Logic-Induced Bisimulations

Jim de Groot

The Australian National University

Helle Hvid Hansen

Delft University of Technology

Alexander Kurz

Chapman University

Abstract

We define a new logic-induced notion of bisimulation (called ρ-bisimulation) for coalgebraic modal logics given by a logical connection, and investigate its properties. We show that it is structural in the sense that it is defined only in terms of the coalgebra structure and the one-step modal semantics and, moreover, can be characterised by a form of relation lifting. Furthermore we compare ρ-bisimulations to several well-known equivalence notions, and we prove that the collection of bisimulations between two models often forms a complete lattice. The main technical result is a Hennessy-Milner type theorem which states that, under certain conditions, logical equivalence implies ρ-bisimilarity. In particular, the latter does *not* rely on a duality between functors T (the type of the coalgebras) and L (which gives the logic), nor on properties of the logical connection ρ.

Keywords: Modal Logic, Coalgebraic Logic, Bisimulation.

1 Introduction

In this paper, we investigate when logical equivalence for a given modal language can be captured by a structural semantic equivalence notion, understood as a form of bisimulation. Our investigation is carried out in the setting of coalgebraic modal logic [21], where semantic structures are given by coalgebras for a functor $\mathsf{T}\colon \mathbf{C} \to \mathbf{C}$ [27]. This allows for a uniform treatment of a wide variety of modal logics [21,25,28]. Coalgebras come with general notions of *behavioural equivalence* and *bisimilarity*, and a logic is said to be *expressive* if logical equivalence implies behavioural equivalence, in which case we have a generalisation of the classic Hennessy-Milner theorem [15].

For **Set**-coalgebras, i.e., when $\mathbf{C} = \mathbf{Set}$, it has been shown that a coalgebraic modal logic is expressive if the language has sufficiently large conjunctions and the set Λ of modalities is *separating*, meaning that they separate points in $\mathsf{T}X$

[24,26,28]. In the more abstract setting of coalgebraic modal logic, where a logic is given by a functor and its semantics by a natural transformation ρ [9,19], a sufficient condition for a logic being expressive is that the so-called mate of ρ is monic [19, Thm. 4.2].

In this line of research, modal logics are often viewed as specification languages for coalgebras. Therefore behavioural equivalence is a given, and the aim is to find expressive logics. However, sometimes the modal language is of primary interest [5] and the relevant modalities need not be separating, see e.g. [12]. This leads us to consider the following question:

Given a possibly non-expressive coalgebraic modal logic, can we characterise logical equivalence by a notion of bisimulation?

Such investigations have been carried out earlier in [4] where a notion of Λ-bisimulation was proposed for **Set**-coalgebras and coalgebraic modal logics with a classical propositional base.

Here we generalise and extend the work of [4] beyond **Set** using the formulation of coalgebraic modal logic via dual adjunctions [9,19,23]. Examples include coalgebras over ordered and topological spaces and modal logics on different propositional bases. After recalling basic definitions of coalgebraic modal logic in Section 2, we define the concept of a ρ-*bisimulation* in Section 3. For **Set**-coalgebras, this is a relation B between coalgebras for which the so-called B-coherent pairs [14,5] give rise to a congruence between complex algebras.

The definition of ρ-bisimulation is structural in the sense that it is defined in terms of the coalgebra structure and the one-step modal semantics ρ. Moreover, it can often be characterised as a greatest fixpoint via relation lifting. We also prove results concerning truth-preservation, composition and lattice structure, and we show that ρ-bisimilarity, like T-bisimilarity for coalgebras. For coalgebras on finite sets, this means that ρ-bisimilarity can be computed by a partition refinement algorithm.

The main technical results are found in Section 4 and concern the distinguishing power of ρ-bisimulations. We first compare ρ-bisimulations with other coalgebraic equivalence notions. Subsequently, we prove a Hennessy-Milner style theorem (Thm. 4.4) in which we give conditions that guarantee that logical equivalence is a ρ-bisimulation. We emphasise that the logic is *not* assumed to be expressive and ρ-bisimilarity will generally differ from bisimilarity for T-coalgebras. Finally, we define a notion of translation between logics and show that if the language of ρ' is a propositional extension of the language of ρ, then ρ-bisimulations are also ρ'-bisimulations (Prop. 4.10).

By instantiating Prop. 4.10, we obtain that for labelled transition systems the ρ-bisimilarity notions for Hennessy-Milner logic [15] and trace logic [19] coincide and are equal to the standard notion of bisimilarity even without assuming image-finiteness. These two logics have the same modalities, which are separating, but trace logic has $\top$ as the only propositional connective.

Due to lack of space, some proofs are omitted here. They will be made available in an extended version of this paper online.

2 Coalgebraic modal logic

We review some background on coalgebraic logic, categorical algebra, and Stone duality. For more see e.g. [27,21,2,3,17]. We write **Set** for the category of sets and functions.

Coalgebraic modal logic generalises modal logic from Kripke frames to coalgebras for a functor T.

Coalgebras can be understood as generalised, state-based systems defined parametrically in the system type T. Formally, we require T to be an endofunctor on a category $\mathbf{C}$. A T-*coalgebra* is then a pair (X, γ) such that $\gamma : X \to \mathsf{T}X$ is a morphism in $\mathbf{C}$. The object X is the state space, and the arrow γ is the coalgebra structure map. A T-*coalgebra morphism* from (X, γ) to (X', γ') is a $\mathbf{C}$-morphism $f : X \to X'$ satisfying $\gamma' \circ f = \mathsf{T}f \circ \gamma$. Together, T-coalgebras and T-coalgebra morphisms form a category which we write as $\mathbf{Coalg}(\mathsf{T})$.

An **algebra** for a functor is the dual notion of a coalgebra. Given an endofunctor $\mathsf{L}\colon \mathbf{A} \to \mathbf{A}$, an L-*algebra* is a pair (A, α) such that $\alpha : \mathsf{L}A \to A$ is a morphism in $\mathbf{A}$. An L-*algebra morphism* from (A, α) to (A', α') is an $\mathbf{A}$-morphism $h : A \to A'$ such that $h \circ \alpha = \alpha' \circ \mathsf{L}h$. We write $\mathbf{Alg}(\mathsf{L})$ for the category of L-algebras and L-algebra morphisms.

Example 2.1 A Kripke frame $(W, R \subseteq X \times X)$ is a coalgebra for the covariant powerset functor $\mathcal{P}\colon \mathbf{Set} \to \mathbf{Set}$ which maps a set to its set of subsets, and a function $f\colon X \to Y$ to the direct image map $f[-]\colon \mathcal{P}X \to \mathcal{P}Y$, by defining $\gamma\colon X \to \mathcal{P}X$ as $\gamma(x) = R[x] = \{y \in X \mid xRy\}$. Similarly, a Kripke model (W, R, V) where V is a valuation of a set P_0 of atomic propositions is a coalgebra for the $\mathbf{Set}$-functor $\mathcal{P}(-) \times \mathcal{P}(P_0)$ (which is constant in its second component) by taking $\gamma(x) = (R[x], V'(x))$ where $V'(x) = \{p \in P_0 \mid x \in V(p)\}$. It can be verified that the ensuing notion of coalgebra morphism coincides with the usual notion of bounded morphism for Kripke frames and Kripke models, respectively.

Example 2.2 Labelled transition systems (LTSs) are coalgebras for the **Set**-functor $\mathsf{T} = \mathcal{P}(-)^A$ where $\mathcal{P}$ is the covariant powerset functor and A is the set of labels. A coalgebra $\gamma\colon X \to \mathcal{P}(X)^A$ specifies for each state $x \in X$ and label $a \in A$, the set $\gamma(x)(a)$ of a-successors of x. In other words, an LTS is an A-indexed multi-relational Kripke frame, and one verifies that coalgebra morphisms are A-indexed bounded morphisms.

Logical connections To investigate logics for T-coalgebras in this generality, we use the Stone duality approach to modal logic [13,1], but rather than a full duality, here one requires only a dual adjunction $\mathsf{P} : \mathbf{C} \rightleftarrows \mathbf{A} : \mathsf{S}$ (sometimes called a *logical connection*) between a category $\mathbf{C}$ of state spaces and a category $\mathbf{A}$ of algebras that encode a propositional base logic. We emphasise that the functors P and S are contravariant. The classic example is then the instance $\mathcal{Q}_{\mathsf{BA}} : \mathbf{Set} \rightleftarrows \mathbf{BA} : \mathsf{Uf}$ where $\mathcal{Q}_{\mathsf{BA}}$ maps a set to its Boolean algebra of predicates (i.e. subsets), and Uf maps a Boolean algebra to its set of ultrafilters.

We denote the units of a dual adjunction $\mathsf{P} : \mathbf{C} \rightleftarrows \mathbf{A} : \mathsf{S}$ by $\eta^{\mathbf{C}}\colon \mathrm{Id}_{\mathbf{C}} \to \mathsf{SP}$

and $\eta^{\mathbf{A}}\colon \mathrm{Id}_{\mathbf{A}} \to \mathsf{PS}$, and the bijection of Hom-sets $\mathbf{C}(C, \mathsf{S}A) \cong \mathbf{A}(A, \mathsf{P}C)$ in both directions by $(-)^\sharp$. Recall that for $f\colon C \to \mathsf{S}A$, the adjoint transpose of f is $f^\sharp = \mathsf{P}f \circ \eta^{\mathbf{A}}_A$, and for $g\colon A \to \mathsf{P}C$, the adjoint is $g^\sharp = \mathsf{S}g \circ \eta^{\mathbf{C}}_C$.

Coalgebraic Modal Logic Given a dual adjunction $\mathsf{P} : \mathbf{C} \rightleftarrows \mathbf{A} : \mathsf{S}$ and an endofunctor T on $\mathbf{C}$, a *modal logic for* T*-coalgebras* is a pair (L, ρ) consisting of an endofunctor $\mathsf{L}\colon \mathbf{A} \to \mathbf{A}$ (defining modalities) and a natural transformation $\rho : \mathsf{LP} \to \mathsf{PT}$, (defining the *one-step modal semantics*). This data gives rise to a functor $\mathbf{Coalg}(\mathsf{T}) \to \mathbf{Alg}(\mathsf{L})$ which sends a coalgebra (X, γ) to its *complex algebra* $(\mathsf{P}X, \gamma^*)$, where $\gamma^* = \mathsf{P}\gamma \circ \rho_X$. Assuming that $\mathbf{Alg}(\mathsf{L})$ has an initial algebra $\alpha : \mathsf{L}\Phi \to \Phi$, which generalises the Lindenbaum-Tarski algebra, the semantics of (equivalence classes of) formulas is obtained as the unique $\mathbf{Alg}(\mathsf{L})$-morphism $[\![-]\!]_\gamma\colon (\Phi, \alpha) \to (\mathsf{P}X, \gamma^*)$. Viewing the semantics as an $\mathbf{A}$-morphism $[\![-]\!]_\gamma\colon \Phi \to \mathsf{P}X$, its adjoint $\mathrm{th}_\gamma = [\![-]\!]^\sharp_\gamma\colon X \to \mathsf{S}\Phi$, is called the *theory map*, since in the classic case it maps a state in X to the ultrafilter of L-formulas it satisfies. By their definitions, the semantics and the theory map make the following diagrams commute:

$$\begin{tikzcd}
\mathsf{L}\Phi \arrow[rr, "\alpha"] \arrow[d, "\mathsf{L}[\![\cdot]\!]_\gamma"'] & & \Phi \arrow[d, "{[\![\cdot]\!]_\gamma}"] \\
\mathsf{LP}X \arrow[r, "\rho_X"] & \mathsf{PT}X \arrow[r, "\mathsf{P}\gamma"] & \mathsf{P}X
\end{tikzcd}
\qquad
\begin{tikzcd}
X \arrow[rr, "\mathrm{th}_\gamma"] \arrow[d, "\gamma"] & & \mathsf{S}\Phi \arrow[d, "\mathsf{S}\alpha"] \\
\mathsf{T}X \arrow[r, "\mathsf{T}\,\mathrm{th}_\gamma"] & \mathsf{TS}\Phi \arrow[r, "\rho^\flat_\Phi"] & \mathsf{SL}\Phi
\end{tikzcd}$$

where $\rho^\flat\colon \mathsf{TS} \to \mathsf{SL}$ is the natural transformation, the so-called mate of ρ, obtained (component-wise) as the adjoint of $\rho\mathsf{S} \circ \mathsf{L}\eta^{\mathbf{A}}$.

Example 2.3 Consider the self-dual adjunction $\mathcal{Q} : \mathbf{Set} \rightleftarrows \mathbf{Set} : \mathcal{Q}$ given in both directions by the contravariant powerset functor $\mathcal{Q}$, which maps a set to its powerset 2^X, and a function $f\colon X \to Y$ to its inverse image map $f^{-1}\colon 2^Y \to 2^X$. In this case, the adjoints are given by transposing. For $f\colon X \to 2^Y$, $f^\sharp\colon Y \to 2^X$ is defined by $f^\sharp(y)(x) = f(x)(y)$.

Considering LTSs as $\mathcal{P}(-)^A$-coalgebras over $\mathbf{Set}$ (cf. Exm. 2.2), we obtain *trace logic for LTSs* [19, Exm. 3.2] by taking $\mathsf{L}^{\mathsf{tr}}\colon \mathbf{Set} \to \mathbf{Set}$ to be the functor $\mathsf{L}^{\mathsf{tr}} = 1 + A \times (-)$ (where $1 = \{*\}$) which encodes a modal signature with a constant modality $\top$ and a unary modality for each $a \in A$. Since $\mathbf{A} = \mathbf{Set}$, trace logic has no other connectives. The initial L^{tr}-algebra consists of finite sequences over A with the empty word as constant, and prefixing with elements from A as the unary operations. That is, L^{tr}-formulas are of the form $\langle a_1\rangle \cdots \langle a_k\rangle\top$, $k \geq 0$.

We obtain the usual semantics of $\top$ and A-labelled diamonds by defining the modal semantics $\rho^{\mathsf{tr}}\colon 1 + A \times \mathcal{Q}(-) \to \mathcal{Q}(\mathcal{P}(-)^A)$ as $\rho^{\mathsf{tr}}_X(*) = \mathcal{P}(X)^A$ and $\rho^{\mathsf{tr}}_X(a, U) = \{t \in \mathcal{P}(X)^A \mid t(a) \cap U \neq \emptyset\}$. Hence for an LTS (X, γ), $[\![\langle a_1\rangle \cdots \langle a_k\rangle\top]\!]_\gamma$ is the subset of X consisting of states x that can execute the trace $a_1 \cdots a_k$.

Example 2.4 Consider again LTSs as $\mathcal{P}(-)^A$-coalgebras over $\mathbf{Set}$ (cf. Exm. 2.2), but take now the classic adjunction $\mathcal{Q}_{\mathsf{BA}} : \mathbf{Set} \rightleftarrows \mathbf{BA} : \mathsf{Uf}$. Hennessy-Milner logic [15] (or equivalently, normal multi-modal logic) is here defined as classical propositional logic extended with join-preserving diamonds. This is achieved by defining $\mathsf{L}^{\mathsf{hm}}\colon \mathbf{BA} \to \mathbf{BA}$ as follows. For a Boolean algebra B, $\mathsf{L}^{\mathsf{hm}}B$ is the free Boolean algebra generated by the set $\{\langle a\rangle b \mid b \in B, a \in A\}$ modulo the congruence generated by the usual diamond equations: $\langle a\rangle\bot = \bot$ and $\langle a\rangle(\varphi_1 \vee \varphi_2) = \langle a\rangle\varphi_1 \vee \langle a\rangle\varphi_2$ for all $a \in A$. The modal semantics $\rho^{\mathsf{hm}}\colon \mathsf{L}^{\mathsf{hm}}\mathcal{Q}_{\mathsf{BA}} \to \mathcal{Q}_{\mathsf{BA}}(\mathcal{P}(-)^A)$ is essentially the Boolean extension of ρ^{tr}. In particular, $\rho^{\mathsf{hm}}_X(\langle a\rangle U) = \{t \in \mathcal{P}(X)^A \mid t(a) \cap U \neq \emptyset\}$.

The above description of Hennessy-Milner logic is a special case of a more general approach described in the next example.

Example 2.5 If $\mathbf{A}$ in the dual adjunction is a variety of algebras, we can define a logic (L, ρ) for $\mathsf{T} : \mathbf{C} \to \mathbf{C}$ by *predicate liftings and axioms* as in [20, Def. 4.2] and [23, Thms 4.7 and 8.8]. An *n-ary predicate lifting* is a natural transformation $\lambda : \mathsf{UP}^n \to \mathsf{UPT}$, where $\mathsf{P}^n X$ is the n-fold product of $\mathsf{P}X$ in $\mathbf{A}$ and $\mathsf{U} : \mathbf{A} \to \mathbf{Set}$ is the forgetful functor. Together with a suitable notion of *axioms*, a collection Λ of such predicate liftings yields a functor $\mathsf{L} : \mathbf{A} \to \mathbf{A}$ sending $A \in \mathbf{A}$ to the free algebra generated by $\{\underline{\lambda}(a_1, \dots, a_n) \mid \lambda \in \Lambda, a_i \in A\}$ modulo (instantiations of) the axioms. Define $\rho : \mathsf{LP} \to \mathsf{PT}$ on generators by $\rho_X(\underline{\lambda}(a_1, \dots, a_n)) = \lambda_X(a_1, \dots, a_n) \in \mathsf{PT}X$. If ρ is well-defined then (L, ρ) is a logic for $\mathbf{Coalg}(\mathsf{T})$. All logics in e.g. [4,7,21] are instances hereof.

Next, we interpret positive modal logic [11,10], whose coalgebraic semantics over posets is given in [18, Example 2.4], in topological spaces:

Example 2.6 Consider the dual adjunction $\Omega : \mathbf{Top} \rightleftarrows \mathbf{DL} : \mathsf{pf}$, where Ω takes open subsets of a topological space, viewed as a distributive lattice, and pf takes prime filters of a distributive lattice topologised by the subbase $\{\widetilde{a} \mid a \in A\}$, where $\widetilde{a} = \{p \in \mathsf{pf}A \mid a \in p\}$. The Vietoris functor $\mathsf{V} : \mathbf{Top} \to \mathbf{Top}$ takes $X \in \mathbf{Top}$ to its collection of compact subsets topologised by the subbase consisting of $⧈a = \{c \in \mathsf{V}X \mid c \subseteq a\}$ and $\diamondsuit a = \{c \in \mathsf{V}X \mid c \cap a \neq \emptyset\}$, where a ranges over the opens of X. For a continuous map $f : X \to X'$ the map $\mathsf{V}f$ takes direct images. The logic functor $\mathsf{N} : \mathbf{DL} \to \mathbf{DL}$ is defined as in [18, Exm. 2.4]. The natural transformation $\rho : \mathsf{N}\Omega \to \Omega\mathsf{V}$ given on generators by $\Box a \mapsto ⧈a$ and $\Diamond a \mapsto \diamondsuit a$ then gives rise to semantics for positive modal logic.

We now recall *linear weighted automata*, see e.g. [8, Section 3.2].

Example 2.7 Let $\Bbbk$ be a field and $\mathbf{Vec}_\Bbbk \rightleftarrows \mathbf{Vec}_\Bbbk$ the dual adjunction between vector spaces over $\Bbbk$ given in both directions by taking dual space via the contravariant functor $(-)^\vee = \mathrm{Hom}(-, \Bbbk) : \mathbf{Vec}_\Bbbk \to \mathbf{Vec}_\Bbbk$. Linear weighted automata for a set A of labels are coalgebras for the endofunctor $\mathsf{W} = \Bbbk \times (-)^A$ on $\mathbf{Vec}_\Bbbk$, where $(-)^A$ is simply the collection of maps $A \to (-)$ with a pointwise vector space structure. We work with the language given by the grammar

$$\varphi ::= 0 \mid p \mid r \cdot \varphi \mid \varphi + \varphi \mid \langle a\rangle\varphi,$$

where $a \in A$, $r \in \Bbbk$, and p is a single proposition letter (the termination predicate). Note that, contrary to *loc. cit.*, we also include connectives corresponding to the signature of vector spaces (because we will use vector spaces as algebraic semantics). The interpretation of a formula φ in this (many-valued) setting is a linear map $[\![\varphi]\!] : X \to \Bbbk$. The connectives 0, + and r are interpreted via the corresponding operations in vector spaces. For p, we use the nullary predicate lifting $\lambda^p \in \mathsf{U}(\mathsf{W}-)^\vee$, where $\mathsf{U} : \mathbf{Vec}_\Bbbk \to \mathbf{Set}$ is the forgetful functor, given by $\lambda^p_X : \mathsf{W}X \to \Bbbk : (r,t) \mapsto r$. That is $[\![p]\!]_\gamma = \lambda^p_X \circ \gamma : X \to \Bbbk$. As for the diamonds, we use $\lambda^{\langle a \rangle} : \mathsf{U}(-)^\vee \to \mathsf{U}(\mathsf{W}-)^\vee$ defined by

$$\lambda^{\langle a \rangle}_X(m) : \mathsf{W}X \to \Bbbk : (r,t) \mapsto m(t(a)).$$

Concretely, this means that if $[\![p]\!]_\gamma(y) = r \in \Bbbk$ and there is an a-transition $x \xrightarrow{a} y$, then $[\![\langle a \rangle p]\!](x) = r$. Together with the axioms $\langle a \rangle(\varphi + \psi) = \langle a \rangle\varphi + \langle a \rangle\psi$ and $r \cdot \langle a \rangle\varphi = \langle a \rangle(r \cdot \varphi)$ this gives rise to an endofunctor $\mathsf{L} : \mathbf{Vec}_\Bbbk \to \mathbf{Vec}_\Bbbk$, and a logic (L, ρ) for linear weighted automata. One can show that logical equivalence coincides with language semantics if the state-space is finite-dimensional.

Relations as jointly mono spans We are interested in giving certain relations a special status. In **Set**, a binary relation $B \subseteq X \times X$ corresponds to an injective map $B \hookrightarrow X \times X$. This generalises to an arbitrary category (possibly lacking products) via the following notion. A span $X_1 \xleftarrow{\pi_1} B \xrightarrow{\pi_2} X_2$ in a category $\mathbf{C}$ is called *jointly mono* if for all $\mathbf{C}$-arrows h, h' with codomain B it satisfies: if $\pi_1 \circ h = \pi_1 \circ h'$ and $\pi_2 \circ h = \pi_2 \circ h'$ then $h = h'$. We sometimes write the above span as (B, π_1, π_2), leaving codomains implicit. If $\mathbf{C}$ has products, then (B, π_1, π_2) is a jointly mono span if and only if the pairing $\langle \pi_1, \pi_2 \rangle : B \to X_1 \times X_2$ is monic.

The collection of jointly mono spans between two objects $X_1, X_2 \in \mathbf{C}$ can be ordered as follows: $(B, \pi_1, \pi_2) \leq (B', \pi'_1, \pi'_2)$ if there exists a (necessarily monic) map $k : B \to B'$ such that $\pi_i = \pi'_i \circ k$. If $(B, \pi_1, \pi_2) \leq (B', \pi'_1, \pi'_2)$ *and* $(B', \pi'_1, \pi'_2) \leq (B, \pi_1, \pi_2)$, then the two spans must be isomorphic. We write $\mathbf{Rel}(X_1, X_2)$ for the poset of jointly mono spans between X_1 and X_2 up to isomorphism.

Image factorisations and regular epis We will also need a generalisation of image factorisation. A category $\mathbf{C}$ is said to have $(\mathcal{E}, \mathcal{M})$-*factorisations* for some classes $\mathcal{E}$ and $\mathcal{M}$ of $\mathbf{C}$-morphisms, if every morphism $f \in \mathbf{C}$ factorises as $f = m \circ e$ with $e \in \mathcal{E}$ and $m \in \mathcal{M}$. We say that $\mathbf{C}$ has an $(\mathcal{E}, \mathcal{M})$-*factorisation system* [2, Def. 14.1] if moreover both $\mathcal{E}$ and $\mathcal{M}$ are closed under composition, and whenever $g \circ e = m \circ f$, with $e \in \mathcal{E}$ and $m \in \mathcal{M}$, there exists a unique diagonal fill-in d such that $f = d \circ e$ and $g = m \circ d$.

An epi e is *regular* if it is a coequalizer. In a variety, the regular epis are precisely the surjective morphisms. $\mathbf{Set}, \mathbf{Pos}, \mathbf{Top}, \mathbf{Vec}, \mathbf{Stone}$ all have a $(\mathcal{R}eg\mathcal{E}pi, \mathcal{M}ono)$-factorisation system.

3 Logic-induced bisimulations

We are now ready to define our logic-induced notion of bisimulation. Throughout this section, we assume we are given a dual adjunction $\mathsf{P} : \mathbf{C} \rightleftarrows \mathbf{A} : \mathsf{S}$, an endofunctor T on $\mathbf{C}$, and a logic (L, ρ) for T-coalgebras. Moreover, we assume that **C** has pullbacks and, in addition, that **A** has pullbacks or **C** has pushouts. Both conditions hold in our examples. In particular, if **A** is variety of algebras then pullbacks exist and are computed as in **Set**.

3.1 Definition and first examples

The basic ingredient for the definition of ρ-bisimulation is the notion of a *dual span*: A jointly mono span $X_1 \xleftarrow{\pi_1} B \xrightarrow{\pi_2} X_2$ in **C** is mapped by P to a cospan $\mathsf{P}X_1 \xrightarrow{\mathsf{P}\pi_1} \mathsf{P}B \xleftarrow{\mathsf{P}\pi_2} \mathsf{P}X_2$ in **A**. Taking its pullback we obtain a jointly mono span in **A**, which we denote by $(\overline{B}, \overline{\pi}_1, \overline{\pi}_2)$ and refer to as the *dual span* of (B, π_1, π_2). In case **C** has pushouts, dual spans exist because dual adjoints send pushouts to pullbacks. In the classic case where $\mathsf{P} = \mathcal{Q}_{\mathsf{BA}} \colon \mathbf{Set} \to \mathbf{BA}$ maps a set to its Boolean algebra of subsets, the dual span $(\overline{B}, \overline{\pi}_1, \overline{\pi}_2)$ consists of *B-coherent pairs (of subsets of X_1 and X_2)*, used in the definition of Λ-bisimulation [4], neighbourhood bisimulation [14], and conditional bisimulation [5]. We proceed to the definition of a ρ-bisimulation.

Definition 3.1 Let $\gamma_1 \colon X_1 \to \mathsf{T}X_1$ and $\gamma_2 \colon X_2 \to \mathsf{T}X_2$ be T-coalgebras. A jointly mono span $X_1 \xleftarrow{\pi_1} B \xrightarrow{\pi_2} X_2$ is a *ρ-bisimulation between γ_1 and γ_2* if

$$\mathsf{P}\pi_1 \circ \gamma_1^* \circ \mathsf{L}\overline{\pi}_1 = \mathsf{P}\pi_2 \circ \gamma_2^* \circ \mathsf{L}\overline{\pi}_2, \tag{1}$$

Definition 3.1 is structural in the sense that it is defined in terms of the coalgebra structure and the one-step modal semantics ρ (via the complex algebras γ_i^*). In particular, it does not refer to the set of all formulas nor to the initial L-algebra. Equation (1) provides a coherence condition that can be checked in concrete settings. We provide examples below. First, we give a more conceptual characterisation in terms of dual spans.

Proposition 3.2 *A jointly mono span (B, π_1, π_2) is a ρ-bisimulation between γ_1 and γ_2 iff its dual span $(\overline{B}, \overline{\pi}_1, \overline{\pi}_2)$ is a congruence between γ_1^* and γ_2^*.*

Proof. ($\Rightarrow$) Equation (1) says that the outer shell of the diagram on the right commutes. The universal property of the pullback $\overline{B}$ then yields a morphism $\beta : \mathsf{L}\overline{B} \to \overline{B}$ such that all squares in (2) commute. ($\Leftarrow$) The existence of such a β implies commutativity of (the outer shell of) the diagram. $\square$

$$\begin{tikzcd}
& \mathsf{L}\overline{B} \arrow[dl, "\mathsf{L}\overline{\pi}_1"'] \arrow[dr, "\mathsf{L}\overline{\pi}_2"] \arrow[dd, dashed, "\beta"] & \\
\mathsf{LP}X_1 \arrow[dd, "\gamma_1^*"'] & & \mathsf{LP}X_2 \arrow[dd, "\gamma_2^*"] \\
& \overline{B} \arrow[dl, "\overline{\pi}_1"'] \arrow[dr, "\overline{\pi}_2"] & \\
\mathsf{P}X_1 \arrow[dr, "\mathsf{P}\pi_1"'] & & \mathsf{P}X_2 \arrow[dl, "\mathsf{P}\pi_2"] \\
& \mathsf{P}B &
\end{tikzcd} \tag{2}$$

We instantiate the definition for some of the examples of Section 2.

Example 3.3 Consider the setting of Example 2.5 where **A** is a variety and (L, ρ) is given by predicate liftings and axioms. If **C** is concrete, then (B, π_1, π_2) is a ρ-bisimulation if for all $(x_1, x_2) \in B$, $\lambda \in \Lambda$ and $(a_1, a_2) \in \overline{B} \subseteq \mathsf{P}X_1 \times \mathsf{P}X_2$, we have:

$$\gamma_1(x_1) \in \lambda_{X_1}(a_1) \quad \text{iff} \quad \gamma_2(x_2) \in \lambda_{X_2}(a_2).$$

The notion of a ρ-bisimulation thus generalises that of a Λ-bisimulation from [4,7], where Λ denotes a collection of (open) predicate liftings. Examples 3.4 and 3.5 below are instances hereof.

Example 3.4 In the setting of positive modal logic from Exm. 2.6, a ρ-bisimulation between (X_1, γ_1) and (X_2, γ_2) is a subspace $B \subseteq X_1 \times X_2$ with projections $\pi_i : B \to X_i$ satisfying for all $(x_1, x_2) \in B$ and all B-coherent pairs of opens $(a_1, a_2) \in \Omega X_1 \times \Omega X_2$:

$$\gamma_1(x_1) \subseteq a_1 \text{ iff } \gamma_2(x_2) \subseteq a_2 \quad \text{and} \quad \gamma_1(x_1) \cap a_1 \neq \emptyset \text{ iff } \gamma_2(x_2) \cap a_2 \neq \emptyset.$$

Example 3.5 In the setting of modal vector logic from Exm. 2.7, jointly mono spans are linear subspaces, and the dual span of (B, π_1, π_2) consists of those pairs of $\Bbbk$-valued, linear predicates $(h_1, h_2) \in X_1^\vee \times X_2^\vee$ such that $(x_1, x_2) \in B$ implies $h_1(x_1) = h_2(x_2)$. Unravelling the definitions shows that a linear subspace (B, π_1, π_2) of $X_1 \times X_2$ is a ρ-bisimulation between (X_1, γ_1) and (X_2, γ_2), if for all $(x_1, x_2) \in B$, we have $[\![p]\!]_\gamma(x_1) = [\![p]\!]_\gamma(x_2)$, and:

$$\text{if } x_1 \xrightarrow{a} y_1 \text{ and } x_2 \xrightarrow{a} y_2, \text{ then } h_1(y_1) = h_2(y_2) \text{ for all } (h_1, h_2) \in \overline{B}.$$

We say that a span (B, π_1, π_2) between (X_1, γ_1) and (X_2, γ_2) is *truth preserving* if $\mathrm{th}_{\gamma_1} \circ \pi_1 = \mathrm{th}_{\gamma_2} \circ \pi_2$. If $\mathbf{C}$ is concrete, this means that if $(x_1, x_2) \in B$ then x_1 and x_2 have the same theory, i.e., satisfy the same formulas. As desired, ρ-bisimulations are truth-preserving:

Proposition 3.6 *If* $X_1 \xleftarrow{\pi_1} B \xrightarrow{\pi_2} X_2$ *is a* ρ*-bisimulation between* T*-coalgebras* (X_1, γ_1) *and* (X_2, γ_2)*, then* $\mathrm{th}_{\gamma_1} \circ \pi_1 = \mathrm{th}_{\gamma_2} \circ \pi_2$.

Proof. Let $\beta : \mathsf{L}\overline{B} \to \overline{B}$ be given as in (2), and let $h_\beta \colon \Phi \to \overline{B}$ be the unique morphism from the initial L-algebra. By construction of β, $\overline{\pi}_i \colon (\overline{B}, \beta) \to (\mathsf{P}X_i, \gamma_i^*)$ are L-algebra morphisms. By uniqueness of initial morphisms, $[\![-]\!]_{\gamma_i} = \overline{\pi}_i \circ h_\beta$, and hence $\mathsf{S}[\![-]\!]_{\gamma_i} = \mathsf{S}h_\beta \circ \mathsf{S}\overline{\pi}_i$, $i = 1, 2$. Combining this with $\mathsf{S}\overline{\pi}_1 \circ \mathsf{SP}\pi_1 = \mathsf{S}\overline{\pi}_2 \circ \mathsf{SP}\pi_2$ (obtained by applying S to the pullback square of $(\overline{B}, \overline{\pi}_1, \overline{\pi}_2)$), it follows that $\mathsf{S}[\![\cdot]\!]_{\gamma_1} \circ \mathsf{SP}\pi_1 = \mathsf{S}[\![\cdot]\!]_{\gamma_2} \circ \mathsf{SP}\pi_2$. Recall that the theory map is the adjoint of the semantic map, i.e., $\mathrm{th}_{\gamma_i} = \mathsf{S}[\![-]\!]_{\gamma_i} \circ \eta^{\mathbf{C}}_{X_i}$ where $\eta^{\mathbf{C}} : \mathrm{Id}_{\mathbf{C}} \to \mathsf{SP}$ is a unit of the logical connection $\mathsf{P} : \mathbf{C} \rightleftarrows \mathbf{A} : \mathsf{S}$. Together with naturality of $\eta^{\mathbf{C}}$, it now follows that:

$$\begin{aligned} \mathrm{th}_{\gamma_1} \circ \pi_1 = \mathsf{S}[\![\cdot]\!]_{\gamma_1} \circ \eta^{\mathbf{C}}_{X_1} \circ \pi_1 \quad &= \mathsf{S}[\![\cdot]\!]_{\gamma_1} \circ \mathsf{SP}\pi_1 \circ \eta^{\mathbf{C}}_B \\ = \mathsf{S}[\![\cdot]\!]_{\gamma_2} \circ \mathsf{SP}\pi_2 \circ \eta^{\mathbf{C}}_B = \mathsf{S}[\![\cdot]\!]_{\gamma_2} \circ \eta^{\mathbf{C}}_{X_2} \circ \pi_2 \quad &= \mathrm{th}_{\gamma_2} \circ \pi_2 \end{aligned}$$

□

3.2 Lattice structure and composition of ρ-bisimulations

In the remainder of Section 3 we assume that $\mathbf{C}$ is finitely complete and well-powered, hence $\mathbf{Rel}(X_1, X_2)$ is simply the poset of subobjects of $X_1 \times X_2$. Besides, assume that $\mathbf{C}$ has an $(\mathcal{E}, \mathcal{M})$-factorisation system with $\mathcal{M} = \mathcal{M}ono$.

It is well known that bisimulations for **Set**-based coalgebras are closed under composition if and only if the coalgebra functor preserves weak pullbacks [27]. We know from [4, Exm. 3.3] that Λ-bisimulations do not always compose, even for weak pullback-preserving functors, so the same failure occurs for ρ-bisimulations (cf. Exm. 3.3). However, in special cases we *can* compose.

The composition of two jointly mono spans (B, π_1, π_2) in $\mathbf{Rel}(X_1, X_2)$ and (B', π_2', π_3) in $\mathbf{Rel}(X_2, X_3)$ is given as follows: The pullback (C, c_1, c_3) of π_2 and

π_2' yields projections $\pi_i \circ c_i : C \to X_i$, and we define $B \circ B'$ via the $(\mathcal{E}, \mathcal{M}ono)$-factorisation of $\langle \pi_1 \circ c_1, \pi_3 \circ c_3 \rangle : C \to X_1 \times X_3$ as $C \twoheadrightarrow B \circ B' \hookrightarrow X_1 \times X_3$.

Call a ρ-bisimulation *full* if both projections are split epi, that is, they have a section. For **Set**-based coalgebras this means that the projections are surjective, i.e., each state in (X_1, γ_1) is ρ-bisimilar to some state in (X_2, γ_2), and vice versa.

Lemma 3.7 *Pullbacks preserve split epimorphisms.*

Lemma 3.8 *Let (f', v) be a pullback of (w, f). If w, f are regular epic and v, f' are split epic then (w, f) is a pushout of (f', v).*

Lemma 3.9 *Let $X_1 \xleftarrow{\zeta_1} S \xrightarrow{\zeta_2} X_2$ and $X_1 \xleftarrow{\pi_1} B \xrightarrow{\pi_2} X_2$ be spans between T-coalgebra (X_1, γ_1) and (X_2, γ_2) and suppose $e : S \to B$ is an epi such that $\zeta_i = \pi_i \circ e$. Then (S, ζ_1, ζ_2) satisfies (1) if and only if (B, π_1, π_2) does.*

Now we can show that full bisimulations compose.

Proposition 3.10 *Full bisimulations are closed under composition.*

Proof. By Lemma 3.9 it suffices to show that $(C, \pi_1 c_1, \pi_3 c_3)$ satisfies the ρ-bisimulation condition. Since all the π_i are split epic, so are c_1 and c_3 (cf. Lemma 3.7). According to Lemma 3.8 this implies that the square is also a pushout. Therefore 4 below is a pullback, while 1, 2 and 3 are pullbacks by definition. It follows that the outer square is a pullback.

$$\begin{tikzcd}
 & & & & \overline{C} \arrow[dll, "\overline{c}_1"'] \arrow[drr, "\overline{c}_3"] \arrow[dd, phantom, "1"] & & & & \\
 & & \overline{B}_1 \arrow[dll, "\overline{\pi}_1"'] \arrow[drr, "\overline{\pi}_2"] \arrow[dd, phantom, "2"] & & & & \overline{B}_2 \arrow[dll, "\overline{\pi}_2'"'] \arrow[drr, "\overline{\pi}_3"] \arrow[dd, phantom, "3"] & & \\
\mathsf{P}X_1 \arrow[drr, "\mathsf{P}\pi_1"'] & & & & \mathsf{P}X_2 \arrow[dll, "\mathsf{P}\pi_2"'] \arrow[drr, "\mathsf{P}\pi_2'"] \arrow[dd, phantom, "4"] & & & & \mathsf{P}X_3 \arrow[dll, "\mathsf{P}\pi_3"] \\
 & & \mathsf{P}B_1 \arrow[drr, "\mathsf{P}c_1"'] & & & & \mathsf{P}B_2 \arrow[dll, "\mathsf{P}c_2"] & & \\
 & & & & \mathsf{P}C & & & &
\end{tikzcd} \tag{3}$$

As a consequence $(\overline{\pi}_1 \overline{c}_1, \overline{\pi}_3 \overline{c}_3)$ is jointly monic. Furthermore, using the fact that B_1 and B_2 are ρ-bisimulations, a straightforward computation shows that $(C, \pi_1 c_1, \pi_3 c_3)$ is a ρ-bisimulation. □

Another well-known result for bisimulations on **Set**-coalgebras is that they form a complete lattice [27]. We now show that, provided **C** has all coproducts, this also holds for ρ-bisimulations. Recall that the empty coproduct $\coprod \emptyset =: \mathbf{0}$ is an initial object, i.e., for all $C \in \mathbf{C}$ there is a unique morphism $!_C : \mathbf{0} \to C$.

Definition 3.11 The *join* of a family $(B_i, \pi_{i,1}, \pi_{i,2})$, $i \in I$, in $\mathbf{Rel}(X_1, X_2)$, is the jointly mono span $\bigcup_{i \in I} B_i$ that arises from the factorisation

$$\begin{tikzcd}
\coprod_i B_i \arrow[r, two heads] \arrow[rr, bend left, "\coprod_i \langle \pi_{i,1}, \pi_{i,2} \rangle"] & \bigcup_i B_i \arrow[r, hook] & X_1 \times X_2
\end{tikzcd}$$

The *bottom element* (I, ι_1, ι_2) in $\mathbf{Rel}(X_1, X_2)$ is defined by the factorisation of the initial morphism: $\mathbf{0} \twoheadrightarrow I \xhookrightarrow{\langle \iota_1, \iota_2 \rangle} X_1 \times X_2$.

Indeed, $\bigcup_i B_i$ is an upper bound in $\mathbf{Rel}(X_1, X_2)$. Suppose $(B_i, \pi_{i,1}, \pi_{i,2}) \leq (S, s_1, s_2)$ for all i, then there are $t_i : B_i \to S$ such that $\pi_{i,j} = s_j \circ t_i$. From the coproduct we get $t : \coprod_{i \in I} B_i \to S$ and this makes the diagram on the right commute. The factorisation system now gives a diagonal $d : \bigcup_{i \in I} B_i \to S$ witnessing that S is bigger than $\bigcup_{i \in I} B_i$ in $\mathbf{Rel}(X_1, X_2)$.

$$\begin{tikzcd}
 & \bigcup_i B_i \arrow[dr, hook] \arrow[d, dashed, "d"] & \\
\coprod_i B_i \arrow[ur, two heads] \arrow[r, "t"] & S \arrow[r, hook] & X_1 \times X_2
\end{tikzcd}$$

Proposition 3.12 *If* $\mathbf{C}$ *has an* $(\mathcal{E}, \mathcal{M}ono)$*-factorisation system, binary products and all coproducts, then the poset* $\mathbf{Rel}(X_1, X_2)$ *is a complete join-semilattice with join* $\bigcup$ *and bottom element* (I, ι_1, ι_2)*.*

Proof. Commutativity and associativity of the join follows from the fact that coproducts are commutative and associative. For idempotency note that for every (B, π_1, π_2) in $\mathbf{Rel}(X_1, X_2)$ we have an $(\mathcal{E}, \mathcal{M}ono)$-factorisation $B + B \xrightarrow{\nabla} B \hookrightarrow X_1 \times X_2$, where ∇ is the codiagonal, so $B \cup B = B$.

Next, we show that (I, ι_1, ι_2) is the bottom element in $(\mathbf{Rel}(X_1, X_2), \cup)$. That is, for all (B, π_1, π_2) in $\mathbf{Rel}(X_1, X_2)$, $B \cup I$ is isomorphic to B. By the definition of a coproduct, $\begin{tikzcd} B \arrow[r, "\cong"'] \arrow[rr, bend left, "i"] & \mathbf{0} + B \arrow[r, two heads, "!_I + \mathrm{id}_B"'] & I + B \end{tikzcd}$ commutes, where i is the inclusion that arises from the coproduct, and because $\mathcal{E}$ is closed under composition, the map $i : B \to I + B$ is in $\mathcal{E}$. By definition of the join, the following commutes:

$$\begin{tikzcd}
B \arrow[r, two heads] \arrow[rrr, hook, bend left, "\langle \pi_1, \pi_2 \rangle"] & I + B \arrow[r, two heads] \arrow[rr, bend right, "{[\langle \iota_1, \iota_2 \rangle, \langle \pi_1, \pi_2 \rangle]}"'] & I \cup B \arrow[r, hook] & X_1 \times X_2
\end{tikzcd}$$

Since factorisation systems are unique up to isomorphism, we get an isomorphism $B \cong B \cup I$. $\square$

We define *ρ-bisimilarity* as the join of all ρ-bisimulations in $\mathbf{Rel}(X_1, X_2)$. The following proposition tells us that ρ-bisimilarity is itself a ρ-bisimulation. Given two T-coalgebras (X_1, γ_1) and (X_2, γ_2), we denote by ρ-$\mathbf{Bis}(\gamma_1, \gamma_2)$ the sub-poset of $\mathbf{Rel}(X_1, X_2)$ of ρ-bisimulations between (X_1, γ_1) and (X_2, γ_2).

Proposition 3.13 *Under the assumptions of Proposition 3.12,* ρ-$\mathbf{Bis}(\gamma_1, \gamma_2)$ *is closed under joins and bottom element in* $\mathbf{Rel}(X_1, X_2)$*. Consequently,* ρ-$\mathbf{Bis}(\gamma_1, \gamma_2)$ *is a complete join semilattice, and hence also a complete lattice.*

While ρ-$\mathbf{Bis}(\gamma_1, \gamma_2)$ is a complete sub-semilattice of $\mathbf{Rel}(X_1, X_2)$, it need not inherit the meets. This resembles the situation for Kripke bisimulations, which are generally not closed under intersections.

Example 3.14 The categories $\mathbf{Set}$, $\mathbf{Top}$ and $\mathbf{Vec}_{\Bbbk}$ from Examples 3.4, 3.3 and 3.5 are well-powered, complete and cocomplete, and as mentioned in Section 2 have a $(\mathcal{R}eg\mathcal{E}pi, \mathcal{M}ono)$-factorisation system. Hence ρ-bisimulations for positive modal logic, modal vector logic and coalgebraic geometric logic form complete lattices, and we recover the similar result for Λ-bisimulations in [4].

3.3 Characterisation via relation lifting

Another property of bisimulations for **Set**-coalgebras is that they can be characterised via relation lifting (see e.g. [29, Sec. 2.2]), and that bisimilarity on a coalgebra (X, γ) is a greatest fixpoint of a monotone operator on the lattice of relations $\mathcal{P}(X \times X)$. In this subsection and the following, we show that these results generalise to ρ-bisimulations.

Given X_1, X_2 in $\mathbf{C}$, we shall define a monotone map

$$\mathsf{T}^\rho : \mathbf{Rel}(X_1, X_2) \to \mathbf{Rel}(\mathsf{T}X_1, \mathsf{T}X_2)$$

which lifts (B, π_1, π_2) in $\mathbf{Rel}(X_1, X_2)$ to $(\mathsf{T}^\rho B, \mathsf{T}^\rho \pi_1, \mathsf{T}^\rho \pi_2)$ in $\mathbf{Rel}(\mathsf{T}X_1, \mathsf{T}X_2)$. In order to do so, consider the composition, for $i = 1, 2$,

$$\sigma_i : \mathsf{T}X_i \xrightarrow{\eta^{\mathbf{C}}_{\mathsf{T}X_i}} \mathsf{SPT}X_i \xrightarrow{\mathsf{S}\rho_{X_i}} \mathsf{SLP}X_i \xrightarrow{\mathsf{SL}\overline{\pi}_i} \mathsf{SL}\overline{B} \tag{4}$$

For a concrete example of σ_i, see Example 3.17 below.

Definition 3.15 Given (B, π_1, π_2) in $\mathbf{Rel}(X_1, X_2)$, we define $\mathsf{T}^\rho(B, \pi_1, \pi_2) = (\mathsf{T}^\rho B, \mathsf{T}^\rho \pi_1, \mathsf{T}^\rho \pi_2)$ as the pullback of $\mathsf{T}X_1 \xrightarrow{\sigma_1} \mathsf{SL}\overline{B} \xleftarrow{\sigma_2} \mathsf{T}X_2$.

Observe that $(\mathsf{T}^\rho B, \mathsf{T}^\rho \pi_1, \mathsf{T}^\rho \pi_2)$ is a jointly mono span because it is a pullback. Monotonicity of T^ρ follows from unravelling the definitions. We can now characterise ρ-bisimulations as in [16] using the relation lifting T^ρ.

Theorem 3.16 *A jointly mono span (B, π_1, π_2) between two T-coalgebras (X_1, γ_1) and (X_2, γ_2) is a ρ-bisimulation if and only if there exists a morphism $\delta : B \to \mathsf{T}^\rho B$ in $\mathbf{C}$ making diagram (5) commute.*

$$\begin{tikzcd}
X_1 & B \arrow[l, "\pi_1"'] \arrow[r, "\pi_2"] \arrow[d, dashed, "\delta"] & X_2 \arrow[d, "\gamma_2"] \\
\mathsf{T}X_1 \arrow[from=u, "\gamma_1"'] & \mathsf{T}^\rho B \arrow[l, "\mathsf{T}^\rho\pi_1"'] \arrow[r, "\mathsf{T}^\rho\pi_2"] & \mathsf{T}X_2
\end{tikzcd} \tag{5}$$

Proof. *If δ exists, then B is a ρ-bisimulation.* Suppose such a δ exists. In order to show that B is a ρ-bisimulation, we need to show that the outer shell of the left diagram below commutes. Recall that $\eta^{\mathbf{C}}$ and $\eta^{\mathbf{A}}$ are the units of the dual adjunction $\mathsf{P} : \mathbf{C} \rightleftarrows \mathbf{A} : \mathsf{S}$.

$$\begin{tikzcd}
 & \mathsf{L}\overline{B} \arrow[dl, bend right, "\mathsf{L}\overline{\pi}_1"'] \arrow[dr, bend left, "\mathsf{L}\overline{\pi}_2"] \arrow[d, "\eta^{\mathbf{A}}_{\mathsf{L}\overline{B}}"] & \\
\mathsf{LP}X_1 \arrow[d, "\rho_{X_1}"'] & \mathsf{PSL}\overline{B} \arrow[dl, "\mathsf{P}\sigma_1"'] \arrow[dr, "\mathsf{P}\sigma_2"] & \mathsf{LP}X_2 \arrow[d, "\rho_{X_2}"] \\
\mathsf{PT}X_1 \arrow[r, "\mathsf{PT}^\rho\pi_1"'] \arrow[d, "\mathsf{P}\gamma_1"'] & \mathsf{PT}^\rho B \arrow[d, "\mathsf{P}\delta"] & \mathsf{PT}X_2 \arrow[l, "\mathsf{PT}^\rho\pi_2"] \arrow[d, "\mathsf{P}\gamma_2"] \\
\mathsf{P}X_1 \arrow[r, "\mathsf{P}\pi_1"'] & \mathsf{P}B & \mathsf{P}X_2 \arrow[l, "\mathsf{P}\pi_2"]
\end{tikzcd}
\qquad
\begin{tikzcd}
\mathsf{L}\overline{B} \arrow[r, "\eta^{\mathbf{A}}_{\mathsf{L}\overline{B}}"] \arrow[d, "\mathsf{L}\overline{\pi}_i"] & \mathsf{PSL}\overline{B} \arrow[ddd, bend left, "\mathsf{P}\sigma_i"] \\
\mathsf{LP}X_i \arrow[r, "\eta^{\mathbf{A}}_{\mathsf{LP}X_i}"] \arrow[d, "\rho_{X_i}"] & \mathsf{PSLP}X_i \arrow[from=u, "\mathsf{PSL}\overline{\pi}_i"'] \\
\mathsf{PT}X_i \arrow[r, "\eta^{\mathbf{A}}_{\mathsf{PT}X_i}"] \arrow[dr, equal] & \mathsf{PSPT}X_i \arrow[from=u, "\mathsf{PS}\rho_{X_i}"'] \arrow[d, "\mathsf{P}\eta^{\mathbf{C}}_{\mathsf{T}X_i}"'] \\
 & \mathsf{PT}X_i
\end{tikzcd} \tag{6}$$

Commutativity of the bottom two squares follows from applying P to the diagram in (5). The middle square commutes because of the definition of $\mathsf{T}^\rho B$. The top two squares commute because they are the outer shell of the right diagram in (6). In (6), the right square commutes by definition of σ_i (Eq. 4).

The other two squares commute by naturality of $\eta^{\mathbf{A}}$ and the lower triangle is a triangle identity of the dual adjunction. Therefore the outer shell commutes.

If B is a ρ-bisimulation, then we can find δ. Suppose (B, π_1, π_2) is a ρ-bisimulation. If we can prove that $\sigma_1 \circ \gamma_1 \circ \pi_1 = \sigma_2 \circ \gamma_2 \circ \pi_2$ then we obtain δ as the mediating map induced by the pullback which defines $\mathsf{T}^\rho B$, as shown in the diagram (7).

$$
\begin{tikzcd}
 & B \arrow[dl, "\pi_1"'] \arrow[dr, "\pi_2"] \arrow[dd, dashed, "\delta"] & \\
X_1 \arrow[dd, "\gamma_1"'] & & X_2 \arrow[dd, "\gamma_2"] \\
 & \mathsf{T}^\rho B \arrow[dl, "\mathsf{T}^\rho\pi_1"'] \arrow[dr, "\mathsf{T}^\rho\pi_2"] & \\
\mathsf{T}X_1 \arrow[dr, "\sigma_1"'] & & \mathsf{T}X_2 \arrow[dl, "\sigma_2"] \\
 & \mathsf{SL}Q &
\end{tikzcd}
\tag{7}
$$

We claim that the following diagram commutes. Since its outer shell is the same as the outer shell of (7), this proves the proposition. So consider:

$$
\begin{tikzcd}
 & X_1 \arrow[dl, bend right, "\gamma_1"'] \arrow[d, "\eta_{X_1}"] & B \arrow[l, "\pi_1"'] \arrow[r, "\pi_2"] \arrow[d, "\eta_B"] & X_2 \arrow[dr, bend left, "\gamma_2"] \arrow[d, "\eta_{X_2}"] & \\
 & \mathsf{SP}X_1 \arrow[d, "\mathsf{SP}\gamma_1"] & \mathsf{SP}B \arrow[l, "\mathsf{SP}\pi_1"] \arrow[r, "\mathsf{SP}\pi_2"'] & \mathsf{SP}X_2 \arrow[d, "\mathsf{SP}\gamma_1"] & \\
\mathsf{T}X_1 \arrow[r, "\eta_{\mathsf{T}X_1}"] \arrow[drr, bend right, "\sigma_1"'] & \mathsf{SPT}X_1 \arrow[d, "\mathsf{S}\rho_{X_1}"] & & \mathsf{SPT}X_2 \arrow[d, "\mathsf{S}\rho_{X_2}"] & \mathsf{T}X_2 \arrow[l, "\eta_{\mathsf{T}X_2}"'] \arrow[dll, bend left, "\sigma_2"] \\
 & \mathsf{SLP}X_1 \arrow[dr, "\mathsf{SL}\overline{\pi}_1"] & & \mathsf{SLP}X_2 \arrow[dl, "\mathsf{SL}\overline{\pi}_2"'] & \\
 & & \mathsf{SL}Q & &
\end{tikzcd}
$$

Commutativity of the middle part follows from the fact that B is a ρ-bisimulation. The four top squares commute because η is a natural transformation. The two remaining squares commute by definition of σ_i. □

We work out the explicit description of T^ρ in a special case:

Example 3.17 Suppose we work with the classic dual adjunction $\mathcal{Q}_{\mathsf{BA}} : \mathbf{Set} \rightleftarrows \mathbf{BA} : \mathsf{Uf}$, T is an endofunctor on $\mathbf{Set}$, and the logic (L, ρ) is given by predicate liftings and axioms (cf. Example 2.5). Then the type of σ_i is $\mathsf{T}X_i \to \mathsf{UfL}\overline{B}$ and the ultrafilter $\sigma_i(t_i)$ is determined by the elements of the form $\underline{\lambda}(a_1, a_2)$ it contains, where $\lambda \in \Lambda$ and $(a_1, a_2) \in \overline{B}$. Therefore the action of T^ρ on (B, π_1, π_2) is given by

$$
\begin{aligned}
\mathsf{T}^\rho B = \{(t_1, t_2) \in \mathsf{T}X_1 \times \mathsf{T}X_2 \mid \forall \lambda \in \Lambda \text{ and } B\text{-coherent } (a_1, a_2) \\
\text{we have } t_1 \in \lambda_{X_1}(a_1) \Leftrightarrow t_2 \in \lambda_{X_2}(a_2)\}.
\end{aligned}
$$

Informally, these are the pairs in $\mathsf{T}X_1 \times \mathsf{T}X_2$ that cannot be distinguished by lifted B-coherent predicates.

3.4 Characterisation as a (post)fixpoint

As for $\mathbf{Set}$-coalgebras, given a relation lifting of T and T-coalgebras (X_1, γ_1), (X_2, γ_2), we can define a map $\mathsf{T}^\rho_{\gamma_1,\gamma_2} : \mathbf{Rel}(X_1, X_2) \to \mathbf{Rel}(X_1, X_2)$ by, essentially, taking inverse images under the γ_i. This is a relational version of a predicate transformer on a coalgebra.

Definition 3.18 Given T-coalgebras (X_1, γ_1) and (X_2, γ_2), define $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}(B, \pi_1, \pi_1) = (\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B, \mathsf{T}^{\rho}_{\gamma_1,\gamma_2}\pi_1, \mathsf{T}^{\rho}_{\gamma_1,\gamma_2}\pi_2) \in \mathbf{Rel}(X_1, X_2)$ via the pullback on the right. This is well defined because pullbacks are jointly mono spans.

$$\begin{tikzcd}
\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B \arrow[rr, "\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}\pi_2"] \arrow[dd, "\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}\pi_1"] & & X_2 \arrow[d, "\gamma_2"] \\
& & \mathsf{T}X_2 \arrow[d, "\sigma_2"] \\
X_1 \arrow[r, "\gamma_1"] & \mathsf{T}X_1 \arrow[r, "\sigma_1"] & \mathsf{SL}\overline{B}
\end{tikzcd}$$

Lemma 3.19 *The map* $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2} \colon \mathbf{Rel}(X_1, X_2) \to \mathbf{Rel}(X_1, X_2)$ *is monotone.*

Proof. If $(B, \pi_1, \pi_2) \leq (B', \pi'_1, \pi'_2)$ then there exists an $m : B \to B'$ such that $\pi_i = \pi'_i \circ m$. As a consequence the pullback $\overline{B}'$ is a cone for $\overline{B}$ and we have a mediating map $k : \overline{B}' \to \overline{B}$ satisfying $\overline{\pi}'_i = \overline{\pi}_i \circ k$. Unravelling the definitions reveals that $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B$ with its projections is a cone for $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B'$, hence there is a (unique) map $t : \mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B \to \mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B'$ such that $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}\pi_i = \mathsf{T}^{\rho}_{\gamma_1,\gamma_2}\pi'_i \circ t$ which witnesses that $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}(B, \pi_1, \pi_2) \leq \mathsf{T}^{\rho}_{\gamma_1,\gamma_2}(B', \pi'_1, \pi'_2)$. □

As announced, ρ-bisimulations are precisely the post-fixpoints of $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}$.

Theorem 3.20 *A relation* $X_1 \xleftarrow{\pi_1} B \xrightarrow{\pi_2} X_2$ *is a ρ-bisimulation between* (X_1, γ_1) *and* (X_2, γ_2) *if and only if* $(B, \pi_1, \pi_2) \leq \mathsf{T}^{\rho}_{\gamma_1,\gamma_2}(B, \pi_1, \pi_2)$.

Proof. If (B, π_1, π_2) is a ρ-bisimulation, then by Theorem 3.16 there is a map $\beta \colon B \to \mathsf{T}^{\rho}B$ such that diagram (5) commutes. We then get a map $\beta' \colon B \to \mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B$ from the pullback property of $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B$. Conversely, given $\beta' \colon B \to \mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B$, we obtain $\beta \colon B \to \mathsf{T}^{\rho}B$ from the pullback property of $\mathsf{T}^{\rho}B$. □

Monotonicity of $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}$ and the Knaster-Tarski fixpoint theorem imply:

Corollary 3.21 *Under the assumptions of Prop. 3.12,* $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}$ *has a greatest fixpoint, and this greatest fixpoint is ρ-bisimilarity.*

Example 3.22 We return to the classic setting of Example 3.17. Let (B, π_1, π_2) be a relation between T-coalgebras (X_1, γ) and (X_2, γ_2). Then

$$\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B = \{(x_1, x_2) \in X_1 \times X_2 \mid (\gamma_1(x_1), \gamma_2(x_2)) \in \mathsf{T}^{\rho}B\}.$$

Informally, $\mathsf{T}^{\rho}_{\gamma_1,\gamma_2}B$ consists of all pairs of worlds whose one-step behaviours are indistiguishable by lifted B-coherent predicates.

4 Distinguishing power

In this section we compare the distinguishing power of ρ-bisimulations with that of other semantic equivalence notions and logical equivalence. We make the same assumptions here as at the start of Section 3. Given a cospan $(X_1, \gamma_1) \to (Y, \delta) \leftarrow (X_2, \gamma_2)$ in $\mathbf{Coalg}(\mathsf{T})$, we call (Y, δ) a *congruence* (of T-coalgebras).

4.1 Comparison with known equivalence notions

We briefly recall three coalgebraic equivalence notions, in descending order of distinguishing power. For more details, see e.g. [4, Def. 3.9].

Definition 4.1 A jointly mono span (B, π_1, π_2) between (X_1, γ_1) and (X_2, γ_2) is a: (i) T-*bisimulation* if there is $t \colon B \to \mathsf{T}B$ such that the π_i become coalgebra morphisms; (ii) *precocongruence* if its pushout $\widehat{\pi}_1 : X_1 \to \widehat{B} \leftarrow X_2 : \widehat{\pi}_1$ can be turned into a congruence between (X_1, γ_1) and (X_2, γ_2), more precisely, if there is $t \colon \widehat{B} \to \mathsf{T}\widehat{B}$ such that $\widehat{\pi}_1$ and $\widehat{\pi}_2$ become coalgebra morphisms; (iii) *behavioural equivalence* if it is a pullback in $\mathbf{C}$ of some cospan $(X_1, \gamma_1) \to (Y, \delta) \leftarrow (X_2, \gamma_2)$.

When T preserves weak pullbacks, all three notions coincide (when considering associated "bisimilarity" notions), but in general, they may differ. In particular, expressive logics can generally only capture behavioural equivalence [14]. The next proposition can be proved in the same way as [4, Prop. 3.10].

Proposition 4.2 *(i) Every T-bisimulation is a ρ-bisimulation. (ii) Every precocongruence is a ρ-bisimulation.*

The converse direction requires additional assumptions.

Proposition 4.3 *If $\mathbf{C}$ has pushouts, P is faithful, and either (i) ρ is pointwise epic or (ii) $\rho^\flat$ is pointwise monic and T preserves monos, then every ρ-bisimulation is a precocongruence. If, in addition, T preserves weak pullbacks, then ρ-bisimilarity coincides with all three notions in Def. 4.1.*

Proof. Suppose $X_1 \xleftarrow{\pi_1} B \xrightarrow{\pi_2} X_2$ is a ρ-bisimulation with pushout $(\widehat{B}, \widehat{\pi}_1, \widehat{\pi}_2)$ be the pushout. We need to find a coalgebra structure $\zeta : \widehat{B} \to \mathsf{T}\widehat{B}$ which turns $\widehat{\pi}_1$ and $\widehat{\pi}_2$ into coalgebra morphisms. It suffices to show that $\mathsf{T}\widehat{\pi}_1 \circ \gamma_1 \circ \pi_1 = \mathsf{T}\widehat{\pi}_2 \circ \gamma_2 \circ \pi_2$, because then the universal property of the pushout yields the desired ζ. If P is faithful and ρ is pointwise epic, then it suffices to prove that $\mathsf{P}\pi_1 \circ \mathsf{P}\gamma_1 \circ \mathsf{PT}\widehat{\pi}_1 \circ \rho_{\widehat{B}} = \mathsf{P}\pi_2 \circ \mathsf{P}\gamma_2 \circ \mathsf{PT}\widehat{\pi}_2 \circ \rho_{\widehat{B}}$. This follows from the left diagram below, where the outer shell commutes because (B, π_1, π_2) is a ρ-bisimulation and the top two squares commute by naturality of ρ.

$$
\begin{tikzcd}
 & \mathsf{LP}\widehat{B} \arrow[dl, "\mathsf{LP}\widehat{\pi}_1"'] \arrow[dr, "\mathsf{LP}\widehat{\pi}_2"] \arrow[dd, "\rho_{\widehat{B}}"] & \\
\mathsf{LP}X_1 \arrow[dd, "\rho_{X_1}"'] & & \mathsf{LP}X_2 \arrow[dd, "\rho_{X_2}"] \\
 & \mathsf{PT}\widehat{B} \arrow[dl, "\mathsf{PT}\widehat{\pi}_1"'] \arrow[dr, "\mathsf{PT}\widehat{\pi}_2"] & \\
\mathsf{PT}X_1 \arrow[dd, "\mathsf{P}\gamma_1"'] & & \mathsf{PT}X_2 \arrow[dd, "\mathsf{P}\gamma_2"] \\
 & \mathsf{P}\widehat{B} \arrow[dl, "\mathsf{P}\widehat{\pi}_1"'] \arrow[dr, "\mathsf{P}\widehat{\pi}_2"] & \\
\mathsf{P}X_1 \arrow[dr, "\mathsf{P}\pi_1"'] & & \mathsf{P}X_2 \arrow[dl, "\mathsf{P}\pi_2"] \\
 & \mathsf{P}B &
\end{tikzcd}
\qquad
\begin{tikzcd}
 & \mathsf{SLP}\widehat{B} & \\
\mathsf{SLP}X_1 \arrow[ur, "\mathsf{SLP}\widehat{\pi}_1"] & & \mathsf{SLP}X_2 \arrow[ul, "\mathsf{SLP}\widehat{\pi}_2"'] \\
 & \mathsf{T}\widehat{B} \arrow[uu, "\rho^\sharp_{\widehat{B}}"] & \\
\mathsf{T}X_1 \arrow[uu, "\rho^\sharp_{X_1}"] \arrow[ur, "\mathsf{T}\widehat{\pi}_1"] & & \mathsf{T}X_2 \arrow[uu, "\rho^\sharp_{X_2}"'] \arrow[ul, "\mathsf{T}\widehat{\pi}_2"'] \\
 & \widehat{B} & \\
X_1 \arrow[uu, "\gamma_1"] \arrow[ur, "\widehat{\pi}_1"] & & X_2 \arrow[uu, "\gamma_2"'] \arrow[ul, "\widehat{\pi}_2"'] \\
 & B \arrow[ul, "\pi_1"] \arrow[ur, "\pi_2"'] &
\end{tikzcd}
$$

Alternatively, suppose P is faithful (hence $\eta^{\mathbf{C}} \colon \mathrm{Id}_{\mathbf{C}} \to \mathsf{SP}$ is pointwise monic), $\rho^\flat$ is pointwise monic and T preserves monos. Then the transpose $\rho^\sharp_{\widehat{B}} : \mathsf{T}\widehat{B} \to \mathsf{SLP}\widehat{B}$ of $\rho_{\widehat{B}}$ is monic, because $\rho^\sharp_{\widehat{B}} = \mathsf{S}\rho_{\widehat{B}} \circ \eta^{\mathbf{C}}_{\mathsf{T}\widehat{B}} = \rho^\flat_{\mathsf{P}\widehat{B}} \circ \mathsf{T}\eta^{\mathbf{C}}_{\widehat{B}}$, so it suffices to show that $\rho^\sharp_{\widehat{B}} \circ \mathsf{T}\widehat{\pi}_1 \circ \gamma_1 \circ \pi_1 = \rho^\sharp_{\widehat{B}} \circ \mathsf{T}\widehat{\pi}_2 \circ \gamma_2 \circ \pi_2$. But this follows from transposing the left diagram above, which yields the diagram to the right.

When T preserves weak pullbacks, T-bisimilarity coincides with behavioural equivalence [27], and hence also with the largest precocongruence and ρ-bisimilarity. $\square$

We note that condition (ii) in Proposition 4.3 entails that (L, ρ) is expressive [19, Thm. 4.2], i.e., that logical equivalence implies behavioural equivalence. In our abstract setting, *logical equivalence* with respect to (L, ρ) is the kernel pair (B, π, π') of the theory map $\mathrm{th} : X \to \mathsf{S}\Phi$. Hence, (L, ρ) is *expressive* if (B, π, π') is below a behavioural equivalence in $\mathbf{Rel}(X, X)$.

4.2 Hennessy-Milner type theorem

We now prove a partial converse to Proposition 3.6 (truth-preservation). We show that under certain conditions logical equivalence implies ρ-bisimilarity.

Theorem 4.4 *Let $\mathbf{C}' \rightleftarrows \mathbf{A}'$ be the dual equivalence induced by the dual adjunction $\mathbf{C} \rightleftarrows \mathbf{A}$. Suppose that*

- $\mathbf{C}$ *has $(\mathcal{R}eg\mathcal{E}pi, \mathcal{M}ono)$-factorisations for morphisms with domain $\in \mathbf{C}'$;*
- $\mathbf{C}'$ *is closed under regular epimorphic images;*
- S *is faithful and* L *preserves epis.*

Then for all T-coalgebras (X, γ) with $X \in \mathbf{C}'$, logical equivalence, i.e., the kernel pair (B, π, π') of $\mathrm{th} : X \to \mathsf{S}\Phi$, is a ρ-bisimulation.

Proof. In order to prove that (B, π, π') is a ρ-bisimulation, we need to show that the outer shell of the diagram on the right commutes. From B being the kernel pair of th we have that $(\Phi, [\![\cdot]\!]_\gamma, [\![\cdot]\!]_\gamma)$ is a cone for the pullback $\overline{B}$. Hence we get a morphism $h : \Phi \to \overline{B}$ such that the triangles left and right of h commute, and it is easy to see that all the inner squares and triangles in the diagram on the right commute. Thus, in order to show that the outer shell commutes, it suffices to show that $\mathsf{L}h$ is epic. By the assumption that L preserves epis, it suffices to show that $h : \Phi \to \overline{B}$ is epic. Let $m \circ e$ be the $(\mathcal{R}eg\mathcal{E}pi, \mathcal{M}ono)$-factorisation of th. Then the left diagram in (8) commutes. Since m is monic the upper square is a pullback, and by [2, Proposition 11.33] it is also a pushout. As a consequence, the lower square in the right diagram of (8), obtained from dualising the left one, is a pullback.

$$\begin{tikzcd}
& \mathsf{L}\overline{B} \arrow[dl, bend right, "\mathsf{L}\overline{\pi}"'] \arrow[dr, bend left, "\mathsf{L}\overline{\pi}'"] & \\
\mathsf{LP}X \arrow[dd, "\gamma^*"'] & \mathsf{L}\Phi \arrow[l, "\mathsf{L}[\![\cdot]\!]_\gamma"] \arrow[r, "\mathsf{L}[\![\cdot]\!]_\gamma"'] \arrow[u, dashed, "\mathsf{L}h"] \arrow[d, "\alpha"] & \mathsf{LP}X \arrow[dd, "\gamma^*"] \\
& \Phi \arrow[dl, "[\![\cdot]\!]_\gamma"'] \arrow[dr, "[\![\cdot]\!]_\gamma"] \arrow[d, dashed, "h"] & \\
\mathsf{P}X \arrow[dr, bend right, "\mathsf{P}\pi"'] & \overline{B} \arrow[l, "\overline{\pi}"] \arrow[r, "\overline{\pi}'"'] & \mathsf{P}X \arrow[dl, bend left, "\mathsf{P}\pi'"] \\
& \mathsf{P}B &
\end{tikzcd}$$

$$\begin{tikzcd}
& B \arrow[dl, "\pi"'] \arrow[dr, "\pi'"] & \\
X \arrow[r, two heads, "e"] \arrow[dr, bend right, "\mathrm{th}"'] & A \arrow[d, hook, "m"] & X \arrow[l, two heads, "e"'] \arrow[dl, bend left, "\mathrm{th}"] \\
& \mathsf{S}\Phi &
\end{tikzcd}
\qquad
\begin{tikzcd}
& \Phi \arrow[ddl, bend right, "[\![\cdot]\!]_\gamma"'] \arrow[ddr, bend left, "[\![\cdot]\!]_\gamma"] \arrow[d, "h"] & \\
& \mathsf{P}A \arrow[dl, hook', "\mathsf{P}e"'] \arrow[dr, hook, "\mathsf{P}e"] & \\
\mathsf{P}X \arrow[dr, hook, "\mathsf{P}\pi"'] & & \mathsf{P}X \arrow[dl, hook', "\mathsf{P}\pi'"] \\
& \mathsf{P}B &
\end{tikzcd}
\tag{8}$$

Here h denotes the adjoint transpose of m. Applying S to h gives the morphism $\mathsf{S}h : \mathsf{SP}A \to \mathsf{S}\Phi$ which by assumption is isomorphic to m (because $A \cong \mathsf{SP}A$). Since S is faithful and m is monic, h and therefore $\mathsf{L}h$ are epic. □

Example 4.5 In the classic case, $\mathbf{Set} \rightleftarrows \mathbf{BA}$ restricts to the full duality between finite sets and finite Boolean algebras. $\mathbf{Set}$ has $(\mathcal{R}eg\mathcal{E}pi, \mathcal{M}ono)$-

factorisations [2, Exm. 14.2(2)]. In **Set** and **BA**, all epis are regular and coincide with surjections [2,6], and finite sets are closed under surjective images. The ultrafilter functor S is faithful. If the logic functor L is given by predicate liftings and relations, then by [23, Remark 4.10] it preserves regular epis, and since all epis are regular, L preserves epis. Applying Theorem 4.4, we recover [4, Theorem 4.5], and thereby all examples given there. In particular, taking (L, ρ) to be Hennessy-Milner logic (Example 2.4), then we recover from Theorem 4.4 that over finite labelled transition systems, logical equivalence implies ρ-bisimilarity for Hennessy-Milner logic.

Remark 4.6 For positive modal logic from Examples 2.6 and 3.4, we have not been able to show that the logic functor $\mathsf{N} : \mathbf{DL} \to \mathbf{DL}$ preserves epis.

Example 4.7 We return to modal vector logic from Examples 2.7 and 3.5. The dual adjunction $\mathbf{Vec}_{\Bbbk} \rightleftarrows \mathbf{Vec}_{\Bbbk}$ restricts to the well-known self-duality of finite-dimensional vector spaces $\mathbf{FinVec}_{\Bbbk}$. The category $\mathbf{Vec}_{\Bbbk}$ has $(\mathcal{R}eg\mathcal{E}pi, \mathcal{M}ono)$-factorisations [2, Ex. 14.2] and the regular epis in both $\mathbf{Vec}_{\Bbbk}$ and $\mathbf{FinVec}_{\Bbbk}$ are the surjections [2, Exm. 7.72]. Moreover, the surjective image of a finite-dimensional vector space is again finite-dimensional, and the functor $(-)^{\vee}$ is faithful. Finally, since L is generated by predicate liftings and axioms it preserves surjections, so we can apply Theorem 4.4 to conclude that logical equivalence and ρ-bisimilarity coincide on W-coalgebras state-spaces in $\mathbf{FinVec}_{\Bbbk}$.

Example 4.8 An example where logical equivalence does not imply ρ-bisimilarity is given by trace logic for labelled transitions systems (Example 2.3). The conditions for Theorem 4.4 hold for trace logic, but the induced dual equivalence is in this case trivial, i.e., $\mathbf{C}'$ and $\mathbf{A}'$ are the empty category, hence Theorem 4.4 does not tell us anything.

4.3 Invariance under translations

In this section we assume that $\mathbf{C}$ has pushouts. The example of Hennessy-Milner logic (Exm. 2.4) and trace logic (Exm. 2.3 and 4.8) is a situation where one logic is a reduct of the other. This can be considered a special case of translating a logic into another. We will show under which conditions ρ-bisimilarity is preserved under translations. To make this formal, we first generalise [22, Def 4.1].

Definition 4.9 Assume we are given a "triangle situation" as in diagram (9(a)) such that $\mathsf{P} = \mathsf{U}\mathsf{P}'$, and we have modal semantics $\rho' : \mathsf{L}'\mathsf{P}' \to \mathsf{P}'\mathsf{T}$ and $\rho : \mathsf{L}\mathsf{P} \to \mathsf{P}\mathsf{T}$. A *translation* from (L', ρ') to (L, ρ) is a natural transformation $\tau : \mathsf{L}\mathsf{P} \to \mathsf{U}\mathsf{L}'\mathsf{P}'$ such that $\rho = \mathsf{U}\rho' \circ \tau$, see diagram (9(b)).

$$
\begin{tikzcd}
 & \mathbf{A}' \arrow[loop right, "\mathsf{L}'"] \arrow[d, shift left, "\mathsf{U}"] \\
\mathbf{C} \arrow[loop left, "\mathsf{T}"] \arrow[ur, "\mathsf{P}'"] \arrow[dr, "\mathsf{P}"'] & \mathbf{A} \arrow[loop right, "\mathsf{L}"] \arrow[u, shift left, "\mathsf{F}"] \arrow[u, phantom, "\dashv"]
\end{tikzcd}
\qquad
\begin{tikzcd}
\mathsf{LP} \arrow[r, "\tau"] \arrow[d, "\rho"] & \mathsf{UL'P'} \arrow[d, "\mathsf{U}\rho'"] \\
\mathsf{PT} \arrow[r, "="] & \mathsf{UP'T}
\end{tikzcd}
\qquad
\begin{tikzcd}
\mathsf{FLP} \arrow[r, "\tau^{\sharp}"] \arrow[d, "\mathsf{F}\rho"'] & \mathsf{L'P'} \arrow[d, "\rho'"] \\
\mathsf{FPT} \arrow[r, "\varepsilon_{\mathsf{P'T}}"] & \mathsf{P'T}
\end{tikzcd}
\qquad (9)
$$

(a) (b) (c)

In (c), ε is the counit of $\mathsf{F} \dashv \mathsf{U}$ (which is adjoint to the identity) because $\mathsf{P} = \mathsf{UP}'$, and $\tau^\sharp$ is the $(\mathsf{F} \dashv \mathsf{U})$-adjoint of τ. In the presence of such a translation, every ρ'-bisimulation is also a ρ-bisimulation. We leave the straightforward proof to the reader. A sufficient condition for the converse is that the transpose $\tau^\sharp$ of τ is epic, see diagram (9(c)). Note that due to the adjunction $\mathsf{F} \dashv \mathsf{U}$, diagram (b) commutes if and only if (c) does. Intuitively, $\tau^\sharp : \mathsf{FLP} \to \mathsf{L}'\mathsf{P}'$ being epic formalises that every modality in L' is a propositional combination of a modal formula of L.

Proposition 4.10 *Suppose that $\tau^\sharp$ is pointwise epic. Then every ρ-bisimulation is a ρ'-bisimulation.*

Proof. Commutativity of the outer shell of the following diagram will prove that $X_1 \xleftarrow{\pi_1} B \xrightarrow{\pi_2} X_2$ is a ρ'-bisimulation:

$$\begin{tikzcd}
 & \mathsf{L}'\mathsf{P}'X_1 \arrow[dl, bend right, "\rho'_{X_1}"'] \arrow[dr, phantom, "2"] & \mathsf{L}'\mathsf{P}'\widehat{B} \arrow[l, "\mathsf{L}'\mathsf{U}\widehat{\pi}_1"'] \arrow[r, "\mathsf{L}'\mathsf{U}\widehat{\pi}_2"] & \mathsf{L}'\mathsf{P}'X_2 \arrow[dr, bend left, "\rho'_{X_2}"] \arrow[dl, phantom, "3"] & \\
1 & \mathsf{FLP}X_1 \arrow[u, "\tau^\sharp_{X_1}"] \arrow[d, "\mathsf{F}\rho_{X_1}"'] & \mathsf{FLP}\widehat{B} \arrow[u, "\tau^\sharp_{\widehat{B}}"] \arrow[l, "\mathsf{FLP}\widehat{\pi}_1"] \arrow[r, "\mathsf{FLP}\widehat{\pi}_2"'] \arrow[d, phantom, "5"] & \mathsf{FLP}X_2 \arrow[u, "\tau^\sharp_{X_2}"'] \arrow[d, "\mathsf{F}\rho_{X_2}"] & 4 \\
\mathsf{P}'\mathsf{T}X_1 \arrow[ddr, bend right, "\mathsf{P}'\gamma_1"'] & \mathsf{FPT}X_1 \arrow[l, "\varepsilon_{\mathsf{P}'\mathsf{T}X_1}"] \arrow[d, "\mathsf{FP}\gamma_1"'] & \mathsf{FP}\widehat{B} \arrow[dl, bend right, "\mathsf{FP}\widehat{\pi}_1"'] \arrow[dr, bend left, "\mathsf{FP}\widehat{\pi}_2"] \arrow[d, phantom, "6"] & \mathsf{FPT}X_2 \arrow[r, "\varepsilon_{\mathsf{P}'\mathsf{T}X_2}"'] \arrow[d, "\mathsf{FP}\gamma_2"] & \mathsf{P}'\mathsf{T}X_2 \arrow[ddl, bend left, "\mathsf{P}'\gamma_2"] \\
7 & \mathsf{FP}X_1 \arrow[r, "\mathsf{FP}\pi_1"'] \arrow[d, "\varepsilon_{\mathsf{P}'X_1}"'] \arrow[dr, phantom, "8"] & \mathsf{FP}B \arrow[d, "\varepsilon_{\mathsf{P}'B}"] & \mathsf{FP}X_2 \arrow[l, "\mathsf{FP}\pi_2"] \arrow[d, "\varepsilon_{\mathsf{P}'X_1}"] \arrow[dl, phantom, "9"] & 10 \\
 & \mathsf{P}'X_1 \arrow[r, "\mathsf{P}'\pi_1"'] & \mathsf{P}'B & \mathsf{P}'X_2 \arrow[l, "\mathsf{P}'\pi_2"] &
\end{tikzcd}$$

Cells 1 and 4 commute by diagram (c) in (9), and cells 2 and 3 by naturality of $\tau^\sharp$. Commutativity of 5 and 6 together follows from applying F to the diagram witnessing the fact that $X_1 \longleftarrow B \longrightarrow X_2$ is a ρ-bisimulation. Commutativity of the remaining cells follows from the naturality of the counit ε. $\square$

In the setting of Examples 2.4, 2.5 and 3.3, where $\mathbf{A}'$ is a variety of algebras and the logic (L, ρ) is given by predicate liftings and axioms, we can consider the special case of (9) where (L, ρ) is the "modal reduct" of (L', ρ').

Example 4.11 Suppose $\mathbf{A}$ is a variety of algebras with free-forgetful adjunction $\mathsf{F} \dashv \mathsf{U}$. Let (L, ρ) be a logic for T-coalgebras given by a collection Λ of predicate liftings and axioms (Example 2.5). Then we can define $\mathsf{P}_0 = \mathsf{U} \circ \mathsf{P}$, which has dual adjoint $\mathsf{S}_0 = \mathsf{SF}$, where S is the dual adjoint of P. Define the logic functor $\mathsf{L}_0 : \mathbf{Set} \to \mathbf{Set}$ by $\mathsf{L}_0 X = \{\underline{\lambda}_0(a_1, \ldots, a_n) \mid \lambda \in \Lambda a_i \in X\}$ and $\mathsf{L}_0 f(\underline{\lambda}_0(a_1, \ldots, a_n)) = \underline{\lambda}_0(fa_1, \ldots, fa_n)$. Define

$$\begin{tikzcd}
 & \mathbf{A} \arrow[loop right, "\mathsf{L}"] \arrow[dd, shift left, "\mathsf{U}"] \arrow[dd, phantom, "\dashv"] \\
\mathbf{C} \arrow[loop left, "\mathsf{T}"] \arrow[ur, "\mathsf{P}"] \arrow[dr, "\mathsf{P}_0"'] & \\
 & \mathbf{Set} \arrow[loop right, "\mathsf{L}_0"] \arrow[uu, shift left, "\mathsf{F}"]
\end{tikzcd}$$

$\tau : \mathsf{L}_0\mathsf{P}_0 \to \mathsf{ULP}$ by $\tau_X(\underline{\lambda}_0(a_1, \ldots, a_n)) = \underline{\lambda}(a_1, \ldots, a_n) \in \mathsf{ULP}X$. The logic (L, ρ) gives rise to the logic (L_0, ρ_0), where $\rho_0 = \mathsf{U}\rho \circ \tau : \mathsf{L}_0\mathsf{P}_0 \to \mathsf{P}_0\mathsf{T}$. Then τ is a translation. One can verify that, in this situation, $\tau^\sharp$ is pointwise epic.

Therefore a jointly mono span $X_1 \longleftarrow B \longrightarrow X_2$ in $\mathbf{C}$ is a ρ-bisimulation if and only if it is a ρ_0-bisimulation. Hence it suffices to look at the underlying sets when verifying whether a jointly mono span is a ρ-bisimulation.

Hennessy-Milner logic and trace logic are a specific instance of Example 4.11.

Example 4.12 Recall trace logic (Exm. 2.3) and Henessy-Milner logic (Exm. 2.4) for LTSs. We gave a concrete definition of the Hennessy-Milner logic functor L^{hm} in Example 2.4. It can be equivalently defined as follows in terms of the free-forgetful adjunction $\mathsf{F} \dashv \mathsf{U}$ of $\mathbf{BA}$ over $\mathbf{Set}$. The join-preservation of the diamond modalities is encoded in L^{hm} by factoring the free-forgetful adjunction $\mathsf{F} \dashv \mathsf{U}$ via the category $\mathbf{JSL}$ of join-semilattices as shown in diagram (10(d)). The logic functor L^{hm} is then defined as $\mathsf{L}^{\mathsf{hm}} = \mathsf{F}\mathsf{L}^{\mathsf{tr}}\mathsf{U} = \mathsf{F}_{\mathsf{BA}}\mathsf{F}_{\mathsf{JSL}}\mathsf{L}^{\mathsf{tr}}\mathsf{U}_{\mathsf{JSL}}\mathsf{U}_{\mathsf{BA}}$. Letting $\mathsf{U} = \mathsf{U}_{\mathsf{JSL}}\mathsf{U}_{\mathsf{BA}}$ and $\mathsf{F} = \mathsf{F}_{\mathsf{BA}}\mathsf{F}_{\mathsf{JSL}}$, we have $\mathcal{Q} = \mathsf{U}\mathcal{Q}_{\mathsf{BA}}$, and in particular $\mathsf{L}^{\mathsf{hm}} = \mathsf{F}\mathsf{L}^{\mathsf{tr}}\mathsf{U}$. The semantics of Hennessy-Milner logic coincides with taking ρ^{hm} to be the adjoint of ρ^{tr}, so trace logic is the modal reduct of Hennessy-Milner logic.

Defining τ as in Exm. 4.11 implies that $\tau^\sharp \colon \mathsf{F}\mathsf{L}^{\mathsf{tr}}\mathcal{Q} \to \mathsf{L}^{\mathsf{hm}}\mathcal{Q}_{\mathsf{BA}}$ is the identity (this follows from $\mathsf{L}^{\mathsf{hm}} = \mathsf{F}\mathsf{L}^{\mathsf{tr}}\mathsf{U}$). Concretely, $\tau^\sharp$ is identity, because formulas of Hennessy-Milner logic are precisely the Boolean combinations of trace logic formulas. Hence, in particular, $\tau^\sharp$ is epic. It now follows from Proposition 4.10 that a ρ^{tr}-bisimulation is a ρ^{hm}-bisimulation (and the converse also holds).

$$
\begin{tikzcd}
 & & \mathbf{BA} \arrow[loop right, "\mathsf{L}^{\mathsf{hm}}"] \arrow[dll, bend left, "\mathsf{Uf}"'] \arrow[d, shift left, "\mathsf{U}_{\mathsf{BA}}"] \arrow[d, phantom, "\dashv"] \\
\mathbf{Set} \arrow[loop left, "\mathsf{T}"] \arrow[urr, bend left, "\mathcal{Q}_{\mathsf{BA}}"] \arrow[drr, bend left, "\mathcal{Q}"] & & \mathbf{JSL} \arrow[u, shift left, "\mathsf{F}_{\mathsf{BA}}"] \arrow[d, shift left, "\mathsf{U}_{\mathsf{JSL}}"] \arrow[d, phantom, "\dashv"] \\
 & & \mathbf{Set} \arrow[loop right, "\mathsf{L}^{\mathsf{tr}}"] \arrow[ull, bend left, "\mathcal{Q}"] \arrow[u, shift left, "\mathsf{F}_{\mathsf{JSL}}"]
\end{tikzcd}
\qquad
\begin{tikzcd}
 & \mathbf{Frm} \arrow[loop right, "\mathsf{N}'"] \arrow[d, "\mathsf{U}'"] \\
\mathbf{Top} \arrow[loop left, "\mathsf{V}"] \arrow[ur, bend left, "\Omega'"] \arrow[r, "\Omega"] \arrow[dr, bend right, "\Omega_0"'] & \mathbf{DL} \arrow[loop right, "\mathsf{N}"] \arrow[d, "\mathsf{U}"] \\
 & \mathbf{Set} \arrow[loop right, "\mathsf{N}_0"]
\end{tikzcd}
\tag{10}
$$

(d) (e)

Example 4.13 We squeeze the topological semantics for positive modal logic from Example 2.6 between two other logics with varying base logics, see diagram (10(e)). Here $\mathbf{Frm}$ is the category of frames and $\Omega' : \mathbf{Top} \to \mathbf{Frm}$ is the functor that sends a topological space to its frame of opens. Let $\mathsf{N}' : \mathbf{Frm} \to \mathbf{Frm}$ be the functor given as in [17, Section III4.3] (known also as the *Vietoris locale*) and define $\rho' : \mathsf{N}'\Omega' \to \Omega'\mathsf{V}$ on generators by $\Box a \mapsto \boxbar a$ and $\Diamond a \mapsto \diamondbar a$. The translation $\tau : \mathsf{N}\Omega \to \mathsf{U}'\mathsf{N}'\Omega'$ given by $\Box a \mapsto \Box a$ and $\Diamond a \mapsto \Diamond a$ is such that $\tau^\sharp$ is epic, thus satisfies the assumptions of Proposition 4.10.

The bottom triangle is an instance of Exm. 4.11. We conclude that a jointly mono span between V-coalgebras (X_1, γ_1) and (X_2, γ_2) is a ρ-bisimulation if and only if it is a ρ'-bisimulation if and only if it is a ρ_0-bisimulation.

5 Conclusion

Our main question was whether we can characterise logical equivalence for (possibly non-expressive) coalgebraic logics by a notion of bisimulation. Towards this goal, we generalised the logic-induced bisimulations in [4] for coalgebraic logics for **Set**-coalgebras to coalgebraic logics parameterised by a dual adjunction. We identified sufficient conditions for when logical equivalence coincides with logic-induced bisimilarity (Thm. 4.4). These are conditions on the categories in the dual adjunction, and *not* on the natural transformation ρ defining (the semantics of) the logic. In particular, we do not require the logic to be expressive.

We found that the distinguishing power of ρ-bisimulations depends on the modalities of the language but not on the propositional connectives. More generally, we showed that certain translations between logics preserve ρ-bisimilarity (Prop. 4.10). Furthermore, as in the expressivity result of [19], ρ-bisimilarity agrees with behavioural equivalence if the mate of ρ is pointwise monic (Prop. 4.3). However, Example 4.12 shows that this is not a necessary condition which raises the question whether one can characterise, purely in terms of ρ, when ρ-bisimilarity coincides with behavioural equivalence.

There are many other avenues for further research. When is a congruence on complex algebras induced by a ρ-bisimulation? Can we drop in Theorem 4.4 the restriction to the subcategory if T is finitary? Can we take quotients with respect to (the largest) ρ-bisimulation on a T-coalgebra?

Moreover, the definition of ρ-bisimulation has a natural generalisation to the order-enriched setting. This gives rise to *ρ-simulations*. Can one prove an ordered Hennessy-Milner theorem where "logical inequality" is recognised by ρ-simulations? Since this question naturally falls into the realm of order-enriched category theory, we will also seek a generalisation to the quantale-enriched setting, accounting for metric versions of simulation.

References

[1] Abramsky, S., *Domain theory in logical form*, Ann. Pure Appl. Logic **51** (1991).

[2] Adámek, J., H. Herrlich and G. E. Strecker, *Abstract and concrete categories* (1990).

[3] Adámek, J., J. Rosický and E. M. Vitale, *Algebraic theories: a categorical introduction to general algebra* (2010).

[4] Bakhtiari, Z. and H. H. Hansen, *Bisimulation for weakly expressive coalgebraic modal logics*, in: *7th Conference on Algebra and Coalgebra in Computer Science, CALCO*, Leibniz International Proceedings in Informatics (LIPIcs), 2017.

[5] Baltag, A. and G. Cinà, *Bisimulation for conditional modalities*, Stud. Log. **106** (2018).

[6] Banaschewski, B., *On the strong amalgamation of Boolean algebras*, Algebra Univers **63** (2010), pp. 235–238.

[7] Bezhanishvili, N., J. de Groot and Y. Venema, *Coalgebraic Geometric Logic*, in: *8th Conference on Algebra and Coalgebra in Computer Science (CALCO 2019)*, Leibniz International Proceedings in Informatics (LIPIcs), 2019.

[8] Bonchi, F., M. M. Bonsangue, M. Boreale, J. Rutten and A. Silva, *A coalgebraic perspective on linear weighted automata*, Information and Computation **211** (2012).

[9] Bonsangue, M. M. and A. Kurz, *Duality for logics of transition systems*, in: *Foundations of Software Science and Computational Structures (FoSSaCS'05)*, LNCS **3441** (2005).

[10] Celani, S. and R. Jansana, *Priestley duality, a Sahlqvist theorem and a Goldblatt-Thomason theorem for positive modal logic*, Logic Journal of the IGPL **7** (1999).

[11] Dunn, J. M., *Positive modal logic*, Studia Logica **55** (1995).

[12] Fan, J., Y. Wang and H. van Ditmarsch, *Almost necessary*, in: *Advances in Modal Logic (AiML'14)* (2014), pp. 178–196.

[13] Goldblatt, R., *Metamathematics of modal logic I*, Rep. on Math. Log. **6** (1976).

[14] Hansen, H. H., C. Kupke and E. Pacuit, *Neighbourhood structures: bisimilarity and basic model theory*, Log. Meth. Comp. Sci. **5** (2009).

[15] Hennessy, M. and R. Milner, *Algebraic laws for nondeterminism and concurrency*, Journal of the ACM (1985).

[16] Hermida, C. and B. Jacobs, *Structural induction and coinduction in a fibrational setting*, Inform. and Comput. **145** (1998).

[17] Johnstone, P. T., "Stone Spaces," Cambridge University Press, 1982.

[18] Kapulkin, K., A. Kurz and J. Velebil, *Expressiveness of positive coalgebraic logic*, Advances in Modal Logic **9** (2014).

[19] Klin, B., *Coalgebraic modal logic beyond sets*, in: *Mathematical Foundations of Programming Semantics (MFPS XXIII)*, ENTCS **173**, 2007, pp. 177 – 201.

[20] Kupke, C., A. Kurz and D. Pattinson, *Algebraic semantics for coalgebraic logics*, in: *Coalgebraic Methods in Computer Science (CMCS'04)*, ENTCS **106**, 2004, pp. 219–241.

[21] Kupke, C. and D. Pattinson, *Coalgebraic semantics of modal logics: An overview*, Theor. Comp. Sci. **412** (2011).

[22] Kurz, A. and R. Leal, *Modalities in the stone age: A comparison of coalgebraic logics*, Theor. Comp. Sci. **430** (2012).

[23] Kurz, A. and J. Rosický, *Strongly complete logics for coalgebras*, Log. Meth. Comp. Sci. **8** (2012).

[24] Moss, L., *Coalgebraic logic*, Annals of Pure and Applied Logic **96** (1999).

[25] Pattinson, D., *Coalgebraic modal logic: soundness, completeness and decidability of local consequence*, Theor. Comp. Sci. **309** (2003).

[26] Pattinson, D., *Expressive logics for coalgebras via terminal sequence induction*, Notre Dame Journal of Formal Logic **45** (2004).

[27] Rutten, J. J. M. M., *Universal coalgebra: a theory of systems*, Theor. Comp. Sci. **249** (2000).

[28] Schröder, L., *Expressivity of coalgebraic modal logic: The limits and beyond*, Theor. Comp. Sci. **390** (2008).

[29] Staton, S., *Relating coalgebraic notions of bisimulation*, Log. Meth. Comp. Sci. **7** (2011).