory of coalgebras and coalgebraic modal logic is the Barr extension $\overline{T}$ of a functor T. It is defined uniformly for any set functor T and has been used in the theory of coalgebras and coalgebraic modal logic to:

(i) Define a notion of bisimilarity between states in coalgebras.

(ii) Define a semantics for the nabla modality.

This only works properly for set functors that have the property that they preserve weak pullbacks. Otherwise the Barr extension is not well-behaved. Most set functors preserve weak pullbacks so this is not a strong restriction. Important exceptions, however, are the functors $\mathcal{N}$ and $\mathcal{M}$ that yield neighborhood frames as their coalgebras.

It has been observed, see for example [4], that there is a relation lifting, we call it $\widetilde{\mathcal{M}}$, for the functor $\mathcal{M}$, that yields an adequate notion of bisimilarity between states in monotone neighborhood frames and is distinct from the Barr extension $\overline{\mathcal{M}}$ of $\mathcal{M}$. Moreover, Santocanale and Venema use $\widetilde{\mathcal{M}}$ in [17] to define a semantics for a well-behaved nabla modality on neighborhood frames. So the relation lifting $\widetilde{\mathcal{M}}$ of $\mathcal{M}$ fulfills the same roles (i) and (ii) that, as explained above, the Barr extensions plays for weak pullback preserving functors. This triggers the question under which conditions a relation lifting can be used for (i) and (ii).

The major contribution of this thesis is to show that relation liftings which are lax extensions and satisfy the further condition that they preserve diagonal relations can be used to fulfill the tasks (i) and (ii). Lax extensions that preserve diagonals are like functors in the category of relations with the difference that only one inclusion of the composition of relations is preserved. The relation lifting $\widetilde{\mathcal{M}}$ is a lax extension and preserves diagonals. As a negative result we show in Proposition 3.7 that, in a sense we will make more precise, there is no relation lifting for $\mathcal{N}$ that fulfills task (i). We also give a partial characterization of the functors that have a lax extension that preserves diagonals. So we prove in Theorem 3.26 that a finitary functor T has a lax extension that preserves diagonals iff it has a separating set of monotone predicate lifting. This theorem establishes a connection between the nabla logic of a lax extension and the other flavor of coalgebraic modal logic that uses predicate liftings. For the nabla logic of a lax extension we show that bisimulation quantifiers are definable in the logic if the lax extension satisfies an additional property, that we call quasi-functoriality. This generalizes the work done by Santocanale and Venema in [17] for the monotone neighborhood functor $\mathcal{M}$. An consequence from the definability of bisimulation quantifiers in the logic is that the logic has uniform interpolation.

The structure of this thesis is as follows. In section 2 we fix the notation and introduce the basic mathematical concepts that we use later. Section 3.2 is organized in three parts. In the first two subsections we define what a relation liftings and a bisimulation for a relation lifting is. In the third, subsection 3.3, we introduce lax extensions, prove some of their basic properties, and show that they can be used to define an adequate notion of bisimilarity if they preserve diagonals. We use the whole subsection 3.4 to prove Theorem 3.26. In section 4 we develop the nabla logic that has a semantics defined with help of a lax extension L along the lines of how this is done in [17] for $\widetilde{\mathcal{M}}$. In the first subsection 4.1 we define the semantics and show that it is adequate with respect to L-bisimilarity. In subsection 4.2 we show how one can eliminate conjunctions and negations from nabla formulas. This result we use in subsection 4.3 to show that bisimulation quantifiers are definable in the nabla logic of a lax extension.

2 Preliminaries

This section contains some of the preliminaries and fixes the notation. We presuppose that the reader has made contact with very basic concepts from category theory before. For example we presuppose the notions of a category, a commutative diagram, an isomorphism, an inverse or a functor between categories.

In order to see the motivation behind the concepts introduced later and to understand the examples the reader needs to know some modal logic on Kripke frames. An extensive introduction into modal logic is given for example in [3].

2.1 Sets, Functions and Relations

We will mainly work in the category Set that has sets as its objects and functions between sets as arrows. It is assumed that the reader is familiar with the usual constructions on sets so the following explanations are here to fix notation. We usually use capital Latin letters $X, Y, Z, \ldots, U, V, W, \ldots$ for sets and small Latin letters $f, g, \ldots$ for functions between sets. The notation $f : X \to Y$ means denote that f is a function with domain X and codomain Y. The identity element for a set X is the identity function $\mathsf{id}_X : X \to X$. The composition of two functions $f : X \to Y$ and $g : Y \to Z$ is the usual composition of functions written as $g \circ f : X \to Z$. An isomorphism in Set is a bijective function $f : X \to Y$ and its inverse is written as $f^{-1} : Y \to X$. Given a function $f : X \to Y$ and a set $X' \subseteq X$ we define the restriction of f to X' as $f{\restriction_{X'}}: X' \to Y, x \mapsto f(x)$. For sets $X' \subseteq X$ the inclusion of X' into X is the map $\iota_{X',X} : X' \to X, x \mapsto x$.

Another category that we will use a lot is the category Rel of relations between sets. Its arrows from a set X to a set Y are all the relations between X and Y. We use capital letters $R, S, \ldots$ for relations and write $R : X \nrightarrow Y$ to indicate that R is a relation between X and Y. A relation $R : X \nrightarrow Y$ as an arrow in the category Rel is not just a set of pairs, that is a subset of $X \times Y$, but it also contains information abouts its domain and codomain. We write R^{gr} for the set of pairs that encodes a relation $R : X \nrightarrow Y$. Note that $R : X \nrightarrow Y$ is an arrow in the category Rel whereas $R^{gr} \subseteq X \times Y$ is an object in Set or Rel. At some places, especially once we use relation liftings later, it matters what the domain and codomain of a relation are. Nevertheless we are often a bit sloppy with the notation and for example use $=$ and $\subseteq$ between relations that do not have the same domain or codomain.

The graph of any function $f : X \to Y$ is a relation between X and Y for which we write again $f : X \nrightarrow Y$. It will be clear from the context in which a symbol f occurs whether it is meant as the function $f : X \to Y$ in Set or as the relation $f : X \nrightarrow Y$ in Rel.

Identity elements in the category Rel are the diagonal relations $\Delta_X : X \nrightarrow X$ with $(x, x') \in \Delta_X$ iff $x = x'$. Note that $\Delta_X = \mathsf{id}_X$ if we consider $\mathsf{id}_X : X \nrightarrow X$ as a relation. The composition of two relations $R : X \nrightarrow Y$ and $S : Y \nrightarrow Z$ is written as $R\,;S : X \nrightarrow Z$ and defined by

$$R\,;S = \{(x,z) \in X \times Z \mid (x,y) \in R \text{ and } (y,z) \in S \text{ for a } y \in Y\}.$$

The composition of relations is written the other way round than the composition of functions. So we have, using the identification of functions with the relation of its graph, that $g \circ f = f\,;g$ for functions $f : X \to Y$ and $g : Y \to Z$.

For every relation $R : X \nrightarrow Y$ its converse $R^\circ : Y \nrightarrow X$ with $(y, x) \in R^\circ$ iff $(x, y) \in R$ is again a relation. The projections of a relation $R : X \nrightarrow Y$ are denoted by $\pi_X : R^{gr} \to X$ and $\pi_Y : R^{gr} \to Y$. It holds that $R = \pi_X^\circ\,;\pi_Y$. For any relation $R : X \nrightarrow Y$ we use $R^e : X \uplus Y \nrightarrow X \uplus Y$ for the smallest equivalence relation on the disjoint union $X \uplus Y$ of X and Y that contains all the pairs that are also in R.

There is an order $\subseteq$ on the relations between X and Y that is defined for relations $R', R : X \nrightarrow Y$ such that $R' \subseteq R$ iff $R'^{gr} \subseteq R^{gr}$. Every set $\mathcal{R}$ of relations between X and Y has an infimum $\bigcap \mathcal{R} : X \nrightarrow Y$ and supremum $\bigcup \mathcal{R} : X \nrightarrow Y$ with respect to the order $\subseteq$. They are just the usual intersection and union of the graphs.

For a relation $R : X \nrightarrow Y$ we define the sets

$$\begin{aligned} \mathsf{preimg}(R) &= \{x \in X \mid \exists y \in Y.(x,y) \in R\} \subseteq X, \\ \mathsf{img}(R) &= \{y \in Y \mid \exists x \in X.(x,y) \in R\} \subseteq Y. \end{aligned}$$

The relation $R : X \nrightarrow Y$ is *left-total* if $\mathsf{preimg}(R) = X$ and *right-total* if $\mathsf{img}(R) = Y$. Given sets $X' \subseteq X$ and $Y' \subseteq Y$ we define the restriction $R{\restriction_{X' \times Y'}} : X' \nrightarrow Y'$ of the relation $R : X \nrightarrow Y$ as $R{\restriction_{X' \times Y'}} = R \cap (X' \times Y')$.

For any set X let $\in_X : X \nrightarrow \mathcal{P}X$ be the membership relation between elements of X and subsets of X.

We are going to use some universal constructions in the category Set.

For every family $\mathcal{X}$ of sets we use $\prod \mathcal{X}$ to denote the product of all the sets in $\mathcal{X}$ with projections $\pi_X : \prod \mathcal{X} \to X$ for all $X \in \mathcal{X}$. The product of two sets X and Y is denoted by $X \times Y$ with projections $\pi_X : X \times Y \to X$ and $\pi_Y : X \times Y \to Y$.

For every family $\mathcal{X}$ of sets we use $\coprod \mathcal{X}$ to denote the coproduct of all the sets in $\mathcal{X}$ with injections $i_X : X \to \coprod \mathcal{X}$ for all $X \in \mathcal{X}$. One can think of the coproduct of $\mathcal{X}$ as the disjoint union $\biguplus \mathcal{X}$ of the sets in $\mathcal{X}$. The coproduct of two sets X and Y is denoted by $X + Y = X \uplus Y$ with injections $i_X : X \to X + Y$ and $i_Y : Y \to X + Y$.

Given two functions $f : X \to Z$ and $g : Y \to Z$ the *pullback* of f and g is the set

$$\mathsf{pb}(f, g) = (f \,;\, g^{\circ})^{gr} = \{(x, y) \in X \times Y \mid f(x) = g(y)\}$$

together with the projections $\pi_X : \mathsf{pb}(f, g) \to X$ and $\pi_Y : \mathsf{pb}(f, g) \to Y$. For these it holds that $f \circ \pi_X = g \circ \pi_Y$. The pullback of f and g is determined up-to-isomorphism by the universal property that for any other set Z' and functions $h : Z' \to X$ and $k : Z' \to Y$ that satisfy $f \circ h = g \circ k$ there is a unique function $m : Z' \to \mathsf{pb}(f, g)$ such that $h = \pi_X \circ m$ and $k = \pi_Y \circ m$. This universal property is depicted in the following diagram:

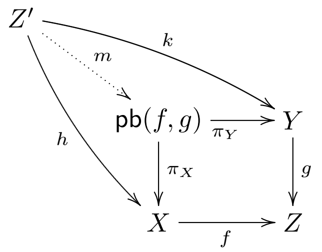

A *weak pullback* of two functions $f : X \to Z$ and $g : Y \to Z$ is any set P together with morphisms $p_X : P \to X$ and $p_Y : P \to Y$ that satisfies the same universal property as the pullback $\mathsf{pb}(f, g)$ with functions π_X and π_Y except the arrow m is not required to be unique. It is possible that two weak pullbacks P and P' of f and g are not isomorphic.

Given two functions $f : X \to Z$ and $g : Y \to Z$ the *pushout* of f and g is the set $\mathsf{po}(f, g) = (X + Y)/R^e$, that is the disjoint union of X and Y modulo the equivalence relation R^e, together with the projections $p_X : X \to \mathsf{po}(f, g), x \mapsto [i_X(x)]$ and $p_Y : Y \to \mathsf{po}(f, g), y \mapsto [i_Y(y)]$. For these it holds that $p_X \circ f = p_Y \circ g$. The pushout of f and g is determined up-to-isomorphism by a universal property that is dual of the universal property of the pullback. For any other set Z' and functions $h : X \to Z'$ and $k : Y \to Z'$ that satisfy $h \circ f = k \circ g$ there is a unique function $m : \mathsf{po}(f, g) \to R'$ such that $h = m \circ p_X$ and $k = m \circ p_Y$. This universal property is depicted in the following diagram:

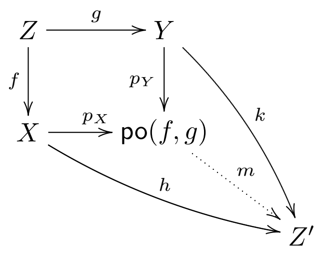

2.2 Set Functors

We will work with various set functors, that are functors from Set to Set or to Set^{op}. The category Set^{op} is like Set as it has sets as objects but an arrow f from X to Y in Set^{op} is a function $f : Y \to X$ from Y to X. A set functor with codomain Set is called covariant whereas one with codomain Set^{op} is contravariant. By default we assume that set functors a covariant and mention it explicitly if they are not.

A natural transformation $\lambda : F \Rightarrow G$ from a set functor F to a set functor G provides a function $\lambda_X : FX \to GX$ for every set X such that for all functions $f : X \to Y$ the following diagram commutes:

$$\begin{array}{ccc} FX & \xrightarrow{\lambda_X} & GX \\ \downarrow{\scriptstyle Ff} & & \downarrow{\scriptstyle Gf} \\ FY & \xrightarrow{\lambda_Y} & GY \end{array}$$

It is possible that F and G are both contravariant. In that case the arrows at the left and right sides of the above square are reversed. If each $\lambda_X : FX \to GX$ is an isomorphism, that means in Set that it is a bijective function, then λ is called a natural isomorphism and the functors F and G are said to be naturally isomorphic. In this case we can define $\lambda^{-1} : G \Rightarrow F$ with $(\lambda^{-1})_X = (\lambda_X)^{-1}$ which is automatically a natural transformation as well. One can think of two naturally isomorphic functors F and G as being the same in the sense that FX is always isomorphic to GX and for every function $f : X \to Y$ there are isomorphisms $\lambda_X : FX \to GY$ and $\lambda_Y : FY \to GY$ such that $Ff = \lambda_Y^{-1} \circ Gf \circ \lambda_X$.

Example 2.1. (i) The *powerset functor* $\mathcal{P}$ maps a set X to $\mathcal{P}X$, the set of all its subsets. A function $f : X \to Y$ is sent to

$$\begin{aligned} &\mathcal{P}f : \mathcal{P}X \to \mathcal{P}Y, \\ &\qquad U \mapsto f[U] = \{f(x) \in Y \mid x \in X\}. \end{aligned}$$

(ii) Similarly to the powerset functor the *contravariant powerset functor* $\breve{\mathcal{P}}$ maps a set X to $\breve{\mathcal{P}}X = \mathcal{P}X$. On functions $\breve{\mathcal{P}}$ is the inverse image map, that is for an $f : X \to Y$

$$\begin{aligned} &\breve{\mathcal{P}}f : \breve{\mathcal{P}}Y \to \breve{\mathcal{P}}X, \\ &\qquad V \mapsto f^{-1}[V] = \{x \in X \mid f(x) \in V\}. \end{aligned}$$

(iii) The identity functor Id is defined on sets as $\mathsf{Id}X = X$ and on functions as $\mathsf{Id}f = f$.

(iv) For every set C there is a constant functor C that sends a a set X to $CX = C$ and every morphism to id_C.

(v) For any collection of functors F_i for $i \in I$ where I is any index set the product $\prod_{i \in I} F_i-$ is again a functor that maps a set X to $\prod_{i \in I} F_i X$. It maps a function $f : X \to Y$ on $\prod_{i \in I} F_i f = (F_i f)_{i \in I}$ which is defined as

$$\begin{aligned} &(F_i f)_{i \in I} : \prod_{i \in I} F_i X \to \prod_{i \in I} F_i Y, \\ &\qquad (\xi_i)_{i \in I} \mapsto (F_i f(\xi_i))_{i \in I}. \end{aligned}$$

(vi) For any collection of functors F_i for $i \in I$ where I is any index set the coproduct $\coprod_{i \in I} F_i-$ is a functor. It maps a set X on the coproduct $\coprod_{i \in I} F_i X$ and a morphism $f : X \to Y$ such that

$$\begin{aligned} &\coprod_{i \in I} F_i f : \coprod_{i \in I} F_i X \to \coprod_{i \in I} F_i Y, \\ &\qquad \xi \mapsto i_{F_i Y} \circ F_i f(\xi'). \qquad \text{where } \xi = i_{F_i X}(\xi') \end{aligned}$$

(vii) The composition of two functors F and G is a functor written as $F \circ G$ or just FG which maps X to $F(GX)$ and $f : X \to Y$ to $F(Gf) : FGX \to FGY$.

(viii) The *neighborhood functor* or double contravariant powerset functor $\mathcal{N} = \breve{\mathcal{P}}\breve{\mathcal{P}}$ maps a set X to $\mathcal{N}X = \breve{\mathcal{P}}\breve{\mathcal{P}}X$ and a function $f : X \to Y$ to $\mathcal{N}f = \breve{\mathcal{P}}\breve{\mathcal{P}}f : \mathcal{N}X \to \mathcal{N}Y$ or more concretely for all $\xi \in \mathcal{N}X = \breve{\mathcal{P}}\breve{\mathcal{P}}X$

$$\mathcal{N}f(\xi) = \{V \subseteq Y \mid f^{-1}[V] \in \xi\}.$$

For any cardinal α there is an α-ary variant ${}^{\alpha}\mathcal{N}$ of $\mathcal{N}$ that maps a set X to ${}^{\alpha}\mathcal{N}X = \breve{\mathcal{P}}((\breve{\mathcal{P}}X)^{\alpha})$. So we have that the elements of ${}^{\alpha}\mathcal{N}X$ are sets of α-tuples of subsets of X.

For a $U \in \xi \in {}^{\alpha}\mathcal{N}X = (\breve{\mathcal{P}}(\breve{\mathcal{P}}X)^{\alpha})$ we write U_{β} for $U(\beta)$ that is the β-th component of U. So if α is a finite number, that is $\alpha = n \in \omega$, then then we have that $U = (U_0, U_1, \ldots, U_{n-1})$ for $U \in \xi$. A function $f : X \to Y$ is mapped by ${}^{\alpha}\mathcal{N}$ to ${}^{\alpha}\mathcal{N}f : {}^{\alpha}\mathcal{N}X \to {}^{\alpha}\mathcal{N}Y$ such that for all $\xi \in {}^{\alpha}\mathcal{N}X = \breve{\mathcal{P}}((\breve{\mathcal{P}}X)^{\alpha})$

$${}^{\alpha}\mathcal{N}f(\xi) = \{V \in (\breve{\mathcal{P}}Y)^{\alpha} \mid (f^{-1}[V_{\beta}])_{\beta \in \alpha} \in \xi\}.$$

(ix) A restriction of the neighborhood functor $\mathcal{N}$ from (viii) is the *monotone neighborhood functor* $\mathcal{M}$. It maps a set X to $\mathcal{M}X \subseteq \mathcal{N}X$ with the additional requirement that all $\xi \in \mathcal{M}X$ are upsets. That means that for all $U, U' \subseteq X$ if $U' \subseteq U$ and $U' \in \xi$ then also $U \in \xi$. On functions $\mathcal{M}$ is defined in the same way as $\mathcal{N}$. So we have for $f : X \to Y$ that

$$\begin{aligned} &\mathcal{M}f : \mathcal{M}X \to \mathcal{M}Y, \\ &\qquad \xi \mapsto \{V \subseteq Y \mid f^{-1}[V] \in \xi\}. \end{aligned}$$

One has to check that this is well defined. So we need that $\mathcal{M}f(\xi)$ is an upset if ξ is. So consider $V' \subseteq V \subseteq Y$ such that $V' \in \mathcal{M}f(\xi)$. That means that $f^{-1}[V'] \in \xi$ and so $f^{-1}[V] \in \xi$, hence $V \in \mathcal{M}f(\xi)$, because $f^{-1}[V'] \subseteq f^{-1}[V]$ and ξ is an upset.

There is also an α-ary version ${}^{\alpha}\mathcal{M}$ of $\mathcal{M}$ that is defined analogously to ${}^{\alpha}\mathcal{N}$ where the monotonicity requirement becomes that if $U'_{\beta} \subseteq U_{\beta}$ for all $\beta \in \alpha$ and $U' \in \xi$ then also $U \in \xi$. Similarly to the 1-ary case ${}^{1}\mathcal{M} = \mathcal{M}$ one can check that ${}^{\alpha}\mathcal{M}f$ is well defined.

(x) The functor F_2^3 maps a set X to

$$F_2^3 X = \{(x_0, x_1, x_2) \in X^3 \mid |\{x_0, x_1, x_2\}| \leq 2\}$$

the set of all triples over X that consist of at most two distinct elements. A function $f : X \to Y$ is mapped by F_2^3 as follows

$$\begin{aligned} &F_2^3 f : F_2^3 X \to F_2^3 Y, \\ &(x_0, x_1, x_2) \mapsto (f(x_0), f(x_1), f(x_2)). \end{aligned}$$

This functor is important since it is relatively simple and can often be used to construct counterexamples to seemingly obvious claims, due to its rather particular properties.

A very important property of set functors for the theory of coalgebras is weak pullback preservation. A functor T *preserves weak pullbacks* if it maps every weak pullback P with projections $\pi_X : P \to X$ and $\pi_Y : P \to Y$ of function $f : X \to Z$ and $g : Y \to Z$ onto a weak pullback TP with projections $T\pi_X : TP \to TX$ and $T\pi_Y : TP \to TY$ of $Tf : TX \to TZ$ and $Tg : TY \to TZ$. In diagrams that means that every weak pullback diagram on the left side is mapped to a weak pullback on the right side:

$$
\begin{array}{ccc}
P & \xrightarrow{\pi_Y} & Y \\
\downarrow{\scriptstyle \pi_X} & & \downarrow{\scriptstyle g} \\
X & \xrightarrow{f} & Z
\end{array}
\qquad\qquad
\begin{array}{ccc}
TP & \xrightarrow{T\pi_Y} & TY \\
\downarrow{\scriptstyle T\pi_X} & & \downarrow{\scriptstyle Tg} \\
TX & \xrightarrow{Tf} & TZ
\end{array}
$$

Example 2.2. From the functors introduced in Example 2.1 only F_2^3, ${}^\alpha\mathcal{N}$ and its monotone variant ${}^\alpha\mathcal{M}$ do not preserve weak pullbacks. Products, coproducts and composition of functors preserve weak pullbacks if all of their components do. A way to find out whether a functor F preserves weak pullbacks is to check whether its Barr extension $\overline{F}$ is functorial. This will be explained in Example 3.2 (vii).

A set functor T *restricts to finite sets* if TX is finite whenever X is. All the functors mentioned here, expect infinite products, restrict to finite sets.

A set functor T is *finitary* if it satisfies for all sets X

$$TX = \bigcup\{T\iota_{X',X}[TX'] \subseteq TX \mid X' \subset X, X' \text{ is finite}\}.$$

The idea behind this definition is that finitary functors have the property that in order to describe an element $\xi \in TX$ one has to use only a finite amount of information from the possibly infinite set X. This is important in the context of modal logic because usually formulas are defined to be finite objects. Finite functors have the property that every element in TX can be fully specified by one single finite formula.

Example 2.3. (i) Examples of finitary functors are: the identity functor, the constant functor C for any possibly infinite set C, finite products of finitary functors, any coproduct of finitary functors or the F_2^3 functor. The powerset functor $\mathcal{P}$ and neighborhood functors $\mathcal{N}$ and $\mathcal{M}$ are not finitary.

(ii) Every set functor T has a finitary version T_ω that is defined such that it maps a set X to

$$T_\omega X = \bigcup\{T\iota_{X',X}[TX'] \subseteq TX \mid X' \subseteq X, X' \text{ is finite}\}.$$

A function $f : X \to Y$ is mapped by T_ω to the function

$$
\begin{aligned}
T_\omega f : T_\omega X &\to T_\omega Y, \\
\xi &\mapsto T\iota_{f[X'],Y} \circ Tf_{X'}(\xi'),
\end{aligned}
$$

where $\xi' \in TX'$ is such that $\xi = \iota_{X',X}(\xi')$ for a finite $X' \subseteq X$ and $f_{X'}$ is the function $f_{X'} : X' \to f[X'], x' \mapsto f(x')$. This is well-defined because the following diagrams commutes for all $X', X'' \subseteq X$:

$$
\begin{array}{ccccc}
X' & \xrightarrow{\iota_{X',X}} & X & \xleftarrow{\iota_{X'',X}} & X'' \\
\downarrow{\scriptstyle f_{X'}} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f_{X''}} \\
f[X'] & \xrightarrow{\iota_{f[X'],Y}} & Y & \xleftarrow{\iota_{f[X''],Y}} & f[X'']
\end{array}
$$

It is immediate from the definition that $T_\omega X \subseteq TX$ for all sets X and that a functor T that is already finitary is identical to T_ω. In section 4 we will use $\mathcal{P}_\omega$ a lot. On can see by instantiating the above definition that this functor maps a set X to the set of all its finite subsets.

A set functor T is called *standard* if it preserves inclusions and all its distinguished points are standard. That T preserves inclusions means that $T\iota_{X',X} = \iota_{TX',TX}$ for all sets $X' \subseteq X$. Note that this in particular implies that $TX' \subseteq TX$ if $X' \subseteq X$. We do not explain the second condition, that all distinguished points are standard. For a precise definition of distinguished points consult [1, Chapter III, Definition 4.4] or [8, Appendix A]. All the functors we are considering do not have any distinguished points that are not standard.

In [1, Chapter III, page 132] it is proved that for every set functor T there is a standard functor T_s that is naturally isomorphic to it with the only possible exception of the empty set. If one examines the proof more carefully, one finds that the functor T_s that is constructed there is actually really isomorphic to T in the case where T has no distinguished points that are not standard. Since we are only looking at functors without non-standard distinguished points we can assume that there is always a standard functor T_s naturally isomorphic to T.

The basic idea behind the construction of T_s in [1] is to associate with any element $\xi \in TX$ the pair (X, ξ). Then we identify pairs by the equivalence relation

$$(X, \xi) \sim (Y, \upsilon) \quad \text{iff} \quad T\iota_{X,X\cup Y}(\xi) = T\iota_{Y,X\cup Y}(\upsilon).$$

The functor T_s is defined such that $T_s X = \{[X, \xi] \mid \xi \in TX\}$ where $[X, \xi]$ is the equivalence class of the pair (X, ξ) under the relation $\sim$. For the details of this construction consult [1, Chapter III, pages 132-134]

It is also shown in [1, Chapter III, Proposition 4.6] that standard functors distribute over finite intersections. That means that for all sets X and Y

$$T(X \cap Y) = TX \cap TY.$$

For a standard functor T we define for every set X the function

$$\begin{aligned}\mathsf{Base} : T_\omega X &\to \mathcal{P}_\omega X, \\ \xi &\mapsto \bigcap\{X' \subseteq X \mid \xi \in TX'\}.\end{aligned}$$

This is well defined because $\xi \in T_\omega X$ which means that there is a finite $X'' \subseteq X$ such that $\xi \in TX''$. The definition is useful because for all $\xi \in T_\omega X$ we have that $\mathsf{Base}(\xi) \in \mathcal{P}_\omega X$ is the least set $U \in \mathcal{P}_\omega X$ such that $\xi \in TU$. To see that $\xi \in T\mathsf{Base}(\xi)$ note that since $\xi \in TX''$ for a finite $X'' \subseteq X$ we can write $\mathsf{Base}(\xi)$ as the intersection of all the finitely many $X' \subseteq X''$ such that $\xi \in TX'$. Now because T preserves finite intersections and ξ is in all the finitely many TX' it is also in the T of the intersection of those X', which is $T\mathsf{Base}(\xi)$. It is clear from the definition of $\mathsf{Base}(\xi)$ that if $\xi \in T\mathsf{Base}(\xi)$ then $\mathsf{Base}(\xi)$ must also be the least set with this property.

Example 2.4. One can check that all the functors we are using except ${}^\alpha\mathcal{N}$ and ${}^\alpha\mathcal{M}$ are standard.

(i) The neighborhood functor is not standard. To verify this consider any non-empty set X and an $X' \subsetneq X$. Now take a $\xi \in \mathcal{N}X'$ such that $X' \in \xi$.

Clearly we have that $\iota^{-1}_{X',X}[X] = X' \in \xi$. So it follows from the definition of $\mathcal{N}$ on morphisms that $X \in \mathcal{N}\iota_{X',X}(\xi)$. But $X \notin \xi$ because $\xi \subseteq \mathcal{P}X'$ and $X' \subsetneq X$. So $\xi \neq \mathcal{N}\iota_{X,X'}(\xi)$ which shows that $\mathcal{N}\iota_{X,X'} \neq \iota_{\mathcal{N}X,\mathcal{N}X'}$.

(ii) One can use the same example as in (i) to see that $\mathcal{M}$ is not standard. As we have mentioned above there is a standard functor $\mathcal{M}_s$ that is naturally isomorphic to $\mathcal{M}$.

2.3 Coalgebras

In the remaining parts of this section we define the basic notions from the theory of coalgebras that we will use later. For a detailed introduction into coalgebras see for example [16] and for the coalgebraic modal logic [15] or [21].

Fix a covariant set functor T. A *T-coalgebra* on a set X is a function $\xi : X \to TX$. The elements of X are called the *states* of ξ and the function ξ is called the *transition function*. A *T-coalgebra morphism* from a T-coalgebra $\xi : X \to TX$ to a T-coalgebra $\zeta : Z \to TZ$ is a function $f : X \to Z$ that makes the following diagram commute:

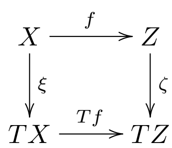

The T-coalgebras together with the T-coalgebra morphisms are a category where the identity arrow on one coalgebra $\xi : X \to TX$ is just the coalgebra morphism $\mathsf{id}_X : X \to X$ and the composition of two arrows is the composition of the underlying set functions.

Consider two states x_0 in a T-coalgebra $\xi : X \to TX$ and y_0 in $\upsilon : Y \to TY$. The states x_0 in ξ and y_0 in υ are *behaviorally equivalent* if there exists a T-coalgebra ζ and coalgebra morphisms f from ξ to ζ and g from υ to ζ such that $f(x_0) = g(y_0)$.

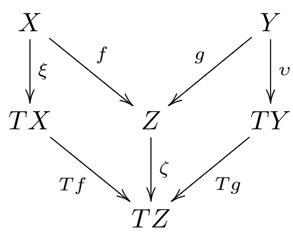

Example 2.5. Many different structures from automata theory and modal logic can be presented as coalgebra for some set functor. The following are some particularly important examples.

(i) Kripke frames are $\mathcal{P}$-coalgebras. This works because every relation $R : X \nrightarrow X$ can be presented as a function $R[\{-\}] : X \to \mathcal{P}X$ that maps every point to the set of its R-successors. One can also check that $\mathcal{P}$-coalgebra morphisms are exactly the bounded morphisms between Kripke frames and that two states are bisimilar iff they are behaviorally equivalent.

Similarly one can represent Kripke models as coalgebras for the functor $\mathcal{P}(\mathsf{P}) \times \mathcal{P}-$ where P is a set of propositional letters.

(ii) Deterministic automata are coalgebras for the functor $2 \times (-)^C$ where C is an alphabet. This functor associates with every state a truth value, that is an element from the set 2, which indicates whether the state is terminating, and a function from C into the set of states, which determines which state the automaton moves into after reading a letter from the alphabet.

(iii) Neighborhood frames that are used as a semantics for classical modal logics are coalgebras for the neighborhood functor $\mathcal{N}$. Coalgebras for the monotone neighborhood functor $\mathcal{M}$ are used in the semantics of monotone modal logic. For more on coalgebras and monotone modal logic see [4].

A construction that we us later is the coproduct of coalgebras. Given T-coalgebras $\xi_i : X_i \to TX_i$ for every $i \in I$ of an arbitrary index set I the coproduct $\xi = \coprod_{i \in I} \xi_i$ is defined to be a T-coalgebra $\xi : X \to TX$ where $X = \coprod_{i \in I} X_i = \biguplus_{i \in I} X_i$ with

$$\xi(x) = Ti_{X_i} \circ \xi_i(x'), \quad \text{where } x = i_{X_i}(x').$$

Here the $i_{X_i} : X_i \to X$ are the injections into the coproduct $\coprod_{i \in I} X_i$ in the category of sets. The injection from ξ_i into ξ as the coproduct in the category of T coalgebras is just the underlying set inclusions $i_X : X_i \to X$, which can be shown to be a coalgebra morphisms. One can easily check that ξ has the universal property of the coproduct in the category of T-coalgebras. In fact this is just an instantiation of the more general fact that every category of coalgebras has all colimits and that they are computed as in the category of sets.

A notion form coalgebraic modal logic that we are using later are predicate liftings. Predicate liftings for a functor T were originally introduced in [14] to define a modal logic for T-coalgebras that resembles the standard modal logic with boxes and diamonds on Kripke frames. More about this can be found in [14] but also the introductory texts [15] and [21].

An n-ary *predicate lifting* for a functor T is a natural transformation

$$\lambda : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T.$$

The *transposite* $\lambda^\flat$ of predicate lifting λ for a functor T is the mapping that is defined at a set X as

$$\begin{aligned} \lambda^\flat_X : TX &\to {}^n\mathcal{N}X = \breve{\mathcal{P}}\breve{\mathcal{P}}^n X \\ \xi &\mapsto \{U \in (\breve{\mathcal{P}}X)^n \mid \xi \in \lambda_X(U)\}. \end{aligned}$$

An n-ary predicate lifting $\lambda : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T$ is *monotone* if $U_i \subseteq U_i'$ for all $i \in n$ implies that $\lambda(U) \subseteq \lambda(U')$ for any $U, U' \in (\breve{\mathcal{P}}X)^n$.

Proposition 2.6. *If λ is an n-ary predicate lifting for T then:*

(i) Its transposite $\lambda^\flat : T \Rightarrow {}^n\mathcal{N}$ is a natural transformation.

(ii) If λ is monotone then the codomain of its transposite can be restricted to ${}^n\mathcal{M}$. That means $\lambda^\flat : T \Rightarrow {}^n\mathcal{M}$ defined as above is well defined.

Proof. For (i) we need to show that the following diagram commutes for all $f : X \to Y$:

$$\begin{array}{ccc} TX & \xrightarrow{\lambda^\flat_X} & {}^n\mathcal{N}X \\ \downarrow{\scriptstyle Tf} & & \downarrow{\scriptstyle {}^n\mathcal{N}f} \\ TY & \xrightarrow{\lambda^\flat_Y} & {}^n\mathcal{N}Y \end{array}$$

Because λ is a predicate lifting that is a natural transformation $\lambda : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T$ we know that

$$\begin{array}{ccc} \breve{\mathcal{P}}^n X & \xrightarrow{\lambda_X} & \breve{\mathcal{P}}TX \\ {\scriptstyle \breve{\mathcal{P}}^n f} \uparrow & & \uparrow {\scriptstyle \breve{\mathcal{P}}Tf} \\ \breve{\mathcal{P}}^n Y & \xrightarrow{\lambda_Y} & \breve{\mathcal{P}}TY \end{array}$$

commutes. So we need that $\breve{\mathcal{P}}\breve{\mathcal{P}}^n f \circ \lambda^\flat_X(\xi) = \lambda^\flat \circ Tf(\xi)$ for all $\xi \in TX$. The two sets are the same because we have for all $V \in \breve{\mathcal{P}}^n Y$ That

$$\begin{array}{rlll} V \in \breve{\mathcal{P}}\breve{\mathcal{P}}^n f \circ \lambda^\flat_X(\xi) & \text{iff} & V \in (\breve{\mathcal{P}}^n f)^{-1}[\lambda^\flat_X(\xi)] & \text{Definition of } \breve{\mathcal{P}} \\ & \text{iff} & \breve{\mathcal{P}}^n f(V) \in \lambda^\flat_X(\xi) & \text{Definition of } (\breve{\mathcal{P}}^n f)^{-1}[-] \\ & \text{iff} & \xi \in \lambda_X(\breve{\mathcal{P}}^n f(V)) & \text{Definition of } \lambda^\flat \\ & \text{iff} & \xi \in \breve{\mathcal{P}}Tf(\lambda_Y(V)) & \lambda \text{ natural transformation} \\ & \text{iff} & Tf(\xi) \in \lambda_Y(V) & \text{Definition of } \breve{\mathcal{P}} \\ & \text{iff} & V \in \lambda^\flat_Y(Tf(\xi)). & \text{Definition of } \lambda^\flat \end{array}$$

For (ii) we need to check for an arbitrary $\xi \in TX$ that the set

$$\lambda^\flat_X(\xi) = \{U \in (\breve{\mathcal{P}}X)^n \mid \xi \in \lambda_X(U)\} \in \breve{\mathcal{P}}\breve{\mathcal{P}}^n X$$

is upwards-closed if λ is monotone. This is fairly obvious: Take $U, U' \in \breve{\mathcal{P}}^n X$ such that $U_i \subseteq U'_i$ for all $i \in n$ and assume that $U \in \lambda^\flat_X(\xi)$. That means that $\xi \in \lambda_X(U) \subseteq \lambda_X(U')$ where the inclusion holds by the monotonicity of λ. But $\xi \in \lambda(U')$ entails by the definition of $\lambda^\flat$ that $U' \in \lambda^\flat_X(\xi)$. $\square$

In order to avoid tiresome compatibility issues when dealing with multiple monotone predicate liftings of possibly different finite arity it will be handy to compose them with the natural transformation $e^n : {}^n\mathcal{M} \Rightarrow {}^\omega\mathcal{M}$ defined by

$$\begin{aligned} e^n_X : {}^n\mathcal{M}X &\to {}^\omega\mathcal{M} \\ U &\mapsto (U_0, U_1, \ldots, U_{n-1}, \emptyset, \emptyset, \ldots). \end{aligned}$$

It is very straightforward to check that this indeed defines a natural transformation. With this we are going to write $e \circ \lambda^\flat : T \Rightarrow {}^\omega\mathcal{M}$ for $e^n \circ \lambda^\flat$ where $\lambda : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T$ is an n-ary natural transformation.

Next we define what a separating set of predicate liftings is. Intuitively a set of natural transformations for a functor T is separating if it is expressive enough to recognize every difference between elements in TX for any set X.

A family $\mathcal{F}$ of functions from X to Y is *jointly injective* if for any $x, x' \in X$ we have that $f(x) = f(x')$ for all $f \in \mathcal{F}$ implies that $x = x'$. A set Λ of predicate liftings for a functor T is *separating* if the set of functions $\{e \circ \lambda : TX \to {}^\omega\mathcal{M}X\}_{\lambda \in \Lambda}$ is jointly injective for every set X. That means that for all $\xi, \xi' \in TX$ if $e \circ \lambda(\xi) = e \circ \lambda(\xi')$ for all $\lambda \in \Lambda$ or equivalently, because e is injective, if $\lambda(\xi) = \lambda(\xi')$ for all $\lambda \in \Lambda$ then $\xi = \xi'$.

3 Relation Liftings and Bisimulations

In this section we use relation liftings for a functor T to define a very general notion of bisimulation for T coalgebras. It turns out that relation liftings that are lax extensions of T are particularly well-behaved. We discuss them in greater detail in the third and forth part of this section.

3.1 Relation Liftings

Definition 3.1 (Relation Lifting)**.** A *relation lifting* L for a set functor T is a collection of relations LR for every relation R, such that $LR : TX \nrightarrow TY$ if $R : X \nrightarrow Y$. We require relation liftings to preserve converses, this means that $L(R^\circ) = (LR)^\circ$ for all relations R.

The restriction that L preserves converses is not essential because all the notions we are considering are symmetrical. Note that in the above definition it matters what the domain and codomain of a relation are. It is possible to have a relation lifting that sends two relations, that have the same graph but different domain or codomain, to two different relations.

Example 3.2. (i) The *Egli-Milner lifting* $\overline{\mathcal{P}}$ is a relation lifting for the covariant powerset functor $\mathcal{P}$ that is defined such that $\overline{\mathcal{P}}R : \mathcal{P}X \nrightarrow \mathcal{P}Y$ for any $R : X \nrightarrow Y$ and $(U, V) \in \overline{\mathcal{P}}R$ iff

- for all $u \in U$ there is a $v \in V$ such that $(u, v) \in R$ (forth condition), and
- for all $v \in V$ there is a $u \in U$ such that $(u, v) \in R$ (back condition).

A more concise way to write this is $\overline{\mathcal{P}}R = \overrightarrow{\mathcal{P}}R \cap \overleftarrow{\mathcal{P}}R$ where we use the abbreviations

$$\overrightarrow{\mathcal{P}}R = \{(U, V) \in \mathcal{P}X \times \mathcal{P}Y \mid \forall u \in U.\exists v \in V.(u, v) \in R\},$$
$$\overleftarrow{\mathcal{P}}R = \{(U, V) \in \mathcal{P}X \times \mathcal{P}Y \mid \forall v \in V.\exists u \in U.(u, v) \in R\},$$

(ii) For the constant functor C of a fixed set C define a relation lifting $\overline{C}$ such that for any $R : X \nrightarrow Y$

$$\overline{C}R : C \nrightarrow C,$$
$$\overline{C}R = \Delta_C.$$

(iii) Let I be an arbitrary index set. If for all $i \in I$ T_i is a functor with relation lifting L_i then $\prod_{i\in I} L_i-$ defined on an $R : X \nrightarrow Y$ as $\prod_{i\in I} L_iR : \prod_{i\in I} T_iX \nrightarrow \prod_{i\in I} T_iY$ with

$$(\xi, \upsilon) \in \prod_{i\in I} L_iR \quad \text{iff} \quad (\xi_i, \upsilon_i) \in L_iR \text{ for all } i \in I$$

is a relation lifting for the functor $\prod_{i\in I} T_i-$.

(iv) Let I be an arbitrary index set. If for all $i \in I$ T_i is a functor with relation lifting L_i then $\coprod_{i\in I} L_i-$ defined on an $R : X \nrightarrow Y$ as $\coprod_{i\in I} L_iR : \coprod_{i\in I} T_iX \nrightarrow \coprod_{i\in I} T_iY$ with

$$(\xi, \upsilon) \in \coprod_{i\in I} L_iR \quad \text{iff} \quad (\xi, \upsilon) \in L_iR \text{ for some } i \in I$$

is a relation lifting for the functor $\coprod_{i\in I} T_i-$.

(v) If T is a functor with relation lifting L and T' is a functor with relation lifting L' then $L' \circ L = L'L-$ defined on a $R : X \nrightarrow Y$ as $L'LR : T'TX \nrightarrow T'TY$ is a relation lifting for $T' \circ T$.

(vi) Recall the notation $\overrightarrow{\mathcal{P}}R$ and $\overleftarrow{\mathcal{P}}R$ from item (i). We can define a relation lifting $\widetilde{\mathcal{M}}$ for the monotone neighborhood functor $\mathcal{M}$ on a relation $R : X \nrightarrow Y$ as

$$\begin{aligned}
&\widetilde{\mathcal{M}}R : \mathcal{M}X \nrightarrow \mathcal{M}Y \\
&\widetilde{\mathcal{M}}R = \overrightarrow{\mathcal{P}}\overleftarrow{\mathcal{P}}R \cap \overleftarrow{\mathcal{P}}\overrightarrow{\mathcal{P}}R.
\end{aligned}$$

One can also define the α-ary version of $\widetilde{\mathcal{M}}$ that maps an $R : X \nrightarrow Y$ on

$$\begin{aligned}
&\widetilde{{}^{\alpha}\mathcal{M}}R : {}^{\alpha}\mathcal{M}X \nrightarrow {}^{\alpha}\mathcal{M}Y \\
&\widetilde{{}^{\alpha}\mathcal{M}}R = \{(\xi, \upsilon) \mid \forall U \in \xi.\exists V \in \upsilon.\forall \beta \in \alpha.(U_\beta, V_\beta) \in \overleftarrow{\mathcal{P}}R\} \cap \\
&\qquad\qquad \{(\xi, \upsilon) \mid \forall V \in \upsilon.\exists U \in \xi.\forall \beta \in \alpha.(U_\beta, V_\beta) \in \overrightarrow{\mathcal{P}}R\}.
\end{aligned}$$

Similar to $\widetilde{\mathcal{M}}$ there is also a relation lifting $\widetilde{\mathcal{M}_s}$ for $\mathcal{M}_s$, the standardized version of $\mathcal{M}$. To define it, let $i : \mathcal{M}_s \Rightarrow \mathcal{M}$ be the natural isomorphism that witnesses that $\mathcal{M}$ and $\mathcal{M}_s$ are isomorphic. Now define $\widetilde{\mathcal{M}_s}R = i_X \,;\widetilde{\mathcal{M}}R\,; i_Y^\circ$ for any relation $R : X \nrightarrow Y$.

(vii) Items (i) and (ii) are instances of a relation lifting that is definable for arbitrary functors T. The *Barr extension* $\overline{T}$ of a functor T is a relation lifting for T that defined on a relation $R : X \nrightarrow Y$ with projections $\pi_X : R \to X$ and $\pi_Y : R \to Y$ such that

$$\overline{T}R = \{(T\pi_X(\rho), T\pi_Y(\rho)) \mid \rho \in TR^{gr}\}.$$

(viii) Another relation lifting that is definable for an arbitrary functor T is the lifting $\widehat{T}$, the introduction of which is attributed to Alexander Kurz in [5]. To see how it is defined consider a relation $R : X \nrightarrow Y$ with projections $\pi_X : R^{gr} \to X$ and $\pi_Y : R^{gr} \to Y$. Let $\mathsf{po}(\pi_X, \pi_Y)$ be the pushout of π_X and π_Y with projections $p_X : X \to \mathsf{po}(\pi_X, \pi_Y)$ and $p_Y : Y \to \mathsf{po}(\pi_X, \pi_Y)$. With these define

$$\begin{aligned}
&\widehat{T}R : TX \nrightarrow TY, \\
&\widehat{T}R = \{(\xi, \upsilon) \in TX \times TY \mid Tp_X(\xi) = Tp_Y(\upsilon)\}.
\end{aligned}$$

3.2 Bisimulations

An important use of relation liftings is to yield a notion of bisimilarity.

Definition 3.3 (Bisimilarity)**.** Let L be a relation lifting for the functor T and $\xi : X \to TX$ and $\upsilon : Y \to TY$ be two T-coalgebras. An *L-bisimulation between* ξ and υ is a relation $R : X \nrightarrow Y$ such that $(\xi(x), \upsilon(y)) \in LR$ for all $(x, y) \in R$. An L-bisimulation between ξ and ξ is also called an *L-bisimulation on* ξ. An *L-bisimulation equivalence* is a bisimulation on a single coalgebra that is also an equivalence relation.

A state x of ξ is *L-bisimilar* to a state y of υ if there is an $R : X \nrightarrow Y$ that is an L-bisimulation between ξ and υ with $(x, y) \in R$. We also write $\leftrightarroweq_L$ for the notion of L-bisimilarity between two fixed coalgebras that are given by the context.

Remark 3.4. One can check that a relation $R : X \nrightarrow Y$ is an L-bisimulation between coalgebra $\xi : X \to TX$ and $\upsilon : Y \to TY$ iff it satisfies the inequality

$$R \subseteq \xi \,;LR;\upsilon^\circ.$$

A motivation to define a notion of L-bisimulation is to get a simpler characterization of behavioral equivalence. Often it is easier to check whether there is a bisimulation between two states than to find two coalgebra morphisms into a third coalgebra that identify the states. Of course this only works for relation liftings for which the notion of bisimilarity is the same as behavioral equivalence.

Definition 3.5. A relation lifting L for a functor T *characterizes behavioral equivalence* if for any states x_0 in a T-coalgebra $\xi : X \to TX$ and y_0 in a T-coalgebra $Y \to TY$ it holds that $x_0 \leftrightarrows_L y_0$ iff x_0 and y_0 are behaviorally equivalent.

Example 3.6. (i) The Egli-Milner lifting $\overline{\mathcal{P}}$ for the powerset functor $\mathcal{P}$ characterizes behavioral equivalence and $\overline{\mathcal{P}}$-bisimulations are exactly the usual bisimulations between Kripke frames.

(ii) The Barr extension $\overline{T}$ for a functor T from Example 3.2 (vii) characterizes behavioral equivalence if T preserves weak pullbacks. This is a well-known fact and a consequence of the more general Proposition 3.15 that we are proving later.

There is also another definition of relations that are $\overline{T}$-bisimulations that does not make use of the notion of a relation lifting. One can check that a relation $R : X \nrightarrow Y$ with projections $\pi_X : R \to X$ and $\Pi_Y : R \to Y$ is a $\overline{T}$-bisimulation between coalgebra $\xi : X \to TX$ and $\upsilon : Y \to TY$ iff there is a map $\rho : R \to TR$ such that the following diagram commutes:

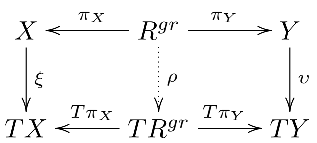

(iii) One can construct a counterexample that shows that the Barr extension $\overline{F_2^3}$ of the functor F_2^3 does not characterizes behavioral equivalence. This entails that F_2^3 does not preserve weak pullbacks.

(iv) The relation lifting $\widehat{T}$ characterizes behavioral equivalence for all the functors we discuss except the neighborhood functor $\mathcal{N}$. As for the Barr extension $\overline{\mathcal{P}}$ there is an alternative definition of a $\widehat{T}$-bisimulation. Let $R : X \nrightarrow Y$ be any relation with projections $\pi_X : R^{gr} \to X$ and $\pi_Y : R^{gr} \to Y$ and let $Z = \mathsf{po}(\pi_X, \pi_Y)$ be the pushout of π_X and π_Y with projections $p_X : X \to Z$ and $p_Y : Y \to Z$. Then one can verify that R is a $\widehat{T}$-bisimulation iff there is a function $\zeta : Z \to TZ$ such that the following diagram commutes:

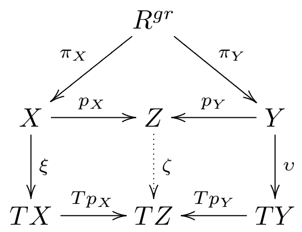

(v) The relation lifting $\widetilde{\mathcal{M}}$ for the monotone neighborhood functor $\mathcal{M}$ characterizes behavioral equivalence. For a detailed discussion of the different relation liftings for $\mathcal{M}$ check [4, Section 4]. There one can also find an example that shows that the Barr extension $\overline{\mathcal{M}}$ does not characterize behavioral equivalence, which entails that $\mathcal{M}$ does not preserve weak pullbacks.

(vi) There is no relation lifting for the neighborhood functor $\mathcal{N}$ that characterizes behavioral equivalence. This is shown by the counterexample in Proposition 3.7 below. The argument there only shows that there is no relation lifting that characterizes behavioral equivalence between two distinct $\mathcal{N}$-coalgebras. It is proved in [6] that on one single $\mathcal{N}$-coalgebra the relation liftings $\overline{\mathcal{N}}$ and $\widehat{\mathcal{N}}$ characterize behavioral equivalence.

Proposition 3.7. *There is no relation lifting for the neighborhood functor $\mathcal{N}$ that characterizes behavioral equivalence.*

Proof. For the proof we need the fact that for any two functions $f : X \to Z$ and $g : Y \to Z$ we have that $\mathcal{N}f(\{\emptyset\}) \neq \mathcal{N}g(\emptyset)$. This holds because otherwise we would get by unfolding the definition of $\mathcal{N}$ on functions that

$$\begin{aligned}
\emptyset &\in \{W \subseteq Z \mid f^{-1}[W] \in \{\emptyset\}\} && f^{-1}[\emptyset] = \emptyset \\
&= \mathcal{N}f(\{\emptyset\}) && \text{definition of } \mathcal{N} \\
&= \mathcal{N}g(\emptyset) && \text{assumption} \\
&= \{W \subseteq Z \mid g^{-1}[W] \in \emptyset\} && \text{definition of } \mathcal{N} \\
&= \emptyset. && V \notin \emptyset \text{ for all } V
\end{aligned}$$

This is clearly impossible.

Now assume for a contradiction that there is a relation lifting L such that L characterizes behavioral equivalence of states in $\mathcal{N}$ coalgebras. Consider the following example of a behavioral equivalence between coalgebras $\xi : X \to \mathcal{N}X$ where $X = \{x_1, x_2, x_3\}$ with $x_1 \mapsto \{\{x_2\}\}, x_2, x_3 \mapsto \{\emptyset\}$, $\upsilon : Y \to \mathcal{N}Y$ where $Y = \{y_1\}$ with $y_1 \mapsto \emptyset$ and $\zeta : Z \to \mathcal{N}Z$ with $Z = \{z_1, z_2\}$ with $z_1 \mapsto \emptyset, z_2 \mapsto \{\emptyset\}$. For these coalgebras the functions $f : X \to Z, x_1 \mapsto z_1, x_2, x_3 \mapsto z_2$ and $g : Y \to Z, y_1 \mapsto z_1$ are coalgebra morphisms from ξ to ζ and from υ to ζ. One can easily check this by verifying that $\mathcal{N}f \circ \xi = \zeta \circ f$ and $\mathcal{N}g \circ \upsilon = \zeta \circ g$. Because $f(x_1) = g(y_1)$ this shows that x_1 and y_1 are behaviorally equivalent.

The situation is depicted in the following figure:

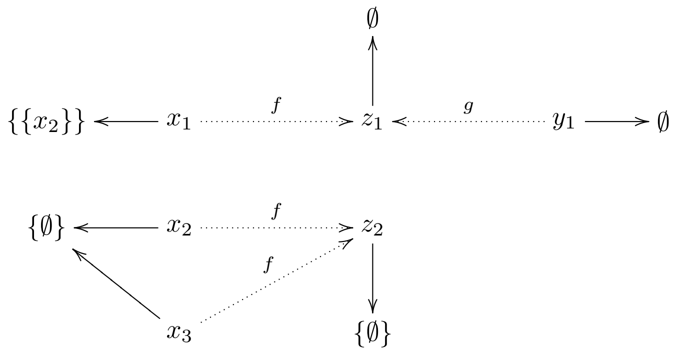

It follows from the assumption that L characterizes behavioral equivalence that there is an L-bisimulation $R : X \nrightarrow Y$ such that $(x_1, y_1) \in R$. Moreover we can show that $(x_2, y_1), (x_3, y_1) \notin R$. We do this for (x_2, y_1) since the argument for (x_3, y_1) is similar. So suppose $(x_2, y_1) \in R$. This means that x_2 and y_1 are L-bisimilar. Hence because L characterizes behavioral equivalence there are coalgebra morphisms h from ξ to ζ and l from υ to ζ such that $h(x_2) = l(y_1)$. Because h and l are coalgebra morphisms we get that

$$\mathcal{N}h(\{\emptyset\}) = \mathcal{N}h \circ \xi(x_2) = \zeta \circ h(x_2) = \zeta \circ l(y_1) = \mathcal{N}l \circ \upsilon(y_1) = \mathcal{N}l(\emptyset).$$

But we showed that this can not be the case. So it follows that $R = \{(x_1, y_1)\}$ and because R is an L-bisimulation that $(\{\{x_2\}\}, \emptyset) = (\xi(x_1), \upsilon(y_1)) \in LR$.

Next we modify the example a little by replacing ξ with the coalgebra $\xi' : X \to \mathcal{N}X, x_1 \mapsto \{\{x_2\}\}, x_2 \mapsto \{\emptyset\}, x_3 \mapsto \emptyset$. We still have that $(\xi'(x_1), \upsilon(y_1)) = (\{\{x_2\}\}, \emptyset) \in LR$ which entails that $R = \{(x_1, y_1)\}$ is an L-bisimulation between x_1 in ξ' and y_1 in υ. Because L characterizes behavioral equivalence it follows that there is a coalgebra $\zeta : Z \to \mathcal{N}Z$ and there are coalgebra morphisms h from ξ to ζ and k from υ to ζ such that $h(x_1) = k(y_1)$. Because h and k are coalgebra morphism this implies that

$$\mathcal{N}h(\{\{x_2\}\}) = \mathcal{N}h \circ \xi(x_1) = \zeta \circ h(x_1) = \zeta \circ k(y_1) = \mathcal{N}k \circ \upsilon(y_1) = \mathcal{N}k(\emptyset).$$

By writing out the definition of $\mathcal{N}$ one can see that this means

$$h^{-1}[C] \in \{\{x_2\}\} \quad \text{iff} \quad k^{-1}[C] \in \emptyset \quad \text{for all } C \subseteq Z.$$

Because the right hand side is never true it follows that $h^{-1}[C] \neq \{x_2\}$ for all $C \subseteq Z$. In the special case $C = \{h(x_2)\}$ this becomes $h^{-1}[\{h(x_2)\}] \neq \{x_2\}$. Certainly $x_2 \in h^{-1}[\{h(x_2)\}]$ so it follows that also $x_1 \in h^{-1}[\{h(x_2)\}]$ or $x_3 \in h^{-1}[\{h(x_2)\}]$. Thus $h(x_2) = h(x_1)$ or $h(x_2) = h(x_3)$. Using that h and k are coalgebra morphisms we can calculate in the former case that

$$\begin{aligned}\mathcal{N}h(\{\emptyset\}) &= \mathcal{N}h \circ \xi'(x_2) = \zeta \circ h(x_2) = \zeta \circ h(x_1) = \zeta \circ k(y_1) = \mathcal{N}k \circ \upsilon(y_1) \\ &= \mathcal{N}k(\emptyset)\end{aligned}$$

and in the latter case that

$$\mathcal{N}h(\{\emptyset\}) = \mathcal{N}h \circ \xi'(x_2) = \zeta \circ h(x_2) = \zeta \circ h(x_3) = \mathcal{N}h \circ \xi'(x_3) = \mathcal{N}k(\emptyset).$$

Hence $\mathcal{N}h(\{\emptyset\}) = \mathcal{N}k(\emptyset)$, which, as argued above, can not hold. □

The next Proposition describes the construction of bisimulation quotients for a wide class of relation liftings.

Proposition 3.8 (Bisimulation Quotient)**.** *Let L be a relation lifting for a functor T satisfying the condition that*

$$L(p\,;p^\circ) \subseteq Tp\,;(Tp)^\circ \text{ for all surjective } p : X \to Y \tag{1}$$

and let the relation $E : X \nrightarrow X$ be an L-bisimulation equivalence on a coalgebra $\xi : X \to TX$. Then there is a transition function $\delta : X/E \to T(X/E)$ such that the projection $p : Z \to Z/E$ is a coalgebra morphism from ξ to δ.

Proof. The projection is defined as

$$\begin{aligned} p : X &\to X/E, \\ x &\mapsto [x] \end{aligned}$$

where $[x]$ is the equivalence class of $x \in X$ under the equivalence relation E. This map is clearly surjective and we have that $E = p\,;p^\circ$.

We intend to define the transition function δ on X/E as

$$\begin{aligned} \delta : X/E &\to T(X/E), \\ [x] &\mapsto Tp \circ \xi(x). \end{aligned}$$

With this definition of δ it holds that $\delta \circ p = Tp \circ \xi$ which means that p is a coalgebra morphism from ξ to δ as required. But we have to show that δ is well defined. To prove this we need that $Tp \circ \xi(x) = Tp \circ \xi(x')$ for arbitrary $x, x' \in X$ with $(x, x') \in E$. Because E is an L-bisimulation it follows that $(\xi(x), \xi(x')) \in LE$ and moreover

$$\begin{aligned} LE &= L(p\,;p^\circ) && E = p\,;p^\circ \\ &\subseteq Tp\,;(Tp)^\circ. && (1) \end{aligned}$$

So we get $(\xi(x), \xi(x')) \in Tp\,;(Tp)^\circ$ which entails $Tp \circ \xi(x) = Tp \circ \xi(x')$. □

3.3 Lax Extensions

In this part introduce lax extensions. These are relation liftings satisfying certain conditions that make them well-behaved. We prove some general properties of lax extensions and show that they characterize behavioral equivalence. For some additional discussion of lax extensions we refer to [19]. Lax extension have also been studied under the name 'monotone relator' in [20, Section 2.1] and very recently in [12, Definition 6], where they are just called 'relators'. In [20] it is additionally required that composition of relation is preserved, that means $=$ instead of $\subseteq$ in our condition (L2) of Definition 3.9, but it is noted in [20] that the $\supseteq$-inclusion can be omitted for most of the proofs. Both [20] and [12] use a different set of conditions in their definitions, but it can be checked that they are equivalent to our Definition 3.9.

Definition 3.9. A relation lifting L for a functor T is a *lax extension* of T if it satisfies the following conditions for all relations $R, R' : X \nrightarrow Z$ and $S : Z \nrightarrow Y$, and functions $f : X \to Z$:

(L1) $R' \subseteq R$ implies $LR' \subseteq LR$,

(L2) $LR\,;LS \subseteq L(R\,;S)$,

(L3) $Tf \subseteq Lf$.

We say that a lax extension L *preserves diagonals* if it additionally satisfies:

(L4) $L\Delta_X \subseteq \Delta_{TX}$.

We require only the inclusion of (L4) for a lax extension to preserve diagonals. This is justified because condition (L4) implies together with condition (L3) that $L\Delta_X = \Delta_{TX}$. The proof of this is in the following Proposition which states some basic properties of lax extensions.

Proposition 3.10. *If L is a lax extension of T then for all functions $f : X \to Z$, $g : Y \to Z$ and relations $R : X \nrightarrow Z$, $S : Z \nrightarrow Y$:*

(i) $\Delta_{TX} \subseteq L\Delta_X$,

(ii) $Tf\,;LS = L(f\,;S)$ *and* $LR\,;(Tg)^\circ = L(R\,;g^\circ)$,

and if L preserves diagonals then

(iii) $L\Delta_X = \Delta_{TX}$ *and* $Lf = Tf$.

(iv) $Tf\,;(Tg)^\circ = L(f\,;g^\circ)$,

Proof. For (i) recall that we identify a function with the relation of its graph. So we have that $\Delta_X = \mathsf{id}_X$ and we can calculate

$$\begin{aligned} \Delta_{TX} = \mathsf{id}_{TX} &= T\mathsf{id}_X && T \text{ functor} \\ &\subseteq L\mathsf{id}_X = L\Delta_X. && \text{(L3)} \end{aligned}$$

The $\subseteq$-inclusion of $Tf\,;LS = L(f\,;S)$ in (ii) holds because $Tf\,;LS \subseteq Lf\,;LS \subseteq L(f\,;S)$ where the first inclusion is condition (L3) and the second inclusion is (L2). For the $\supseteq$-inclusion consider

$$\begin{aligned} L(f\,;S) &\subseteq Tf\,;(Tf)^\circ\,;L(f\,;S) && \Delta_{TX} \subseteq Tf\,;(Tf)^\circ \\ &\subseteq Tf\,;(Lf)^\circ\,;L(f\,;S) && \text{(L3)} \\ &\subseteq Tf\,;Lf^\circ\,;L(f\,;S) && \text{preservation of converses} \\ &\subseteq Tf\,;L(f^\circ\,;f\,;S) && \text{(L2)} \\ &\subseteq Tf\,;LS. && f^\circ\,;f \subseteq \Delta_Y \text{ and (L1)} \end{aligned}$$

For $LR\,;(Tg)^\circ = L(R\,;g^\circ)$ we can use the same argument and the fact that L preserves converses.

For (iv) and (iii) first notice that if L preserves diagonals then $L\Delta_X = \Delta_{TX}$ because of (L4) and (i).

The equation $Tf = Lf$ from (iii) holds because of

$$\begin{aligned} Lf &= L(f\,;\Delta_X) \\ &= Tf\,;L\Delta_X && \text{(ii)} \\ &= Tf. && L\Delta_X = \Delta_{TX} \end{aligned}$$

The claim (iv) holds because

$$
\begin{aligned}
Tf\,;(Tg)^\circ &= Tf\,;L\Delta_X\,;(Tg)^\circ && \Delta_{TX} = L\Delta_X \\
&= L(f\,;\Delta_X\,;g^\circ) && \text{(ii) twice} \\
&= L(f\,;g^\circ).
\end{aligned}
$$

□

Example 3.11. (i) It is easy to see that the Barr extension $\overline{T}$ of a functor T satisfies (L1). We can also show that $\overline{T}f = Tf$ for all function $f : X \to Y$. This means that $\overline{T}$ satisfies (L3) and (L4). For a proof consider the following commutative diagram and note that the projection $\pi_X : f^{gr} \to X$ is an isomorphism.

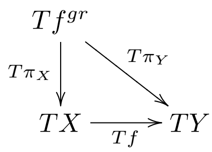

Condition (L2) for Barr extensions is more difficult. One can show that that $\overline{T}R\,;\overline{T}S = \overline{T}(R\,;S)$ for all relations $R : X \nrightarrow Z$ and $S : Z \nrightarrow Y$ iff T preserves weak pullbacks. See for example [10, Fact 3.6] for a proof of this claim. The $\subseteq$-inclusion of this proof really uses that T preserves weak pullbacks and we do not have an example of a functor T that does not preserve weak pullbacks for which $\overline{T}$ still satisfies (L2).

Also note that the condition $\overline{T}R\,;\overline{T}S = \overline{T}(R\,;S)$ for all relations $R : X \nrightarrow Z$ and $S : Z \nrightarrow Y$ is very strong. Together with $\overline{T}f = Tf$ for all function $f : X \to Y$ it means that $\overline{T}$ is a functor from Rel to Rel that extends T. Such an extension of a functor T is unique if it exists because for every relation $R : X \nrightarrow Y$ with projections $\pi_X : R^{gr} \to X$ and $\pi_Y : R^{gr} \to Y$ we have that

$$
\begin{aligned}
\overline{T}R &= \overline{T}(\pi_X^\circ\,;\pi_Y) && R = \pi_X^\circ\,;\pi_Y \\
&= \overline{T}\pi_X^\circ\,;\overline{T}\pi_Y && \overline{T}R\,;\overline{T}S = \overline{T}(R\,;S) \\
&= T(\pi_X^\circ)\,;T\pi_Y. && \overline{T}f = Tf
\end{aligned}
$$

(ii) The relation lifting $\widetilde{\mathcal{M}}$ as defined in Example 3.2 (vi) for the monotone neighborhood functor $\mathcal{M}$ is a lax extension that preserves diagonals. It is easy to check the conditions (L1) and (L2). To check (L3) we show that $(\xi, \mathcal{M}f(\xi)) \in \widetilde{\mathcal{M}}f$ for all functions $f : X \nrightarrow Y$ and $\xi \in \mathcal{M}X$. First pick any $U \in \xi$. Then we clearly have that $(U, f[U]) \in \overleftarrow{\mathcal{P}}f$ which shows $(\xi, \mathcal{M}f(\xi)) \in \overrightarrow{\mathcal{P}}\overleftarrow{\mathcal{P}}f$. Now take any $V \in \mathcal{M}f(\xi)$. That means that $f^{-1}[V] \in \xi$ and for this we have $(f^{-1}[V], V) \in \overrightarrow{\mathcal{P}}f$, and hence $(\xi, \mathcal{M}f(\xi)) \in \overleftarrow{\mathcal{P}}\overrightarrow{\mathcal{P}}f$. To check condition (L4) we prove that $\xi \subseteq \xi'$ for any $(\xi, \xi') \in \widetilde{\mathcal{M}}\Delta_X$. A similar argument shows $\xi \supseteq \xi'$ and hence $(\xi, \xi') \in \Delta_{\mathcal{M}X}$. So take any $U \in \xi$. It follows that there is a $U' \in \xi'$ such that $(U, U') \in \overleftarrow{\mathcal{P}}\Delta_X$. This means that $U \supseteq U'$ and because ξ' is an upset we get that $U \in \xi'$.

One can also check that $\widetilde{\mathcal{M}_s}$ is a lax extension for $\mathcal{M}_s$ and preserves diagonals. One can verify this fact directly but it is also a consequence of Proposition 3.18 that we prove later and the fact that, already using the terminology from

Definition 3.17, the relation lifting $\widetilde{\mathcal{M}_s} = \widetilde{\mathcal{M}}^{\{i\}}$ is the initial lift of $\widetilde{\mathcal{M}}$ along $\{i\}$ where $i : \mathcal{M}_s \to \mathcal{M}$ is the natural isomorphism between $\mathcal{M}_s$ and $\mathcal{M}$.

(iii) The product of lax extensions as defined in Example 3.2 (iii) is a lax extension. It preserves diagonals if all of its factors preserve diagonals.

(iv) The coproduct of lax extensions as defined in Example 3.2 (iv) is a lax extension. It preserves diagonals if all of its summands preserve diagonals.

(v) The compositions of two lax extensions as defined in Example 3.2 (v) is a lax extension. It preserves diagonals if the two composed lax extensions preserve diagonals.

(vi) The F_2^3 functor has a lax extension L_2^3 that preserves diagonals. L_2^3 is defined componentwise. That means for a relation $R : X \nrightarrow Y$ we have that

$$\begin{aligned}
&L_2^3 R : F_2^3 X \nrightarrow F_2^3 Y,\\
&L_2^3 R = \{((x_0, x_1, x_2), (y_0, y_1, y_2)) \mid (x_0, y_0), (x_1, y_1), (x_2, y_2) \in R\}.
\end{aligned}$$

It is straightforwards to verify that L_2^3 satisfies conditions (L1), (L2), (L3) and (L4).

The next proposition shows that for a lax extension for a standard functor it does not really matter what the domain and codomain of a relation are.

Proposition 3.12. *For any lax extension L of a standard functor T we have that for all relations $R : X \nrightarrow Y$ and sets $X' \subseteq X$ and $Y' \subseteq Y$*

$$L(R\restriction_{X' \times Y'}) = (LR)\restriction_{TX \times TY} .$$

Proof. We can rewrite the restriction as $R\restriction_{X' \times Y'} = \iota_{X',X} \,;R\,; \iota^\circ_{Y',Y}$ where $\iota_{X',X} : X' \to X$ and $\iota_{Y',Y} : Y' \to Y$ are inclusions. Then it follows for a lax extension L of a standard functor T that

$$\begin{aligned}
L(R\restriction_{X' \times Y'}) &= L(\iota_{X',X} \,; R\,; \iota^\circ_{Y',Y}) \\
&= T\iota_{X',X} \,; LR\,; (T\iota_{Y',Y})^\circ && \text{Proposition 3.10 (ii)}\\
&= \iota_{TX',TX} \,; LR\,; \iota^\circ_{TY'TY} && T \text{ standard}\\
&= (LR)\restriction_{TX' \times TY'} .
\end{aligned}$$

□

The conditions (L1), (L2) and (L3) of a lax extension L directly entail useful properties of L-bisimulations. The condition (L1) ensures that the union of L-bisimulations is again an L-bisimulation, (L2) yields that the composition of L-bisimulations is an L-bisimulation and because of (L3) coalgebra morphisms are L-bisimulations. This facts are summarized in the following Proposition.

Proposition 3.13. *For a lax extension L of T and T-coalgebras $\xi : X \to TX$, $\upsilon : Y \to TY$ and $\zeta : Z \to TZ$ it holds that*

(i) The graph of every coalgebra morphism f from ξ to υ is an L-bisimulation between ξ and υ.

(ii) If $R : X \nrightarrow Z$ respectively $S : Z \nrightarrow Y$ are L-bisimulations between ξ and ζ respectively ζ and υ then $R;S : X \nrightarrow Y$ is an L-bisimulation between ξ and υ.

(iii) Every arbitrary union of L-bisimulations between ξ and υ is again an L-bisimulation between ξ and υ.

Proof. For claim (i) we have to show that $(\xi(x), \upsilon(y)) \in Lf$ for arbitrary $(x, y) \in f$. Since f is a function $(x, y) \in f$ means just $y = f(x)$. Applying υ on both sides yields $\upsilon(y) = \upsilon \circ f(x) = Tf \circ \xi(x)$ where the latter equality holds because f is a coalgebra morphism. It follows that $(\xi(x), \upsilon(y)) \in Tf \subseteq Lf$ by (L3).

For claim (ii) we show $(\xi(x), \upsilon(y)) \in L(R\,;S)$ for any $(x, y) \in R\,;S$. From the choice of x and y we get a $z \in Z$ such that $(x, z) \in R$ and $(z, y) \in S$. Because R and S are L-bisimulations it follows that $(\xi(x), \zeta(z)) \in LR$ and $(\zeta(z), \upsilon(y)) \in LS$. Hence $(\xi(x), \upsilon(y)) \in LR\,;LS \subseteq L(R\,;S)$ where the inclusion is condition (L2) from the definitions of lax extensions.

For claim (iii) let $\mathcal{R}$ be a set of L-bisimulations between ξ and υ. Now pick an arbitrary $(x, y) \in \bigcup\mathcal{R}$. We need to show that $(\xi(x), \upsilon(y)) \in L(\bigcup\mathcal{R})$. From $(x, y) \in \bigcup\mathcal{R}$ it follows that there is an $R \in \mathcal{R}$ such that $(x, y) \in R$. Because R is an L-bisimulation we get that $(\xi(x), \upsilon(y)) \in LR \subseteq L(\bigcup\mathcal{R})$ where the inclusions holds by the monotonicity (L1) of lax extensions and the fact that $R \subseteq \bigcup\mathcal{R}$. □

Corollary 3.14. *Let L be a lax extension for T and $\xi : X \to TX$ and $\upsilon : Y \to TY$ be two T-coalgebras. The relation of L-bisimilarity $\leftrightarroweq_L : X \nrightarrow Y$ between two coalgebras is an L-bisimulation between ξ and υ. Moreover the relation of L-bisimilarity $\leftrightarroweq_L : X \nrightarrow X$ on one single coalgebra ξ is an equivalence relation.*

Proof. The relation $\leftrightarroweq_L : X \nrightarrow Y$ can in the following way be written as a union of L-bisimulations between ξ and υ:

$$\leftrightarroweq_L = \bigcup\{R : X \nrightarrow Y \mid R \text{ is an } L\text{-bisimulation between } \xi \text{ and } \upsilon\}.$$

Hence it follows by Proposition 3.13 (iii) that $\leftrightarroweq_L$ is an L-bisimulation between ξ and υ.

To check that L-bisimilarity $\leftrightarroweq_L : X \nrightarrow X$ on one single coalgebra ξ is an equivalence relation we need that it is reflexive, symmetric and transitive. For reflexivity observe that the graph of the coalgebra morphism id_X from ξ to ξ is an L bisimulation by claim (i) of Proposition 3.13. Symmetry follows from the assumption that all relation liftings we consider preserve converses. Transitivity follows by claim (ii) of Proposition 3.13. □

We can now prove that lax extensions that preserve diagonals characterize behavioral equivalence.

Proposition 3.15. *If L is a lax extension for T that preserves diagonals then a state x_0 in a T-coalgebra $\xi : X \to TX$ and a state y_0 in a T-coalgebra $\upsilon : Y \to TY$ are behaviorally equivalent iff they are L-bisimilar.*

Proof. For the direction from left to right assume that x_0 and y_0 are behaviorally equivalent. That means that there are T-coalgebra $\zeta : Z \to TZ$ and coalgebra morphisms f from ξ to ζ and g from υ to ζ such that $f(x_0) = g(y_0)$. To see that x_0 and y_0 are L-bisimilar observe that by Proposition 3.13 (i) (ii) the relation $f\,;g^\circ : X \nrightarrow Y$ is an L-bisimulation between ξ and υ because it is the composition of graphs from morphisms. This implies that x_0 and y_0 are L-bisimilar because $(x_0, y_0) \in f\,;g^\circ$.

For the other direction we show that given any L-bisimulation $R : X \nrightarrow Y$ between ξ and υ and $(x, y) \in R$ then there is a coalgebra $\zeta : Z \to TZ$ and coalgebra morphisms f from ξ to ζ and g from υ to ζ such that $f(x) = g(y)$.

Consider first the coproduct $\xi + \upsilon : X + Y \to T(X + Y)$ of ξ and υ with $i_X : X \to X + Y$ and $i_Y : Y \to X + Y$ as injections. Next define the relation $R' : (X + Y) \nrightarrow (X + Y)$ with $R' = i_X^\circ ; R ; i_Y$. By Proposition 3.13 we know that the graphs of the coalgebra morphisms i_X and i_Y are L-bisimulations and that the composition of L-bisimulations is an L-bisimulation. Therefore R' is an L-bisimulation on $\xi + \upsilon$ and witnesses the bisimilarity of $i_X(x)$ and $i_Y(y)$.

Now let $\Leftrightarrow_L : (X + Y) \nrightarrow (X + Y)$ be L-bisimilarity on $\xi + \upsilon$. By Corollary 3.14 $\Leftrightarrow_L$ is an L-bisimulation equivalence on $\xi + \upsilon$. Because L is a lax extension we have by Proposition 3.10 (iv) that $L(p ; p^\circ) \subseteq Tp ; (Tp)^\circ$ for all surjective $p : X \to Y$ and so can apply Proposition 3.8 to a $\zeta : Z \to TZ$ where $Z = (X + Y)/\Leftrightarrow_L$ that is the bisimulation quotient of $\xi + \upsilon$ by the relation $\Leftrightarrow_L$ such that the projection $p : (X + Y) \to Z$ is a coalgebra morphism from $\xi + \upsilon$ to ζ.

Let $f = p \circ i_X : X \to Z$ and $g = p \circ i_Y$. Clearly f is a coalgebra morphism from ξ to ζ and g is one from υ to ζ. Moreover they identify x and y because

$$\begin{aligned} f(x) &= p \circ i_X(x) = [i_X(x)] = [i_Y(y)] \qquad && i_X(x) \Leftrightarrow_L i_Y(y) \\ &= p \circ i_Y(y) = g(y). && \end{aligned}$$

□

Remark 3.16. From Proposition 3.15 and Proposition 3.7 it follows that there is no lax extension that preserves diagonals for the neighborhood functor $\mathcal{N}$.

3.4 Lax Extensions of Finitary Functors

The goal of this subsection is to prove Theorem 3.26 which says that a finitary functor has a lax extension that preserves diagonals iff it has a separating set of monotone predicate liftings. For the right to left direction we use the following construction from [19].

Definition 3.17. Let L be a relation lifting for a set functor T, and a set $\Lambda = \{\lambda : T' \to T\}_{\lambda \in \Lambda}$ of natural transformations from an other set functor T' to T. Then we can define a relation lifting L^Λ for T called the *initial lift of L along* Λ as

$$L^\Lambda R = \bigcap_{\lambda \in \Lambda} (\lambda_X ; LR ; \lambda_Y^\circ), \qquad \text{for all sets } X, Y \text{ and } R : X \nrightarrow Y.$$

Another way to write the above definition of $L^\Lambda R : T'X \nrightarrow T'Y$ is

$$L^\Lambda R = \{(\xi, \upsilon) \in T'X \times T'Y \mid (\lambda_X(\xi), \lambda_Y(\upsilon)) \in LR \text{ for all } \lambda \in \Lambda\}.$$

The good thing about the initial lift construction is that it preserves lax extensions.

Proposition 3.18. *Let $\Lambda = \{\lambda : T' \Rightarrow T\}_{\lambda \in \Lambda}$ be a set of natural transformations from a set functor T' to a set functor T and let L be a relation lifting for T. Then L^Λ is a lax extension for T' if L is a lax extension for T. Moreover if $\{\lambda_X : T'X \to TX\}_{\lambda \in \Lambda}$ is jointly injective at every set X and L preserves diagonals then L^Λ preserves diagonals.*

Proof. We need to show that the conditions of the definition of a lax extension are preserved along initial lifts. For (L1) take two relations $R, R' : X \nrightarrow Y$ with $R' \subseteq R$. From the assumption that L satisfies (L1) we get $LR' \subseteq LR$. Because composition and joins of relations clearly preserve order we can calculate

$$L^\Lambda R' = \bigcap_{\lambda \in \Lambda} (\lambda_X \,; LR' \,; \lambda_Y^\circ) \subseteq \bigcap_{\lambda \in \Lambda} (\lambda_X \,; LR \,; \lambda_Y^\circ) = L^\Lambda R.$$

For condition (L2) of the definition of lax extensions assume that $LR \,; LS \subseteq L(R \,; S)$. Then consider

$$\begin{aligned}
L^\Lambda R \,; L^\Lambda S &= \bigcap_{\lambda \in \Lambda} (\lambda_X \,; LR \,; \lambda_Y^\circ) \,;\, \bigcap_{\lambda \in \Lambda} (\lambda_Y \,; LS \,; \lambda_Z^\circ) \\
&\subseteq \bigcap_{\lambda \in \Lambda} (\lambda_X \,; LR \,; \lambda_Y^\circ \,; \lambda_Y \,; LS \,; \lambda_X^\circ) && \text{basic set theory} \\
&\subseteq \bigcap_{\lambda \in \Lambda} (\lambda_X \,; LR \,; LS \,; \lambda_X^\circ) && \lambda_Y^\circ \,; \lambda_Y \subseteq \Delta_{TY} \\
&\subseteq \bigcap_{\lambda \in \Lambda} (\lambda_X \,; L(R \,; S) \,; \lambda_X^\circ) && \text{assumption} \\
&= L^\Lambda(R \,; S).
\end{aligned}$$

For preservation of condition (L3) assume that $Tf \subseteq Lf$ for a function $f : X \to Y$. We show that then $T'f \subseteq L^\Lambda f$. Because the $\lambda \in \Lambda$ are natural transformations we have that $T'f \,; \lambda_Y = \lambda_X \,; Tf$. Composing with λ_Y° from right yields $T'f \subseteq \lambda_X \,; Tf \,; \lambda_Y^\circ$ because $\Delta_{T'Y} \subseteq \lambda_Y \,; \lambda_Y^\circ$. From the assumption it follows that $T'f \subseteq \lambda_X \,; Lf \,; \lambda_Y^\circ$ for all $\lambda \in \Lambda$. Hence

$$T'f \subseteq \bigcap_{\lambda \in \Lambda} (\lambda_X \,; LR \,; \lambda_Y^\circ) = L^\Lambda f.$$

In order to prove that L^Λ preserves diagonals, if L does and $\{\lambda_X : T'X \to TX\}_{\lambda \in \Lambda}$ is jointly injective at every set X, we show that $(L4)$ is preserved from L to L^Λ. For this we first show that if $\{\lambda_X : T'X \to TX\}_{\lambda \in \Lambda}$ is jointly injective at every set X then

$$\bigcap_{\lambda \in \Lambda} (\lambda_X \,; \lambda_X^\circ) = \Delta_{T'X}. \tag{2}$$

For the $\subseteq$-inclusion take $\xi, \xi' \in T'X$ with $(\xi, \xi') \in \bigcap_{\lambda \in \Lambda} (\lambda_X \,; \lambda_X^\circ)$. That means that $\lambda_X(\xi) = \lambda_X(\xi')$ for every $\lambda \in \Lambda$. Because the predicate liftings in Λ are jointly injective this implies that $\xi = \xi'$ and hence $(\xi, \xi') \in \Delta_{TX}$. The $\supseteq$-inclusion follows from the general fact that $f \,; f^\circ \supseteq \Delta_{TX}$ for any function f.

Now assume $L\Delta_X \subseteq \Delta_{TX}$. It follows that $L^\Lambda \Delta_X \subseteq \Delta_{T'X}$ because

$$\begin{aligned}
L^\Lambda \Delta_X &= \bigcap_{\lambda \in \Lambda} (\lambda_X \,; L\Delta_X \,; \lambda_X^\circ) && \text{definition} \\
&\subseteq \bigcap_{\lambda \in \Lambda} (\lambda_X \,; \Delta_{TX} \,; \lambda_X^\circ) && \text{assumption} \\
&= \bigcap_{\lambda \in \Lambda} (\lambda_X \,; \lambda_X^\circ) && \Delta_{TX} \text{ neutral element} \\
&= \Delta_{T'X}. && (2)
\end{aligned}$$

□

For the direction from left to right from Theorem 3.26 we are going to define the so called Moss liftings. It is shown in [11] that if we consider the Barr extension of a weak pullback preserving functor then the Moss liftings are monotone predicate liftings. Here we check that the argument also works for lax extensions.

The first ingredient that is needed to define the Moss liftings is the following natural transformation.

Definition 3.19. Given a lax extension L of a functor T we define for every set X the map

$$\begin{aligned} \lambda^L_X : T\breve{\mathcal{P}}X &\to \breve{\mathcal{P}}TX, \\ \Xi &\mapsto \{\xi \in TX \mid (\xi, \Xi) \in L{\in_X}\}, \end{aligned}$$

Proposition 3.20. *For a lax extension L the mapping $\lambda^L : T\breve{\mathcal{P}} \Rightarrow \breve{\mathcal{P}}T$ is a natural transformation.*

Proof. We have to verify that the following diagram commutes for any function $f : X \to Y$:

$$\begin{array}{ccc} T\breve{\mathcal{P}}X & \xrightarrow{\lambda^L_X} & \breve{\mathcal{P}}TX \\ {\scriptstyle T\breve{\mathcal{P}}f}\uparrow & & \uparrow{\scriptstyle \breve{\mathcal{P}}Tf} \\ T\breve{\mathcal{P}}Y & \xrightarrow{\lambda^L_Y} & \breve{\mathcal{P}}TY \end{array} \tag{3}$$

First observe that

$$L{\in_X} \,;(T\breve{\mathcal{P}}f)^\circ = Tf \,; L{\in_Y}. \tag{4}$$

This is shown by the calculation

$$\begin{aligned} L{\in_X} \,;(T\breve{\mathcal{P}}f)^\circ &= L\left(\in_X ;(\breve{\mathcal{P}}f)^\circ\right) && \text{Proposition 3.10 (ii)} \\ &= L(f \,; \in_Y) && (*) \\ &= Tf \,; L{\in_Y}. && \text{Proposition 3.10 (ii)} \end{aligned}$$

Here $(*)$ follows from the equality $\in_X ;(\breve{\mathcal{P}}f)^\circ = f \,; \in_Y$, which holds because by the definition of $\breve{\mathcal{P}}$ we have that $x \in_X \breve{\mathcal{P}}f(V) = f^{-1}[Y]$ iff $f(x) \in_Y V$ for all $x \in X$ and $V \subseteq Y$.

In order to check the commutativity of (3) take an $\Upsilon \in T\breve{\mathcal{P}}Y$. We need that $\breve{\mathcal{P}}Tf \circ \lambda^L_Y(\Upsilon) = \lambda^L_X \circ T\breve{\mathcal{P}}f(\Upsilon)$. This holds because for any $\xi \in TX$ we have that

$$\begin{aligned} \xi \in \lambda^L_X \circ T\breve{\mathcal{P}}f(\Upsilon) \quad &\text{iff} \quad (\xi, T\breve{\mathcal{P}}f(\Upsilon)) \in L{\in_X} && \text{definition of } \lambda^L \\ &\text{iff} \quad (\xi, \Upsilon) \in L{\in_X} \,;(T\breve{\mathcal{P}}f)^\circ && \text{basic set theory} \\ &\text{iff} \quad (\xi, \Upsilon) \in Tf \,; L{\in_Y} && (4) \\ &\text{iff} \quad (Tf(\xi), \Upsilon) \in L{\in_Y} && \text{basic set theory} \\ &\text{iff} \quad Tf(\xi) \in \lambda^L_Y(\Upsilon) && \text{definition of } \lambda^L \\ &\text{iff} \quad \xi \in (Tf)^{-1}[\lambda^L_Y(\Upsilon)] = \breve{\mathcal{P}}Tf \circ \lambda^L_Y(\Upsilon). && \text{definition of } \breve{\mathcal{P}} \end{aligned}$$

□

The second mathematical object we need to define the Moss liftings is a finitary presentation of the functor T.

Definition 3.21. A *finitary presentation* (Σ, E) *of* a functor T is a functor Σ of the form

$$\Sigma X = \coprod_{n \in \omega} \Sigma_n \times X^n$$

together with a surjective natural transformation $E : \Sigma \Rightarrow T$.

One can show, as we do in Example 3.22 (ii), that every finitary functor has a finitary presentation. A finitary presentation of T allows us to capture all the information in the sets TX for a possibly very complex functor T by means of a relatively simple polynomial functor Σ. This is, because for every $\xi \in TX$ there is at least one $(r, u) \in \Sigma_n \times X^n$ for an $n \in \omega$ for which $\xi = E_X(r, u)$ and that behaves in a similar way as ξ, since E is a natural transformation.

Example 3.22. (i) The standard presentation of the finitary powerset functor $\mathcal{P}_\omega$ is defined as $\Sigma = \coprod_{n \in \omega}(-)^n$ and

$$\begin{aligned} E_X : \coprod_{n \in \omega} X^n &\to \mathcal{P}_\omega X, \\ U &\mapsto \{U_i \in X \mid i \in n\}. \qquad \text{where } U \in X^n \text{ for an } n \in \omega \end{aligned}$$

It is obvious that E_X is surjective for every set X and one can easily verify that E is a natural transformation.

(ii) The next example shows that every finitary functor has a finitary presentation. The *canonical presentation* of a finitary functor T is defined such that $\Sigma_n = Tn$ for every cardinal $n \in \omega$ and E is defined at a set X as

$$\begin{aligned} E_X : \coprod_{n \in \omega} Tn \times X^n &\to TX, \\ (\nu, U) &\mapsto TU(\nu). \qquad \text{where } \nu \in Tn \text{ and } U \in X^n \text{ for an } n \in \omega \end{aligned}$$

In this definition we take $U \in X^n$ to be a map $U : n \to X$. To show that this is indeed a finitary presentation of T we have to check that E is a natural transformation and surjective for every set X.

For the naturality of E take any function $f : X \to Y$ and any element $(\nu, U) \in Tn \times X^n$ for any $n \in \omega$. Then we calculate

$$\begin{aligned} Tf \circ E_X(\nu, U) &= Tf \circ TU(\nu) && \text{definition of } E \\ &= T(f \circ U)(\nu) && T \text{ functor} \\ &= E_Y(\nu, f \circ U) && \text{definition of } E \\ &= E_Y \circ \Sigma f(\nu, U). && \text{definition of } \Sigma \text{ on functions} \end{aligned}$$

To see that E_X is surjective pick any $\xi \in TX$. Because T is finitary that means that there is a finite $X' \subseteq_\omega X$ and a $\xi' \in TX'$ such that $\xi = T\iota_{X',X}(\xi')$. Because X' is finite there is an $n \in \omega$ with a bijection $b : X' \to n$. Now we claim that $(Tb(\xi'), \iota_{X',X} \circ b^{-1}) \in Tn \times X^n$ is mapped by E_X to ξ. This is proved by the calculation

$$\begin{aligned} E_X(Tb(\xi'), \iota_{X',X} \circ b^{-1}) &= T(\iota_{X',X} \circ b^{-1})(Tb(\xi')) && \text{definition of } E \\ &= T(\iota_{X',X} \circ b^{-1} \circ b)(\xi') && T \text{ functor} \\ &= T\iota_{X',X}(\xi') && b^{-1} \circ b = \mathsf{id}_{X'} \\ &= \xi. && T\iota_{X',X}(\xi') = \xi \end{aligned}$$

The next Lemma shows how a lax extension for T interacts with a finitary presentation of T. This Lemma is similar to one direction of [11, Lemma 6.3] where this result is proved for the Barr extension. One can use the lax extension L_2^3 of F_2^3 to construct an example which shows that the back direction of [11, Lemma 6.3] does not hold for lax extensions in general.

Lemma 3.23. *Let (Σ, E) be a presentation of a finitary functor T, let L be a lax extension for T and let $R : X \nrightarrow Y$ be any relation. Then for all $n \in \omega$, $r \in \Sigma_n$, $u \in X^n$ and $v \in Y^n$ we have that if $u_i R v_i$ for all $i \in n$ then $(E_X(r, u), E_Y(r, v)) \in LR$.*

Proof. Let $\pi_Y : R \to X$ and $\pi_Y : R \to Y$ be the projections of R. For these it holds that $R = \pi_X^\circ \,; \pi_Y$. Because $(u_i, v_i) \in R$ for all $i \in n$ we have that $\rho = (r, ((u_0, v_0), (u_1, v_1), \ldots, (u_{n-1}, v_{n-1}))) \in \Sigma R^{gr}$. With the definition of Σ on morphisms it holds that $\Sigma\pi_X(\rho) = (r, u)$ and $\Sigma\pi_Y(\rho) = (r, v)$. The following two diagrams commute because $E : \Sigma \Rightarrow T$ is a natural transformation

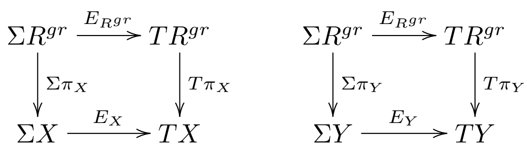

We can use this to get that $E_X(r, u) = E_X(\Sigma\pi_X(\rho)) = T\pi_X(E_R(\rho))$ and $E_Y(r, v) = E_Y(\Sigma\pi_Y(\rho)) = T\pi_Y(E_R(\rho))$. It is entailed by this identities that $(E_X(r, u), E_R(\rho)) \in (T\pi_X)^\circ$ and that $(E_R(\rho), E_Y(r, v)) \in T\pi_Y$. So we obtain

$$\begin{aligned} (E_X(r, u), E_Y(r, v)) \in (T\pi_X)^\circ \,; (T\pi_Y) &\subseteq L\pi_X^\circ \,; L\pi_Y && \text{(L3)} \\ &\subseteq L(\pi_X^\circ \,; \pi_Y) = LR. && \text{(L2)} \end{aligned}$$

□

Definition 3.24. Given a finitary functor T and a lax extension L for T take any finitary presentation (Σ, E) of T according to Definition 3.21 and let λ^L be the natural transformation of Definition 3.19. For every $r \in \Sigma_n$ of any $n \in \omega$ the *Moss lifting* of r is an n-ary predicate lifting for T that is defined as $\mu^p : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T, \mu^r = \lambda^L \circ E_{\breve{\mathcal{P}}}(r, -)$. This definition yields the following diagram for every set X:

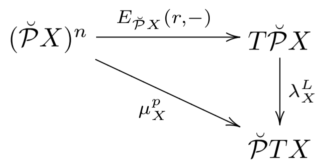

Proposition 3.25. *The Moss liftings of a functor T with finitary presentation (Σ, E) and lax extension L are monotone.*

Proof. Take any Moss lifting $\mu^r = \lambda^L \circ E_{\breve{\mathcal{P}}}(p, -) : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T$ of a $r \in \Sigma_n$ for an $n \in \omega$. Now assume we have $U, U' \in (\breve{\mathcal{P}}X)^n$ for any set X such that $U_i \subseteq U_i'$ for all $i < n$. To prove that μ^r is monotone we need to show that $\mu_X^r(U) \subseteq \mu_X^r(U')$.

So pick any $\xi \in \mu_X^r(U) = \lambda_X^L \circ E_{\breve{\mathcal{P}}X}(r, U)$. By the definition of λ^L this means that $(\xi, E_{\breve{\mathcal{P}}X}(r, U)) \in L{\in_X}$. Moreover we get from the assumption that

$U_i \subseteq U'_i$ for all $i \in n$ with Lemma 3.23 that $(E_{\breve{\mathcal{P}}X}(r,U), E_{\breve{\mathcal{P}}X}(r,U')) \in L(\subseteq)$. Putting this together yields

$$\begin{aligned}(\xi, E_{\breve{\mathcal{P}}X}(r,U')) \in L{\in_X} \,;\, L(\subseteq) &\subseteq L(\in_X \,;\, (\subseteq)) && \text{(L2)}\\ &\subseteq L{\in_X}. && \text{(L1)}\end{aligned}$$

For the last inequality we need that $\in_X ; (\subseteq) \subseteq \in_X$ which is immediate from the definition of subsets. So we have that $(\xi, E_{\breve{\mathcal{P}}X}(r,U')) \in L{\in_X}$ and hence by the definition of λ^L that $\xi \in \lambda^L_X \circ E_{\breve{\mathcal{P}}X}(r,U') = \mu^r_X(U')$. □

Theorem 3.26. *Let T be a finitary functor. Then T has a lax extension that preserves diagonals iff there is a separating set of monotone predicate liftings with finite arity for T.*

Proof. We first show the direction from right to left. So assume we have a separating set Λ of monotone predicate liftings with finite arity for T. By Proposition 2.6 (ii) the monotonicity of $\lambda \in \Lambda$ entails that we can consider $\lambda^\flat : T \to {}^n\mathcal{N}$ to have codomain ${}^n\mathcal{M}$. That Λ is separating, means that the set of functions $\{e_X \circ \lambda^\flat_X : TX \Rightarrow {}^\omega\mathcal{M}X\}_{\lambda\in\Lambda}$, where $e : {}^n\mathcal{M} \Rightarrow {}^\omega\mathcal{M}$ is the embedding as defined in section 2.3, is jointly injective at every set X. Therefore we can apply Proposition 3.18 for the set $\Gamma = \{e \circ \lambda^\flat : T \Rightarrow {}^\omega\mathcal{M}\}_{\lambda\in\Lambda}$ and get that the initial lift $(\widetilde{{}^\omega\mathcal{M}})^\Gamma$ of $\widetilde{{}^\omega\mathcal{M}}$ along Γ as defined in Definition 3.17 is a lax extension for T that preserves diagonals.

For the other direction we assume that T has a lax extension L. Because T is finitary it has a finitary presentation (Σ, E) as demonstrated in Example 3.22 (ii). With this we can consider the set of all the Moss liftings as defined in Definition 3.24:

$$M = \{\mu^r = \lambda^L \circ E(r,-) \mid r \in \Sigma_n, n \in \omega\}.$$

By Proposition 3.25 we know that the Moss liftings are monotone. So it only remains to show that the set M is separating. To do this take $\xi, \xi' \in TX$ for an arbitrary set X such that $(\mu^r)^\flat_X(\xi) = (\mu^r)^\flat_X(\xi')$ for all $r \in \Sigma_n$ of all $n \in \omega$. We have to show that $\xi = \xi'$. By the definition of the transposite of a natural transformation it follows that for all $n \in \omega$ and $r \in \Sigma_n$

$$\{U \in (\breve{\mathcal{P}}X)^n \mid \xi \in \mu^r_X(U)\} = \{U \in (\breve{\mathcal{P}}X)^n \mid \xi' \in \mu^r_X(U)\}.$$

This is equivalent to

$$\xi \in \mu^r_X(U) \quad \text{iff} \quad \xi' \in \mu^r_X(U), \text{ for all } U \in (\breve{\mathcal{P}}X)^n.$$

Unfolding the definitions of $\mu^r = \lambda^L \circ E_{\breve{\mathcal{P}}}(r,-)$ and $\lambda^L(\Xi) = \{\xi \in TX \mid (\xi,\Xi) \in L{\in_X}\}$ yields that for all $n \in \omega$, $r \in \Sigma_n$ and $U \in (\breve{\mathcal{P}}X)^n$

$$(\xi, E_{\breve{\mathcal{P}}X}(r,U)) \in L{\in_X} \quad \text{iff} \quad (\xi', E_{\breve{\mathcal{P}}X}(r,U)) \in L{\in_X}.$$

Because $E_{\breve{\mathcal{P}}X}$ is surjective, and the variables n, r and U quantify over the whole domain of $E_{\breve{\mathcal{P}}X} : \coprod_{n\in\omega} \left(\Sigma_n \times (\breve{\mathcal{P}}X)^n\right) \to T\breve{\mathcal{P}}X$, it follows that for all $\Xi \in T\mathcal{P}X$

$$(\xi, \Xi) \in L{\in_X} \quad \text{iff} \quad (\xi', \Xi) \in L{\in_X}. \tag{5}$$

In order to use (5) consider the map

$$\begin{aligned} s_X : X &\to \mathcal{P}X, \\ x &\mapsto \{x\}. \end{aligned}$$

Because of (L3) we have that $(\xi, Ts_X(\xi)) \in Ts_X \subseteq Ls_X$. Moreover we clearly have that $s_X \subseteq \in_X$ and because of (L1) it follows that $(\xi, Ts_X(\xi)) \in L{\in_X}$. With (5) we get that $(\xi', (T\mu_X)\xi) \in L{\in_X}$. Then we compute

$$\begin{aligned} (\xi, \xi') \in Ls_X \,;\, L(\ni_X) &\subseteq L(s_X \,;\, (\ni_X)) && \text{(L2)} \\ &= L(\Delta_X) && s_X \,;\, \ni_X = \Delta_X \\ &= \Delta_{TX}. && \text{Proposition 3.10 (iii)} \end{aligned}$$

From this it follows that $\xi = \xi'$. □

4 The Nabla Logic of a Lax Extension

In this section we study the nabla logic $\mathcal{L}_T$ for a fixed lax extension L of a fixed standard functor T. The assumption that T is standard is needed to get a well-behaved syntax. As we observed in section 2.2 on page 9 this is not an essential restriction because every set functor is almost isomorphic to a standard functor. We also fix an arbitrary set P of propositional letters.

4.1 Syntax and Semantics

In this subsection we define the syntax and semantics of the nabla logic for the lax extension L and prove that it is adequate and expressive with respect to L-bisimilarity. This shows that the logic is strong enough to describe properties of states in coalgebras up to L-bisimilarity.

In order to give a semantics for the language $\mathcal{L}_T$ on T-coalgebras we also have to give an interpretation for the propositional letters. This is done by adding a valuation to T-coalgebras yielding T-models.

Definition 4.1. A *T-model* $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ is a T-coalgebra $\xi : X \to TX$ together with a *valuation* $V_{\mathbb{X}}$ that is a function $V_{\mathbb{X}} : \mathsf{P} \to \mathcal{P}X$.

For a $C \subseteq \mathsf{P}$ define the functor $T_C = (T-) \times \mathcal{P}C$ and the relation lifting $L_C = L - \times \overline{\mathcal{P}C}$ for the functor T_C as in Example 3.2 (iii). Here $\mathcal{P}C$ is a constant functor and $\overline{\mathcal{P}C}$ the Barr extension of the constant functor. So we have that $((c, \xi), (c', \xi')) \in L_C R$ iff $c = c'$ and $(\xi, \xi') \in LR$.

There is an one-to-one correspondence between T-models and T_{P}-coalgebras. A T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ corresponds to the T_{P}-coalgebra

$$\begin{aligned} \widehat{\mathbb{X}} : X &\to T_{\mathsf{P}}X = TX \times \mathcal{P}\mathsf{P}, \\ x &\mapsto (\xi(x), \{p \in \mathsf{P} \mid x \in V_{\mathbb{X}}(p)\}). \end{aligned}$$

In the other direction there is for every T_{P}-coalgebra $\sigma : X \to T_{\mathsf{P}}X$ a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ defined as $\xi = \pi_{TX} \circ \sigma : X \to TX$ and $V_{\mathbb{X}} : \mathsf{P} \to \mathcal{P}X, p \mapsto \{x \in X \mid p \in \pi_{\mathcal{P}\mathsf{P}} \circ \sigma(x)\}$ where $\pi_{TX} : \mathcal{P}TX \times \mathcal{P}\mathsf{P} \to TX$ and $\pi_{\mathcal{P}\mathsf{P}} : TX \times \mathcal{P}\mathsf{P} \to \mathcal{P}\mathsf{P}$ are the projections.

A morphism f from a T-model $\mathbb{X}$ to a T-model $\mathbb{Y}$ is defined to be a T_{P}-coalgebra morphism from $\widehat{\mathbb{X}}$ to $\widehat{\mathbb{Y}}$.

For $C' \subseteq C \subseteq \mathsf{P}$ we can define the natural transformation $r^{C,C'} : T_C \Rightarrow T_{C'}$ at a set X as

$$\begin{aligned} r_X^{C,C'} &: T_C X \to T_{C'} X, \\ &(\alpha, c) \mapsto (\alpha, c \cap C'). \end{aligned}$$

The *projection* of a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ to a $C \subseteq \mathsf{P}$ is the T_C-coalgebra $r_X^{\mathsf{P},C} \circ \widehat{\mathbb{X}} : X \to T_C X$.

An L_C-bisimulation between a T-model $\mathbb{X}$ and a T-model $\mathbb{Y}$ is defined to be an L_C-bisimulation between the T_C-coalgebras $r_X^{\mathsf{P},C} \circ \widehat{\mathbb{X}} : X \to T_C X$ and $r_Y^{\mathsf{P},C} \circ \widehat{\mathbb{Y}} : Y \to T_C Y$.

Because T-models are just T_C-coalgebras and bisimulations between them are the usual L_C-bisimulations between them we can take the product of T-models and use the constructions of Proposition 3.13 on L_C-bisimulations between T-models.

It follows directly from the definitions that a relation $R : X \nrightarrow Y$ is a L_C-bisimulation between T-models $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ and $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ iff R is an L-bisimulation between the T-coalgebras $\xi : X \to TX$ and $\upsilon : Y \to TY$ and R preserves the truth of all propositional letters in C, that is for all $(x, y) \in R$ we have that

$$x \in V_{\mathbb{X}}(p) \quad \text{iff} \quad y \in V_{\mathbb{Y}}(p), \qquad \text{for all } p \in C.$$

With this it is easy to see that for any $C' \subseteq C$ if a relation R is an L_C-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$ then it is also an $L_{C'}$-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$.

Next we define the syntax of the language $\mathcal{L}_T(C)$ for a $C \subseteq \mathsf{P}$.

Definition 4.2. For any $C \subseteq \mathsf{P}$ define the language $\mathcal{L}_T(C)$ by the grammar

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

where $p \in C$, $A \in \mathcal{P}_\omega \mathcal{L}_T$ and $\alpha \in T_\omega \mathcal{L}_T$.

We use $\mathcal{L}_T$ as an abbreviation for $\mathcal{L}_T(\mathsf{P})$.

Set $\top = \bigwedge \emptyset$ and $\bot = \bigvee \emptyset$.

Later we use the maps $\neg : \mathcal{L}_T \to \mathcal{L}_T, a \mapsto \neg a$ and $\wedge : \mathcal{P}_\omega \mathcal{L}_T \to \mathcal{L}_T, A \mapsto \bigwedge A$.

Observe that for $C' \subseteq C \subseteq \mathsf{P}$ we have that $\mathcal{L}_T(C') \subseteq \mathcal{L}_T(C)$. This can be proved by induction on the complexity of formulas in $\mathcal{L}_T(C')$.

We use arbitrary finitary conjunctions and disjunctions instead of the usual binary ones. This approach yields an equivalent logic but it facilitates the notation.

The recursive definition of the formulas in $\mathcal{L}_T(C)$ allows us to define functions or relations for formulas by recursion and prove properties about them by induction on their complexity. In the case of $\nabla \alpha$ with $\alpha \in T_\omega \mathcal{L}_T(C)$ this means that if we want to make a definition for $\nabla \alpha$ we can assume that we already have a definition for all the formulas in the set $\mathsf{Base}(\alpha)$ and if we want to prove something by induction for $\nabla \alpha$ we can assume that the claim already holds for all formulas in $\mathsf{Base}(\alpha)$. Recall from section 2.2 page 9 that $\mathsf{Base}(\alpha)$ is the smallest set U such that $\alpha \in TU$. An example of a recursive definition is the following.

Definition 4.3. The *modal rank* $\mathsf{rank}(a) \in \omega$ of a formula $a \in \mathcal{L}_T$ is defined recursively by

$$\begin{aligned} \mathsf{rank}(p) &= 0, & p &\in \mathsf{P} \\ \mathsf{rank}(\neg a) &= \mathsf{rank}(a), & a &\in \mathcal{L}_T \\ \mathsf{rank}(\bigwedge A) &= \max\{\mathsf{rank}(a) \mid a \in A\}, & A &\in \mathcal{P}_\omega \mathcal{L}_T \\ \mathsf{rank}(\bigvee A) &= \max\{\mathsf{rank}(a) \mid a \in A\}, & A &\in \mathcal{P}_\omega \mathcal{L}_T \\ \mathsf{rank}(\nabla \alpha) &= 1 + \max\{\mathsf{rank}(a) \mid a \in \mathsf{Base}(\alpha)\}. & \alpha &\in T_\omega \mathcal{L}_T \end{aligned}$$

The next definition fixes the satisfaction conditions of formulas in $\mathcal{L}_T$ on T-models.

Definition 4.4. Using the fixed lax extension L for the functor T we can define the *semantics* of L for the language $\mathcal{L}_T(\mathsf{P})$ on T-models. For a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ define the satisfaction relation $\Vdash_{\mathbb{X}} : X \nrightarrow \mathcal{L}_T(\mathsf{P})$ by recursion as

$$\begin{aligned} x \Vdash_{\mathbb{X}} p &\quad \text{if} \quad x \in V_{\mathbb{X}}(p) & p &\in \mathsf{P} \\ x \Vdash_{\mathbb{X}} \neg a &\quad \text{if} \quad \text{not } x \Vdash_{\mathbb{X}} a, & a &\in \mathcal{L}_T \\ x \Vdash_{\mathbb{X}} \bigwedge A &\quad \text{if} \quad x \Vdash_{\mathbb{X}} a \text{ for all } a \in A, & A &\in \mathcal{P}_\omega \mathcal{L}_T \\ x \Vdash_{\mathbb{X}} \bigvee A &\quad \text{if} \quad x \Vdash_{\mathbb{X}} a \text{ for some } a \in A, & A &\in \mathcal{P}_\omega \mathcal{L}_T \\ x \Vdash_{\mathbb{X}} \nabla \alpha &\quad \text{if} \quad (\xi(x), \alpha) \in L{\Vdash_{\mathbb{X}}}. & \alpha &\in T_\omega \mathcal{L}_T \end{aligned}$$

Remark 4.5. Strictly speaking are the recursive clauses in Definition 4.4 not stated in a correct recursive way. For conjunction and disjunction we can only assume that ${\Vdash_{\mathbb{X}}}\restriction_{X \times \mathsf{Base}(A)} = {\Vdash_{\mathbb{X}}}\restriction_{X \times A}$ is already defined. That is not an issue because the conditions $x \Vdash_{\mathbb{X}} a$ for all (some) $a \in A$ and $x \Vdash_{\mathbb{X}} \restriction_{X \times A} \; a$ for all (some) $a \in A$ are equivalent. In the recursive clause for the nabla we can only presuppose that ${\Vdash_{\mathbb{X}}}\restriction_{X \times \mathsf{Base}(\alpha)}$ is already defined. So the actual recursive definition is that $x \Vdash_{\mathbb{X}} \nabla\alpha$ iff $(\xi(x), \alpha) \in L({\Vdash_{\mathbb{X}}}\restriction_{X \times \mathsf{Base}(\alpha)})$ and we need a little argument why this is equal to the clause given above. Because T is assumed to be standard we have that $\alpha \in T\mathsf{Base}(\alpha)$ and can use Proposition 3.12 to get that $(\xi(x), \alpha) \in L({\Vdash_{\mathbb{X}}} \restriction_{X \times \mathsf{Base}(\alpha)}) = (L{\Vdash_{\mathbb{X}}}) \restriction_{TX \times T\mathsf{Base}(\alpha)}$ is equivalent to $(\xi(x), \alpha) \in L{\Vdash_{\mathbb{X}}}$.

Example 4.6. With the Egli-Milner lifting $\overline{\mathcal{P}}$ of the $\mathcal{P}$ functor one can define the logic $\Lambda_{\mathcal{P}}$ for Kripke frames. A modal formula $\nabla\alpha$ is in this case of the shape $\nabla\{a_0, a_1, \ldots, a_{n-1}\}$ where $A = \{a_0, a_1, \ldots, a_{n-1}\}$ is a finite set of formulas. The satisfaction condition for nabla becomes $x \Vdash \nabla\{a_0, a_1, \ldots, a_{n-1}\}$ iff for every successor x' of x there is a formula a_i such that $x' \Vdash a_i$ and for every formula $a_i \in \{a_0, a_1, \ldots, a_{n-1}\}$ there is a successor x' of x such that $x' \Vdash a_i$. One can show that ∇A is equivalent to the formula $\Box \bigvee A \wedge \bigwedge \Diamond A$ in standard modal logic. Conversely $\Diamond a$ is equivalent to $\nabla\{a, \top\}$ and $\Box a$ is equivalent to $\nabla\emptyset \vee \nabla\{a\}$.

Remark 4.7. There is another way to define the semantics of $\mathcal{L}_T$ that uses only the natural transformation λ^L from Definition 3.19. See for example [15] how this works in detail. The trick is to consider $\mathcal{L}_T$ to be an initial object in suitable category of algebras. Then one can apply the $\breve{\mathcal{P}}$ functor to any coalgebra and compose with the natural transformation λ^L to get an algebra corresponding to the original coalgebra. The unique arrow from $\mathcal{L}_T$ into the algebra corresponding coalgebra gives an semantics that one can show to be equivalent to the one defined here. It is also noteworthy that in this approach the fact that λ^L is a natural transformation immediately entails that the semantics is adequate for behavioral equivalence.

Definition 4.8. Define the relation of logical consequence $\models : \mathcal{L}_T \nrightarrow \mathcal{L}_T$ by

$$a \models a' \quad \text{iff} \quad x \Vdash_{\mathbb{X}} a \text{ implies } x \Vdash_{\mathbb{X}} a' \text{ for all states } x \text{ in any } T\text{-model } \mathbb{X}.$$

The relation of logical equivalence $\equiv : \mathcal{L}_T \nrightarrow \mathcal{L}_T$ is defined as

$$a \equiv a' \quad \text{iff} \quad x \Vdash_{\mathbb{X}} a \text{ iff } x \Vdash_{\mathbb{X}} a' \text{ for all states } x \text{ in any } T\text{-model } \mathbb{X}.$$

Remark 4.9. There are translations between the nabla modalities of the lax extension L and the modalities associated to the Moss liftings that are defined as in Definition 3.24 for the lax extension L and any finitary presentation (Σ, E) of T_ω. Here we give a very short sketch of how this works. For a much more detailed treatment consider [11]. For any $\alpha \in T_\omega\mathcal{L}$ let $(r, A) \in \Sigma_n \times \mathcal{L}^n$ be such that $E_{\mathcal{L}}(r, A) = \alpha$. This always exists because E is surjective. Now let $\Box^r$ be the n-ary modality of the Moss lifting μ^r. That means that the formula $\Box^r A$ where $A = (A_0, A_1, \ldots, A_{n-1})$ has the following satisfaction condition at any state x_0 in some T-model $\mathbb{X} = (X, \xi, V_\xi)$

$$x_0 \Vdash_{\mathbb{X}} \Box^r A \quad \text{iff} \quad \xi(x_0) \in \mu^r_X([\![A_0]\!], [\![A_1]\!], \ldots, [\![A_{n-1}]\!]),$$

where $[\![-]\!]$ gives the extension of a formula that is

$$\begin{aligned} [\![-]\!] : \mathcal{L} &\to \mathcal{P}X, \\ a &\mapsto \{x \in X \mid x \Vdash_{\mathbb{X}} a\}. \end{aligned}$$

We can show that $x_0 \Vdash_{\mathbb{X}} \Box^r A$ is equivalent to $x_0 \Vdash_{\mathbb{X}} \nabla\alpha$. For this consider

$$\begin{array}{lll}
 & x_0 \Vdash_{\mathbb{X}} \Box^r A & \\
\text{iff} & \xi(x_0) \in \mu^r_X([\![A_0]\!], [\![A_1]\!], \ldots, [\![A_{n-1}]\!]) & \text{semantics of } \Box^r \\
\text{iff} & \xi(x_0) \in \lambda^L_X \circ E_{\breve{\mathcal{P}}X}(r, ([\![A_0]\!], [\![A_1]\!], \ldots, [\![A_{n-1}]\!])) & \text{Definition 3.24} \\
\text{iff} & (\xi(x_0), E_{\breve{\mathcal{P}}X}(r, ([\![A_0]\!], [\![A_1]\!], \ldots, [\![A_{n-1}]\!]))) \in L(\in_X) & \text{Definition 3.19} \\
\text{iff} & (\xi(x_0), (T_\omega[\![-]\!])(E_{\mathcal{L}}(r, A))) \in L(\in_X) & \text{E natural} \\
\text{iff} & (\xi(x_0), (T_\omega[\![-]\!])(\alpha)) \in L(\in_X) & E_{\mathcal{L}}(r, A) = \alpha \\
\text{iff} & (\xi(x_0), (T[\![-]\!])(\alpha)) \in L(\in_X) & \alpha \in T_\omega\mathcal{L} \subseteq T\mathcal{L} \\
\text{iff} & (\xi(x_0), \alpha) \in L(\in_X) \,;\, (T[\![-]\!])^\circ & \text{set theory} \\
\text{iff} & (\xi(x_0), \alpha) \in L(\in_X \,;\, [\![-]\!]^\circ) & \text{Proposition 3.10 (ii)} \\
\text{iff} & (\xi(x_0), \alpha \in L\Vdash_{\mathbb{X}} & \in_X \,;\, [\![-]\!]^\circ = \Vdash_{\mathbb{X}} \\
\text{iff} & x_0 \Vdash_{\mathbb{X}} \nabla\alpha. & \text{semantics of } \nabla
\end{array}$$

Next we show that the nabla logic of a lax extension is adequate with respect to L_C-bisimulation. That means that every two bisimilar states satisfy the same formulas.

Definition 4.10. Two states x_0 in a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ and y_0 in a T-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ are *modally equivalent* for formulas in $\mathcal{L}_T(C)$ iff

$$x_0 \Vdash_{\mathbb{X}} a \quad \text{iff} \quad y_0 \Vdash_{\mathbb{Y}} a, \quad \text{for all } a \in \mathcal{L}_T(C).$$

We write $x_0 \leftrightsquigarrow_C y_0$ if x_0 and y_0 are modally equivalent for formulas in $\mathcal{L}_T(C)$.

Proposition 4.11 (Adequacy)**.** *Given a state x_0 in a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ and a state y_0 in a T-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ if x_0 and y_0 are L_C-bisimilar then $x_0 \leftrightsquigarrow_C y_0$.*

Proof. Let R be an L_C-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$ with $(x_0, y_0) \in R$ and let $\Phi \subseteq \mathcal{L}_T$ be the set of formulas on which bisimilar points agree, that is

$$\Phi := \{a \in \mathcal{L}_T(C) \mid x \Vdash_{\mathbb{X}} a \text{ iff } y \Vdash_{\mathbb{Y}} a, \text{ for all } (x, y) \in R\}.$$

With this definition of Φ it is obvious that

$$R \,;\, \Vdash_{\mathbb{Y}}\restriction_{Y \times \Phi} \subseteq \Vdash_{\mathbb{X}}\restriction_{X \times \Phi},$$

and in the other direction

$$R^\circ \,;\, \Vdash_{\mathbb{X}}\restriction_{X \times \Phi} \subseteq \Vdash_{\mathbb{Y}}\restriction_{Y \times \Phi}. \tag{6}$$

We are now going to prove that $\Phi = \mathcal{L}_T(C)$. This entails that for all $(x, y) \in R$ that x satisfies the same formulas of the language $\mathcal{L}_T(C)$ in $\mathbb{X}$ as y does in $\mathbb{Y}$. So in particular x_0 and y_0 satisfy the same formulas because $(x_0, y_0) \in R$.

We show with induction on the complexity of a formula $a \in \mathcal{L}_T(C)$ that $a \in \Phi$. The base case $a = p \in C \subseteq \mathsf{P}$ follows directly from the semantics of propositional letters and the fact that R is an L_C-bisimulation between the T-models $\mathbb{X}$ and $\mathbb{Y}$. The Boolean cases are as usual. So let us focus on the case where $a = \nabla\alpha$ for some $\alpha \in T_\omega \mathcal{L}_T$. The induction hypothesis is that $\alpha \in T_\omega \Phi$. We have to show that $x \Vdash_{\mathbb{X}} \nabla\alpha$ iff $y \Vdash_{\mathbb{Y}} \nabla\alpha$ for all $(x, y) \in R$.

So assume that $x \Vdash_{\mathbb{X}} \nabla\alpha$. By the definition of the satisfaction relation that means $(\xi(x), \alpha) \in L\Vdash_{\mathbb{X}}$ and because $\alpha \in T_\omega\Phi \subseteq T\Phi$ in particular that $(\xi(x), \alpha) \in (L\Vdash_{\mathbb{X}})\restriction_{TX \times T\Phi} = L(\Vdash_{\mathbb{X}}\restriction_{X \times \Phi})$ where the last equality holds by Proposition 3.12. Because R is an L-bisimulation we have that $(\upsilon(y), \xi(x)) \in LR^\circ$, and so we get

$$\begin{aligned}
(\upsilon(y), \alpha) \in LR^\circ \,;\, L(\Vdash_{\mathbb{X}}\restriction_{X \times \Phi}) &\subseteq L(R^\circ \,;\, \Vdash_{\mathbb{X}}\restriction_{X \times \Phi}) && \text{(L2)} \\
&\subseteq L(\Vdash_{\mathbb{Y}}\restriction_{Y \times \Phi}) && \text{(6) and (L1)} \\
&= (L\Vdash_{\mathbb{Y}})\restriction_{TY \times T\Phi} && \text{Proposition 3.12} \\
&\subseteq L\Vdash_{\mathbb{Y}}.
\end{aligned}$$

This shows that $y \Vdash_{\mathbb{Y}} \nabla\alpha$. The other direction from $y \Vdash_{\mathbb{Y}} \nabla\alpha$ to $x \Vdash_{\mathbb{X}} \nabla\alpha$ is proved analogously. □

There is a partial converse of Proposition 4.11. If two states in T_ω-models satisfy the same formulas then they are bisimilar. This is shown by the next Proposition. The proof is similar to the one given in [2, Theorem 4.3] for the Barr extension. An intuitive explanation why this only works for finitary functors is that every formula of $\mathcal{L}_T$, since it is a finite object, can only capture a finite amount of information. A counterexample in the case of the $\mathcal{P}$-functor is given in [3, Section 2.2, Example 2.23].

Proposition 4.12 (Expressivity). *Given a state x_0 in a T_ω-model $\mathbb{X} = (X, \xi, V_\mathbb{X})$ and a state y_0 in a T_ω-model $\mathbb{Y} = (Y, \upsilon, V_\mathbb{Y})$ then x_0 and y_0 are L_C-bisimilar if $x_0 \leftrightsquigarrow_C y_0$.*

Proof. We first reduce the problem to the L_C-bisimilarity of two modally equivalent states in one single T_ω-model $\mathbb{Z} = (Z, \zeta, V_\mathbb{Z})$. For this let $\mathbb{Z} = \mathbb{X} + \mathbb{Y}$ be the coproduct of $\mathbb{X}$ and $\mathbb{Y}$ with injections $i_X : X \to Z$ and $i_Y : Y \to Z$. The injections are T-model morphisms and hence by Proposition 3.13 (i) their graphs are L_P-bisimulations. By Proposition 4.11 we know that L_P-bisimulations preserve truth and so we have that $i_X(x_0)$ and $i_Y(y_0)$ satisfy the same formulas of the language $\mathcal{L}_T(C)$ in $\mathbb{Z}$ because by assumption x_0 and y_0 satisfy the same formulas of the language $\mathcal{L}_T(C)$ as states of $\mathbb{X}$ and $\mathbb{Y}$. In the proof we construct an L_C-bisimulation $\leftrightsquigarrow : Z \not\to Z$ on $\mathbb{Z}$ that connects any two states that satisfy the same formulas of $\mathcal{L}_T(C)$. So we will have that $i_X(x_0) \leftrightsquigarrow i_Y(y_0)$ and hence by Proposition 3.13 (ii) the relation $i_X \,; \leftrightsquigarrow ; i_Y^\circ : X \not\to Y$ is an L_C-bisimulation between ξ and υ with $(x_0, y_0) \in i_X \,; \leftrightsquigarrow ; i_Y^\circ$.

In the following we consider the relation $\leftrightsquigarrow : Z \not\to Z$ to be modal equivalence for formulas in $\mathcal{L}_T(C)$ between states of Z. So we have $z \leftrightsquigarrow z'$ iff $z \leftrightsquigarrow_C z'$ for $z, z' \in Z$. By assumption we have $x_0 \leftrightsquigarrow y_0$ so it only remains to show that $\leftrightsquigarrow$ is an L_C-bisimulation. For this we have to check for arbitrarily chosen $(z_0, z_1) \in \leftrightsquigarrow$ that $(\zeta(z_0), \zeta(z_1)) \in L\!\leftrightsquigarrow$ and that

$$z_0 \in V_\mathbb{Z}(p) \quad \text{iff} \quad z_1 \in V_\mathbb{Z}(p), \qquad \text{for all } p \in C.$$

The latter holds because for all $p \in C \subseteq \mathcal{L}_T(C)$ we have by assumption that $z_0 \Vdash_\mathbb{Z} p$ iff $z_1 \Vdash_\mathbb{Z} p$. So it remains to be proved that $(\zeta(z_0), \zeta(z_1)) \in L\!\leftrightsquigarrow$.

Consider the set S of successors of z_0 and z_1, that is

$$S = \mathsf{Base}(\zeta(z_0)) \cup \mathsf{Base}(\zeta(z_1)).$$

This set exists and is finite because $\zeta(z_0), \zeta(z_1) \in T_\omega Z$. Because T is standard we have that $\zeta(z_0) \in T\mathsf{Base}(\zeta(z_0)) \subseteq TS$ and similarly $\zeta(z_1) \in TS$.

Note that for all states $z, z' \in Z$ with $z \not\leftrightsquigarrow z'$ there is a formula $d_z^{z'} \in \mathcal{L}_T$ such that $z \Vdash_\mathbb{Z} d_z^{z'}$ but $z' \not\Vdash_\mathbb{Z} d_z^{z'}$. We can use this to define a function $f : S \to \mathcal{L}_T$ that maps every state $z \in S$ to a formula $f(z)$ that is true at z and false at all other states of S. This is done as follows:

$$\begin{aligned} f : S &\to \mathcal{L}_T, \\ z &\mapsto \bigwedge\{d_z^{z'} \mid z' \in S, z \not\leftrightsquigarrow z'\}. \end{aligned}$$

The conjunction $f(z)$ is finite for all $z \in S$ because S is finite. For this definition of f we claim that

$$\Vdash_\mathbb{Z}\restriction_{S \times \mathcal{L}_T} ; f^\circ = \leftrightsquigarrow\restriction_{S \times S}. \tag{7}$$

For the $\subseteq$-inclusion we argue by contraposition. Assume $z' \not\leftrightsquigarrow z$ for $z, z' \in S$. Then $z' \not\Vdash_\mathbb{Z} f(z)$ because $z' \not\Vdash_\mathbb{Z} d_z^{z'}$ and $f(z) = \bigwedge\{d_z^{z'} \mid z' \in S, z \not\leftrightsquigarrow z'\}$. For the $\supseteq$-inclusion take $z, z' \in S$ with $z' \leftrightsquigarrow z$. Because $z \Vdash_\mathbb{Z} d_z^{z''}$ for all $z'' \in S$ with $z \not\leftrightsquigarrow z''$ we have that $z \Vdash_\mathbb{Z} f(z)$ since $f(z)$ is the conjunction $\bigwedge\{d_z^{z''} \mid z'' \in S, z \not\leftrightsquigarrow z''\}$. It follows that $z' \Vdash_\mathbb{Z} f(z)$ because $z' \leftrightsquigarrow z$.

Because $\zeta(z_1) \in TS$ we can calculate

$$\begin{aligned}
(\zeta(z_1), \zeta(z_1)) \in \Delta_{TS} &\subseteq T\Delta_S && \text{Proposition 3.10 (i)}\\
&\subseteq L(\leftrightsquigarrow\restriction_{S\times S}) && \Delta_S \subseteq \leftrightsquigarrow\restriction_{S\times S} \text{ and (L1)}\\
&= L(\Vdash_{\mathbb{X}}\restriction_{S\times\mathcal{L}_T} ; f^\circ) && (7)\\
&= L(\Vdash_{\mathbb{X}}\restriction_{S\times\mathcal{L}_T}) ; (Tf)^\circ. && \text{Proposition 3.10 (ii)}
\end{aligned}$$

This yields an $\alpha \in T\mathcal{L}_T$ with $(\zeta(z_1), \alpha) \in L(\Vdash_{\mathbb{Z}}\restriction_{S\times\mathcal{L}_T})$ and $(\alpha, \zeta(z_1)) \in (Tf)^\circ$. From the latter it follows that $\alpha = (Tf)(\zeta(z_1)) = (T_\omega f)(\zeta(z_1)$ because $\zeta(z_1) \in T_\omega Z$. So we get that $\alpha \in T_\omega\mathcal{L}_T$ hence it follows that $\nabla\alpha \in \mathcal{L}_T$ and $z_1 \Vdash_{\mathbb{Z}} \nabla\alpha$ because $(\zeta(z_1), \alpha) \in L(\Vdash_{\mathbb{Z}}\restriction_{S\times\mathcal{L}_T}) = (L\Vdash_{\mathbb{Z}})\restriction_{TS\times T\mathcal{L}_T} \subseteq L\Vdash_{\mathbb{Z}}$ by Proposition 3.12.

It follows that $z_0 \Vdash_{\mathbb{Z}} \nabla\alpha$ because by assumption $z_0 \leftrightsquigarrow z_1$. Therefore $(\zeta(z_0), \alpha) \in L\Vdash_{\mathbb{Z}}$. Because $\zeta(z_0) \in TS$ this gives $(\zeta(z_0), \alpha) \in (L\Vdash_{\mathbb{Z}})\restriction_{TS\times T\mathcal{L}_T}$ and by Proposition 3.12 $(\zeta(z_0), \alpha) \in L(\Vdash_{\mathbb{Z}}\restriction_{S\times\mathcal{L}_T})$. This yields what we want to show because

$$\begin{aligned}
(\zeta(z_0), \zeta(z_1)) \in L(\Vdash_{\mathbb{Z}}\restriction_{S\times\mathcal{L}_T}) ; (Tf)^\circ &= L(\Vdash_{\mathbb{Z}}\restriction_{D\times\mathcal{L}_T} ; f^\circ) && \text{Proposition 3.10 (ii)}\\
&= L(\leftrightsquigarrow\restriction_{S\times S}) && (7)\\
&= (L\leftrightsquigarrow)\restriction_{TS\times TS} && \text{Proposition 3.12}\\
&\subseteq L\leftrightsquigarrow.
\end{aligned}$$

□

4.2 Disjunctive Nabla Normal Form

The goal of this subsection is to prove Theorem 4.18 which states that every formula in $\mathcal{L}_T$ is equivalent to a formula in which negation and conjunction occur only over propositional letters. We say that such formulas are in disjunctive nabla normal form.

Definition 4.13. A formula $a \in \mathcal{L}_T$ is a *literal* if it is either a propositional letter or the negation of a propositional letter, that is $a = p$ or $a = \neg p$ for some $p \in C$. The set of formulas in *disjunctive nabla normal form* $\mathcal{L}_T^{\mathsf{d}}(C) \subseteq \mathcal{L}_T(C)$ is generated by the grammar:

$$a ::= \bigwedge\Pi \mid \bigvee A \mid \bigwedge\Pi \wedge \nabla\alpha,$$

where Π is a finite set of literals, $A \in \mathcal{P}_\omega\mathcal{L}_T^{\mathsf{d}}(C)$ and $\alpha \in T_\omega\mathcal{L}_T^{\mathsf{d}}(C)$.

The two main ingredients in the proof of Theorem 4.18 are Proposition 4.14 which says how we can distribute negations over nablas and Proposition 4.17 which shows how we can put conjunctions into nablas.

Proposition 4.14 shows that we can define the Boolean dual of nabla as a disjunction of nablas. This shows how we can push negations inside nablas and thereby decrease the modal rank at which negations occur in a formula. This result first appeared in [7] for nablas defined for a Barr extension of a weak pullback preserving functor.

Proposition 4.14. *If T restricts to finite sets then we have for all $\alpha \in T_\omega\mathcal{L}_T(C)$*

$$\neg\nabla(T\neg)(\alpha) \equiv \bigvee\{\nabla T\wedge(\Omega) \mid \Omega \in T_\omega\mathcal{P}_\omega\mathsf{Base}(\alpha), (\alpha, \Omega) \notin L(\not\in_{\mathsf{Base}(\alpha)})\}.$$

Proof. Note that the disjunction on the right hand side is finite and therefore it is a well defined formula in the language $\mathcal{L}_T(C)$. To see this observe that $\mathsf{Base}(\alpha)$ and hence $\mathcal{P}_\omega\mathsf{Base}(\alpha)$ is finite. Because T, and hence T_ω, is assumed to restrict to finite sets this entails that there are only finitely many choices for an $\Omega \in T_\omega\mathcal{P}_\omega\mathsf{Base}(\alpha)$.

For the direction from left to right assume we have a state x_0 in a T-model $\mathbb{X} = (X, \xi, V_\mathbb{X})$ such that $x_0 \Vdash_\mathbb{X} \neg\nabla(T\neg)(\alpha)$. That means that $(\xi(x_0), (T\neg)(\alpha)) \notin L{\Vdash_\mathbb{X}}$ and with Proposition 3.10 (ii) that $(\xi(x_0), \alpha) \notin L{\Vdash_\mathbb{X}} ; (T\neg)^\circ = L(\Vdash_\mathbb{X} ; \neg^\circ)$. We want to find an $\Omega \in T_\omega\mathcal{P}_\omega\mathsf{Base}(\alpha)$ such that $x_0 \Vdash_\mathbb{X} \nabla T\wedge(\Omega)$ and $(\alpha, \Omega) \notin L(\notin_{\mathsf{Base}(\alpha)})$. For this consider the function

$$f : X \to \mathcal{P}_\omega\mathsf{Base}(\alpha),$$
$$x \mapsto \{a \in \mathsf{Base}(\alpha) \mid x \Vdash_\mathbb{X} a\}.$$

Set $\Omega = Tf(\xi(x_0)) \in T\mathcal{P}_\omega\mathsf{Base}(\alpha) = T_\omega\mathcal{P}_\omega\mathsf{Base}(\alpha)$. We have to check that $(\alpha, Tf(\xi(x_0))) \notin L(\notin_{\mathsf{Base}(\alpha)})$ and $x_0 \Vdash_\mathbb{X} \nabla T\wedge(Tf(\xi(x_0)))$.

To verify that $(\alpha, Tf(\xi(x_0))) \notin L(\notin_{\mathsf{Base}(\alpha)})$ we need that

$$f ; \not\ni_{\mathsf{Base}(\alpha)} \subseteq \Vdash_\mathbb{X} ; (\neg\restriction_{\mathsf{Base}(\alpha)})^\circ. \tag{8}$$

This inequality means that if a formula $a \in \mathsf{Base}(\alpha)$ is not in $f(x)$ for a state x then the negation of a is true at x. This holds because if $a \notin f(x)$ for an $a \in \mathsf{Base}(\alpha)$ then we must have $x \not\Vdash_\mathbb{X} a$ by the definition of f and hence $x \Vdash_\mathbb{X} \neg a$.

Now assume for a contradiction that $(\alpha, Tf(\xi(x_0))) \in L(\notin_{\mathsf{Base}(\alpha)})$. This entails that $(\alpha, \xi(x_0)) \in L(\notin_{\mathsf{Base}(\alpha)}) ; (Tf)^\circ$. Then we can compute

$$\begin{aligned}
(\xi(x_0), \alpha) \in Tf ; L(\not\ni_{\mathsf{Base}(\alpha)}) &= L(f ; \not\ni_{\mathsf{Base}(\alpha)}) && \text{Proposition 3.10 (ii)}\\
&\subseteq L(\Vdash_\mathbb{X} ; (\neg\restriction_{\mathsf{Base}(\alpha)})^\circ) && \text{(8) and (L1)}\\
&= L((\Vdash_\mathbb{X} ; \neg^\circ)\restriction_{X\times\mathsf{Base}(\alpha)}) && \text{set theory}\\
&= (L(\Vdash_\mathbb{X} ; \neg^\circ))\restriction_{TX\times T\mathsf{Base}(\alpha)} && \text{Proposition 3.12}\\
&\subseteq L(\Vdash_\mathbb{X} ; \neg^\circ).
\end{aligned}$$

But this is a contradiction to $(\xi(x_0), \alpha) \notin L(\Vdash_\mathbb{X} ; \neg^\circ)$.

The other thing we have to check is that $x_0 \Vdash_\mathbb{X} \nabla T\wedge(Tf(\xi(x_0)))$. For this first observe that

$$\Delta_X \subseteq \Vdash_\mathbb{X} ; (\wedge\restriction_{\mathsf{Base}(\alpha)})^\circ ; f^\circ. \tag{9}$$

This holds because the conjunction of formulas that are true at one state is again true at this state. Now consider

$$\begin{aligned}
(\xi(x_0), \xi(x_0)) \in \Delta TX &\subseteq L\Delta_X && \text{Proposition 3.10 (i)}\\
&\subseteq L(\Vdash_\mathbb{X} ; (\wedge\restriction_{\mathsf{Base}(\alpha)})^\circ ; f^\circ) && \text{(9) and (L1)}\\
&= L{\Vdash_\mathbb{X}} ; \left(T\wedge\restriction_{\mathsf{Base}(\alpha)}\right)^\circ ; (Tf)^\circ. && \text{Proposition 3.10 (ii)}
\end{aligned}$$

So it follows that $(\xi(x_0), T(\wedge\restriction_{\mathsf{Base}(\alpha)})(Tf(\xi(x_0)))) \in L{\Vdash_\mathbb{X}}$. This means that $(\xi(x_0), T\wedge(\Omega)) \in L{\Vdash_\mathbb{X}}$ because $T(\wedge\restriction_{\mathsf{Base}(\alpha)})(Tf(\xi(x_0))) = T(\wedge\restriction_{\mathsf{Base}(\alpha)})(\Omega) = T\wedge(\Omega)$. Hence $x_0 \Vdash_\mathbb{X} \nabla T\wedge(\Omega)$.

For the other direction assume that there is an $\Omega \in T_\omega\mathcal{P}_\omega\mathsf{Base}(\alpha)$ with $(\alpha, \Omega) \notin L(\notin_{\mathsf{Base}(\alpha)})$ such that $x_0 \Vdash_\mathbb{X} \nabla T\wedge(\Omega)$ for an x_0 in a T-model $\mathbb{X} =$

$(X, \xi, V_{\mathbb{X}})$. We have to show that $x_0 \Vdash_{\mathbb{X}} \neg\nabla(T\neg)(\alpha)$. Assume for a contradiction that $x_0 \Vdash_{\mathbb{X}} \nabla(T\neg)(\alpha)$ which means that $(\xi(x_0), (T\neg)(\alpha)) \in L{\Vdash_{\mathbb{X}}}$ and equivalently $(\xi(x_0), \alpha) \in L{\Vdash_{\mathbb{X}}} ; (T\neg)^{\circ}$. By Proposition 3.10 (ii) we can reformulate this as $(\alpha, \xi(x_0)) \in T\neg ; L{\Vdash_{\mathbb{X}}}^{\circ} = L(\neg ; {\Vdash_{\mathbb{X}}}^{\circ})$.

From the assumption that $x_0 \Vdash_{\mathbb{X}} \nabla T\wedge(\Omega)$ we get that $(\xi(x_0), T\wedge(\Omega)) \in L{\Vdash_{\mathbb{X}}}$ and so by Proposition 3.10 (ii) that $(\xi(x_0), \Omega) \in L({\Vdash_{\mathbb{X}}} ; \wedge^{\circ})$. Together with $(\alpha, \xi(x_0)) \in L(\neg ; {\Vdash_{\mathbb{X}}}^{\circ})$ it follows by (L1) that $(\alpha, \Omega) \in L(\neg ; {\Vdash_{\mathbb{X}}}^{\circ}) ; L({\Vdash_{\mathbb{X}}} ; \wedge^{\circ}) \subseteq L(\neg ; {\Vdash_{\mathbb{X}}}^{\circ} ; {\Vdash_{\mathbb{X}}} ; \wedge^{\circ})$. Define $\Gamma = (\neg ; {\Vdash_{\mathbb{X}}}^{\circ} ; {\Vdash_{\mathbb{X}}} ; \wedge^{\circ})$. So $(\alpha, \Omega) \in L\Gamma$.

Now observe that

$$\Gamma{\restriction_{\mathsf{Base}(\alpha) \times \mathcal{P}_{\omega}\mathsf{Base}(\alpha)}} = (\neg ; {\Vdash_{\mathbb{X}}}^{\circ} ; {\Vdash_{\mathbb{X}}} ; \wedge^{\circ}){\restriction_{\mathsf{Base}(\alpha) \times \mathcal{P}_{\omega}\mathsf{Base}(\alpha)}} \subseteq \notin_{\mathsf{Base}(\alpha)}. \tag{10}$$

This holds because a formula whose negation is true at a state can not be a conjunct of a conjunction that is true at that state. Moreover we have that $\alpha \in T\mathsf{Base}(\alpha)$ and $\Omega \in T\mathcal{P}_{\omega}\mathsf{Base}(\alpha)$. Hence by Proposition 3.12 $(\alpha, \Omega) \in (L\Delta){\restriction_{T\mathsf{Base}(\alpha) \times T\mathcal{P}_{\omega}\mathsf{Base}(\alpha)}} = L(\Gamma{\restriction_{\mathsf{Base}(\alpha) \times \mathcal{P}_{\omega}\mathsf{Base}(\alpha)}})$. But by (10) and (L1) we have that $L(\Gamma{\restriction_{\mathsf{Base}(\alpha) \times \mathcal{P}_{\omega}\mathsf{Base}(\alpha)}}) \subseteq L(\notin_{\mathsf{Base}(\alpha)})$. So $(\alpha, \Omega) \in L(\notin_{\mathsf{Base}(\alpha)})$ which is a contradiction to the assumption that $(\alpha, \Omega) \notin L(\notin_{\mathsf{Base}(\alpha)})$. □

Next we want to show how we can get rid of conjunctions over nablas and replace them with disjunctions of nablas. For this we first need to prove the following easy Proposition that states that the nabla logic is in a sense monotone.

Proposition 4.15. *It holds for all $\alpha, \alpha' \in T_{\omega}\mathcal{L}_T$ that*

$$(\alpha, \alpha') \in L{\models} \text{ implies } \nabla\alpha \models \nabla\alpha'.$$

Proof. Assume that $(\alpha, \alpha') \in L{\models}$ and take any state x_0 in a T-coalgebra $\xi : X \to TX$ such that $x_0 \Vdash_{\mathbb{X}} \nabla\alpha$. We have to show that $x_0 \Vdash_{\mathbb{X}} \nabla\alpha'$. By the semantics of the nabla $x_0 \Vdash_{\mathbb{X}} \nabla$ means that $(\xi(x_0), \alpha) \in L{\Vdash_{\mathbb{X}}}$.

By (L2) it follows that $(\xi(x_0), \alpha') \in L{\Vdash_{\mathbb{X}}} ; L{\models} \subseteq L({\Vdash_{\mathbb{X}}} ; {\models})$. Now it is easy to see that ${\Vdash_{\mathbb{X}}} ; {\models} \subseteq {\Vdash_{\mathbb{X}}}$ by the definition of $\models$. So we get with (L1) that $(\xi(x_0), \alpha') \in L{\Vdash_{\mathbb{X}}}$ which means that $x_0 \Vdash_{\mathbb{X}} \nabla\alpha'$. □

Definition 4.16. For a set of formulas $A \subseteq \mathcal{L}_T$ let $\mathsf{Conj}(A) \in \mathcal{P}\mathcal{L}_T$ be the set of all finite conjunctions of formulas in A that is

$$\begin{aligned} \mathsf{Conj} : \mathcal{P}\mathcal{L}_T &\to \mathcal{P}\mathcal{L}_T, \\ A &\mapsto \{\textstyle\bigwedge F \in \mathcal{L}_T \mid F \in \mathcal{P}_{\omega}A\}. \end{aligned}$$

Proposition 4.17. *Assume that T preserves finite sets, take any $A \in \mathcal{P}_{\omega}T\mathcal{L}_T$ and set $S = \bigcup(\mathcal{P}_{\omega}\mathsf{Base})(A)$. Then it holds that*

$$\bigwedge_{\alpha \in A} \nabla\alpha \equiv \bigvee\{\nabla\beta \mid \beta \in T\mathsf{Conj}(S) \text{ with } (\beta, \alpha) \in L{\models} \text{ for all } \alpha \in A\}.$$

Proof. First note that the right hand side is a formula, because the disjunction is finite. This holds because T is assumed to restrict to finite sets and there are only finitely many conjunctions over the finite set $S = \bigcup(\mathcal{P}_{\omega}\mathsf{Base})(A)$.

For the direction from left to right assume we are given a state x_0 in a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ with $x_0 \Vdash_{\mathbb{X}} \nabla\alpha$ for all $\alpha \in A$. In order to witness the right hand side we have to find a $\beta \in T\mathsf{Conj}(S)$ with $(\beta, \alpha) \in L{\models}$ for all $\alpha \in A$ such that $x_0 \Vdash_{\mathbb{X}} \nabla\beta$. For this consider the map

$$\begin{aligned} f : X &\to \mathsf{Conj}(S), \\ x &\mapsto \bigwedge\{a \in S \mid x \Vdash_{\mathbb{X}} a\}. \end{aligned}$$

The map f is well defined because S is finite. Moreover we claim that

$$f^{\circ} ; {\Vdash_{\mathbb{X}}}\restriction_{X \times S} \subseteq {\models}\restriction_{\mathsf{Conj}(S) \times S}. \tag{11}$$

The inequality (11) is equivalent to the claim that $f(x) \models a$ for all states $x \in X$ and $a \in S$ such that $x \Vdash_{\mathbb{X}} a$. This holds because $f(x) = \bigwedge\{a \in S \mid x \Vdash_{\mathbb{X}} a\}$, Whence $a \in S$ is one of the conjuncts of the conjunction $f(x)$ if we assume that $x \Vdash_{\mathbb{X}} a$. Clearly we then have $f(x) \models a$ because a conjunction entails all its conjuncts.

Set $\beta = Tf(\xi(x_0)) \in T\mathsf{Conj}(S)$. Since $\mathsf{Conj}(S) \subseteq \mathcal{L}_T$ is a finite set and T is standard it follows that $\beta \in T\mathsf{Conj}(S) \subseteq T_\omega\mathcal{L}_T$.

To verify $(\beta, \alpha) \in L{\models}$ for all $\alpha \in A$ we use that $(\beta, \xi(x_0)) \in (Tf)^{\circ}$ by definition of β and $(\xi(x_0), \alpha) \in L{\Vdash_{\mathbb{X}}}$ because $x_0 \Vdash_{\mathbb{X}} \nabla\alpha$. So we can compute

$$\begin{aligned}
(\beta, \alpha) \in (Tf)^{\circ} ; (L{\Vdash_{\mathbb{X}}})\restriction_{TX \times TS} &= (Tf)^{\circ} ; L({\Vdash_{\mathbb{X}}}\restriction_{X \times S}) && \text{Proposition 3.12} \\
&\subseteq Lf^{\circ} ; L({\Vdash_{\mathbb{X}}}\restriction_{X \times S}) && \text{(L3)} \\
&\subseteq L(f^{\circ} ; {\Vdash_{\mathbb{X}}}\restriction_{X \times S}) && \text{(L2)} \\
&\subseteq L({\models}\restriction_{\mathsf{Conj}(S) \times S}) && \text{(11) and (L1)} \\
&= (L{\models})\restriction_{T\mathsf{Conj}(S) \times TS} && \text{Proposition 3.12} \\
&\subseteq L{\models}.
\end{aligned}$$

It remains to show that $x_0 \Vdash_{\mathbb{X}} \nabla\beta$. For this observe that from the definition of f we get that $f^{\circ} \subseteq {\Vdash_{\mathbb{X}}}\restriction_{X \times \mathsf{Conj}(S)}$. So we calculate

$$\begin{aligned}
(\xi(x_0), \beta) \in (Tf)^{\circ} &\subseteq Lf^{\circ} && \text{(L3)} \\
&\subseteq L({\Vdash_{\mathbb{X}}}\restriction_{X \times \mathsf{Conj}(S)}) && f^{\circ} \subseteq {\Vdash_{\mathbb{X}}}\restriction_{X \times \mathsf{Conj}(S)} \text{ and (L1)} \\
&= (L{\Vdash_{\mathbb{X}}})\restriction_{TX \times T\mathsf{Conj}(S)} && \text{Proposition 3.12} \\
&\subseteq L{\Vdash_{\mathbb{X}}}.
\end{aligned}$$

By the semantics of ∇ this means that $x_0 \Vdash_{\mathbb{X}} \nabla\beta$.

For the direction from right to left assume that we have a state x_0 in a T-coalgebra $\xi : X \to TX$ such that $x_0 \Vdash_{\mathbb{X}} \nabla\beta$ for a $\beta \in T\mathsf{Conj}(S)$ such that $(\beta, \alpha) \in L{\models}$ for all $\alpha \in A$. By Proposition 4.15 it follows that $x_0 \Vdash_{\mathbb{X}} \nabla\alpha$ for all $\alpha \in A$ which gives the left hand side. □

Theorem 4.18. *For every formula $a \in \mathcal{L}_T(C)$ there is a formula $a^{\mathsf{d}} \in \mathcal{L}_T^{\mathsf{d}}(C)$ that is in disjunctive nabla normal form such that*

$$a \equiv a^{\mathsf{d}}.$$

Proof. This is proved by an induction on the modal rank of the formula a. For formulas with $\mathsf{rank}(a) = 0$, that do not contain any nablas a^{d} is just the usual disjunctive normal form of a as a formula of propositional logic.

In the induction step we first rewrite a in a disjunctive normal form such that it is a disjunction of conjunctions of literals and possibly negated nablas. That this is possible is a basic fact about propositional logic. With Proposition 4.14 we can then get rid of the negated nablas by replacing them with conjunctions of nablas. Now the formula is a disjunction of conjunctions of literals and nablas. Then we use Proposition 4.17 to replace conjunctions of nablas with disjunctions of nablas. So the formula is now a disjunction of conjunctions that contain any number of literals and at most one disjunction of nablas. Using the distributivity of $\wedge$ over $\vee$ we get the whole formula into a form such that it is a disjunction of conjunctions of literals and at most one nabla. Since none of the transformations of Propositions 4.14 and 4.17 increases the modal rank of the formula we have that all the direct subformulas of the nablas occurring in the formula are of smaller modal rank. So we can use the induction hypothesis to replace all the subformulas that are directly under nablas with equivalent formulas in disjunctive nabla normal form. After this the whole formula is in disjunctive nabla normal form and we are done. □

4.3 Bisimulation Quantifiers and Uniform Interpolation

The goal of this part is to prove uniform interpolation for $\mathcal{L}_T$. In order to do this we introduce the notion of a bisimulation quantifier and show that bisimulation quantifiers are definable in the language $\mathcal{L}_T$. Our proofs follows the proof in [17] which shows a similar result in the case of $\mathcal{M}$-models. We generalize their result to arbitrary standard functors that restrict to finite sets and have a lax extension that satisfies the following extra condition.

Definition 4.19. A relation lifting L of T is *quasi-functorial* if it satisfies the following condition for all relations $R : X \nrightarrow Z$ and $S : Z \nrightarrow Y$:

$$LR\,;LS = L(R\,;S) \cap (\mathsf{preimg}(LR) \times \mathsf{img}(LS)). \tag{12}$$

Recall from the definition of $\mathsf{preimg}(LR) \subseteq TX$ that $\xi \in \mathsf{preimg}(LR)$ iff there is a $\zeta_R \in TZ$ such that $(\xi, \zeta_R) \in LR$. Similarly $\upsilon \in \mathsf{img}(LS) \subseteq TY$ iff there is a $\zeta_S \in TZ$ such that $(\zeta_S, \upsilon) \in L$. So the $\subseteq$-inclusion of (12) holds for any lax extension because of (L2). The $\supseteq$-inclusion is the actual substantial requirement. It is equivalent to condition that for all $(\xi, \upsilon) \in L(R\,;S)$ if there is a $\zeta_R \in TZ$ such that $(\xi, \zeta_R) \in LR$ and there is a $\zeta_S \in TZ$ such that $(\zeta_S, \upsilon) \in LS$ then there is a $\zeta \in TZ$ such that $(\xi, \zeta) \in LR$ and $(\zeta, \upsilon) \in LS$.

Example 4.20. (i) Recall from Example 3.11 (i) that the Barr extension $\overline{T}$ for a functor T that preserves weak pullbacks is functorial. Hence it satisfies $\overline{T}R\,;\overline{T}S = \overline{T}(R\,;S)$ for all relation R and S. Clearly this implies that $\overline{T}$ is also quasi-functorial.

(ii) We can prove that the lax extension $\widetilde{\mathcal{M}}$ of $\mathcal{M}$ as defined in Example 3.2 (vi) is quasi-functorial. This implies that the lax extension $\widetilde{\mathcal{M}_s}$ of $\mathcal{M}_s$ is also quasi-functorial. To see this recall that $\widetilde{\mathcal{M}_s}$ was defined on a relation $R : X \nrightarrow Y$ as $\widetilde{\mathcal{M}_s}R = i_X\,;\widetilde{\mathcal{M}}R\,;i_Y^\circ$ for any relation $R : X \nrightarrow Y$ where $i : \mathcal{M}_s \Rightarrow \mathcal{M}$ is a natural isomorphism.

For the proof that $\mathcal{M}$ is quasi-functorial take any two relations $R : X \nrightarrow Z$ and $S : Z \nrightarrow Y$. We have to show that for all $(\zeta, \upsilon) \in \widetilde{\mathcal{M}}(R\,;S)$ if there are

$\zeta_R, \zeta_S \in \mathcal{M}Z$ with $(\xi, \zeta_R) \in \widetilde{\mathcal{M}}R$ and $(\zeta_S, \upsilon) \in \widetilde{\mathcal{M}}S$ then there is a $\zeta \in \mathcal{M}Z$ with $(\xi, \zeta) \in \widetilde{\mathcal{M}}R$ and $(\zeta, \upsilon) \in \widetilde{\mathcal{M}}S$.

From the assumptions $(\xi, \zeta_R) \in \widetilde{\mathcal{M}}R \subseteq \overrightarrow{\mathcal{P}}\overleftarrow{\mathcal{P}}R$ and $(\xi, \upsilon) \in \widetilde{\mathcal{M}}(R\,;S) \subseteq \overrightarrow{\mathcal{P}}\overleftarrow{\mathcal{P}}(R\,;S)$ we get that for every $U \in \xi$ there are elements $W_U \in \zeta_R$ and $V_U \in \upsilon$ such that $(U, W_U) \in \overleftarrow{\mathcal{P}}R$ and $(U, V_U) \in \overleftarrow{\mathcal{P}}(R\,;S)$. Similarly by $(\zeta_S, \upsilon) \in \widetilde{\mathcal{M}}S \subseteq \overleftarrow{\mathcal{P}}\overrightarrow{\mathcal{P}}R$ and $(\xi, \upsilon) \in \widetilde{\mathcal{M}}(R\,;S) \subseteq \overleftarrow{\mathcal{P}}\overrightarrow{\mathcal{P}}(R\,;S)$ we get for every $V \in \upsilon$ elements $W_V \in \zeta_S$ and $U_V \in \xi$ such that $(W_V, V) \in \overrightarrow{\mathcal{P}}R$ and $(U_V, V) \in \overrightarrow{\mathcal{P}}(R\,;S)$.

From $(U, V_U) \in \overleftarrow{\mathcal{P}}(R\,;S)$ it follows that for every $v \in V_U$ there is a $u_v \in U$ such that $(u_v, v) \in R\,;S$. Hence there is a $w_v \in Y$ such that $(u_v, w_v) \in R$ and $(w_v, v) \in S$. With this define for every $U \in \xi$

$$W'_U = W_U \cup \{w_v \in V \mid v \in V_U\}.$$

For this definition we can show that $(U, W'_U) \in \overleftarrow{\mathcal{P}}R$. For, take any $w \in W'_U$. Then it is either in W_U or from $\{w_v \in V \mid v \in V_U\}$. In the former case the claim follows from $(U, W_U) \in \overleftarrow{\mathcal{P}}R$. In the latter case we have by the definition of w_v that there is a $u_v \in U$ such that $(u_v, w_v) \in R$. Moreover we have that $(W'_U, V_U) \in \overleftarrow{\mathcal{P}}S$ because for any $v \in V_U$ there is the element $w_v \in W'_U$ with $(w_v, v) \in S$.

A symmetric argument shows that for every $V \in \upsilon$ we can find a $W'_V \in \zeta_S$ with the properties $(W'_V, V) \in \overrightarrow{\mathcal{P}}S$ and $(U_V, W'_V) \in \overrightarrow{\mathcal{P}}R$.

Set

$$\zeta = \{W \subseteq Z \mid W \supseteq W'_U \text{ for a } U \in \xi \text{ or } W \supseteq W'_V \text{ for a } V \in \upsilon\}.$$

From the definition its clear that ζ is upwards-closed and so we have that $\zeta \in \mathcal{M}Y$. It remains to show that $(\xi, \zeta) \in \widetilde{\mathcal{M}}R$ and $(\zeta, \upsilon) \in \widetilde{\mathcal{M}}S$. We only do the former since the latter is analogous.

For $(\xi, \zeta) \in \overrightarrow{\mathcal{P}}\overleftarrow{\mathcal{P}}R$ note that for any $U \in \xi$ there is the set $W'_U \in \zeta$ with $(U, W'_U) \in \overleftarrow{\mathcal{P}}R$.

For $(\xi, \zeta) \in \overleftarrow{\mathcal{P}}\overrightarrow{\mathcal{P}}R$ pick any $W \in \zeta$. By the definition of ζ it follows that either $W \supseteq W'_U$ for a $U \in \xi$ or $W \supseteq W'_V$ for a $V \in \upsilon$. In the former case consider $W_U \in \zeta_R$. From the assumption $(\xi, \zeta_R) \in \widetilde{\mathcal{M}}R \subseteq \overleftarrow{\mathcal{P}}\overrightarrow{\mathcal{P}}R$ we obtain a $U' \in \xi$ with $(U', W_U) \in \overrightarrow{\mathcal{P}}R$. Because $W_U \subseteq W'_U \subseteq W$ this entails that $(U', W) \in \overrightarrow{\mathcal{P}}R$ as required. In the other case where $W \supseteq W'_V$ for a $V \in \upsilon$ we have from the above that $(U_V, W'_V) \in \overrightarrow{\mathcal{P}}R$ for the set $U_V \in \xi$. Because $W'_V \subseteq W$ it follows that $(U_V, W) \in \overrightarrow{\mathcal{P}}R$.

(iii) One can show that quasi-functoriality of relation liftings is preserved under taking products or coproducts of quasi-functorial relation liftings.

(iv) The lax extension L_3^2 for the functor F_2^3 is not quasi-functorial. This is a consequence of Example 4.26 and Theorem 4.25 but it can also be shown by a direct counterexample.

We define the bisimulation quantifier as a syntactical transformation on formulas in $\mathcal{L}_T$.

Definition 4.21. By recursion on the complexity of formulas in disjunctive nabla normal form we define for all $p \in \mathsf{P}$ the map

$$\begin{aligned} e^p : \mathcal{L}_T^{\mathsf{d}} &\to \mathcal{L}_T, \\ \bigwedge \Pi &\mapsto \begin{cases} \bot, & \text{if } \{p, \neg p\} \subseteq \Pi, \\ \bigwedge (\Pi \setminus \{p, \neg p\}), & \text{otherwise,} \end{cases} \\ \bigvee A &\mapsto \bigvee \mathcal{P}_\omega e^p(A), \\ \bigwedge \Pi \wedge \nabla \alpha &\mapsto \begin{cases} \bot, & \text{if } \nabla\alpha \equiv \bot, \\ e^p(\bigwedge \Pi) \wedge \nabla T_\omega e^p(\alpha), & \text{otherwise,} \end{cases} \end{aligned}$$

where $A \in \mathcal{P}_\omega \mathcal{L}_T^{\mathsf{d}}$, Π is a set of literals and $\alpha \in T_\omega \mathcal{L}_T^{\mathsf{d}}$.

If T restricts to finite sets we can use Theorem 4.18 to extend this Definition to the set of all formulas as follows

$$\begin{aligned} \exists p : \mathcal{L}_T(C \cup \{p\}) &\to \mathcal{L}_T(C), \\ a &\mapsto e^p(a^{\mathsf{d}}). \end{aligned}$$

We will often write $\exists p.a$ for the formula $\exists p(a)$. The operator $\exists p.$ is called *bisimulation quantifier.*

The recursive clauses in the definition of e^p just distribute to the subformulas. So we have for example for $A = \{a_0, a_1, \ldots, a_{n-1}\} \in \mathcal{P}_\omega \mathcal{L}_T$ that

$$e^p(\bigvee A) = \bigvee \mathcal{P}_\omega e^p(A) = \bigvee \{e^p(a_0), e^p(a_1), \ldots, e^p(a_{n-1})\}.$$

Similarly we have that the formula $\nabla T_\omega e^p(\alpha)$ is just the formula $\nabla \alpha$ with e^p applied to all of the immediate subformulas.

Remark 4.22. The function $\exists p$ removes all occurrences of the propositional letter p from its argument. This means that it restrict to an mapping $\exists p : \mathcal{L}_T(C) \to \mathcal{L}_T(C \setminus \{p\})$ for any set of propositional letters $C \subseteq \mathsf{P}$.

The formula $\exists p$ is called bisimulation quantifier because the formula $\exists p.a$ is intended to have special satisfaction conditions that uses the following notion of up-to-p-bisimulations.

Definition 4.23. A relation $R : X \nrightarrow Y$ is an *up-to-p L_{P}-bisimulation* between T-models $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ and $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ if it is an $L_{\mathsf{P}\setminus\{p\}}$-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$. Call two states x in $\mathbb{X}$ and y in $\mathbb{Y}$ *up-to-p L_{P}-bisimilar* if there is an up-to-p L_{P}-bisimulation R between $\mathbb{X}$ and $\mathbb{Y}$ with $(x, y) \in R$. We write $x \leftrightarrows_p y$ if x and y are up-to-p L_{P}-bisimilar.

For any T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ we define the relation $\Vdash_{\mathbb{X}}^{\exists p} : X \nrightarrow \mathcal{L}_T$ by

$$\begin{aligned} x_0 \Vdash_{\mathbb{X}}^{\exists p} a \quad \text{iff} \quad & \text{there is a state } y_0 \text{ in some } T\text{-model } \mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}}) \\ & \text{such that } x_0 \leftrightarrows_p y_0 \text{ and } y_0 \Vdash_\upsilon a. \end{aligned}$$

It is our goal to show in Theorem 4.25 that $\exists p.a$ has the following satisfaction conditions at state x_0 of a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$:

$$x_0 \Vdash_{\mathbb{X}} \exists p.a \quad \text{iff} \quad x_0 \Vdash_{\mathbb{X}}^{\exists p} a.$$

For this we first prove the following Lemma.

Lemma 4.24. *For any T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ there is a T-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ and a left-total up-to-p L_{P}-bisimulation $R : X \nrightarrow Y$ between $\mathbb{X}$ and $\mathbb{Y}$ such that*

$$\Vdash_{\mathbb{X}}^{\exists p} = R\,;\Vdash_{\mathbb{Y}} . \tag{13}$$

Proof. Fix the T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$. We have to construct a T-model $\mathbb{Y} = (\upsilon, V_{\mathbb{Y}})$ that has the required properties. Consider the pairs $(x, a) \in \Vdash_{\mathbb{X}}^{\exists p}$. For any such pair there is a state $y_{x,a}$ in a T-model $\mathbb{Y}_{x,a} = (Y_{x,a}, \upsilon_{x,a}, V_{\mathbb{Y}_{x,a}})$ such that $y_{x,a} \Vdash_{\mathbb{Y}_{x,a}} a$ and there is an up-to-p L_{P}-bisimulation $R_{x,a} : X \nrightarrow Y_{x,a}$ between $\mathbb{X}$ and $\mathbb{Y}_{x,a}$ such that $(x, y_{x,a}) \in R_{x,a}$.

Define $\mathbb{Y} = (\upsilon : Y \to TY, V_{\mathbb{Y}}) = \coprod_{(x,a)\in\Vdash_{\mathbb{X}}^{\exists p}} \mathbb{Y}_{x,a}$ to be the coproduct of the coalgebras $\upsilon_{x,a}$ for every $(x, a) \in \Vdash_{\mathbb{X}}^{\exists p}$, with injections $i_{Y_{x,a}} : Y_{x,a} \to Y$. Because the injections $i_{Y_{x,a}}$ are coalgebra morphisms the relations of their graphs are by Proposition 3.10 (iii) L_{P}-bisimulations between $\mathbb{Y}_{x,a}$ and $\mathbb{Y}$. By Proposition 4.11 it follows that they preserve truth and so we have for $i_{Y_{x,a}}(y_{x,a}) \in Y$ as a state of υ with $i_{Y_{x,a}}(y_{x,a}) \Vdash_{\mathbb{Y}} a$.

The relation $R : X \nrightarrow Y$ is defined by

$$R = \bigcup_{(x,a)\in\Vdash_{\mathbb{X}}^{\exists p}} \left(R_{x,a}\,; i_{Y_{x,a}}\right) .$$

The relations $R_{x,a}\,; i_{Y_{x,a}}$ are up-to-p L_{P}-bisimulations for all $(x, a) \in \Vdash_{\mathbb{X}}^{\exists p}$ because they are a composition of two up-to-p L_{P} bisimulations which is by Proposition 3.13 (ii) also an up-to-p L_{P}-bisimulation. By Proposition 3.13 (iii) we know that an arbitrary union of up-to-p L_{P}-bisimulations is again an up-to-p L_{P}-bisimulation. Hence R is an up-to-p L_{P}-bisimulation.

The relation R is left total because for all $x \in X$ we have that $(x, \top) \in \Vdash_{\mathbb{X}}^{\exists p}$ and hence that $(x, i_{Y_{x,\top}}(y_{x,\top})) \in R_{x,\top}\,; i_{Y_{x,\top}} \subseteq R$.

It remains to check that these definitions satisfy (13). For the $\subseteq$-inclusion, take any $(x, a) \in \Vdash_{\mathbb{X}}^{\exists p}$. Then $(x, a) \in R\,;\Vdash_{\mathbb{Y}}$ follows because $(x, i_{Y_{x,a}}(y_{x,a})) \in R_{x,a}\,; i_{Y_{x,a}} \subseteq R$ and $i_{Y_{x,a}}(y_{x,a}) \Vdash_{\mathbb{Y}} a$. For the $\supseteq$-inclusion take any $(x, a) \in R\,;\Vdash_{\mathbb{Y}}$. So there is a $y \in Y$ such that $(x, y) \in R$ and $y \Vdash_{\mathbb{Y}} a$. Since R is an up-to-p L_{P}-bisimulations it follows by the definition of $\Vdash_{\mathbb{X}}^{\exists p}$ that $(x, a) \in \Vdash_{\mathbb{X}}^{\exists p}$. □

Theorem 4.25. *Assume that T restricts to finite sets and L is quasi-functorial. Then it holds for the function $\exists p : \mathcal{L}_T \to \mathcal{L}_T$, as defined in Definition 4.21, that for any state x_0 in a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ and any formula $a \in \mathcal{L}_T$*

$$x_0 \Vdash_{\mathbb{X}} \exists p.a \quad \textit{iff} \quad \textit{there is a state } y_0 \textit{ in some } T\textit{-model } \mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}}) \textit{ such that } x_0 \leftrightarrows_p y_0 \textit{ and } y_0 \Vdash_{\mathbb{Y}} a.$$

Proof. Because $\exists p = e^p((-)^{\mathsf{d}})$ and we know from Theorem 4.18 that $(-)^{\mathsf{d}}$ preserves truth it is enough to show that for all $a \in \mathcal{L}_T^{\mathsf{d}}$

$$x_0 \Vdash_{\mathbb{X}} e^p(a) \quad \text{iff} \quad x_0 \Vdash_{\mathbb{X}}^{\exists p} a. \tag{14}$$

This is done by induction on the complexity of $a \in \mathcal{L}_T^{\mathsf{d}}$.

We omit the case where $a = \bigwedge \Pi$. It is similar to the case $a = \bigwedge \Pi \wedge \nabla\alpha$ that we prove below but much easier because it only involves propositional letters..

For the case $a = \bigvee A$ where $A \in \mathcal{P}_\omega \mathcal{L}_T^{\mathsf{d}}$ first recall that $e^p(\bigvee A) = \bigvee \mathcal{P}_\omega e^p(A)$. For the left to right direction assume that $x_0 \Vdash_{\mathbb{X}} \bigvee \mathcal{P}_\omega e^p(A)$. So there is an $a \in A$

such that $x_0 \Vdash_{\mathbb{X}} e^p(a)$. By the induction hypothesis this means that there is a state y_0 in a T-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ such that $x_0 \leftrightarroweq_p y_0$ and $y_0 \Vdash_{\mathbb{Y}} a$. Hence also $y_0 \Vdash_{\mathbb{Y}} \bigvee A$ and because $x_0 \leftrightarroweq_p y_0$ it follows that $x_0 \Vdash_{\mathbb{X}}^{\exists p} \bigvee B$.

For the other direction assume that there is a state y_0 in a T-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ such that $x_0 \leftrightarroweq_p y_0$ and $y_0 \Vdash_{\upsilon} \bigvee A$. So $y_0 \Vdash_{\mathbb{Y}} a$ for an $a \in A$. Since $x_0 \leftrightarroweq_p y_0$ it follows that $x_0 \Vdash_{\mathbb{X}}^{\exists p} b$ and by induction hypothesis that $x_0 \Vdash_{\mathbb{X}} e^p(b)$. By the satisfaction conditions of the disjunction we conclude that $x_0 \Vdash_{\mathbb{X}} \bigvee \mathcal{P}_\omega e^p(A)$.

The last case is where $a = \bigwedge \Pi \wedge \nabla\alpha$ for a set Π of literals and an $\alpha \in T_\omega \mathcal{L}_T$. If $p, \neg p \in \Pi$ or $\nabla\alpha \equiv \bot$, then it is not the case that $x_0 \Vdash_{\mathbb{X}}^{\exists p} a$ and the claim follows immediately from the definition of $e^p(\bigwedge \Pi \wedge \nabla\alpha)$. So assume that not $p, \neg p \in \Pi$ and that $\nabla\alpha$ is satisfiable. The induction hypothesis is that (14) holds for all formulas $a \in \mathsf{Base}(\alpha)$. This can be expressed as

$$\Vdash_{\mathbb{X}} ; (e^p\restriction_{\mathsf{Base}(\alpha)})^\circ = \Vdash_{\mathbb{X}}^{\exists p}\restriction_{X \times \mathsf{Base}(\alpha)} . \tag{15}$$

We first do the easier direction from right to left. So assume that there is state y_0 in a T-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ with $y_0 \Vdash_{\mathbb{Y}} \bigwedge \Pi \wedge \nabla\alpha$ and there is an up-to-p L_{P}-bisimulation $R : X \nrightarrow Y$ between $\mathbb{X}$ and $\mathbb{Y}$ with $(x_0, y_0) \in R$. From the definition of $\Vdash_{\mathbb{X}}^{\exists p} : X \times \mathcal{L}_T$ it follows that $R ; \Vdash_{\upsilon} \subseteq \Vdash_{\mathbb{X}}^{\exists p}$. Because the up-to-$p$ L_{P}-bisimulation R is in particular an L-bisimulation it holds that $(\xi(x_0), \upsilon(y_0)) \in LR$. Because of $y_0 \Vdash_{\upsilon} \nabla\alpha$ it we have that $(\upsilon(y_0), \alpha) \in L\Vdash_{\upsilon}$. Putting this together we get $(\xi(x_0), \alpha) \in LR ; L\Vdash_{\upsilon}$ and we can calculate

$$\begin{aligned}
(\xi(x_0), \alpha) &\in (LR ; L\Vdash_{\upsilon})\restriction_{TX \times T\mathsf{Base}(\alpha)} && \alpha \in T\mathsf{Base}(\alpha) \\
&\subseteq (L(R ; \Vdash_{\upsilon}))\restriction_{TX \times T\mathsf{Base}(\alpha)} && \text{(L2)} \\
&\subseteq \left(L\Vdash_{\mathbb{X}}^{\exists p}\right)\restriction_{TX \times T\mathsf{Base}(\alpha)} && R ; \Vdash_{\upsilon} \subseteq \Vdash_{\mathbb{X}}^{\exists p} \text{and (L1)} \\
&= L\left(\Vdash_{\mathbb{X}}^{\exists p}\restriction_{X \times \mathsf{Base}(\alpha)}\right) && \text{Proposition 3.12} \\
&= L\left(\Vdash_{\mathbb{X}} ; (e^p\restriction_{\mathsf{Base}(\alpha)})^\circ\right) && (15) \\
&= (L\Vdash_{\mathbb{X}} ; (e^p)^\circ)\restriction_{TX \times T\mathsf{Base}(\alpha)} && \text{Proposition 3.12} \\
&\subseteq (L\Vdash_{\mathbb{X}} ; (e^p)^\circ) \\
&= L\Vdash_{\mathbb{X}} ; (Te^p)^\circ. && \text{Proposition 3.10 (ii)}
\end{aligned}$$

This entails that $(\xi(x_0), Te^p(\alpha)) \in L\Vdash_{\mathbb{X}}$. Hence $x_0 \Vdash_{\mathbb{X}} \nabla Te^p(\alpha)$.

Because R is an up-to-p L_{P}-bisimulation that connects x_0 and y_0 it follows that x_0 makes the same propositional letters true as y_0, with the possible exception of p. Hence $y_0 \Vdash_{\mathbb{Y}} \bigwedge \Pi$ entails that $x_0 \Vdash_X \bigwedge (\Pi \setminus \{p, \neg p\})$. It follows that $x_0 \Vdash_{\mathbb{X}} \bigwedge (\Pi \setminus \{p, \neg p\}) \wedge \nabla T_\omega e^p(\alpha)$ which is what we need to show because $e^p(\bigwedge \Pi \wedge \nabla\alpha) = \bigwedge (\Pi \setminus \{p, \neg p\}) \wedge \nabla T_\omega e^p(\alpha)$.

For the left to right direction assume that $x_0 \Vdash_{\mathbb{X}} e^p(\bigwedge \Pi \wedge \nabla\alpha)$. By the definition of e^p this means that $x_0 \Vdash_{\mathbb{X}} \bigwedge (\Pi \setminus \{p, \neg p\})$ and that $x_0 \Vdash_{\mathbb{X}} \nabla Te^p(\alpha)$. From the latter it follows that $(\xi(x_0), Te^p\restriction_{\mathsf{Base}(\alpha)} (\alpha)) \in L\Vdash_{\mathbb{X}}$ since $Te^p(\alpha) =$

$T(e^p\restriction_{\mathsf{Base}(\alpha)})(\alpha)$. So we get

$$\begin{aligned}(\xi(x_0),\alpha) \in L\Vdash_{\mathbb{X}} ; (Te^p\restriction_{\mathsf{Base}(\alpha)})^\circ &= L\left(\Vdash_{\mathbb{X}} ; (e^p\restriction_{\mathsf{Base}(\alpha)})^\circ\right) && \text{Proposition 3.10 (ii)}\\ &= L\left(\Vdash_{\mathbb{X}}^{\exists p}\restriction_{X\times\mathsf{Base}(\alpha)}\right) && (15)\\ &= \left(L\Vdash_{\mathbb{X}}^{\exists p}\right)\restriction_{TX\times T\mathsf{Base}(\alpha)} && \text{Proposition 3.12}\\ &\subseteq L\Vdash_{\mathbb{X}}^{\exists p}.\end{aligned}$$

We have to show that $x_0 \Vdash_{\mathbb{X}}^{\exists p} \nabla\alpha$. To get this we need a witnessing state y_0 in a T-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ such that $y_0 \Vdash_{\mathbb{Y}} \bigwedge\Pi \wedge \nabla\alpha$, and an up-to-$p$ L_{P}-bisimulation $R : X \nrightarrow Y$ between $\mathbb{X}$ and $\mathbb{Y}$ with $(x_0, y_0) \in R$.

A first approximation is given by Lemma 4.24 from which we get a T-model $\mathbb{Y}'' = (Y'', \upsilon'', V_{\mathbb{Y}''})$ and a right-total up-to-p L_{P}-bisimulation $R'' : X \nrightarrow Y''$ between $\mathbb{X}$ and $\mathbb{Y}''$ such that

$$\Vdash_{\mathbb{X}}^{\exists p} = R'' ; \Vdash_{\mathbb{Y}''}. \tag{16}$$

We already know that $(\xi(x_0), \alpha) \in L\Vdash_{\mathbb{X}}^{\exists p} = L(R'' ; \Vdash_{\mathbb{Y}''})$. The plan is to use the quasi-functoriality of L to get a $\upsilon_0 \in T(Y'')$ such that $(\xi_T(x_0), \upsilon_0) \in LY''$ and $(\upsilon_0, \alpha) \in L\Vdash_{\upsilon''}$. This υ_0 will then function as the unfolding of a witnessing state y_0.

Because $\nabla\alpha$ is satisfiable there is a state z_0 in a T-model $\mathbb{Z} = (Z, \zeta, V_{\mathbb{Z}})$ such that $z_0 \Vdash_{\mathbb{Z}} \nabla\alpha$. Let $\mathbb{Y}' = \mathbb{Y}'' + \mathbb{Z}$ with $\mathbb{Y}' = (Y', \upsilon', V_{\mathbb{Y}'})$ be the coproduct of $\mathbb{Y}''$ and $\mathbb{Z}$ with injections $i_{Y''}$ and i_Z. The injections are coalgebra morphisms and so they preserves truth by Proposition 3.13 (i) and Proposition 4.11. That gives us that $\Vdash_{\mathbb{Y}''} = i_{Y''} ; \Vdash_{\mathbb{Y}'}$. Now define the relation $R' : X \nrightarrow Y', R' = R'' ; i_{Y''}$ which by Proposition 3.13 (ii) is an up-to-p L_{P}-bisimulation between $\mathbb{X}$ and $\mathbb{Y}'$ which is, as one easily verifies, also left-total and satisfies

$$\Vdash_{\mathbb{X}}^{\exists p} = R'' ; \Vdash_{\mathbb{Y}''} = R' ; \Vdash_{\mathbb{Y}'}.$$

Because i_Z preserves truth we get that $i_Z(z_0) \Vdash_{\mathbb{Y}'} \nabla\alpha$. Hence $(\upsilon'(i_Z(z_0)), \alpha) \in L\Vdash_{\mathbb{Y}'}$ and so $\alpha \in \mathsf{img}(\Vdash_{\mathbb{Y}'})$.

Since R' is left-total there is a function $f : X \to Y'$ with $f \subseteq R'$. Using f we get that $(\xi(x_0), Tf(\xi(x_0))) \in Tf \subseteq Lf \subseteq LR$ by (L3) and (L1). It follows that $\xi(x_0) \in \mathsf{preimg}(LR')$.

Now we have that $(\xi(x_0), \alpha) \in L(\Vdash_{\mathbb{X}}^{\exists p} = R' ; \Vdash_{\mathbb{Y}'})$, $\xi(x_0) \in \mathsf{preimg}(LR')$ and $\alpha \in \mathsf{img}(\Vdash_{\mathbb{Y}'})$. So it follows from the quasi-functoriality of L that there exists a $\upsilon_0 \in TY'$ with $(\xi(x_0), \upsilon_0) \in LR'$ and $(\upsilon_0, \alpha) \in L\Vdash_{\mathbb{Y}'}$. This we can use to define the coalgebra $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ on $Y = \{y_0\} \cup Y'$ for a $y_0 \notin Y'$, with transition function

$$\begin{aligned}\upsilon : Y &\to TY,\\ y &\mapsto \begin{cases}\upsilon_0, & \text{if } y = y_0\\ \upsilon'(y), & \text{if } y \in Y',\end{cases}\end{aligned}$$

The transition function υ is well defined for $y \in Y'$ because $TY' \subseteq TY$ since T is standard. The valuation $V_{\mathbb{Y}}$ is defined such that $V_{\mathbb{Y}}(q) = V_{\mathbb{Y}'}(q) \cup \{y_0\}$ for all $p \in \mathsf{P} \setminus \{p\}$ with $x_0 \in V_{\mathbb{X}}(q)$ and $V_{\mathbb{Y}}(q) = V_{\mathbb{Y}'}(q)$ for all $q \in \mathsf{P} \setminus \{p\}$

with $x_0 \notin V_{\mathbb{X}}(q)$. For the proposition letter p we set $V_{\mathbb{Y}}(p) = V_{\mathbb{Y}'}(q) \cup \{y_0\}$ if $p \in \Pi$ and $V_{\mathbb{Y}}(p) = V_{\mathbb{Y}'}(p)$ otherwise. The relation $R : X \not\rightarrow Y$ is defined by $R = R' ; \iota_{Y',Y} \cup \{(x_0, y_0)\}$.

For this we now first show that the inclusion $\iota_{Y',Y} : Y' \not\rightarrow Y$ is an L_{P}-bisimulation. It is clear from the definition of the valuation $V_{\mathbb{Y}}$ that $\iota_{Y',Y}$ preserves all propositional letters. To see that it is an L-bisimulation consider

$$\begin{aligned}
\iota_{Y',Y} &\subseteq \upsilon' ; \iota_{TY',TY} ; \upsilon^\circ && \text{Definition of } \upsilon \\
&= \upsilon' ; T_{\mathsf{P}}\iota_{Y',Y} ; \upsilon^\circ && T \text{ standard} \\
&\subseteq \upsilon' ; L\iota_{Y',Y} ; \upsilon^\circ. && \text{(L3)}
\end{aligned}$$

By Remark 3.4 it follows that $\iota_{Y',Y}$ is an L-bisimulation between υ' and υ.

Because $\iota_{Y',Y}$ is an L_{P}-bisimulation between $\mathbb{Y}'$ and $\mathbb{Y}$ it preserves truth by Proposition 4.11. So we have that $\iota^\circ_{Y',Y} ; \Vdash_{\mathbb{Y}'} \subseteq \Vdash_{\mathbb{Y}}$. We also have that $(\upsilon_0, \upsilon_0) \in \iota^\circ_{TY',TY} = T(\iota^\circ_{Y',Y}) \subseteq L\iota^\circ_{Y',Y}$ because T is standard and by (L3). Moreover we already know that $(\upsilon_0, \alpha) \in L\Vdash_{\mathbb{Y}'}$. Hence we get

$$\begin{aligned}
(\upsilon_0, \alpha) \in L\iota^\circ_{Y',Y} ; L\Vdash_{\mathbb{Y}'} &\subseteq L(\iota^\circ_{Y',Y} ; \Vdash_{\mathbb{Y}'}) && \text{(L2)} \\
&\subseteq L\Vdash_{\mathbb{Y}} && \iota^\circ_{Y',Y} ; \Vdash_{\mathbb{Y}'} \subseteq \Vdash_{\mathbb{Y}} \text{ and (L1)}
\end{aligned}$$

It follows that $(\upsilon(y_0), \alpha) \in L\Vdash_\upsilon$, and therefore $y_0 \Vdash_{\mathbb{Y}} \nabla\alpha$, because $\upsilon(y_0) = \upsilon_0$. The definition of the valuation $V_{\mathbb{Y}}$ entails, together with the fact that $x_0 \Vdash_{\mathbb{X}} \bigwedge(\Pi \setminus \{p, \neg p\})$, that $y_0 \Vdash_{\mathbb{Y}} \bigwedge \Pi$. It follows that $y_0 \Vdash_{\mathbb{Y}} \bigwedge \Pi \wedge \nabla\alpha$.

The last thing we have to prove is that $R = R' ; \iota_{Y',Y} \cup \{(x_0, y_0)\}$ is an up-to-p L_{P}-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$. The relation R preserves the truth of propositional letters, except for p, because as we already observed R' and $\iota_{Y',Y}$ do and the valuation $V_{\mathbb{Y}}$ was defined such that y_0 makes up-to-p the same propositional letters true as x_0. So it remains to check that the relation R is an L-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$. For this we show that $(\xi(x), \upsilon(y)) \in LR$ for an arbitrary $(x, y) \in R$. We first do the case where $(x, y) \in R' ; \iota_{Y',Y}$. We have by Proposition 3.13 (ii) that the composition $R' ; \iota_{Y',Y}$ is an L-bisimulation and so $(\xi(x), \upsilon(y)) \in L(R' ; \iota_{Y',Y})$ because $(x, y) \in R' ; \iota_{Y',Y}$. Since $R' ; \iota_{Y',Y} \subseteq R$ it follows by (L1) that $(\xi(x), \upsilon(y)) \in LR$. In the other case we have that $(x, y) = (x_0, y_0)$ and have to show that $(\xi(x_0), \upsilon(y_0)) \in LR$. For this we use that $(\xi(x_0), \upsilon_0) \in LR'$ and that $(\upsilon_0, \upsilon(y_0)) \in T\iota_{Y,Y'}$, since T is standard and $\upsilon(y_0) = \upsilon_0$. So we can compute

$$\begin{aligned}
(\xi(x_0), \upsilon(y_0)) \in LR' ; T\iota_{Y,Y'} &\subseteq LR' ; L\iota_{Y,Y'} && \text{(L3)} \\
&\subseteq L(R' ; \iota_{Y,Y'}) && \text{(L2)} \\
&\subseteq LR. && R' ; \iota_{Y,Y'} \subseteq R \text{ and (L1)}
\end{aligned}$$

□

Example 4.26. Consider the functor F^3_2 from Example 2.1 (x) and its lax extension L^3_2 from Example 3.11 (vi). We can show that bisimulation quantifiers are not definable in the language $\mathcal{L}_{F^3_2}$ if the semantics of the nabla is given by L^3_2. We show this by proving that there is no formula in $\mathcal{L}_{F^3_2}$ that has $x_0 \Vdash^{\exists p}_{\mathbb{X}} \nabla(p, \neg p, \neg p)$ as its satisfaction condition at a state x_0 of a F^3_2-model $\mathbb{X}$. For this purpose we first prove the claim that for any state x_0 in a F^3_2-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ with $\xi(x_0) = (x_1, x_2, x_3)$:

$$x_0 \Vdash^{\exists p}_{\mathbb{X}} \nabla(p, \neg p, \neg p) \quad \text{iff} \quad x_2 \leftrightarroweq_p x_3. \tag{17}$$

Assume first that $x_0 \Vdash^{\exists p}_{\mathbb{X}} \nabla(p, \neg p, \neg p)$. This means that there is a state y_0 in a F^3_2-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ and there is an up-to-p $L^3_{2\mathsf{P}}$-bisimulation $R : X \nrightarrow Y$ between $\mathbb{X}$ and $\mathbb{Y}$ such that $y_0 \Vdash_{\mathbb{Y}} \nabla(p, \neg p, \neg p)$. Let $\upsilon(y_0) = (y_1, y_2, y_3)$. We know that $y_1 \neq y_2$ and $y_1 \neq y_3$ because $y_1 \Vdash_\upsilon p$ whereas $y_2 \Vdash_\upsilon \neg p$ and $y_3 \Vdash_\upsilon \neg p$ by the semantics of the nabla. Because (y_1, y_2, y_3) can contain at most two distinct elements it follows that $y_2 = y_3$. Because R is an L^3_2-bisimulation between ξ and υ with $(x_0, y_0) \in R$ we have that $(x_2, y_2), (x_3, y_3) \in R$ since L^3_2 is defined componentwise. This entails $(x_2, x_3) \in R\,; R^\circ$ which shows that x_2 and x_3 are up-to-p $L^3_{2\mathsf{P}}$-bisimilar because by Proposition 3.13 (ii) the composition $R\,; R^\circ$ is an up-to-p $L^3_{2\mathsf{P}}$-bisimulation.

For the other direction of (17) assume that $x_2 \leftrightarroweq_p x_3$. Now we want to construct a witnessing F^3_2-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ for the left hand side. The idea is to first identify the points x_2 and x_3 and then make this identified point distinct from x_1. This enables us to make p true at the point corresponding to x_1 and false at the point corresponding to x_2 and x_3.

It follows from Corollary 3.14 that the bisimilarity relation $\leftrightarroweq_p \colon X \times X$ is an $L^3_{2\mathsf{P}\setminus\{p\}}$-bisimulation equivalence on the T_{P}-coalgebra $\widehat{\mathbb{X}} : X \to F^3_{2\,\mathsf{P}}X$ corresponding to $\mathbb{X}$. From Proposition 3.8 we get the bisimulation quotient $\zeta'' : Z'' \to F^3_{2\,\mathsf{P}\setminus\{p\}}Z''$, where $Z'' = X/{\leftrightarroweq_p}$, such that the projection $p : X \to Z''$ is a $F^3_{2\,\mathsf{P}\setminus\{p\}}$-coalgebra morphism from $r_X^{\mathsf{P},\mathsf{P}\setminus\{p\}} \circ \widehat{\mathbb{X}}$ to ζ''. Because $x_2 \leftrightarroweq_p x_3$ it holds that $p(x_2) = p(x_3)$.

Next consider the coproduct $\zeta' = \zeta'' + \zeta'' : Z' \to F^3_{2\,\mathsf{P}\setminus\{p\}}Z'$ where $Z' = Z'' + Z''$ with injections $i_0, i_1 : Z'' \to Z'$. Intuitively ζ' consists of two copies of ζ'' where the first copy is accessed by i_0 and the second by i_1. So for every equivalence class $z \in Z'' = X/{\leftrightarroweq_p}$ there are two identical copies $i_0(z)$ and $i_1(z)$ of it in ζ'. We can then define the relation $R' = p\,; (i_0 \cup i_1) : X \nrightarrow Z'$ that connects any point in x with the two copies of its equivalence class in ζ'. Because of Proposition 3.13 it we have that $R' = p\,; (i_0 \cup i_1) : X \nrightarrow Z'$ is an up-to-p $L^3_{2\mathsf{P}}$-bisimulation between $r_X^{\mathsf{P},\mathsf{P}\setminus\{p\}} \circ \widehat{\mathbb{X}}$ and ζ'. By assuming that the propositional letter p is false at every state we can take the $F^3_{2\,\mathsf{P}\setminus\{p\}}$-coalgebra ζ' to be a $F^3_{2\,\mathsf{P}}$-coalgebra $\widehat{\mathbb{Y}'} : Y' \to F^3_{2\,\mathsf{P}}Y'$ where $Y' = Z'$. So there is an F^3_2-model $\mathbb{Y}' = (Y', \upsilon', V_{\mathbb{Y}'})$ corresponding to $\widehat{\mathbb{Y}}$. For $\mathbb{Y}'$ we have that R' is an up-to-p $L^3_{2\mathsf{P}}$-bisimulation between the F^3_2-models $\mathbb{X}$ and $\mathbb{Y}'$.

In $\mathbb{Y}'$ there is the state $i_0 \circ p(x_1)$ that is bisimilar to x_1 and distinct from the point $i_1 \circ p(x_2) = i_1 \circ p(x_3)$ that is bisimilar to x_2 and to x_3. That is the situation we were aiming for and we use it to define the F^3_2-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ on the set $Y = Y' \cup \{y_0\}$ for a $y_0 \notin Y'$ with transition function

$$\begin{aligned} \upsilon : Y &\to F^3_2 Y, \\ y &\mapsto \begin{cases} (i_0 \circ p(x_1), i_1 \circ p(x_2), i_1 \circ p(x_3)), & \text{if } y = y_0 \\ \upsilon'(y). & \text{if } y \in Y' \end{cases} \end{aligned}$$

The valuation $V_{\mathbb{Y}} : \mathsf{P} \to \mathcal{P}Y$ is defined such that $V_{\mathbb{Y}}(p) = \{i_0 \circ p(x_1)\}$ and $V_{\mathbb{Y}}(q) = V_{\mathbb{Y}'}(q)$ for $q \in \mathsf{P} \setminus \{p\}$.

Because p holds only at the state $i_0 \circ p(x_1)$ and $i_0 \circ p(x_1) \neq i_1 \circ p(x_2) = i_1 \circ p(x_3)$ it follows that $y_0 \Vdash_\upsilon \nabla(p, \neg p, \neg p)$. One can also check that $R = R' \cup \{(x_0, y_0)\} : X \nrightarrow Y$ is an up-to-p $L^3_{2\mathsf{P}}$-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$. Together this yields that $x_0 \Vdash^{\exists p}_{\xi} \nabla(p, \neg p, \neg p)$ and finishes the proof of (17).

From the equivalence (17) it follows that a formula $b \in \mathcal{L}_{F_2^3}$, that has $x_0 \Vdash^{\exists p}_{\mathbb{X}} \nabla(p, \neg p, \neg p)$ as its satisfaction conditions at a state x_0, characterizes the property that the successors x_2 and x_3 of x_0 are up-to-p L_{P}-bisimilar. This is interesting because bisimilarity is a property that depends on arbitrarily remote successors of a state. But every formula $b \in \mathcal{L}_{F_2^3}$ has a finite modal rank and can only characterize properties of successors that are at most n steps away. Hence it is not possible that there is such an $b \in \mathcal{L}_{F_2^3}$.

To make this argument more precise assume for a contradiction that there is a formula $b \in \mathcal{L}_{F_2^3}$ such that $x \Vdash_{\mathbb{X}} b$ is equivalent to $x \Vdash^{\exists p}_{\mathbb{X}} \nabla(p, \neg p, \neg p)$ at every state x in any F_2^3-model $\mathbb{X}$. Let n be the modal rank of b. Now consider two F_2^3-models $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ and $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$. The model $\mathbb{X}$ is defined on a set $X = \{x\}$ such that

$$\xi(x) = (\emptyset, (x, x, x)),$$

and $V_{\mathbb{X}}(q) = \emptyset$ for all $q \in \mathsf{P}$. The F_2^3-model $\mathbb{Y} = (Y, \upsilon, V_{\mathbb{Y}})$ has the set $Y = \{y, y'\} \cup \{y_i \mid i = 0, \dots, n-1\}$ with $n+2$ distinct elements as its states, its transition function $\upsilon : Y \to F_2^3 Y$ is defined such that

$$\begin{aligned} \upsilon(y) &= (y, y, y_{n-1}), \\ \upsilon(y_{i+1}) &= (y_i, y_i, y_i), \\ \upsilon(y_0) &= (y', y', y'), \\ \upsilon(y') &= (y', y', y'), \end{aligned}$$

and its valuation is such that $V_{\mathbb{Y}}(r) = \{y'\}$ for an $r \in \mathsf{P} \setminus \{p\}$ and $V_{\mathbb{Y}}(q) = \emptyset$ for all $q \in \mathsf{P} \setminus \{r\}$. For these models it is easy to show by induction on $d < n$ that $x \Vdash_{\mathbb{X}} a$ iff $y \Vdash_{\mathbb{Y}} a$ and $x \Vdash_{\mathbb{X}} a$ iff $y_d \Vdash_{\mathbb{Y}} a$ for all $a \in \mathcal{L}_{F_2^3}$ of rank at most d. It follows that $x \Vdash_{\mathbb{X}} a$ iff $y \Vdash_{\mathbb{Y}} a$ for every $a \in \mathcal{L}_{F_2^3}$ with rank at most n. Hence $x \Vdash_{\mathbb{X}} b$ iff $y \Vdash_{\mathbb{Y}} b$. But $x \Vdash^{\exists p}_{\mathbb{X}} \nabla(p, \neg p, \neg p)$ and not $y \Vdash^{\exists p}_{\mathbb{Y}} \nabla(p, \neg p, \neg p)$ since we clearly have that $x \leftrightarroweq_p x$ whereas not $y \leftrightarroweq_p y_{n-1}$ because y_{n-1} has a successor where r is true but y does not.

An application of bisimulation quantifiers is the following interpolation result.

Corollary 4.27 (Uniform Interpolation)**.** *Assume that T restricts to finite sets and that L is quasi-functorial. For any finite sets of propositional letters $C_a \subseteq \mathsf{P}$ and $D \subseteq C_a$ and any formula $a \in \mathcal{L}_T(C_a)$ there is a formula $a_D \in \mathcal{L}_T(D)$ such that for all $C_b \subseteq \mathsf{P}$ with $C_a \cap C_b \subseteq D$ and formulas $b \in \mathcal{L}_T(C_b)$ we have that*

$$a \models b \quad \textit{iff} \quad a_D \models b.$$

Proof. Let $\{p_0, p_1, \dots, p_{n-1}\} = C_a \setminus D$. Then set

$$a_D = \exists p_0 . \exists p_1 . \dots \exists p_{n-1} . a.$$

With Remark 4.22 we have that $a_D \in \mathcal{L}_T(C_a \cap D) \subseteq \mathcal{L}_T(D)$.

To check that $a \models b$ iff $a_D \models b$ assume first that $a \models b$. To prove that $a_D \models b$ we have to show $x_0 \Vdash_{\mathbb{X}} b$ for an arbitrary state x_0 in a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ with $x_0 \Vdash_{\mathbb{X}} a_D$. By the semantics of the bisimulation quantifiers we get states y_i in T-models $\mathbb{Y}_i$ for $i = 1, 2, \dots, n$ such that $x_0 \leftrightarroweq_{p_0} y_1$, $y_1 \leftrightarroweq_{p_1} y_2$, $\dots$, $y_{n-1} \leftrightarroweq_{p_{n-1}} y_n$ and $y_n \Vdash_{\mathbb{Y}_n} a$. From the latter fact it follows that $y_n \Vdash_{\mathbb{Y}_n} b$

since we assume that $a \models b$. Because every of the witnessing up-to-p_i L_{P}-bisimulations for $i = 0, 1, \ldots, n-1$ is also an $L_{\mathsf{P}\setminus\{p_0,p_1,\ldots,p_{n-1}\}}$-bisimulation we can apply Proposition 3.13 (ii) to obtain that x_0 and y_n are $L_{\mathsf{P}\setminus\{p_0,p_1,\ldots,p_{n-1}\}}$-bisimilar. It follows from the assumptions that $C_b \subseteq \mathsf{P} \setminus \{p_0, p_1, \ldots, p_{n-1}\}$. So we can use Proposition 4.11 to get $x_0 \Vdash_{\mathbb{X}} b$.

For the other direction we show that $a \models a_D$. Then $a \models b$ follows by transitivity from $a_D \models b$. So take any state x_0 in a T-model $\mathbb{X} = (X, \xi, V_{\mathbb{X}})$ with $x_0 \Vdash_{\mathbb{X}} a$. Then clearly $x_0 \Vdash_{\mathbb{X}} a_D$ because x_0 is up-to-p L_{P}-bisimilar to itself for any $p \in \mathsf{P}$, since Δ_X is an L_{P}-bisimulation. □

5 Conclusions and Further Questions

In this thesis we proved that lax extensions which preserve diagonals give rise to a notion of bisimulations that is adequate for behavioral equivalence and we demonstrated that lax extension can be used to define a well behaved semantics for the nabla modality. For these reasons it is interesting to study lax extension that preserve diagonals in the context of coalgebraic modal logic. Another indication of the importance of lax extensions that preserve diagonals is that, of all the functors we consider, only the neighborhood functor $\mathcal{N}$ does not posses a lax extension that preserves diagonals and for this functor we showed in Proposition 3.7 that there is no relation lifting that characterizes behavioral equivalence.

An interesting goal for further research would be to characterize the functors which have a lax extension that preserves diagonals. Our Theorem 3.26 is a first step into this direction but it only applies to finitary functors and the condition it gives, that the functor has a separating set of monotone predicate liftings, is not more fundamental than what it is supposed to characterize. Furthermore one could try to find a canonical way to obtain a lax extension that preserves diagonals for the functors that posses one, similar to the Barr extension of weak pullback preserving functors. For this it might be helpful to note that for all our examples of functors which have a lax extension that preserves diagonals, the relation lifting $\widehat{T}$ from Example 3.2 (viii) also characterizes behavioral equivalence, though it is not necessarily a lax extension itself. A more general question one could work on is to find characterizing criteria of functors that have a relation lifting that characterizes behavioral equivalence. It might turn out that every functor with a relation lifting that characterizes behavioral equivalence also has a lax extension that preserves diagonals.

Another starting point for future work is our Theorem 4.25, which states that bisimulation quantifiers are definable in the nabla logic of a quasi-functorial lax extension. For example one could investigate the the property that a lax extension is quasi-functorial more carefully. Which property of the functor $\mathcal{M}$ brings about that $\mathcal{M}$ has a quasi-functorial lax extension and hence definable bisimulation quantifiers? It would also be interesting to see whether bisimulation quantifiers are still definable in the nabla logic of a quasi-functorial lax extension if one adds modal fixpoint operators to the logic.

Acknowledgments

Foremost, I thank Yde Venema for his active supervision of this thesis project. I am especially grateful for his patience with me giving wrong proofs and for his encouragement and support when something did not work as expected.

My sincere thanks also go to the MoL-room people: Gabriela Aslı Rino Nesin, Paula Henk, Marta Sznajder, Peter Fritz and Erik Parmann, for support on a daily basis, for substantial discussions and for providing the social environment, with all ups and downs, during the time this thesis was written.

References

[1] Jiří Adámek and Věra Trnková. *Automata and Algebras in Categories.* Kluwer Academic Publishers, Norwell, MA, USA, 1990.

[2] Alexandru Baltag. A logic for coalgebraic simulation. *Electronic Notes in Theoretical Computer Science*, 33:42–60, 2000.

[3] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge: Cambridge University Press, 2001.

[4] Helle Hvid Hansen and Clemens Kupke. A coalgebraic perspective on monotone modal logic. *Electronic Notes in Theoretical Computer Science*, 106:121–143, December 2004.

[5] Helle Hvid Hansen, Clemens Kupke, and Eric Pacuit. Bisimulation for neighbourhood structures. In *Proceedings of the 2nd International Conference on Algebra and Coalgebra in Computer Science*, CALCO'07, pages 279–293, Berlin, Heidelberg, 2007. Springer-Verlag.

[6] Helle Hvid Hansen, Clemens Kupke, and Eric Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Logical Methods in Computer Science*, 5(2), 2009.

[7] Christian Kissig and Yde Venema. Complementation of coalgebra automata. In *Proceedings of the 3rd international conference on Algebra and coalgebra in computer science*, CALCO'09, pages 81–96, Berlin, Heidelberg, 2009. Springer-Verlag.

[8] Clemens Kupke. *Finitary Coalgebraic Logic.* PhD thesis, Universiteit van Amsterdam, 2006.

[9] Clemens Kupke, Alexander Kurz, and Yde Venema. Completeness of the finitary Moss logic. In Carlos Areces and Rob Goldblatt, editors, *Advances in Modal Logic*, 2008.

[10] Clemens Kupke, Alexander Kurz, and Yde Venema. Completeness for the coalgebraic cover modality, to appear, see [9] for an earlier version.

[11] Alexander Kurz and Raul Andres Leal. Equational coalgebraic logic. *Electronic Notes in Theoretical Computer Science*, pages 333–356, 2009.

[12] Paul Blain Levy. Similarity quotients as final coalgebras. In *FOSSACS*, pages 27–41, 2011.

[13] Lawrence S. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96, 1999.

[14] Dirk Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309(1–3):177–193, 2003.

[15] Dirk Pattinson. An introduction to the theory of coalgebras, 2003. Lecture Notes, Second North American Summer School on Logic, Language and Information.

[16] J. J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theoretical Computer Science*, 249(1):3 – 80, 2000.

[17] Luigi Santocanale and Yde Venema. Uniform interpolation for monotone modal logic. In Lev Beklemishev, Valentin Goranko, and Valentin Shehtman, editors, *Advances in Modal Logic, Volume 8*, pages 350–370. College Publications, 2010.

[18] Lutz Schröder. Expressivity of coalgebraic modal logic: The limits and beyond. *Theoretical Computer Science*, 390(2-3):230 – 247, 2008. Foundations of Software Science and Computational Structures.

[19] Christoph Schubert and Gavin J. Seal. Extensions in the theory of lax algebras. *Theories and Applications of Categories*, 21(7):118–151, 2008.

[20] Albert Thijs. *Simulation and Fixpoint Semantics.* PhD thesis, University of Groningen, 1996.

[21] Yde Venema. Algebras and coalgebras. In Patrick Blackburn, Johan van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 331 – 426. Elsevier, 2007.