Lax Extensions of Coalgebra Functors

Johannes Marti and Yde Venema

ILLC, University of Amsterdam

Abstract. We discuss the use of relation lifting in the theory of set-based coalgebra. On the one hand we prove that the neighborhood functor does not extend to a relation lifting of which the associated notion of bisimilarity coincides with behavorial equivalence.
On the other hand we argue that relation liftings may be of use for many other functors that do not preserve weak pullbacks, such as the monotone neighborhood functor. We prove that for any relation lifting L that is a lax extension extending the coalgebra functor T and preserving diagonal relations, L-bisimilarity captures behavioral equivalence. We also show that if T is finitary, it admits such an extension iff there is a separating set of finitary monotone predicate liftings for T.

Keywords: coalgebra, relation lifting, predicate lifting, bisimilarity.

1 Introduction

There are at least two reasons why the notion of relation lifting plays an important role in the theory of (set-based) coalgebras: to characterize bisimulations, and to define the semantics of Moss-type coalgebraic logics. In both cases, coalgebraists generally have the Barr extension $\overline{T}$ in mind, which, for a functor T and a relation $R \subseteq X \times Y$, is the relation given by

$$\overline{T}R := \{(T\pi_X(\rho), T\pi_Y(\rho)) \in TX \times TY \mid \rho \in TR\} ,$$

where $\pi_X : R \to X$ and $\pi_Y : R \to Y$ are the two projections. This relation lifting characterizes a bisimulation between two coalgebras $\xi : X \to TX$ and $\upsilon : Y \to TY$ as a relation $R \subseteq X \times Y$ such that $(\xi(x), \upsilon(y)) \in \overline{T}R$ whenever $(x, y) \in R$. It is well-known, however, that these applications only work properly in case the functor T satisfies the category-theoretic property of preserving weak pullbacks. The key observation here is that $\overline{T}$ distributes over relation composition iff T preserves weak pullbacks. As an example, the above characterization of bisimilarity only coincides with that of behavioral equivalence (that is the relation of identifiability of two states by morphisms sharing their codomain) if T has this property. For this reason relation liftings are often thought to be of interest only in a setting of coalgebras for a weak pullback preserving functor.

On the other hand, the monotone neighborhood functor $\mathcal{M}$ is an important example of a coalgebra functor which does not preserve weak pullbacks, but which has a relation lifting $\widetilde{\mathcal{M}}$ that is essentially different from the Barr extension $\overline{\mathcal{M}}$ and whose notion of bisimilarity exactly captures behavioral equivalence [4].

And recently it has been shown that this notion of relation lifting can also be used to define the semantics of a Moss-style coalgebraic modality [15].

For this reason we study the notions of relation lifting that can be associated with a set functor T from a more general perspective. Here we take a relation lifting for a set functor T to be a collection of relations LR for every relation R, such that $LR \subseteq TX \times TY$ if $R \subseteq X \times Y$ (in the sequel we will give a more precise definition). Such studies have already been undertaken in the past. In [18] Thijs introduced a class of relation liftings, which he calls 'relators', to generalize different notions of coalgebraic simulation. Later, Baltag used Thijs' framework in [2] to give a semantics for the coalgebraic cover modality nabla. In [7] Hughes and Jacobs defined a generalization of the Barr extension for functors that carry an order. Very recently, Levy investigated the relation between the concept of similarity given by a relation liftings and final coalgebras [11].

In this paper, which forms part of the MSc thesis [12] authored by the first author and supervised by the second, we focus on the question when such a relation lifting captures behavioral equivalence, in the sense that L-bisimilarity (defined in the obvious way) coincides with behavioral equivalence for any pair of T-coalgebras. Our work concerns similar notions as Levy's paper [11]. The difference is that whereas Levy looks for endofunctors in some suitable order-category such that the notion of behavioral equivalence for its coalgebras (in his case: identification in the final coalgebra) coincides with similarity for a fixed relation lifting, we go the other way round and try to find relation liftings, whose notion of bisimilarity captures behavioral equivalence for a fixed functor.

Our main results can be summarized as follows. On the negative side, we prove that there is no way to capture behavioral equivalence between coalgebras for the (arbitrary) neighborhood functor $\mathcal{N}$ by means of relation lifting (Theorem 2). On the other hand, an important notion studied here is that of a lax extension of a functor T [17]. We will see that if such a lax extension preserves diagonals, then it captures behavioral equivalence indeed (Theorem 1) — this takes care of all cases known to us. Furthermore, we will provide some additional evidence that this combination of properties (lax extension preserving diagonals) is a natural one: in Theorem 3 we will prove that any finitary functor T has such an extension iff it admits a separating set of finitary monotone predicate liftings, a notion that is familiar from the theory of coalgebraic modal logic [13].

2 Preliminaries

In this paper we presuppose knowledge of the theory of coalgebras [14]. We recall some of the central definitions in this section, mainly to fix the notation.

2.1 Relations

In the following we consider relations to be arrows in the category of sets and relations. That is, we think of a relation $R : X \nrightarrow Y$ between sets X and Y as not just a subset of $X \times Y$ but as also specifying its codomain X and domain

Y. Nevertheless, we often write $R = R'$, $R \subseteq R'$, $R \cup R', R \cap R' : X \nrightarrow Y$ or $(x, y) \in R$ as if the relations $R, R' : X \nrightarrow Y$ were sets. We use $R^{gr} \subseteq X \times Y$ if we want to make explicit that we mean the set of pairs, considered as an object in the category of sets and functions, that stand in a relation $R : X \nrightarrow Y$.

We write $R \,;S : X \nrightarrow Z$ for the composition of two relations $R : X \nrightarrow Y$, $S : Y \nrightarrow Z$, and $R^\circ : Y \nrightarrow X$ for the converse of $R : X \nrightarrow Y$ with $(y, x) \in R^\circ$ iff $(x, y) \in R$. The graph of any function $f : X \to Y$ is a relation $f : X \nrightarrow Y$ between X and Y for which we also use the symbol f. It will be clear from the context in which a symbol f occurs whether it is meant as a arrow in the category of sets and functions or an arrow in the category of relations. The composition of relations is written the other way round than the composition of functions. So we have for functions $f : X \to Y$ and $g : Y \to Z$ that $g \circ f = f \,; g$.

Identity elements in the category of sets and relations 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 the identity function $\mathsf{id}_X : X \to X$ as a relation.

2.2 Set Functors

In the following we assume, if not explicitly stated otherwise, that functors are covariant endofunctors in the category of sets and functions. A functor T is *finitary* if it satisfies for all sets X:

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

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.

We now introduce some of the functors that we are concerned with in this paper. The *powerset functor* $\mathcal{P}$ maps a set X to the set of all its subsets $\mathcal{P}X$. A function $f : X \to Y$ is sent to $\mathcal{P}f : \mathcal{P}X \to \mathcal{P}Y, U \mapsto f[U]$. The *contravariant powerset functor* $\breve{\mathcal{P}}$ also 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 $f : X \to Y$ we have $\breve{\mathcal{P}}f : \breve{\mathcal{P}}Y \to \breve{\mathcal{P}}X, V \mapsto f^{-1}[V]$.

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 $\xi \in \mathcal{N}X = \breve{\mathcal{P}}\breve{\mathcal{P}}X$ we have

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

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

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

A restriction of the neighborhood functor $\mathcal{N}$ is the *monotone neighborhood functor* $\mathcal{M}$. It maps a set X to the collection $\mathcal{M}X$ of objects ξ in $\mathcal{N}X$ that are *upsets*, meaning 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 the same as $\mathcal{N}$. So we have for $f : X \to Y$ that

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

It is straightforward to check that this is well-defined. 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 $U \in \xi$.

The next two functors F^3_2 and $\mathcal{P}_n$ are interesting examples for us, because they, like the monotone neighborhood functor, do not preserve weak pullbacks but still allow for a relation lifting that captures behavioral equivalence.

The functor F^3_2 maps a set X to

$$F^3_2 X = \{(x_0, x_1, x_2) \in X^3 \mid |\{x_0, x_1, x_2\}| \leq 2\}$$

the set of all triples over X that consist of at most two distinct elements. On functions the functor F^3_2 is defined exactly as $(-)^3$, that is a function $f : X \to Y$ is mapped by F^3_2 such that $F^3_2 f(x_0, x_1, x_2) = (f(x_0), f(x_1), f(x_2))$.

The restricted powerset functor $\mathcal{P}_n$ for an $n \in \omega$ maps a set X to the set $\mathcal{P}_n X = \{U \subseteq X \mid |U| < n\}$ of all its subsets of cardinality smaller than n. On functions it has the same definitions as $\mathcal{P}$, that is $\mathcal{P}_n f(U) = f[U]$.

2.3 Coalgebras

A *T-coalgebra* for a covariant functor T 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 structure*. A *T-coalgebra morphism* from a T-coalgebra $\xi : X \to TX$ to a T-coalgebra $\zeta : Z \to TZ$ is a function $f : X \to Z$ such that $\zeta \circ f = Tf \circ \xi$.

The T-coalgebras together with the T-coalgebra morphisms are a category where the identity arrows, and composition of arrows is the same as for the underlying set functions. This category is cocomplete and all colimits are computed as for the underlying sets.

The central notion of equivalence between states in coalgebra is behavioral equivalence. Two states, x_0 in a T-coalgebra $\xi : X \to TX$ and y_0 in T-coalgebra $\upsilon : Y \to TY$, 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 $f(x_0) = g(y_0)$.

2.4 Predicate Liftings

A notion from coalgebraic modal logic that we are using later are predicate liftings. Predicate liftings for a functor T were originally introduced in [13], but see also [16], to define a modal logic for T-coalgebras that resembles the standard modal logic with boxes and diamonds on Kripke frames.

An n-ary *predicate lifting* for T is a natural transformation $\lambda : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T$. The *transposite* $\lambda^\flat : T \Rightarrow {}^n\mathcal{N} = \breve{\mathcal{P}}\breve{\mathcal{P}}^n$ of predicate lifting λ for a functor T is a natural transformation that is defined at a set X as

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

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 $U, U' \in (\breve{\mathcal{P}}X)^n$. The following observation is crucial for the proof of Theorem 3. The routine proof is left to the reader.

Proposition 1. *If $\lambda : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T$ is a monotone n-ary predicate lifting for T then the codomain of its transposite $\lambda^\flat : T \Rightarrow {}^n\mathcal{N}$ can be restricted to ${}^n\mathcal{M}$. That means $\lambda^\flat : T \Rightarrow {}^n\mathcal{M}$ defined as above is well-defined.*

To avoid tiresome compatibility issues when dealing with the transposites of multiple monotone predicate liftings of possibly different finite arity one can compose them with the componentwise injective natural transformation $e^n : {}^n\mathcal{M} \Rightarrow {}^\omega\mathcal{M}$ defined by

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

We just write $e \circ \lambda^\flat : T \Rightarrow {}^\omega\mathcal{M}$ for $e^n \circ \lambda^\flat$, where $n \in \omega$ is the arity of λ.

A family $\mathcal{F}$ of functions from X to Y is *jointly injective* if given 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 monotone 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 at every set X. 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.

2.5 Relation Liftings and Bisimilarity

Fix a covariant set functor T. A *relation lifting* L for T is a collection of relations $LR : TX \nrightarrow TY$ for every relation $R : X \nrightarrow Y$. Throughout this paper we shall require relation liftings to preserve converses, this means that $L(R^\circ) = (LR)^\circ$ for all relations R. This restriction simplifies the presentation and is not essential for our results because behavioral equivalence, the notion we want to capture with relation liftings, is symmetrical.

Given a relation lifting L for a set functor T and two T-coalgebras $\xi : X \to TX$ and $\upsilon : Y \to TY$, an *L-bisimulation between ξ and υ* is a relation $R : X \nrightarrow Y$ such that $(\xi(x), \upsilon(y)) \in LR$ for all $(x, y) \in R$. The relation $\leftrightarrows^L_{\xi,\upsilon} : X \nrightarrow Y$ of L-bisimilarity between ξ and υ is defined as the union of all L bisimulations between ξ and υ. We sometimes omit the subscripts and just write $x \leftrightarrows^L y$ if the coalgebras x and y belong to are clear from the context. We also write $\leftrightarrows^L_\xi = \leftrightarrows^L_{\xi,\xi} : X \nrightarrow X$ for bisimilarity on one single coalgebra $\xi : X \to TX$.

A relation lifting L for T *captures behavioral equivalence* if for any states x and y in T-coalgebras we have $x \leftrightarrows^L y$ iff x and y are behaviorally equivalent.

3 Lax Extensions

In this section we introduce lax extensions. These are relation liftings satisfying certain conditions that make them well-behaved in the context of coalgebra. We summarize some general properties of lax extensions and show that they capture behavioral equivalence if they preserve diagonals. For some additional discussion of lax extensions, although in a different context, we refer to [17]. For more about the general 2-categorical concept of a lax functor consult [8].

Definition 1. *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) $Tf \subseteq Lf$.

A lax extension L preserves diagonals *if it additionally satisfies:*

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

Condition (L3) in [17] additionally requires that $(Tf)^\circ \subseteq L(f^\circ)$. For us this follows automatically from the preservation of converses.

Only one inclusion is needed in (L4) for a lax extension to preserve diagonals. This is enough because, as shown in Proposition 2 below, condition (L4) implies together with condition (L3) that $L\Delta_X = \Delta_{TX}$.

Remark 1. In [7] a generalization of the Barr extension is defined with the name 'lax relation lifting'. This lax relation lifting is in general not a lax extension in our sense, even if we would not require preservation of converses, because it does not satisfy (L2). The lax relation lifting of [7] always satisfies $LR;LS \supseteq L(R;S)$ which is exactly the condition that distinguishes lax extension that preserve diagonals from the Barr extension and makes them useful for functors that do not preserve weak pullbacks.

Lax extensions have already been studied in the context of coalgebra under the name 'monotone relator' in [18, Section 2.1] and very recently in [11, Definition 6], where they are just called 'relators'. In [18] it is additionally required that composition of relation is preserved, that means = instead of $\subseteq$ in our condition (L2) of Definition 1, but it is noted that the $\supseteq$-inclusion can be omitted for most of the proofs. Both [18] and [11] use a different set of conditions in their definitions, but it can be checked that they are equivalent to our Definition 1. Instead of (L3) [18] requires that

(R3) $\Delta_{TX} \subseteq L\Delta_X$,
(R4) $Tf\,;LR\,;(Tg)^\circ \subseteq L(f\,;R\,;g^\circ)$.

In [11] condition (R4) has = instead of just $\subseteq$. This is superfluous, because we can show that (R3) and (R4) imply (L3). Hence every relator is a lax extension and the equality in (R4) follows from Proposition 2 (ii) below. To see that (R3) and (R4) imply (L3) consider for any function $f : X \to Z$

$$\begin{aligned} Tf = Tf \,;\, \Delta_{TZ} \,;\, (T\mathsf{id}_Z)^\circ &\subseteq Tf \,;\, L\Delta_Z \,;\, (T\mathsf{id}_Z)^\circ && \text{(R3)} \\ &\subseteq L(f \,;\, \Delta_Z \,;\, \mathsf{id}_Z^\circ) = Lf \;. && \text{(R4)} \end{aligned}$$

That every lax extension is a relator, that is every lax extension satisfies (R3) and (R4) follows from our next Proposition that summarizes some basic properties of lax extensions.

Proposition 2. *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) $\Delta_{TX} = L\Delta_X$ *and* $Tf = Lf$*.*
(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

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

The other claim $LR\,;(Tg)^\circ = L(R\,;g^\circ)$ follows from $Tf\,;LS = L(f\,;S)$ because L preserves converses.

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

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

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

For claim (iv) consider

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

Example 1. (i) For any functor T there is a trivial lax extension C that maps any relation $R : X \nrightarrow Y$ to the maximal relation $CR = TX \times TY : TX \nrightarrow TY$. For most functors this lax extension does not preserve diagonals.

(ii) The *Egli-Milner lifting* $\overline{\mathcal{P}}$ is a lax extension of the covariant powerset functor $\mathcal{P}$ that preserves diagonals. It 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).

More concisely we can write $\overline{\mathcal{P}}R = \overrightarrow{\mathcal{P}}R \cap \overleftarrow{\mathcal{P}}R$ where we use the abbreviations

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

(iii) The Egli-Milner lifting from item (ii) is an 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

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

It is easy to see that the Barr extension $\overline{T}$ of a functor T satisfies (L1). One 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 proofs of this basic properties of the Barr extension consult for instance [9].

Condition (L2) is more difficult. It is the case 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 [9, Fact 3.6]. So we have that the Barr extension $\overline{T}$ of a weak pullback preserving functor T is a lax extension that preserves diagonals.

(iv) Even though one can show that the Barr extension $\overline{\mathcal{M}}$ of the monotone neighborhood functor does not satisfy (L2), there is a lax extension $\widetilde{\mathcal{M}}$ of $\mathcal{M}$ that preserves diagonals. For the Definition recall the notation $\overrightarrow{\mathcal{P}}R$ and $\overleftarrow{\mathcal{P}}R$ from item (ii). The lax extension $\widetilde{\mathcal{M}}$ is defined on a relation $R : X \nrightarrow Y$ as

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

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

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

It is easy to check the conditions (L1) and (L2) for $\widetilde{\mathcal{M}}$. To check (L3) we show that $(\xi, \mathcal{M}f(\xi)) \in \widetilde{\mathcal{M}}f$ for all functions $f : X \to Y$ and $\xi \in \mathcal{M}X$. For

$(\xi, \mathcal{M}f(\xi)) \in \overrightarrow{\mathcal{P}}\overleftarrow{\mathcal{P}}f$ observe that $(U, f[U]) \in \overleftarrow{\mathcal{P}}f$ and $f[U] \in \mathcal{M}f(\xi)$ for any $U \in \xi$ because ξ is an upset. To get $(\xi, \mathcal{M}f(\xi)) \in \overleftarrow{\mathcal{P}}\overrightarrow{\mathcal{P}}f$ take any $V \in \mathcal{M}f(\xi)$. By the definition of $\mathcal{M}$ on morphisms this means that $f^{-1}[V] \in \xi$ and for this we have $(f^{-1}[V], V) \in \overrightarrow{\mathcal{P}}f$. 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 exists 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 $U \in \xi'$. Completely analogously one can verify that $\widetilde{{}^{\alpha}\mathcal{M}}$ is a lax extension of ${}^{\alpha}\mathcal{M}$ that preserves diagonals.

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

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

There is an easy counterexample to (L2) for the Barr extension $\overline{F_2^3}$ of F_2^3.

(vi) There is a lax extension $\widetilde{\mathcal{P}_n}$ of the restricted powerset functor $\mathcal{P}_n$ that preserves diagonals. It is defined in the same way as the Egli-Milner lifting $\overline{\mathcal{P}}$ of $\mathcal{P}$, that is $\widetilde{\mathcal{P}_n}R = \overrightarrow{\mathcal{P}}R \cap \overleftarrow{\mathcal{P}}R$ for any relation $R : X \nrightarrow Y$. Nevertheless, $\widetilde{\mathcal{P}_n}$ is distinct from the Barr extension $\overline{\mathcal{P}_n}$ of $\mathcal{P}_n$. As for $\overline{F_2^3}$ one can given a counterexample to (L2) for $\overline{\mathcal{P}_n}$ provided that $n > 3$.

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. Note also that our requirement that relation liftings preserve converses immediately implies that the converse of a bisimulation is a bisimulation. This facts are summarized in the following Proposition whose easy proof is left to the reader.

Proposition 3. *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 their composition $R\,;S : X \nrightarrow Y$ is an L-bisimulation between ξ and υ.
(iii) Every union of L-bisimulations between ξ and υ is again an L-bisimulation between ξ and υ.

Corollary 1. *Let L be a lax extension of T and $\xi : X \to TX$ and $\upsilon : Y \to TY$ be two T-coalgebras. The relation of L-bisimilarity $\leftrightarroweq^L_{\xi,\upsilon}$ between ξ and υ is itself an L-bisimulation between ξ and υ. Moreover L-bisimilarity $\leftrightarroweq^L_{\xi} : X \nrightarrow X$ on one single coalgebra ξ is an equivalence relation.*

We are now ready to prove that lax extensions that preserve diagonals capture behavioral equivalence. Note that in the proof the preservation of diagonals is only used for the application of Proposition 2 (iv) at the end of the direction from bisimilarity to behavioral equivalence.

Theorem 1. *If L is a lax extension of T that preserves diagonals then L captures behavioral equivalence.*

Proof. We have to show that 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.

For the direction from left to right assume that x_0 and y_0 are behaviorally equivalent. That means that there are a 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 (i) and (ii) the relation $f\,;g^\circ : X \nrightarrow Y$ is an L-bisimulation between ξ and υ because it is the composition of graphs of coalgebra morphisms. This implies that x_0 and y_0 are L-bisimilar because $(x_0, y_0) \in f\,;g^\circ$.

In the other direction we have to show that for any pair $(x, y) \in R$, for an L-bisimulation $R : X \nrightarrow Y$ between ξ and υ, the states x and y are behaviorally equivalent. Without loss of generality we can consider the case of two states z and z' in one single coalgebra $\zeta : Z \to TZ$ with an L-bisimulation $S : Z \nrightarrow Z$ on ζ such that $(z, z') \in S$. This is because otherwise we let ζ be the coproduct of ξ and υ with injections i_X and i_Y and then consider the relation $S = i_X^\circ\,; R\,; i_Y$ which, using Proposition 3, can be shown to be an L-bisimulation on ζ.

Now consider the relation $\leftrightarroweq^L_\zeta : Z \nrightarrow Z$ of L-bisimilarity on ζ which by Corollary 1 is both an equivalence relation and an L-bisimulation. Our goal is to put a transition structure $\delta : Z/{\leftrightarroweq^L_\zeta} \to T(Z/{\leftrightarroweq^L_\zeta})$ on the quotient $Z/{\leftrightarroweq^L_\zeta}$ such that the projection $p : Z \to Z/{\leftrightarroweq^L_\zeta}, z \mapsto [z]$ becomes a coalgebra morphism from ζ to δ. Since we assume that $z \leftrightarroweq^L_\zeta z'$ it follows then that $p(z) = p(z')$ which witnesses that z and z' are behaviorally equivalent.

We intend to define the transition function δ on $Z/{\leftrightarroweq^L_\zeta}$ such that

$$\delta([z]) \mapsto Tp \circ \zeta(z)\ .$$

This definition clearly satisfies $\delta \circ p = Tp \circ \zeta$ which means that p is a coalgebra morphism from ζ to δ as required. But we have to show that δ is well-defined. To prove this we need that $Tp \circ \xi(z) = Tp \circ \xi(z')$ for arbitrary $z, z' \in Z$ with $z \leftrightarroweq^L_\zeta z'$. Because $\leftrightarroweq^L_\zeta$ is an L-bisimulation it follows that $(\zeta(z), \zeta(z')) \in L{\leftrightarroweq^L_\zeta}$ and moreover

$$\begin{aligned} L{\leftrightarroweq^L_\zeta} &= L(p\,; p^\circ) && {\leftrightarroweq^L_\zeta} = p\,; p^\circ \\ &= Tp\,; (Tp)^\circ\ . && \text{Proposition 2 (iv)} \end{aligned}$$

So we get $(\zeta(z), \zeta(z')) \in Tp\,; (Tp)^\circ$ which entails $Tp \circ \zeta(z) = Tp \circ \zeta(z')$, as required.

4 (No) Bisimulations for Neighborhood Frames

Already the papers [5] and [6] examine different relation liftings for the neighborhood functor $\mathcal{N}$, and the notions of bisimilarity they give rise to. It is found that none of the proposed relation liftings captures behavioral equivalence. In this section we show that this is indeed not possible. Nevertheless, it should be mentioned that, for the simpler case of behavioral equivalence on one single coalgebra, already the Barr extension $\overline{\mathcal{N}}$ of the neighborhood functor captures behavioral equivalence [5, Proposition 4].

Theorem 2. *There is no relation lifting for the neighborhood functor $\mathcal{N}$ that captures 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

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

which is clearly impossible.

Now suppose for a contradiction that there is a relation lifting L for $\mathcal{N}$ that captures behavioral equivalence. Consider an example with the 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, one can verify, that 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. Because $f(x_1) = g(y_1)$ this shows that x_1 and y_1 are behaviorally equivalent. The situation is depicted in the figure:

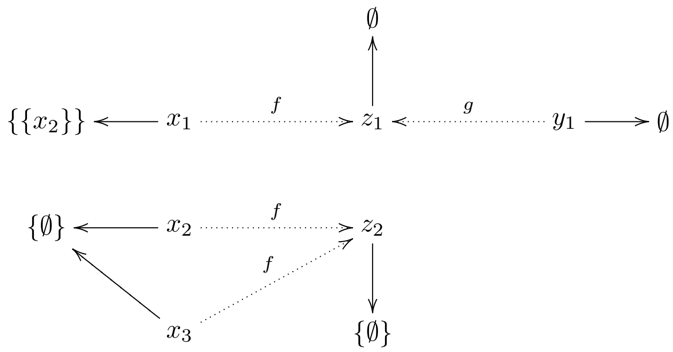

It follows from the assumption that L captures behavioral equivalence that there is an L-bisimulation $R : X \nrightarrow Y$ between ξ and υ such that $(x_1, y_1) \in R$.

Moreover we can show that $(x_2, y_1), (x_3, y_1) \notin R$. We do this only for (x_2, y_1) since the argument for (x_3, y_1) is similar. Suppose for a contradiction that x_2 and y_1 are L-bisimilar. Because L captures behavioral equivalence, it follows that there is a coalgebra $\zeta'' : Z'' \to \mathcal{N}Z''$ and coalgebra morphisms j from ξ to ζ'' and l from υ to ζ'' such that $j(x_2) = l(y_1)$. Using that j and l are coalgebra morphisms we get following contradiction to what we showed above:

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

So it follows that $R = \{(x_1, y_1)\}$ and because R is an L-bisimulation we find that $(\{\{x_2\}\}, \emptyset) = (\xi(x_1), \upsilon(y_1)) \in LR$.

Next we replace 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 linking x_1 in ξ' and y_1 in υ. Because L captures 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 means $h^{-1}[\{h(x_2)\}] \neq \{x_2\}$. Certainly $x_2 \in h^{-1}[\{h(x_2)\}]$ so it must be that $x_1 \in h^{-1}[\{h(x_2)\}]$ 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 it follows in both cases that $\mathcal{N}h(\{\emptyset\}) = \mathcal{N}k(\emptyset)$ which, as argued above, leads to a contradiction.

As a Corollary we obtain that the neighborhood functor has no lax extension that preserves diagonals, since we know from Theorem 1 that such a relation lifting would capture behavioral equivalence.

Corollary 2. *There is no lax extension that preserves diagonals for the neighborhood functor $\mathcal{N}$.*

5 Predicate Liftings and Lax Extensions

In the previous section we saw that the neighborhood functor does not have a lax extension that preserves diagonals. If we add the requirement that the neighborhoods are monotone, that is we look at the monotone neighborhood functor $\mathcal{M}$, then we have the lax extension $\widetilde{\mathcal{M}}$ that preserves diagonals. In this section we show that some sense of monotonicity is exactly what is needed from a functor in order to have a lax extension that preserves diagonals. Our goal is to prove following Theorem:

Theorem 3. *A finitary functor T has a lax extension that preserves diagonals iff there is a separating set of monotone predicate liftings with finite arity for T.*

Proof. This is the overview of the proof that brings together all the results from this Section.

For the direction from left to right assume that T has a lax extension L that preserves diagonals. Because T is finitary it has a finitary presentation (Σ, E) as demonstrated in Example 3. We use this together with the natural transformation $\lambda^L : T\breve{\mathcal{P}} \Rightarrow \breve{\mathcal{P}}T$ from Definition 3 to construct the Moss liftings for T defined as in Definition 5. In Proposition 6 we prove that the Moss liftings are monotone and in Proposition 7 that set of all Moss liftings is separating.

For the direction from right to left assume we have a separating set Λ of monotone predicate liftings with finite arity for T. By Proposition 1 the monotonicity of each $\lambda \in \Lambda$ entails that we can take $\lambda^\flat : T \to {}^n\mathcal{N}$ to have codomain ${}^n\mathcal{M}$. We can then apply the initial lift construction from Definition 2 to the set of natural transformations $\Gamma = \{e \circ \lambda^\flat : T \Rightarrow {}^\omega\mathcal{M}\}_{\lambda\in\Lambda}$, where $e : {}^n\mathcal{M} \Rightarrow {}^\omega\mathcal{M}$ is the embedding as defined in Section 2.4, and obtain a relation lifting $(\widetilde{{}^\omega\mathcal{M}})^\Gamma$ for the functor T. We show in Proposition 4 that the relation lifting $(\widetilde{{}^\omega\mathcal{M}})^\Gamma$ is a lax extension for T that preserves diagonals, since $\widetilde{{}^\omega\mathcal{M}}$ is a lax extension for ${}^\omega\mathcal{M}$ that preserves diagonals and the set of functions $\{e_X \circ \lambda^\flat_X : TX \Rightarrow {}^\omega\mathcal{M}X\}_{\lambda\in\Lambda}$ is jointly injective at every set X because Λ is assumed to be separating.

We now describe the two constructions, initial lift and Moss liftings, that are used in the proof of Theorem 3. The initial lift of a lax extension along a set of natural transformations is taken from [17]. In the proof of Theorem 3 we use it to build a lax extension for T from the lax extension $\widetilde{{}^\omega\mathcal{M}}$ and a separating set of predicate liftings.

Definition 2. *Let L be a relation lifting for T, and $\Lambda = \{\lambda : T' \Rightarrow T\}_{\lambda\in\Lambda}$ a set of natural transformations from another functor T' to T. Then we can define a relation lifting L^Λ for T called the* initial lift *of L along Λ as*

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

Equivalently to the above Definition, one can define $L^\Lambda R : T'X \nrightarrow T'Y$ for an $R : X \nrightarrow Y$ such that

$$(\xi, \upsilon) \in L^\Lambda R \quad \text{iff} \quad (\lambda_X(\xi), \lambda_Y(\upsilon)) \in LR \text{ for all } \lambda \in \Lambda\,.$$

Next we show that the initial lift construction preserves laxness and, which is essential for Theorem 3, it also preserves condition (L4), if the set of natural transformations is jointly injective for every set.

Proposition 4. *Let $\Lambda = \{\lambda : T' \Rightarrow T\}_{\lambda \in \Lambda}$ be a set of natural transformations from a functor T' to a 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 of T. Moreover, L^Λ preserves diagonals, if L preserves diagonals and $\{\lambda_X : T'X \to TX\}_{\lambda \in \Lambda}$ is jointly injective at every set X.*

Proof. It is routine to verify that all the conditions (L1), (L2) and (L3) are preserved by the initial lift construction. That the elements of Λ are natural transformations is only used for the preservation of (L3).

Here we give the proof for the claim 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 first show that if $\{\lambda_X : T'X \to TX\}_{\lambda \in \Lambda}$ is jointly injective at every set X then

$$\bigcap_{\lambda \in \Lambda} (\lambda_X \,;\, \lambda_X^\circ) = \Delta_{T'X} \,. \tag{1}$$

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

Now assume that L satisfies (L4) that is $L\Delta_X \subseteq \Delta_{TX}$ for every set X. It follows that $L^\Lambda \Delta_X \subseteq \Delta_{T'X}$ because

$$\begin{aligned} L^\Lambda \Delta_X &= \bigcap_{\lambda \in \Lambda} (\lambda_X \,;\, L\Delta_X \,;\, \lambda_X^\circ) && \text{definition} \\ &\subseteq \bigcap_{\lambda \in \Lambda} (\lambda_X \,;\, \Delta_{TX} \,;\, \lambda_X^\circ) && \text{assumption} \\ &= \bigcap_{\lambda \in \Lambda} (\lambda_X \,;\, \lambda_X^\circ) && \Delta_{TX} \text{ neutral element} \\ &= \Delta_{T'X} \,. && (1) \end{aligned}$$

This shows that L^Λ satisfies (L4).

Example 2. Consider the natural transformations $\Diamond, \Box : \mathcal{P} \Rightarrow \mathcal{M}$ with

$$\Diamond_X(U) = \{V \subseteq X \mid U \cap V \neq \emptyset\} \,, \qquad \Box_X(U) = \{V \subseteq X \mid U \subseteq V\} \,.$$

These natural transformation are clearly injective at every set X and hence it follows with Proposition 4 that $\widetilde{\mathcal{M}}^{\{\Diamond\}}$ and $\widetilde{\mathcal{M}}^{\{\Box\}}$ are lax extensions of the powerset functor $\mathcal{P}$ that preserve diagonals. Indeed, one can easily verify that they are both equal to the Barr extension $\overline{\mathcal{P}}$ of $\mathcal{P}$.

For the left-to-right direction of Theorem 3 we use the so called Moss liftings. It is shown in [10] 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 arbitrary lax extensions.

The first step in the construction of the Moss liftings is to use the lax extension L of T to define a distributive law between T and the contravariant powerset functor $\breve{\mathcal{P}}$.

Definition 3. *Given a lax extension L of a functor T we define for every set X the function*

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

where $\in_X\colon X \nrightarrow \mathcal{P}X$ denotes the membership relation between elements of X and subsets of X.

Proposition 5. *For a lax extension L the mapping $\lambda^L : T\breve{\mathcal{P}} \Rightarrow \breve{\mathcal{P}}T$ from Definition 3 is a natural transformation.*

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

$$\begin{array}{ccc} T\breve{\mathcal{P}}X & \xrightarrow{\lambda^L_X} & \breve{\mathcal{P}}TX \\ {\scriptstyle T\breve{\mathcal{P}}f}\uparrow & & \uparrow{\scriptstyle \breve{\mathcal{P}}Tf} \\ T\breve{\mathcal{P}}Y & \xrightarrow{\lambda^L_Y} & \breve{\mathcal{P}}TY \end{array} \qquad (2)$$

First observe that

$$L{\in_X} \, ; (T\breve{\mathcal{P}}f)^\circ = Tf \, ; L{\in_Y} \, . \qquad (3)$$

This is shown by the calculation

$$\begin{aligned} L{\in_X} \, ; (T\breve{\mathcal{P}}f)^\circ &= L\left(\in_X ; (\breve{\mathcal{P}}f)^\circ\right) && \text{Proposition 2 (ii)} \\ &= L(f \, ; \in_Y) && \text{direct verification} \\ &= Tf \, ; L{\in_Y} \, . && \text{Proposition 2 (ii)} \end{aligned}$$

To check the commutativity of (2) 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

$$\begin{aligned} \xi \in \lambda^L_X \circ T\breve{\mathcal{P}}f(\Upsilon) \quad &\text{iff} \quad (\xi, T\breve{\mathcal{P}}f(\Upsilon)) \in L{\in_X} && \text{definition of } \lambda^L \\ &\text{iff} \quad (\xi, \Upsilon) \in L{\in_X} \, ; (T\breve{\mathcal{P}}f)^\circ && \text{basic set theory} \\ &\text{iff} \quad (\xi, \Upsilon) \in Tf \, ; L{\in_Y} && (3) \\ &\text{iff} \quad (Tf(\xi), \Upsilon) \in L{\in_Y} && \text{basic set theory} \\ &\text{iff} \quad Tf(\xi) \in \lambda^L_Y(\Upsilon) && \text{definition of } \lambda^L \\ &\text{iff} \quad \xi \in \breve{\mathcal{P}}Tf \circ \lambda^L_Y(\Upsilon) \, . && \text{definition of } \breve{\mathcal{P}} \end{aligned}$$

Apart from the natural transformation $\lambda^L : T\breve{\mathcal{P}} \Rightarrow \breve{\mathcal{P}}T$ we need a finitary presentation of the functor T to define the Moss liftings. For more about presentations of set functors consult [1].

Definition 4. *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, 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. In order to define predicate liftings for an arbitrary finitary functor T it is necessary that we can somehow decompose it into pieces of the form X^n. This is exactly what the polynomial functor of a finitary presentation does.

The availability of a finitary presentation of T is the only part in the proof of Theorem 3 where we need that the functor T is finitary. If we had allowed for predicate liftings of infinite arity then we could generalize the construction of the Moss liftings to arbitrary accessible functors.

Example 3. 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

$$E_X : \coprod_{n\in\omega} Tn \times X^n \to TX\ ,$$

$$(\nu, u) \mapsto Tu(\nu)\ , \quad \text{where } \nu \in Tn \text{ and } u \in X^n \text{ for an } n \in \omega\ .$$

In this definition we take $u \in X^n$ to be a function $u : n \to X$. It is routine to check that this definition indeed provides a finitary presentation of T, meaning that E is a natural transformation and surjective at every set X.

The next Lemma shows how a lax extension of T interacts with a finitary presentation of T. This Lemma is similar to one direction of [10, 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 [10, Lemma 6.3] does not hold for lax extensions in general.

Lemma 1. *Let (Σ, E) be a finitary presentation of a functor T with lax extension L, and let $R : X \nrightarrow Y$ be any relation. Then we have for all $n \in \omega$, $r \in \Sigma_n$, $u \in X^n$ and $v \in Y^n$ 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

$\rho = (r, ((u_0, v_0), (u_1, v_1), \dots, (u_{n-1}, v_{n-1}))) \in \Sigma R^{gr}$. With the definition of Σ on morphisms it holds that $\Sigma\pi_X(\rho) = (r, u)$ and $\Sigma\pi_Y(\rho) = (r, v)$. Since E is a natural transformation from Σ to T we also 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 these 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

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

Which is what we had to show.

We can now define the Moss lifting for a finitary functor T by composing the finitary presentation of T with the natural transformation λ^L.

Definition 5. *Given a finitary functor T and a lax extension L for T take any finitary presentation (Σ, E) of T according to Definition 4 and let λ^L be the natural transformation of Definition 3. 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*

$$\mu^r : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T\ ,$$
$$\mu^r = \lambda^L \circ E_{\breve{\mathcal{P}}}(r, -)\ .$$

This definition yields the following diagram for every set X:

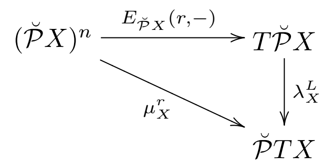

We use Lemma 1 to show that the Moss liftings are monotone.

Proposition 6. *The Moss liftings of a functor T with finitary presentation (Σ, E) and lax extension L are monotone.*

Proof. Take any Moss lifting $\mu^r = \lambda^L \circ E_{\breve{\mathcal{P}}}(r, -) : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T$ of an $r \in \Sigma_n$ 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 $\mu^r_X(U) \subseteq \mu^r_X(U')$.

So pick any $\xi \in \mu^r_X(U) = \lambda^L_X \circ E_{\breve{\mathcal{P}}X}(r, U)$. 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 $i \in n$ and Lemma 1 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 $\xi \in \lambda^L_X \circ E_{\breve{\mathcal{P}}X}(r, U') = \mu^r_X(U')$.

The last thing we have to show is that the set of all Moss liftings is separating. This is the only place in the construction of the Moss liftings where we actually need that the lax extension L preserves diagonals.

Proposition 7. *If L is a lax extension of a finitary functor T that preserves diagonals and let (Σ, E) be a finitary presentation of T. Then the set of all Moss liftings $M = \{\mu^r : \breve{\mathcal{P}}^n \Rightarrow \breve{\mathcal{P}}T \mid r \in \Sigma_n, n \in \omega\}$ is separating.*

Proof. To show that M is separating suppose for arbitrary $\xi, \xi' \in TX$ of any set X that $(\mu^r)^\flat_X(\xi) = (\mu^r)^\flat_X(\xi')$ for all $r \in \Sigma_n$ of all $n \in \omega$. We need to prove 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 full domain of $E_{\breve{\mathcal{P}}X} : \coprod_{n\in\omega}(\Sigma_n \times (\breve{\mathcal{P}}X)^n) \to T\breve{\mathcal{P}}X$, it follows that for all $\Xi \in T\mathcal{P}X$

$$(\xi, \Xi) \in L{\in_X} \quad \text{iff} \quad (\xi', \Xi) \in L{\in_X}. \tag{4}$$

To get $\xi = \xi'$ from (4) 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 (4) we get that $(\xi', Ts_X(\xi)) \in L{\in_X}$. Then we compute

$$\begin{aligned} (\xi, \xi') \in Ls_X \,;\, L{\ni_X} &\subseteq L(s_X \,;\, \ni_X) && \text{(L2)} \\ &= L\Delta_X && s_X \,;\, \ni_X = \Delta_X \\ &\subseteq \Delta_{TX} \ . && \text{(L4)} \end{aligned}$$

From this it follows that $\xi = \xi'$, which finishes the proof.

6 Conclusions and Open Questions

In this paper we showed that lax extension that preserve diagonals can be used in the theory of coalgebra to give a relational characterization of behavioral equivalence. This together with the fact that lax extensions can be used to define the semantics of an adequate cover modality indicates that lax extensions provide an adequate generalization for the role that the Barr extension of weak pullback preserving functors has played so far in the theory of coalgebras and coalgebraic modal logic. In this way the use of relation liftings in the theory of coalgebras can be extended to set functors that do not preserve weak pullbacks but nevertheless admit a lax extension that preserves diagonal relations.

The importance of lax extensions that preserves diagonals would motivate to study their properties on their own right. A pressing question, that we were unable to answer, concerns the uniqueness of such lax extensions. We do not know of an example of a functor with two distinct lax extension that both preserve diagonals. It would be interesting to find such an example or otherwise prove that any set functor has at most one lax extensions that preserve diagonals.

A negative result of this paper is that the neighborhood functor does not allow for a relational lifting that captures behavioral equivalence. This shows that there are limits to the use of relation liftings in the theory of coalgebras. A goal for further research would be to determine which functors have a relation lifting that captures behavioral equivalence. All the examples of such functors we know of also have a lax extension that preserves diagonals. So it might turn out, that whenever a functor allows for a relational characterization of behavioral equivalence it has a lax extension that preserves diagonals.

A further, probably easier, problem would be to characterize the functors that have a lax extension that preserves diagonals. Our Theorem 3 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. It might be interesting to look for a more elementary definition for the kind of monotonicity a functor needs to posses in order to allow for a separating set of monotone predicate liftings or, respectively, for a lax extension preserving diagonals. As a start one could look at the weak limit preservation properties, that are investigated in [3]. Moreover, it would be nice to have a canonical way to obtain a lax extension that preserves diagonals for the functors that posses one, similar to the definition of the Barr extension for weak pullback preserving functors.

We plan to write an other paper about the logic that results when one uses a lax extension to give a semantics to a Moss-style coalgebraic cover modality in the spirit of [2]. Some of this results can already be found in [12].

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. H. Peter Gumm and Tobias Schröder. Types and coalgebraic structure. *Algebra Universalis*, 53:229–252, 2005.

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. Jesse Hughes and Bart Jacobs. Simulations in coalgebra. In *THEOR. COMP. SCI.* Elsevier, 2003.
8. G. M. Kelly. Basic concepts of enriched category theory, 2005.
9. Clemens Kupke, Alexander Kurz, and Yde Venema. Completeness for the coalgebraic cover modality. accepted for publication.
10. Alexander Kurz and Raul Andres Leal. Equational coalgebraic logic. *Electronic Notes in Theoretical Computer Science*, pages 333–356, 2009.
11. Paul Blain Levy. Similarity quotients as final coalgebras. In *FOSSACS*, pages 27–41, 2011.
12. Johannes Marti. Relation liftings in coalgebraic modal logic. Master's thesis, ILLC, University of Amsterdam, 2011.
13. Dirk Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309(1–3):177–193, 2003.
14. J. J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theoretical Computer Science*, 249(1):3 – 80, 2000.
15. 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.
16. 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.
17. Christoph Schubert and Gavin J. Seal. Extensions in the theory of lax algebras. *Theories and Applications of Categories*, 21(7):118–151, 2008.
18. Albert Thijs. *Simulation and Fixpoint Semantics.* PhD thesis, University of Groningen, 1996.