

Rank-1 Modal Logics are Coalgebraic

Lutz Schröder*, DFKI-Lab Bremen
and Dept. of Comput. Sci., Univ. Bremen

Dirk Pattinson†, Depart. of Computing, Imperial College London

Address for Correspondence:

Lutz Schröder
DFKI-Lab Bremen, SKS
Cartesium, Enrique-Schmidt-Str. 5
28359 Bremen
Germany
Phone: +49-421-218-64216
Fax: +49-421-218-9864216
E-Mail: Lutz.Schroeder@dfki.de

Abstract

Coalgebras provide a unifying semantic framework for a wide variety of modal logics. It has previously been shown that the class of coalgebras for an endofunctor can always be axiomatised in rank 1. Here we establish the converse, i.e. every rank 1 modal logic has a sound and *strongly* complete coalgebraic semantics. This is achieved by constructing for a given modal logic a canonical coalgebraic semantics, consisting of a signature functor and interpretations of modal operators, which turns out to be final among all such structures. The canonical semantics may be seen as a coalgebraic reconstruction of neighbourhood semantics, broadly construed. A finitary restriction of the canonical semantics yields a canonical weakly complete semantics which moreover enjoys the Hennessy-Milner property.

As a consequence, the machinery of coalgebraic modal logic, in particular generic decision procedures and upper complexity bounds, becomes applicable to arbitrary rank 1 modal logics, without regard to their semantic status; we thus obtain purely syntactic versions of such results. As an extended example, we apply our framework to recently defined deontic logics. In particular, our methods lead to the new result that these logics are strongly complete.

Keywords: Modal logic, coalgebra, neighbourhood frames, deontic logic, decision procedures

*This work forms part of the DFG project *Generic Algorithms and Complexity Bounds in Coalgebraic Modal Logic* (SCHR 1118/5-1)

†Partially supported by EPSRC grant EP/F031173/1

Introduction

In recent years, coalgebras have received a steadily growing amount of attention as general models of state-based systems [32], encompassing such diverse systems as labelled transition systems, probabilistic systems, game frames, and neighbourhood frames [36]. On the logical side, modal logic has emerged as the adequate specification language for coalgebraically modelled systems. A variety of different frameworks have been proposed [24, 31, 16]. Here, we work with *coalgebraic modal logic* [27], which allows for a high level of generality while retaining a close relationship to the established syntactic and semantic tradition of modal logic.

Reversing the viewpoint that modal logic is a suitable specification language for coalgebra, coalgebraic semantics captures, in a uniform way, the structural similarities of a large variety of different modal logics. In its most traditional form, modal logic comprises a unary operator \Box , read as “necessarily”, in addition to propositional connectives, and is interpreted over relational (Kripke) models. But there are many variations: probabilistic modal logic [21, 13] features operators that speak about the likelihood of events, with relational successors annotated with probabilities on the semantical side. The modalities of Pauly’s coalition logic [29] formalise coalitional power in strategic games and provide the syntactic means to reason about so-called game frames that involve strategies of participating agents. Yet another example is conditional logic [5] that adds a binary non-monotonic conditional to propositional logic and is commonly interpreted over selection function models. In a broad sense, the semantics of all the above logics comprise a notion of state or world and operators that describe the dynamics of state transitions.

The purpose and main idea of a coalgebraic treatment of modal logics is to lift these similarities to a formal level with the goal of deriving meta-theorems that uniformly apply to many structurally different logics. We illustrate this viewpoint by re-visiting the above examples in more detail.

Kripke Frames. The traditional textbook semantics of the modal logic K and its extensions is usually presented in relational form: a Kripke model is a pair (W, R) where W is a set of worlds and $R \subseteq W \times W$ is an accessibility relation. Kripke frames are easily seen to be in 1-1 correspondence to *powerset coalgebras*, i.e. pairs (W, ρ) where $\rho : W \rightarrow \mathcal{P}(W)$ is the transition function that assigns the set of successors $\{w' \in W \mid wRw'\}$ to each world w . In this example, $\rho(w)$ is the unstructured set of successors of a single point, and as we will see shortly, the coalgebraic approach deploys its full power when structured sets of successors are considered.

In the coalgebraic setting, it is most appropriate to think of modal operators as specifying properties of (possibly structured) *successor sets*. For example, if ϕ is a modal formula with extension $\llbracket \phi \rrbracket \subseteq W$, then a world w satisfies $\Box\phi$ if the *successor set* of w is contained in $\llbracket \phi \rrbracket$. In other words,

$$w \models \phi \iff \rho(w) \in \{B \in \mathcal{P}(W) \mid B \subseteq \llbracket \phi \rrbracket\}.$$

Conceptually, this epitomises the interpretation of the modal operator as a *predicate lifting*, that is, an operation that transforms predicates on states to predicates on (for now unstructured) successor sets. As we shall see in the examples to follow, this

re-formulation of modal semantics provides a common denominator for a broad class of structurally different modal logics.

Probabilistic Transition Systems. In its simplest form [21, 13], probabilistic modal logic (PML) extends propositional logic with operators L_p where $p \in [0, 1]$ is a rational number. The intended reading of $L_p\phi$ is ‘ ϕ holds with probability at least p in the next state’. PML is interpreted over probabilistic transition systems (W, P) where W is a set of worlds and $P = (P_w)_{w \in W}$ is a family of probability distributions on W , indexed by the set of worlds. In the spirit of the coalgebraic re-formulation of Kripke frames above, we readily recognise probabilistic transition systems as coalgebras (W, ρ) where $\rho : W \rightarrow \mathbf{D}(W)$ assigns a probability distribution $\rho(w) \in \mathbf{D}(W)$ over W to each world w . The main difference lies in the fact that collections of successors are now *structured*: moving from frames to probabilistic models entails a shift from successor sets to distributions. The classical interpretation of probabilistic formulas, i.e.

$$w \models L_p\phi \iff P_w(\llbracket \phi \rrbracket) \geq p$$

can now be re-phrased in terms of *successor distributions*:

$$w \models L_p\phi \iff \rho(w) \in \{\mu \in \mathbf{D}(W) \mid \mu(\llbracket \phi \rrbracket) \geq p\},$$

i.e. a state w satisfies $L_p\phi$ if its successor distribution assigns probability at least p to the event $\llbracket \phi \rrbracket$. Again, the quintessential nature of a probabilistic modal operator manifests itself as providing a passage from properties of states (subsets of W) to properties of successor distributions.

Conditional Logic. The language of conditional logic [5] extends propositional logic with a binary connective that we write \Rightarrow , using infix notation. The operator \Rightarrow represents a non-monotonic conditional, and the intended reading of $\phi \Rightarrow \psi$ is “ ψ holds under the condition ϕ ”. Note that this operator is in general distinct from implication \rightarrow . For example, the validity of $\phi \Rightarrow \psi$ does *not* imply that of $\phi \wedge \phi' \Rightarrow \psi$. Conditional logic is usually interpreted in so-called (standard) conditional frames (or selection function frames), that is, tuples (W, f) where W is a set of worlds and $f : W \times \mathcal{P}(W) \rightarrow \mathcal{P}(W)$ is a selection function that assigns a proposition $f(w, A) \subseteq W$ to each world w and condition $A \subseteq W$. In coalgebraic parlance, we understand conditional frames as structures (W, ρ) where $\rho : W \rightarrow (\mathcal{P}(W) \rightarrow \mathcal{P}(W))$ maps each world $w \in W$ to a function $\rho(w) : \mathcal{P}(W) \rightarrow \mathcal{P}(W)$ from conditions to propositions, both formalised as subsets of W . That is, successor structures of worlds are now (selection) functions of type $\mathcal{P}(W) \rightarrow \mathcal{P}(W)$.

In a conditional frame (W, f) , the standard semantics of the conditional operator takes the form

$$w \models \phi \Rightarrow \psi \iff f(w, \llbracket \phi \rrbracket) \subseteq \llbracket \psi \rrbracket.$$

Again, the semantics of the conditional operator can be understood as specifying a property of successor structures, i.e. selection functions: under the coalgebraic reading we have

$$w \models \phi \Rightarrow \psi \iff \rho(w) \in \{f : \mathcal{P}(W) \rightarrow \mathcal{P}(W) \mid f(\llbracket \phi \rrbracket) \subseteq \llbracket \psi \rrbracket\}$$

where $\llbracket \phi \rrbracket$ and $\llbracket \psi \rrbracket$ are again the truth-sets of ϕ and ψ , respectively. We note that, as in the other examples above, the semantics of the conditional operator is embodied by an operation, in this case binary, that maps predicates on the set of worlds to predicates on the set of successor structures, in this case selection functions.

The above examples suggest that a great variety of structurally different modal logics can be studied in a uniform and more abstract setting. Models take the form of coalgebras, i.e. pairs $(W, \rho : W \rightarrow T(W))$ where ρ assigns a successor structure $\rho(w) \in T(W)$ to every world $w \in W$. By varying the notion of successor structure, one captures Kripke frames ($T(W)$ consists of the subsets of W), probabilistic transition systems ($T(W)$ are the probability distributions over W) and conditional frames (where $T(W)$ are the selection functions over W).

In all cases, the link to the semantics of modal languages is provided by the abstract notion of *predicate liftings*, i.e. operations $\mathcal{P}(W) \rightarrow \mathcal{P}(TW)$ that single out properties of successor structures in terms of predicates on the state set. This is the starting point of *coalgebraic modal logic*: the study of abstract properties of modal logics based on a coalgebraic formulation in terms of predicate liftings. Rather than dealing with modal languages and their semantics on a case-by-case basis, the goal is to find easy-to-check general *coherence conditions* between a modal logic and its semantics that guarantee properties like completeness or decidability. In other words, the theory abstracts from the particular definition of successor structures and liftings. Results on particular logics can then be obtained by checking these coherence conditions for a specific instance of the general theory. The meta-theory of coalgebraic modal logic in this sense has been expanding rapidly in recent years; prominent results include the Hennessy-Milner property [33], bisimulation-somewhere-else [18], and generic decidability and complexity criteria [34, 36].

It has been shown in [34] that every coalgebraic modal logic, when interpreted over *all* coalgebras of a given type, can be axiomatized by formulas of rank 1, i.e. with nesting depth of modal operators uniformly equal to 1 (logics of arbitrary rank are obtained by restricting the relevant class of coalgebras); such axioms may be regarded as concerning precisely the single next transition step. Here, we establish the converse: given a modal logic \mathcal{L} of rank 1, we construct a functor $M_{\mathcal{L}}$ that provides a sound and *strongly* complete semantics for \mathcal{L} ; i.e. *coalgebraic modal logic subsumes all rank-1 modal logics*. The functor $M_{\mathcal{L}}$ can be viewed as a generalization of the neighbourhood frame functor, so that in traditional terms, our results imply in particular that every rank-1 modal logic is strongly complete for the associated class of neighbourhood frames. The semantics over $M_{\mathcal{L}}$ is moreover canonical in a precise categorical sense: it is final among all possible coalgebraic semantics of \mathcal{L} , i.e. we obtain an adjunction between modal syntax and semantics. A finitary modification of $M_{\mathcal{L}}$ provides a canonical finitely branching semantics, which necessarily (i.e. due to finite branching) fails to be strongly complete, but is still weakly complete and moreover, unlike $M_{\mathcal{L}}$ itself, satisfies the Hennessy-Milner property, which states that logically indistinguishable states are behaviourally equivalent.

Besides rounding off the picture in a pleasant way, these results make the extensive generic machinery of coalgebraic modal logic applicable to arbitrary rank-1 modal logics, even when the latter are given purely syntactically or equipped with a semantics that fails to be, or has not yet been recognized as, coalgebraic. For instance, we obtain a purely syntactic version of the decidability criterion of [34]. As an extended example,

we discuss applications of these results to recently defined variants of deontic logic [11]. In particular, our results immediately imply that these logics are strongly complete, while only weak completeness is proved in loc. cit.

The material is organised as follows. We recall the fundamental concepts of coalgebraic modal logic in Section 1. We then describe the construction of the canonical semantics and its finitely branching variant in Sections 2 and 3, and in Section 4 lay out how these constructions form part of an adjunction between modal syntax and coalgebraic semantics. Applications to decidability and complexity issues are discussed in Section 5, and the extended example is presented in Section 6. We have taken care to explain the required basic categorical terminology at first use, so that Sections 1, 5, 6, and most of Section 2 should be accessible also to readers without previous knowledge of category theory. Section 4 does assume a more extensive categorical background, but is not required for the understanding of the following sections. This work is an extended version of [38].

1 Coalgebraic Modal Logic

We briefly recapitulate the basics of the coalgebraic semantics of modal logic. Coalgebraic modal logic in the form considered here has been introduced in [27], generalising previous frameworks [15, 30, 19, 25]. We work in an extended setting with polyadic modal operators [33], found e.g. in conditional and default logics.

A (*modal*) *similarity type* is a set Λ of modal operators with associated finite arities. The set $\mathcal{F}(\Lambda)$ of Λ -formulas ϕ is defined by the grammar

$$\phi ::= \perp \mid \phi_1 \wedge \phi_2 \mid \neg\phi \mid L(\phi_1, \dots, \phi_n),$$

where L ranges over all modalities in Λ and n is the arity of L . Other Boolean operations are defined as usual; propositional atoms can be expressed as nullary modalities.

Remark 1 Admitting modal operators of arity greater than 1 causes some notational overhead, but otherwise no actual additional technical problems. Many modal logics of interest have polyadic operators, including conditional logic (Example 12.5 below), Presburger modal logic [8], and some forms of probabilistic modal logic [9].

Generally, we denote the set of propositional formulas over a set V by $\text{Prop}(V)$, generated by the basic connectives \neg and \wedge . Moreover, we denote by $\Lambda(V)$ the set $\{L(a_1, \dots, a_n) \mid L \in \Lambda \text{ } n\text{-ary}, a_1, \dots, a_n \in V\}$. A *literal* over V is a formula of the form either a or $\neg a$, with $a \in V$. A (*conjunctive*) *clause* is a finite, possibly empty, disjunction (conjunction) of literals. Although we regard clauses as formulas rather than sets of literals, we shall sometimes use terminology such as ‘a literal is contained in a clause’ or ‘a clause contains another’, with the obvious meaning.

If $V \subseteq \mathcal{F}(\Lambda)$, we also regard propositional formulas over V as Λ -formulas. A *Z-substitution* for a set V is a map $V \rightarrow Z$ into some set Z ; for a formula ϕ over V , we call the result $\phi\sigma$ of applying the substitution σ to ϕ a *Z-instance* of ϕ . If $\Phi \subseteq \text{Prop}(V)$ and $\psi \in \text{Prop}(V)$, we say that ψ is a *propositional consequence* of Φ , written $\Phi \vdash_{PL} \psi$, if there are finitely many $\phi_1, \dots, \phi_n \in \Phi$ such that $\phi_1 \wedge \dots \wedge \phi_n \rightarrow \psi$ is a propositional tautology over V , regarded as a set of atoms; e.g. $\{\Box a, \Box a \rightarrow \Box b\} \vdash_{PL} \Box b$.

Definition 2 (Rank-1 logic) A *one-step rule* R over a set V of propositional variables is a rule ϕ/ψ , where $\phi \in \text{Prop}(V)$ and $\psi \in \text{Prop}(\Lambda(V))$. In an *extended one-step rule* ϕ/ψ , we more generally allow $\psi \in \text{Prop}(\Lambda(\text{Prop}(V)))$. We will refer to extended one-step rules just as *rules* when this is unlikely to cause confusion. An *axiom* or a *one-step formula* is an extended one-step rule with empty premise, i.e. an element of $\text{Prop}(\Lambda(\text{Prop}(V)))$. A *rank-1 (modal) logic* is a pair $\mathcal{L} = (\Lambda, \mathbf{R})$ consisting of a similarity type Λ and a set \mathbf{R} of extended one-step rules. We identify rules up to injective renaming of variables; e.g. if $a/\Box a$ is in \mathbf{R} , then $b/\Box b$ is also in \mathbf{R} and denotes the same rule.

Note that the definition of one-step formula rules out axioms involving propositional variables at the top level, such as $\Box a \rightarrow a$; first results concerning these more general logics have appeared in [28]. The archetypal rank-1 logic is K (Example 12.1 below), whose main axiom is the one-step formula $\Box(a \rightarrow b) \rightarrow \Box a \rightarrow \Box b$ where one has propositional connectives above and below a single layer of boxes.

Given a rank-1 logic $\mathcal{L} = (\Lambda, \mathbf{R})$, we say that a Λ -formula ϕ is \mathcal{L} -*derivable*, and write $\vdash_{\mathcal{L}} \phi$, if ϕ is derivable by the following rules:

$$\begin{aligned} (P) \quad & \frac{\phi_1; \dots; \phi_n}{\psi} (\{\phi_1, \dots, \phi_n\} \vdash_{PL} \psi) \\ (R) \quad & \frac{\phi\sigma}{\psi\sigma} (\phi/\psi \in \mathbf{R}; \sigma \text{ an } \mathcal{F}(\Lambda)\text{-substitution}) \\ (C) \quad & \frac{\phi_1 \leftrightarrow \psi_1; \dots; \phi_n \leftrightarrow \psi_n}{L(\phi_1, \dots, \phi_n) \leftrightarrow L(\psi_1, \dots, \psi_n)} (L \in \Lambda \text{ } n\text{-ary}). \end{aligned}$$

The last rule above is referred to as the *congruence rule*. For a set $\Phi \subseteq \mathcal{F}(\Lambda)$, we write $\Phi \vdash_{\mathcal{L}} \psi$ if $\vdash_{\mathcal{L}} \phi_1 \wedge \dots \wedge \phi_n \rightarrow \psi$ for suitable $\phi_1, \dots, \phi_n \in \Phi$. Of course, the premise of rule (R) is vacuous in the case of axioms. The propositional reasoning rule (P) may be replaced by the more traditional combination of modus ponens and introduction of substitution instances of propositional tautologies.

It has been shown that axioms, one-step rules, and extended one-step rules may equivalently replace each other [34]. We shall thus on occasion assume that a given rank-1 logic is presented only in terms of one-step rules or only in terms of axioms when convenient. In more detail, extended one-step rules are by definition more general than axioms, and in turn may trivially be replaced by equivalent one-step rules in which propositional formulas under modal operators in the conclusion are abbreviated as single propositional variables by introducing suitable additional premises. The non-trivial step is to replace one-step rules by equivalent axioms; this is discussed in detail in [34]. A typical (and easy) example is the monotonicity rule $a \rightarrow b/\Box a \rightarrow \Box b$, which may be replaced by the axiom $\Box(a \wedge b) \rightarrow \Box a$.

Remark 3 We can always assume that every propositional variable a appearing in the premise ϕ of a one-step rule appears also in the conclusion: otherwise, we can eliminate a by passing from ϕ to $\phi[\top/a] \vee \phi[\perp/a]$.

Coalgebraic modal logic interprets modal formulas over coalgebras, which abstract from concrete notions of reactive system:

Definition 4 (Coalgebras) Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a functor, referred to as the *signature functor*, where \mathbf{Set} is the category of sets. Explicitly, T maps sets X to sets TX , and maps $f : X \rightarrow Y$ to maps $Tf : TX \rightarrow TY$, preserving identities and composition. A T -coalgebra is a pair $C = (X, \xi)$ where X is a set (of *states*) and ξ is a function $X \rightarrow TX$ called the *transition function*. A *morphism* $(X_1, \xi_1) \rightarrow (X_2, \xi_2)$ of T -coalgebras is a map $f : X_1 \rightarrow X_2$ such that $\xi_2 \circ f = Tf \circ \xi_1$. States x, y in coalgebras C, D are *behaviourally equivalent* if there exist coalgebra morphisms $f : C \rightarrow E$ and $g : D \rightarrow E$ such that $f(x) = g(y)$.

We view coalgebras as generalised transition systems: the transition function maps states to structured sets of successors and observations, with the structure determined by the signature functor; correspondingly, we refer to elements of TX as *successor structures*.

Example 5 We give examples of functors and the system types modelled by their coalgebras.

1. *Kripke frames*: The *covariant powerset functor* $\mathcal{P} : \mathbf{Set} \rightarrow \mathbf{Set}$ takes a set X to its powerset $\mathcal{P}(X)$, and for $f : X \rightarrow Y$, the map $\mathcal{P}(f) : \mathcal{P}(X) \rightarrow \mathcal{P}(Y)$ takes direct images. Coalgebras $\xi : X \rightarrow \mathcal{P}(X)$ for $\mathcal{P}(X)$ can be identified with Kripke frames (X, R) by putting $R = \{(x, y) \mid y \in \xi(x)\}$. Morphisms of \mathcal{P} -coalgebras are precisely bounded morphisms of Kripke frames. The *finite powerset functor* \mathcal{P}^{fin} takes a set X to the set $\mathcal{P}^{fin}(X)$ of finite subsets of X ; \mathcal{P}^{fin} -coalgebras are finitely branching Kripke frames.

2. *Neighbourhood frames*: The covariant powerset functor is distinguished from the *contravariant powerset functor* $\mathcal{Q} : \mathbf{Set}^{op} \rightarrow \mathbf{Set}$, which also maps a set X to its powerset, denoted $\mathcal{Q}(X)$ in this case, but which takes a map $f : X \rightarrow Y$, i.e. a morphism $Y \rightarrow X$ in \mathbf{Set}^{op} , to the map $\mathcal{Q}(f) : \mathcal{Q}(Y) \rightarrow \mathcal{Q}(X)$ which takes preimages. (Generally, contravariant functors, indicated by the notation $\mathbf{Set}^{op} \rightarrow \mathbf{Set}$, are like functors, but reverse the direction of maps.) Being contravariant, \mathcal{Q} does not itself have coalgebras according to Definition 4. However, composing \mathcal{Q} with itself, we obtain the (covariant) *neighbourhood functor*

$$\mathbf{N} = \mathcal{Q} \circ \mathcal{Q}^{op} : \mathbf{Set} \rightarrow \mathbf{Set}.$$

Coalgebras $\xi : X \rightarrow \mathbf{N}(X)$ for \mathbf{N} are precisely neighbourhood frames: for each state x , ξ determines a set $\xi(x) \subseteq \mathcal{Q}(X) = \mathcal{P}(X)$ of neighbourhoods. Note that while the action of the neighbourhood functor on objects is identical to that of $\mathcal{P} \circ \mathcal{P}$, the action on morphisms differs. This impacts on the definition of modalities, where it turns out that \mathbf{N} is the right choice for modelling neighbourhood semantics.

3. *Multigraphs*: The *finite multiset functor* \mathbf{B} takes a set X to the set of maps $B : X \rightarrow \mathbb{N}$ with finite support, where $B(x) = n$ is read ‘multiset B contains x with multiplicity n ’. We extend B to an integer-valued measure on $\mathcal{P}(X)$ by putting $B(A) = \sum_{x \in A} B(x)$. The action of \mathbf{B} on morphisms is then defined by $\mathbf{B}(f)(y) = B(f^{-1}[\{y\}])$. Coalgebras $\xi : X \rightarrow \mathbf{B}(X)$ for \mathbf{B} are *multigraphs* [7], i.e. directed graphs with \mathbb{N} -weighted edges: the transition weight from x to y in X is $\xi(x)(y)$.

4. *Markov chains*: The finite distribution functor \mathbf{D} takes a set X to the set $\mathbf{D}(X)$ of finitely supported probability distributions over X ; for $f : X \rightarrow Y$, $\mathbf{D}(f) : \mathbf{D}(X) \rightarrow$

$D(Y)$ takes image measures. Coalgebras $\xi : X \rightarrow D(X)$ are Markov chains: for $x, y \in X$, $\xi(x)(\{y\})$ is the transition probability from x to y .

5. *Conditional frames:* The *conditional frame functor* Cf is defined by $\text{Cf}(X) = \mathcal{Q}(X) \rightarrow \mathcal{P}(X)$, where \rightarrow denotes function space. Note that the negative occurrence of the contravariant powerset functor \mathcal{Q} makes the overall action of Cf on functions covariant. Coalgebras $\xi : X \rightarrow \text{Cf}(X)$ associate to each $x \in X$ and each $A \subseteq X$ a set $\xi(x)(A) \subseteq X$; i.e. they are precisely the *conditional frames* appearing in the semantics of conditional logic [5], also known as *selection function models*.

Assumption 6 We can assume w.l.o.g. that $T : \text{Set} \rightarrow \text{Set}$ preserves injective maps ([2], proof of Theorem 3.2). For convenience, we will in fact sometimes assume that $TX \subseteq TY$ if $X \subseteq Y$. Moreover, we assume that T is non-trivial, i.e. $TX = \emptyset \implies X = \emptyset$ (otherwise, $TX = \emptyset$ for all X).

While the type of system underlying the semantics is encapsulated in the chosen signature functor, the interpretation of modalities is given by a choice of predicate liftings, so named because they lift predicates on the state space to predicates on the set of successor structures:

Definition 7 (Predicate liftings) An *n-ary predicate lifting* for T is a natural transformation

$$\lambda : \mathcal{Q}^n \rightarrow \mathcal{Q} \circ T^{\text{op}},$$

where $\mathcal{Q} : \text{Set}^{\text{op}} \rightarrow \text{Set}$ is the contravariant powerset functor as before, and \mathcal{Q}^n denotes its (pointwise) n -th power, i.e. $\mathcal{Q}^n(X) = (\mathcal{Q}(X))^n$. Moreover, $T^{\text{op}} : \text{Set}^{\text{op}} \rightarrow \text{Set}^{\text{op}}$ is the dual of T . Explicitly, λ is a family of maps $\lambda_X : \mathcal{Q}(X)^n \rightarrow \mathcal{Q}(TX)$, indexed over all sets X , satisfying the *naturality* equation

$$\lambda_X(f^{-1}[A_1], \dots, f^{-1}[A_n]) = (Tf)^{-1}[\lambda_Y(A_1, \dots, A_n)]$$

for all $f : X \rightarrow Y$ and all $A_1, \dots, A_n \subseteq Y$.

Example 8 The basic example of a predicate lifting is the following unary lifting for the covariant powerset functor \mathcal{P} inducing the standard box modality on Kripke frames (Example 12.1 below): for a set X , let $\lambda_X : \mathcal{Q}(X) \rightarrow \mathcal{Q}(\mathcal{P}(X))$ be the map defined by

$$\lambda_X(A) = \{B \in \mathcal{P}(X) \mid B \subseteq A\};$$

i.e. λ_X lifts a predicate A on X to the predicate $\lambda_X(A)$ on $\mathcal{P}(X)$ satisfied by $B \in \mathcal{P}(X)$ iff $B \subseteq A$. The family (λ_X) is a predicate lifting for \mathcal{P} : the naturality equation translates into the set-theoretic fact that for $f : X \rightarrow Y$, $B \in \mathcal{P}(X)$, and $A \subseteq Y$,

$$B \subseteq f^{-1}[A] \iff \mathcal{P}(f)(B) = f[B] \subseteq A.$$

More examples of predicate liftings are found in Example 12 below.

The coalgebraic semantics of modal logics is now defined as follows. Given a similarity type Λ , a Λ -*structure* $\mathcal{M} = (T, (\llbracket L \rrbracket^{\mathcal{M}})_{L \in \Lambda})$ consists of a signature functor T and an assignment of an n -ary predicate lifting $\llbracket L \rrbracket^{\mathcal{M}}$ for T to every modal operator $L \in \Lambda$ of arity n . We say that \mathcal{M} is *based on* T , or that T is the *underlying functor* of \mathcal{M} .

If there is no danger of confusion, we drop the superscripts on the liftings. Given a Λ -structure \mathcal{M} , the satisfaction relation $\models_C^{\mathcal{M}}$ between states x of T -coalgebras $C = (X, \xi)$ and Λ -formulas is defined inductively, with the usual clauses for the Boolean operations. The clause for an n -ary modal operator L is

$$x \models_C^{\mathcal{M}} L(\phi_1, \dots, \phi_n) \iff \xi(x) \in \llbracket L \rrbracket^{\mathcal{M}}(\llbracket \phi_1 \rrbracket_C^{\mathcal{M}}, \dots, \llbracket \phi_n \rrbracket_C^{\mathcal{M}})$$

where $\llbracket \phi \rrbracket_C^{\mathcal{M}} = \{x \in X \mid x \models_C^{\mathcal{M}} \phi\}$. Again, we drop sub- and superscripts if they are clear from the context. When we speak of a *coalgebraic modal logic* informally, we formally refer to a Λ -structure.

Remark 9 The reason for restricting the exposition to rank-1 logics is that we are going to construct a coalgebraic semantics for a given logic in which the logic is interpreted over *all* coalgebras for a given functor; it has been shown that all such logics can be completely axiomatised in rank-1 [34]. Extensions of our results to logics outside rank-1, in particular employing axioms with nested modalities, are the subject of ongoing [28] and future work; they will necessarily concern completeness over suitable restricted classes of coalgebras.

Remark 10 Note that a nullary predicate lifting for T is just a family of subsets $\lambda_X \subseteq TX$ for all sets X such that $\lambda_X = (Tf)^{-1}[\lambda_Y]$ for all maps $f : X \rightarrow Y$. If $L \in \Lambda$ is a nullary modal operator, then a state x in a coalgebra (X, ξ) satisfies the formula L iff $\xi(x) \in \llbracket L \rrbracket_X$. E.g. we can interpret a nullary deadlock operator δ over the covariant powerset functor \mathcal{P} by $\llbracket \delta \rrbracket_X = \{\emptyset\} \subseteq \mathcal{P}(X)$, and then a state x in a $\mathcal{P}(X)$ -coalgebra (X, ξ) satisfies δ iff $\xi(x) = \emptyset$, i.e. iff x does not have any successors in the corresponding Kripke frame.

The next definition introduces the semantic consequence relations that we are using in the remainder of the paper.

Definition 11 (Soundness and completeness) Given a Λ -structure \mathcal{M} based on T , a formula $\phi \in \mathcal{F}(\Lambda)$ is a *local semantic consequence* of a set $\Phi \subseteq \mathcal{F}(\Lambda)$, written $\Phi \models_{\mathcal{M}} \psi$, if, for every state x in every T -coalgebra, $x \models \psi$ whenever $x \models \Phi$ (i.e. $x \models \phi$ for all $\phi \in \Phi$). The logic \mathcal{L} is *sound* for \mathcal{M} if $\Phi \models_{\mathcal{M}} \psi$ whenever $\Phi \vdash_{\mathcal{L}} \psi$, *strongly complete* if $\Phi \vdash_{\mathcal{L}} \psi$ whenever $\Phi \models_{\mathcal{M}} \psi$, and *weakly complete* if $\vdash_{\mathcal{L}} \psi$ whenever $\emptyset \models_{\mathcal{M}} \psi$.

Example 12 We give a brief description of some coalgebraic modal logics, illustrating in particular the fact that many interesting modal logics are axiomatised in rank 1.

1. The modal logic K has a single unary modal operator \Box and rank-1 axioms $\Box\top$ and $\Box(a \rightarrow b) \rightarrow \Box a \rightarrow \Box b$. Note that $\Box\top$ together with the congruence rule entails the usual necessitation rule. The standard Kripke semantics of K is obtained as the K -structure based on the covariant powerset functor \mathcal{P} that interprets \Box by the predicate lifting of Example 8, i.e.

$$\llbracket \Box \rrbracket_X(A) = \mathcal{P}(A) \subset \mathcal{P}(X)$$

for $A \subseteq X$. Indeed, under the correspondence between \mathcal{P} -coalgebras and Kripke frames described in Example 5.1, the coalgebraic semantics for \Box is transformed into the usual semantic clause

$$x \models_{(X,R)} \Box \phi \quad \text{iff} \quad \forall y \in X. xRy \implies y \models_{(X,R)} \phi.$$

2. The neighbourhood semantics of the minimal modal logic E , i.e. the modal logic with a single operator \Box and no rules or axioms (except replacement of equivalents, i.e. the congruence rule) is captured as an E -structure based on the neighbourhood functor \mathbf{N} (Example 5.2) by

$$\llbracket \Box \rrbracket_X(A) = \{\mathfrak{A} \in \mathbf{N}(X) \mid A \in \mathfrak{A}\}.$$

3. *Graded modal logic* (GML) [10] has operators \diamond_k for $k \in \mathbb{N}$ of the nature ‘in more than k successor states, it is the case that’. GML has originally been interpreted over Kripke frames by just counting successor states. This semantics fails to be coalgebraic, as it violates the naturality equation: for $f : X \rightarrow Y$, $A \subseteq Y$, and $B \in \mathcal{P}(X)$, $\#(B \cap f^{-1}[A]) > k$ is not in general equivalent to $\#(f[B] \cap A) > k$. Following [7], we may however equip GML with a coalgebraic semantics in the shape of a GML-structure based on the finite multiset functor \mathbf{B} (Example 5.3); the interpretation of the modal operators is then defined by

$$\llbracket \diamond_k \rrbracket_X(A) = \{B \in \mathbf{B}(X) \mid B(A) > k\}$$

This satisfies the naturality equation because $\mathbf{B}f(B)(A) = B(f^{-1}[A])$. This semantics induces the same notion of satisfiability as the original Kripke semantics [34]. The following rank-1 axiomatisation of GML has been given in [10]; it uses the abbreviations $\Box_k \phi := \neg \diamond_k \neg \phi$ and $\diamond_{-1} \phi := \top$:

$$\begin{aligned} \Box_0(a \rightarrow b) &\rightarrow \Box_0 a \rightarrow \Box_0 b & \diamond_k a &\rightarrow \diamond_l a \quad (k > l) \\ \diamond_k a &\leftrightarrow \bigvee_{i=0}^{k+1} (\diamond_{i-1}(a \wedge b) \vee \diamond_{k-i}(a \wedge \neg b)) & \Box_0(a \rightarrow b) &\rightarrow \diamond_k a \rightarrow \diamond_k b \end{aligned}$$

4. *Probabilistic modal logic* (PML) [21, 13] has modal operators L_p for $p \in [0, 1] \cap \mathbb{Q}$, read ‘in the next step, it is with probability at least p the case that’. A PML-structure based on the finite distribution functor \mathbf{D} (Example 5.4) is defined by

$$\llbracket L_p \rrbracket_X(A) = \{P \in \mathbf{D}(X) \mid P(A) \geq p\}.$$

A complete rank-1 axiomatisation of PML has been given in [13].

5. *Conditional logic* [5] has a binary infix modal operator \Rightarrow , with $\phi \Rightarrow \psi$ read e.g. ‘if ϕ , then normally ψ ’ (further readings include e.g. relevant implication); i.e. \Rightarrow is a non-monotonic conditional. The conditional logic CK has axioms $a \Rightarrow \top$ and $(a \Rightarrow (b \rightarrow c)) \rightarrow (a \Rightarrow b) \rightarrow (a \Rightarrow c)$, i.e. behaves like K in the second variable and like E in the first. A CK -structure based on \mathbf{Cf} is defined by

$$\llbracket \Rightarrow \rrbracket_X(A, B) = \{f \in \mathbf{Cf}(X) \mid f(A) \subseteq B\}.$$

This induces precisely the conditional frame semantics described in [5].

In this work, we expand on the fact that there may be different structures for a given similarity type (and indeed for a given logic, as defined further below). We compare structures by means of morphisms: A *morphism* $\mu : \mathcal{N} \rightarrow \mathcal{M}$ between two Λ -structures $\mathcal{N} = (S, (\llbracket L \rrbracket^{\mathcal{N}})_{L \in \Lambda})$, $\mathcal{M} = (T, (\llbracket L \rrbracket^{\mathcal{M}})_{L \in \Lambda})$ is a natural transformation $\mu : S \rightarrow T$ that commutes with predicate liftings, i.e. $\llbracket L \rrbracket^{\mathcal{N}} = \mathcal{Q}\mu \circ \llbracket L \rrbracket^{\mathcal{M}}$ for all $L \in \Lambda$

(explicitly, $\llbracket L \rrbracket_X^{\mathcal{N}}(A_1, \dots, A_n) = \mu_X^{-1}[\llbracket L \rrbracket_X^{\mathcal{M}}(A_1, \dots, A_n)]$ for L n -ary, $A_1, \dots, A_n \subseteq X$):

$$\begin{array}{ccccc}
 \mathcal{Q}^n(X) & \xrightarrow{\llbracket L \rrbracket_X^{\mathcal{M}}} & \mathcal{Q}(TX) & & TX \\
 & \searrow \llbracket L \rrbracket_X^{\mathcal{N}} & \swarrow \mathcal{Q}(\mu_x) & & \nearrow \mu_x \\
 & & \mathcal{Q}(SX) & \xrightarrow{\mu_x} & SX
 \end{array}$$

A morphism $\mu : \mathcal{N} \rightarrow \mathcal{M}$ induces a translation of S -coalgebras $C = (X, \xi)$ into T -coalgebras

$$\mu(C) = (X, \mu_X \circ \xi).$$

It is easy to see that the semantics of formulas is invariant under morphisms of structures:

Lemma 13 Let $\mu : \mathcal{N} \rightarrow \mathcal{M}$ be a morphism of Λ -structures $\mathcal{N} = (S, (\llbracket L \rrbracket^{\mathcal{N}})_{L \in \Lambda})$ and $\mathcal{M} = (T, (\llbracket L \rrbracket^{\mathcal{M}})_{L \in \Lambda})$, and let $C = (X, \xi)$ be an S -coalgebra. Then

$$\llbracket \phi \rrbracket_C^{\mathcal{N}} = \llbracket \phi \rrbracket_{\mu(C)}^{\mathcal{M}}$$

for all $\phi \in \mathcal{F}(\Lambda)$. □

Related to these notions is the concept of *substructure*: recall that a functor S is a *subfunctor* of a functor T if $SX \subseteq TX$ for all sets X , and for every map $f : X \rightarrow Y$, Sf is the restriction of $Tf : TX \rightarrow TY$ to a map $SX \rightarrow SY$; in other words, S consists in a choice of subsets $SX \subseteq TX$ such that $Tf[SX] \subseteq SY$ for all $f : X \rightarrow Y$. If $\mathcal{M} = (T, (\llbracket L \rrbracket^{\mathcal{M}})_{L \in \Lambda})$ is a Λ -structure, then every subfunctor S of T can be uniquely completed to a Λ -structure $\mathcal{N} = (S, (\llbracket L \rrbracket^{\mathcal{N}})_{L \in \Lambda})$ in such a way that the inclusion natural transformation $S \rightarrow T$ becomes a morphism $\mathcal{N} \rightarrow \mathcal{M}$, namely by putting $\llbracket L \rrbracket_X^{\mathcal{N}}(A_1, \dots, A_n) = SX \cap \llbracket L \rrbracket_X^{\mathcal{M}}(A_1, \dots, A_n)$ for $L \in \Lambda$ n -ary and $A_1, \dots, A_n \subseteq X$; we call \mathcal{N} the *substructure of \mathcal{M} induced by S* .

We note that every Λ -structure, via the associated predicate liftings, gives rise to a translation in the above sense between T -coalgebras and neighbourhood frames.

Lemma and Definition 14 (Neighbourhood structure, canonical translation)

Given a similarity type Λ , the Λ -structure \mathcal{N}_Λ consisting of the endofunctor

$$\mathbf{N}_\Lambda = \prod_{L \in \Lambda \text{ } n\text{-ary}} \mathcal{Q} \circ (\mathcal{Q}^{\text{op}})^n$$

(where $\mathcal{Q} : \text{Set}^{\text{op}} \rightarrow \text{Set}$ is contravariant powerset) and the liftings

$$\llbracket L \rrbracket_X^{\mathcal{N}_\Lambda}(A_1, \dots, A_n) = \{\mathfrak{A}_L \in \mathbf{N}_\Lambda(X) \mid (A_1, \dots, A_n) \in \mathfrak{A}_L\}$$

is the *neighbourhood structure of Λ* . For every Λ -structure \mathcal{M} based on T , the *canonical translation* defined by

$$\begin{array}{ccc}
 \mu_X : T(X) & \rightarrow & \mathbf{N}_\Lambda(X) \\
 t & \mapsto & (\{(A_1, \dots, A_n) \in \mathcal{Q}(X)^n \mid t \in \llbracket L \rrbracket_X^{\mathcal{M}}(A_1, \dots, A_n)\})_{L \in \Lambda \text{ } n\text{-ary}}
 \end{array}$$

is a morphism $\mu : \mathcal{M} \rightarrow \mathcal{N}_\Lambda$ of structures.

Proof Straightforward unravelling of the definitions. \square

By virtue of the neighbourhood frame translation, we can convert T -coalgebras into multi-neighbourhood frames, i.e. coalgebras for products of n -ary neighbourhood functors $\mathcal{Q} \circ (\mathcal{Q}^{\text{op}})^n$; by Lemma 13, this translation is compatible with the interpretation of Λ -formulas.

We now introduce a single-step version of the logic induced by a Λ -structure, the *one-step logic*, which is characterised syntactically by excluding nested modal operators, and semantically does not involve state transitions in the sense that it speaks about single successor structures rather than entire coalgebras. Below, we introduce notions of formula and satisfaction for the one-step logic; we shall later consider variants of these notions.

Definition 15 (One-step logic) Given a set X , the set of *one-step formulas over X* is the set $\text{Prop}(\Lambda(\mathcal{P}(X)))$. By means of the interpretation map $\llbracket - \rrbracket : \text{Prop}(\mathcal{P}(X)) \rightarrow \mathcal{P}(X)$ that assigns to $\phi \in \text{Prop}(\mathcal{P}(X))$ its extension $\llbracket \phi \rrbracket \in \mathcal{P}(X)$ arising from the Boolean algebra structure of $\mathcal{P}(X)$, we identify $\text{Prop}(\Lambda(\text{Prop}(\mathcal{P}(X))))$ with $\text{Prop}(\Lambda(\mathcal{P}(X)))$, i.e. we regard also elements of $\text{Prop}(\Lambda(\text{Prop}(\mathcal{P}(X))))$ as one-step formulas over X , but immediately evaluate the inner propositional layer. For a one-step formula ψ over X , we define an interpretation $\llbracket \psi \rrbracket \subseteq TX$ by extending the interpretation $\llbracket L(\phi_1, \dots, \phi_n) \rrbracket = \llbracket L \rrbracket(\llbracket \phi_1 \rrbracket, \dots, \llbracket \phi_n \rrbracket)$ of $\Lambda(\mathcal{P}(X))$ to $\text{Prop}(\Lambda(\mathcal{P}(X)))$ according to the Boolean algebra structure of $\mathcal{P}(TX)$. We then write $X \models \phi$ if $\llbracket \phi \rrbracket = X$, and $TX \models \psi$ if $\llbracket \psi \rrbracket = TX$. For $t \in TX$, we say that t *satisfies* ψ , and write $t \models^X \psi$, if $t \in \llbracket \psi \rrbracket$. The *one-step theory of t* is the set $\{\psi \in \text{Prop}(\Lambda(\mathcal{P}(X))) \mid t \models^X \psi\}$. We say that ψ is *(one-step) satisfiable* if $\llbracket \psi \rrbracket \neq \emptyset$, i.e. if there exists $t \in TX$ such that $t \models^X \psi$.

Note that one-step formulas over X mix syntax and semantics by treating the inner propositional layer semantically while the modal layer and the outer propositional layer are treated syntactically; this will be convenient in the construction of canonical structures below.

The requirement that axioms are of rank 1 means that every axiom makes assertions precisely about the next transition step. This allows us to capture soundness as a property exhibited in a single transition step as follows.

Definition 16 (One-step soundness, \mathcal{L} -structures) Given a set X , a $\mathcal{P}(X)$ -*valuation* for V is just a $\mathcal{P}(X)$ -substitution for V , i.e. a map $V \rightarrow \mathcal{P}(X)$. An extended one-step rule ϕ/ψ is *one-step sound* for a Λ -structure \mathcal{M} based on T if $TX \models \psi\tau$ for each set X and each $\mathcal{P}(X)$ -valuation τ such that $X \models \phi\tau$ (of course, the latter condition is void for axioms). An \mathcal{L} -*structure* for a rank-1 logic $\mathcal{L} = (\Lambda, \mathbf{R})$ is a Λ -structure for which all rules in \mathbf{R} are one-step sound.

It is easy to see that one-step soundness implies soundness, i.e.

Proposition 17 A rank-1 logic \mathcal{L} is sound for all \mathcal{L} -structures.

Additional conditions guarantee weak completeness; see Section 2. In general, this is all one can hope for, as many coalgebraic modal logics fail to be compact [34]. However, it will turn out that \mathcal{L} is indeed *strongly* complete for the canonical \mathcal{L} -structure constructed below.

Example 18 We give some examples of one-step sound and unsound rules for the structures in the running examples.

1. *Kripke frames* (Example 12.1) The rule

$$\frac{a \wedge b \rightarrow c}{\Box a \wedge \Box b \rightarrow \Box c}$$

is one-step sound: let X be a set and let τ be a $\mathcal{P}(X)$ -valuation such that $X \models ((a \wedge b) \rightarrow c)\tau$, i.e. $\tau(a) \cap \tau(b) \subseteq \tau(c)$. We have to show that $\mathcal{P}(X) \models (\Box a \wedge \Box b \rightarrow \Box c)\tau$, i.e. that $\llbracket \Box \rrbracket_X \tau(a) \cap \llbracket \Box \rrbracket_X \tau(b) \subseteq \llbracket \Box \rrbracket_X \tau(c)$. This is clear: if $A \in \mathcal{P}(X)$ satisfies $A \subseteq \tau(a)$ and $A \subseteq \tau(b)$, then $A \subseteq \tau(a) \cap \tau(b) \subseteq \tau(c)$.

Contrastingly, the rule $\neg a / \neg \Box a$ fails to be one-step sound: if $\tau(a) = \emptyset$, then $X \models (\neg a)\tau$, but not $\mathcal{P}(X) \models (\neg \Box a)\tau$, as $\llbracket \neg \Box a \rrbracket \tau = X - \llbracket \Box \rrbracket_X \tau(a) = X - \{\emptyset\}$.

2. *Multigraphs* (Example 12.3) The rule $\phi / \psi =$

$$\frac{\neg(a \wedge b) \wedge (a \vee b \rightarrow c)}{\Diamond_k a \wedge \Diamond_l b \rightarrow \Diamond_{k+l+1} c}$$

is one-step sound: let X be a set and let τ be a $\mathcal{P}(X)$ -valuation such that $X \models \phi\tau$, i.e. $\tau(a) \cap \tau(b) = \emptyset$ and $\tau(a) \cup \tau(b) \subseteq \tau(c)$. Let $B \in \mathbf{B}(X)$ be a multiset such that $B \models (\Diamond_k a \wedge \Diamond_l b)\tau$, i.e. $B(\tau(a)) > k$ and $B(\tau(b)) > l$. Since $\tau(a)$ and $\tau(b)$ are disjoint, it follows that $B(\tau(c)) > k + l + 1$, i.e. $B \models (\Diamond_{k+l+1} c)\tau$ as required.

An obviously one-step unsound variant of the above rule is the rule

$$\frac{a \vee b \rightarrow c}{\Diamond_k a \wedge \Diamond_l b \rightarrow \Diamond_{k+l+1} c},$$

for which the above one-step soundness proof fails as $\tau(a)$ and $\tau(b)$ need not be disjoint; a counterexample to one-step soundness is $\tau(a) = \tau(b) = \tau(c) \neq \emptyset$ and $k = l = B(\tau(a)) - 1$.

2 From Rank-1 Logics to Coalgebraic Models

In this section we construct for a given rank-1 modal logic \mathcal{L} a *canonical \mathcal{L} -structure* $\mathcal{M}_{\mathcal{L}}$ for which \mathcal{L} is (sound and) *strongly complete*. (Recall that by results of [34], this result will not generalise directly to non-rank-1 logics, as every structure can be axiomatised in rank-1). Moreover, we consider a finitely branching substructure $\mathcal{M}_{\mathcal{L}}^{\text{fin}}$ of $\mathcal{M}_{\mathcal{L}}$ which is canonical among the finitely branching \mathcal{L} -structures. For $\mathcal{M}_{\mathcal{L}}^{\text{fin}}$, \mathcal{L} is (sound and) weakly complete and has the Hennessy-Milner property, i.e. states satisfying the same formulas are behaviourally equivalent. This tradeoff is typical: the Hennessy-Milner property holds only over finitely branching systems, while strong completeness will fail over such systems due to the breakdown of compactness. For the remainder of the section, we fix a rank-1 modal logic $\mathcal{L} = (\Lambda, \mathbf{R})$.

The construction of the canonical structure resembles the construction of canonical models using maximally consistent sets, but works, like many concepts of coalgebraic modal logic, at the single step level. We begin by fixing a notion of derivability at the one-step level, i.e. for one-step formulas over a set X . Informally speaking, a one-step formula is one-step derivable over X if it can be propositionally derived

from a set of conclusions of rules whose premises hold in X . That is to say, one-step derivations formalise information that becomes apparent after one transition step. Like the syntax of one-step formulas, one-step deduction works semantically on the inner propositional layer and syntactically on the outer layers.

Definition 19 (One-step derivations) Let X be a set. We say that a one-step formula ψ over X , i.e. $\psi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$, is *one-step derivable* (over X), and write $\vdash_{\mathcal{L}}^X \psi$, if ψ is propositionally entailed by conclusions of $\mathcal{P}(X)$ -instances of rules in \mathbf{R} whose premises hold in X , formally: if

$$\Theta \vdash_{PL} \psi,$$

where

$$\Theta = \{\chi\tau \mid \phi/\chi \in \mathbf{R}; \tau \text{ a } \mathcal{P}(X)\text{-valuation}; X \models \phi\tau\}.$$

The condition $X \models \phi\tau$ in the definition of Θ is, of course, vacuous in the case of axioms. Formulas $\rho \in \text{Prop}(\mathcal{P}(X))$ occurring as subformulas in $\chi\tau \in \Theta$ are implicitly interpreted as their extension $\llbracket \rho \rrbracket \in \mathcal{P}(X)$. We say that ψ is *one-step derivable* from $\Phi \subseteq \text{Prop}(\Lambda(\mathcal{P}(X)))$, $\Phi \vdash_{\mathcal{L}}^X \psi$, if there are $\phi_1, \dots, \phi_n \in \Phi$ such that $\vdash_{\mathcal{L}}^X \phi_1 \wedge \dots \wedge \phi_n \rightarrow \psi$. The set Φ is *one-step consistent* if $\Phi \not\vdash_{\mathcal{L}}^X \perp$, and *maximally one-step consistent* if Φ is maximal w.r.t. \subseteq among the one-step consistent subsets of $\text{Prop}(\Lambda(\mathcal{P}(X)))$.

Remark 20 Of course, the notion of one-step derivation may equivalently be presented in the shape of two rules for propositional entailment and the application of \mathbf{R} ; the formulation above emphasises the trivial fact that all one-step derivations have a normal form where applications of \mathbf{R} precede propositional reasoning. There is no rule for congruence (i.e. replacement of equivalents) in the one-step derivation system presented above. Such a rule is unnecessary, as the ‘formulas’ under modal operators are already subsets. Further below (Definition 49), we will consider a variant of one-step deduction that works with propositional variables instead of sets, and hence does include a congruence rule.

Note that unlike in previous definitions [26, 34, 36], we do not require that a one-step derivation of $\psi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ uses only Boolean combinations of sets occurring in ψ . In the cited contexts, i.e. when a coalgebraic semantics of \mathcal{L} is already given, this requirement is essentially justified by Proposition 3.10 of [36]. For purposes of the present work, we need to establish this form of proof normalisation directly. The crucial ingredient is a fact on solvability of Boolean equations, proved using a strategy recently employed in [22], which ultimately relies on Boole’s expansion method [4]:

Lemma 21 Let $\phi_i, \psi_i \in \text{Prop}(V \cup W)$, $i = 1, \dots, n$, where V and W are disjoint finite sets of variables, let A be a Boolean algebra, and let τ be an A -valuation for V . If the system of equations

$$\phi_i\tau = \psi_i\tau \quad (i = 1, \dots, n)$$

is *solvable* in A , i.e. there exists an A -valuation κ for W such that $\phi_i\tau\kappa = \psi_i\tau\kappa$ for all i , then there exists a $\text{Prop}(V)$ -substitution σ for W such that $\phi_i\sigma\tau = \psi_i\sigma\tau$ for all i .

In other words, if a system of Boolean equations with coefficients in a Boolean algebra A is solvable in A , then it is solvable by Boolean combinations of the coefficients.

Proof It is straightforward to reduce to a single equation of the form

$$\phi\tau = \top.$$

We then proceed by induction on the size of W . For $|W| = 0$, there is nothing to prove. For $x \in W$, we apply Boole's expansion method to obtain

$$\begin{aligned} \phi &\equiv (x \rightarrow \phi[\top/x]) \wedge (\neg x \rightarrow \phi[\perp/x]) \\ &\equiv (x \rightarrow \phi[\top/x]) \wedge (\neg\phi[\perp/x] \rightarrow x) \end{aligned} \quad (*)$$

(with \equiv denoting propositionally equivalent transformation steps). In particular, ϕ propositionally entails $\neg\phi[\perp/x] \rightarrow \phi[\top/x]$, and hence the equation

$$(\neg\phi[\perp/x] \rightarrow \phi[\top/x])\tau = \top$$

over $W - \{x\}$ is solvable in A . By induction, this equation is solvable by a $\text{Prop}(V)$ -substitution σ for $W - \{x\}$, and by $(*)$, the $\text{Prop}(V)$ -substitution

$$\sigma[\phi[\top/x]/x]$$

for W solves $\phi\tau = \top$, i.e. satisfies $\phi\sigma\tau = \top$. \square

For reference, we moreover note the following trivial fact, which states that propositional entailment is invariant under exchanging atoms as long as equality of atoms is respected.

Lemma 22 Let $\Phi \subseteq \text{Prop}(V)$, let $\psi \in \text{Prop}(V)$, let σ be a W -substitution, and let τ be a U -substitution such that $\sigma(a) = \sigma(b)$ implies $\tau(a) = \tau(b)$ for all $a, b \in V$. Then $\Phi\sigma \vdash_{PL} \psi\sigma$ implies $\Phi\tau \vdash_{PL} \psi\tau$.

Proof By the assumptions, one can find a U -substitution κ for W such that $\tau = \sigma\kappa$. Let $\Phi\sigma \vdash_{PL} \psi\sigma$. By the substitution lemma of propositional logic, it follows that $\Phi\sigma\kappa \vdash_{PL} \psi\sigma\kappa$, i.e. $\Phi\tau \vdash_{PL} \psi\tau$. \square

Of this fact, we need the following slight variant:

Lemma 23 Let $\Phi \subseteq \text{Prop}(V)$, let $\psi \in \text{Prop}(W)$, and let σ, τ be W -substitutions such that $\sigma(a) = \sigma(b)$ implies $\tau(a) = \tau(b)$ for all $a, b \in V$ and $\sigma(a) = c$ implies $\tau(a) = c$ for all $a \in V$ and all $c \in W$ occurring in ψ . Then $\Phi\sigma \vdash_{PL} \psi$ implies $\Phi\tau \vdash_{PL} \psi$.

Proof Apply the previous lemma to W -substitutions σ', τ' for the disjoint union $V + W$ which extend σ and τ , respectively, by the identity substitution on W . \square

This allows us to establish the desired normalisation result, essentially a subformula property for the one-step logic (whose proof however does not rely on any form of cut elimination):

Proposition 24 Let $\psi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ such that $\vdash_{\mathcal{L}}^X \psi$, and let $\mathfrak{A} \subseteq \mathcal{P}(X)$ be the set of sets occurring in ψ . Then there exists a one-step derivation of ψ that uses only $\text{Prop}(\mathfrak{A})$ -instances of rules, where we identify $\text{Prop}(\mathfrak{A})$ with a subset of $\mathcal{P}(X)$; formally: $\Theta \vdash_{PL} \psi$, where $\Theta = \{\chi\tau \mid \phi/\chi \in \mathfrak{R}; \tau \text{ a } \text{Prop}(\mathfrak{A})\text{-valuation}; X \models \phi\tau\}$.

Proof By the results of [34], we can assume (purely for ease of notation) that \mathbf{R} consists only of axioms. By the definition of one-step derivation, ψ is propositionally entailed by (finitely many) formulas of the form $\phi_i \sigma_i \in \text{Prop}(\Lambda(\mathcal{P}(X)))$, $i = 1, \dots, n$, where $\text{Prop}(\Lambda(\text{Prop}(V_i))) \ni \phi_i \in \mathbf{R}$ and σ_i is a $\mathcal{P}(X)$ -valuation; by renaming variables and discarding unused variables, we may assume a single finite set V of propositional variables and a single $\mathcal{P}(X)$ -valuation σ . By Lemma 23, it suffices to find a $\text{Prop}(\mathfrak{A})$ -valuation τ for V such that

1. for all formulas $L(\rho_1, \dots, \rho_k), L(\rho'_1, \dots, \rho'_k)$ occurring in the ϕ_i , equality of $L(\rho_1, \dots, \rho_k)\sigma$ and $L(\rho'_1, \dots, \rho'_k)\sigma$ in $\Lambda(\mathcal{P}(X))$ (i.e. with the inner propositional level evaluated in $\mathcal{P}(X)$) implies equality of $L(\rho_1, \dots, \rho_k)\tau$ and $L(\rho'_1, \dots, \rho'_k)\tau$ in $\Lambda(\mathcal{P}(X))$, and
2. for every formula $L(A_1, \dots, A_k) \in \Lambda(\mathfrak{A}) \subseteq \Lambda(\mathcal{P}(X))$ occurring in ψ and every formula $L(\rho_1, \dots, \rho_k)$ occurring in the ϕ_i , equality of $L(\rho_1, \dots, \rho_k)\sigma$ and $L(A_1, \dots, A_k)$ in $\Lambda(\mathcal{P}(X))$ implies equality of $L(\rho_1, \dots, \rho_k)\tau$ and $L(A_1, \dots, A_k)$ in $\Lambda(\mathcal{P}(X))$.

As equality of e.g. $L(\rho_1, \dots, \rho_k)$ and $L(\rho'_1, \dots, \rho'_k)$ in $\Lambda(\mathcal{P}(X))$ under some $\mathcal{P}(X)$ -valuation amounts to equalities $\rho_1 = \rho'_1, \dots, \rho_k = \rho'_k$ in $\mathcal{P}(X)$, the above conditions amount to solving a finite system of Boolean equations in $\mathcal{P}(X)$, with constants from \mathfrak{A} appearing on some of the right-hand sides (due to condition 2). The system is, by construction, solvable in $\mathcal{P}(X)$ (by the given valuation σ), and hence in $\text{Prop}(\mathfrak{A})$ by Lemma 21. \square

We continue our preparations for the definition of the canonical structure:

Notation 25 For $f : X \rightarrow Y$, let σ_f denote the substitution replacing $A \subseteq Y$ with $f^{-1}[A] \subseteq X$.

Lemma 26 Let $f : X \rightarrow Y$. If $\Phi \vdash_{\mathcal{L}}^Y \psi$, then $\Phi \sigma_f \vdash_{\mathcal{L}}^X \psi \sigma_f$.

Proof One immediately reduces to the case $\Phi = \emptyset$. If $\vdash_{\mathcal{L}}^Y \psi$, then ψ is propositionally entailed by formulas $\chi_i \tau_i$, $i = 1, \dots, n$, where $\phi_i / \chi_i \in \mathbf{R}$ for some ϕ_i and τ_i is a $\mathcal{P}(Y)$ -valuation such that $Y \models \phi_i \tau_i$. Because σ_f is a homomorphism $\mathcal{P}(Y) \rightarrow \mathcal{P}(X)$ of Boolean algebras, we have $X \models \phi_i \tau_i \sigma_f$ for all i . Moreover, by the substitution lemma of propositional logic, $\psi \sigma_f$ is propositionally entailed by the $\chi_i \tau_i \sigma_f$; hence $\vdash_{\mathcal{L}}^X \psi \sigma_f$. \square

The canonical \mathcal{L} -structure $\mathcal{M}_{\mathcal{L}}$ for \mathcal{L} will be based on the functor $M_{\mathcal{L}}$ that takes a set X to the set of maximally one-step consistent subsets of $\text{Prop}(\Lambda(\mathcal{P}(X)))$. For a map $f : X \rightarrow Y$, $M_{\mathcal{L}}(f)$ is defined by

$$M_{\mathcal{L}}(f)(\Phi) = \{\phi \in \text{Prop}(\Lambda(\mathcal{P}(Y))) \mid \phi \sigma_f \in \Phi\}.$$

This definition is justified by

Lemma 27 For $\Phi \in M_{\mathcal{L}}(X)$, the set $M_{\mathcal{L}}(f)(\Phi)$ is maximally one-step consistent.

Proof Consistency of $M_{\mathcal{L}}(f)(\Phi)$ is immediate by Lemma 26. To prove maximality, let $\psi \notin M_{\mathcal{L}}(f)(\Phi)$ for some $\psi \in \text{Prop}(\Lambda(\mathcal{P}(Y)))$. Then $\psi \sigma_f \notin \Phi$, hence $\neg \psi \sigma_f \in \Phi$, and thus $\neg \psi \in M_{\mathcal{L}}(f)(\Phi)$. \square

From the perspective of Stone duality, the functor $M_{\mathcal{L}}$ is the imperfect dual of the functorial presentation $L : \mathbf{BA} \rightarrow \mathbf{BA}$ of a rank-1 logic $\mathcal{L} = (\Lambda, \mathbf{R})$. We think of it as imperfect because the (perfect) duality between Stone spaces and Boolean algebras is replaced by a dual adjunction between sets and Boolean algebras. The latter is given by functors

$$\bar{Q} : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{BA} \quad \text{and} \quad S : \mathbf{BA}^{\text{op}} \rightarrow \mathbf{Set}$$

where \bar{Q} takes a set to its Boolean algebra of subsets, and S takes a Boolean algebra to the set of its ultrafilters. We review the terminology of [17]:

Definition 28 (Functorial presentation of a logic) The *functorial presentation* of $\mathcal{L} = (\Lambda, \mathbf{R})$ is the functor $L_{\mathcal{L}} : \mathbf{BA} \rightarrow \mathbf{BA}$ defined by $L_{\mathcal{L}}A = F(\Lambda(A))/\sim$ where $F : \mathbf{Set} \rightarrow \mathbf{BA}$ is the free construction and \sim is the congruence generated by

$$\psi\sigma \sim \top$$

for all $\phi/\psi \in \mathbf{R}$ and all A -substitutions σ such that $\phi\sigma = \top$ (where we assume that applying the A -substitution σ to $\psi \in \text{Prop}(\Lambda(\text{Prop}(V)))$ evaluates propositional formulas over A to elements of A).

Example 29 The standard example of a functorial presentation is the functorial presentation of K : $L_K(A)$ is the quotient of the free Boolean algebra over the set $\{\Box a \mid a \in A\}$ of atoms modulo the congruence generated by the equations $\Box \top = \top$ and $\Box(a \wedge b) = \Box a \wedge \Box b$ for all $a, b \in A$.

Remark 30 It has been shown in [17] that the category of $L_{\mathcal{L}}$ -algebras is isomorphic to the category of Boolean algebras with operators for the similarity type Λ that satisfy all the axioms corresponding to rules of \mathbf{R} . As a consequence, the initial $L_{\mathcal{L}}$ -algebra is precisely the Lindenbaum-Tarski algebra of the logic (Λ, \mathbf{R}) , i.e. the set of closed \mathcal{L} -formulas modulo logical equivalence.

We have the following close relationship between $M_{\mathcal{L}}$ and the functorial presentation of \mathcal{L} :

Proposition 31 Let $L_{\mathcal{L}}$ be the functorial presentation of \mathcal{L} . Then

$$M_{\mathcal{L}} \cong SL_{\mathcal{L}}\bar{Q},$$

where \bar{Q} and S constitute the dual adjunction between \mathbf{BA} and \mathbf{Set} described above.

Proof Given a set X , $SL_{\mathcal{L}}\bar{Q}(X)$ consists of the ultrafilters in the Boolean algebra $L_{\mathcal{L}}\bar{Q}(X)$. We can regard $L_{\mathcal{L}}\bar{Q}(X)$ as consisting of equivalence classes $[\phi]_{\equiv}$ in $\text{Prop}(\Lambda(\mathcal{P}(X)))$ modulo the equivalence $\phi \equiv \psi$ iff $\vdash_{\mathcal{L}}^X \phi \leftrightarrow \psi$. We then define an injective natural transformation $\nu : SL_{\mathcal{L}}\bar{Q} \rightarrow M_{\mathcal{L}}$ by

$$\nu_X(u) = \{\phi \in \text{Prop}(\Lambda(\mathcal{P}(X))) \mid [\phi]_{\equiv} \in u\}.$$

To see that ν_X is also surjective, let $\Phi \in M_{\mathcal{L}}(X)$, and define an ultrafilter u in $L_{\mathcal{L}}\bar{Q}(X)$ by

$$[\phi]_{\equiv} \in u \iff \phi \in \Phi.$$

This is well-defined, as by maximal one-step consistency of Φ , $\phi \in \Phi \iff \psi \in \Phi$ whenever $\phi \equiv \psi$. It is then clear that $\nu_X(u) = \Phi$. \square

We now complete the definition of $\mathcal{M}_{\mathcal{L}}$ as a Λ -structure $(M_{\mathcal{L}}, (\llbracket L \rrbracket^{\mathcal{M}_{\mathcal{L}}})_{L \in \Lambda})$ by

$$\llbracket L \rrbracket(A_1, \dots, A_n) = \{\Phi \in M_{\mathcal{L}}(X) \mid L(A_1, \dots, A_n) \in \Phi\}$$

for $L \in \Lambda$ n -ary; it is easy to check that the above indeed defines a predicate lifting. We then have a single-step version of the truth lemma:

Lemma 32 (One-step truth lemma) For every set X and every one-step formula $\psi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$,

$$\Phi \models^X \psi \quad \text{iff} \quad \psi \in \Phi$$

for all $\Phi \in M_{\mathcal{L}}(X)$.

Proof Induction over ψ , where the steps for Boolean operators are by maximal consistency of Φ and the step for modal operators is by construction. \square

Theorem and Definition 33 (Canonical structure) The Λ -structure $\mathcal{M}_{\mathcal{L}}$ is an \mathcal{L} -structure, the *canonical \mathcal{L} -structure*.

Proof One-step soundness is immediate by the one-step truth lemma. \square

Remark 34 The construction of $M_{\mathcal{L}}$ always gives rise to subclasses of neighbourhood models and can be viewed from a more general perspective, as follows. Given a similarity type Λ and a Λ -structure \mathcal{M} based on T , we can regard a set \mathbf{R} of extended one-step rules as a set of *frame conditions* to be satisfied by T -coalgebras, where we assume w.l.o.g. that \mathbf{R} consists only of axioms. To this end, we interpret \mathcal{L} -formulas with propositional variables over T -models $((X, \xi), \tau)$, consisting of a T -coalgebra $C = (X, \xi)$ and a $\mathcal{P}(X)$ -valuation τ for a set V of propositional variables as usual, thus inducing a satisfaction relation \models_C^τ between states of C and formulas over V . We say that a T -model (C, τ) satisfies an \mathcal{L} -formula ϕ over V ($C, \tau \models \phi$) if $x \models_C^\tau \phi$ for all states x in C , and that C satisfies ϕ ($C \models \phi$) if $C, \tau \models \phi$ for all $\mathcal{P}(X)$ -valuations τ for V . The set \mathbf{R} of axioms *defines* the class of T -coalgebras C such that $C \models \phi$ for all $\phi \in \mathbf{R}$.

The crucial observation is then that frame conditions in rank 1 can always be pushed into the functor: putting

$$T_{\mathbf{R}}(X) = \{t \in TX \mid t \models^X \phi\tau \text{ for all } \phi \in \mathbf{R} \text{ and all } \mathcal{P}(X)\text{-valuations } \tau\}$$

defines, due to naturality of predicate liftings, a subfunctor $T_{\mathbf{R}}$ of T . It then turns out that a T -coalgebra $C = (X, \xi)$ satisfies the frame conditions in \mathbf{R} iff its structure map ξ factors through $T_{\mathbf{R}}(X)$, i.e. iff C is a $T_{\mathbf{R}}$ -coalgebra, the crucial lemma being that

$$x \models_C^\tau \phi \quad \text{iff} \quad \xi(x) \models^X \phi\tau$$

for every one-step formula ϕ over V , every $\mathcal{P}(X)$ -valuation τ for V , and every state x in C . (This lemma implies also that rank-1 axioms are one-step sound for a Λ -structure \mathcal{M} iff they are frame-valid over \mathcal{M} , i.e. satisfied by all T -coalgebras.) In other words, \mathbf{R} *defines the class of $T_{\mathbf{R}}$ -coalgebras*. (However, this construction does not necessarily preserve completeness of logics, see Example 54 below.)

In this setting, the construction of $M_{\mathcal{L}}$ for $\mathcal{L} = (\Lambda, \mathbf{R})$ may be decomposed as follows. To begin, let M_{Λ} denote the functor defined by taking $M_{\Lambda}(X)$ to be the

set of maximally *propositionally* consistent subsets of $\text{Prop}(\Lambda(\mathcal{P}(X)))$, with the action on morphisms as above. Then in the above notation, $(M_\Lambda)_\mathbb{R} = M_{(\Lambda, \mathbb{R})} = M_\mathcal{L}$: given $\Phi \in M_\Lambda(X)$, Φ is one-step \mathcal{L} -consistent iff $\phi\tau \in \Phi$ for all $\phi \in \mathbb{R}$ and all $\mathcal{P}(X)$ -valuations τ iff (by the one-step truth lemma) $\Phi \models^X \phi\tau$ for all $\phi \in \mathbb{R}$, τ iff $\Phi \in (M_\Lambda)_\mathbb{R}(X)$. It is moreover easy to see that the canonical translation $\mu : M_\Lambda \rightarrow \mathbb{N}_\Lambda$ into the neighbourhood structure of Λ (Lemma and Definition 14) is an isomorphism: unravelling the definitions, μ_X maps $\Phi \in M_\Lambda(X)$ to the family of sets $(\{(A_1, \dots, A_n) \in \mathcal{Q}(X)^n \mid L(A_1, \dots, A_n) \in \Phi\})_{L \in \Lambda \text{ } n\text{-ary}}$, and the inverse transformation $\mathbb{N}_\Lambda(X) \rightarrow M_\Lambda(X)$ maps a family of sets $(\mathfrak{A}_L)_{L \in \Lambda}$, where $\mathfrak{A}_L \subseteq \mathcal{Q}(X)^n$ for L n -ary, to its one-step theory (Definition 15), explicitly described as the unique maximally one-step consistent extension of the set $\{L(A_1, \dots, A_n) \mid L \in \Lambda \text{ } n\text{-ary}, (A_1, \dots, A_n) \in \mathfrak{A}_L\} \cup \{\neg L(A_1, \dots, A_n) \mid L \in \Lambda \text{ } n\text{-ary}, (A_1, \dots, A_n) \notin \mathfrak{A}_L\}$. Hence we can, by the above, equivalently describe $M_\mathcal{L}$ as $(\mathbb{N}_\Lambda)_\mathbb{R}$. That is, $M_\mathcal{L}$ -coalgebras are those neighbourhood frames that satisfy all frame conditions prescribed by \mathcal{L} . Of course, satisfaction of rank-1 frame conditions in a neighbourhood frame translates immediately into obvious closure conditions on sets of neighbourhoods.

Example 35 We discuss the concrete shape of the canonical structure by re-visiting some of the logics introduced in Example 12.

1. For the minimal modal logic E (Example 12.2), M_E is the neighbourhood functor $\mathbb{N} = \mathcal{Q} \circ \mathcal{Q}^{\text{op}}$: the natural transformation $M_E \rightarrow \mathbb{N}$ which takes $\Phi \in M_E(X)$ to $\{A \subseteq X \mid \Box A \in \Phi\} \in \mathbb{N}(X)$ is easily seen to be a natural isomorphism. This is an instance of a more general fact on neighbourhood structures noted in Remark 34. In particular, the inverse natural isomorphism takes $\mathfrak{A} \in \mathbb{N}(X)$ to the one-step theory Φ of \mathfrak{A} w.r.t. the structure described in Example 12.2; explicitly, $\Phi \subseteq \text{Prop}(\{\Box\}(\mathcal{P}(X)))$ is the unique maximally one-step consistent extension of the set $\{\Box A \mid A \in \mathfrak{A}\} \cup \{\neg \Box A \mid A \in \mathcal{P}(X) \setminus \mathfrak{A}\}$.

2. For the modal logic K (Example 12.1), M_K is the filter functor [12], i.e. the subfunctor \mathbb{F} of \mathbb{N} where $\mathbb{F}(X)$ consists of all (not necessarily proper, or augmented) filters on X ; this is witnessed by a natural isomorphism $M_K \rightarrow \mathbb{F}$ obtained by restricting the isomorphism $M_E \rightarrow \mathbb{N}$ from the previous example.

3. The canonical structure \mathcal{M}_{CK} for the conditional logic CK (see Example 12.5) is equivalently described by a subfunctor \mathbb{F}_{CK} of $\mathcal{Q} \circ (\mathcal{Q}^{\text{op}})^2$. The elements of $\mathbb{F}_{CK}(X)$ are those subsets \mathfrak{A} of $\mathcal{Q}(X)^2$ such that for every $A \subseteq X$, the set $\{B \mid (A, B) \in \mathfrak{A}\}$ is a filter. This functor is isomorphic to the functor T defined by $T(X) = \mathcal{Q}(X) \rightarrow \mathbb{F}(X)$.

Recall that the fact that $\mathcal{M}_\mathcal{L}$ is a structure for \mathcal{L} implies soundness of \mathcal{L} over $\mathcal{M}_\mathcal{L}$. We now turn to *strong completeness*, which is established by a canonical model construction that generalises the standard notion of canonical neighbourhood frame. As usual, we call a set $\Phi \subseteq \mathcal{F}(\Lambda)$ of Λ -formulas \mathcal{L} -consistent (or just consistent) if there do not exist formulas $\phi_1, \dots, \phi_n \in \Phi$ such that $\vdash_\mathcal{L} \neg(\phi_1 \wedge \dots \wedge \phi_n)$. The set Φ is *maximally \mathcal{L} -consistent* if it is maximal w.r.t. \subseteq among the \mathcal{L} -consistent sets. The carrier of the canonical model is then the set $C_\mathcal{L}$ of maximally \mathcal{L} -consistent sets of Λ -formulas. The key to the construction is the *existence proof* (rather than the explicit construction) of a suitable $M_\mathcal{L}$ -coalgebra structure on $C_\mathcal{L}$, a technique first employed in [34]:

Definition 36 (Coherent coalgebra structure, canonical model) An $M_{\mathcal{L}}$ -coalgebra structure $\zeta : C_{\mathcal{L}} \rightarrow M_{\mathcal{L}}(C_{\mathcal{L}})$ on $C_{\mathcal{L}}$ is *coherent* if

$$\zeta(\Phi) \in \llbracket L \rrbracket(\hat{\psi}_1, \dots, \hat{\psi}_n) \quad \text{iff} \quad L(\psi_1, \dots, \psi_n) \in \Phi$$

for all $\mathcal{F}(\Lambda)$ -formulas $L(\psi_1, \dots, \psi_n)$ and all $\Phi \in C_{\mathcal{L}}$, where $\hat{\psi} = \{\Psi \in C_{\mathcal{L}} \mid \psi \in \Psi\}$. In this case, $(C_{\mathcal{L}}, \zeta)$ is a *canonical model*.

Lemma 37 (Existence lemma) There exists a coherent coalgebra structure on $C_{\mathcal{L}}$.

The proof of the existence lemma relies on a number of additional lemmas.

Lemma 38 (Lindenbaum lemma) Every consistent set of Λ -formulas is contained in a maximally consistent set.

Lemma 39 (One-step Lindenbaum lemma) Every one-step consistent subset of $\text{Prop}(\Lambda(\mathcal{P}(X)))$ is contained in a maximally one-step consistent set.

Both the global and the one-step version of the Lindenbaum lemma are proved by appealing to Zorn's lemma (note that we do not assume that Λ is countable).

Lemma 40 Suppose $(C_{\mathcal{L}}, \gamma)$ is canonical and let $\phi \in \text{Prop}(V)$ and σ be an $\mathcal{F}(\Lambda)$ -substitution. Then $\vdash_{\mathcal{L}} \phi\sigma$ iff $C_{\mathcal{L}} \models \phi\hat{\sigma}$, where $\hat{\sigma}$ is the $\mathcal{P}(C_{\mathcal{L}})$ -valuation given by $\hat{\sigma}(a) = \widehat{\sigma(a)} = \{\Psi \in C_{\mathcal{L}} \mid \sigma(a) \in \Psi\}$.

Proof This is a special case of [34], Lemma 27 (take the closed set Σ of loc. cit. to be all of $\mathcal{F}(\Lambda)$). \square

The following is another piece of propositional logic, an obvious variant of Lemma 22.

Lemma 41 Let $\Phi \subseteq \text{Prop}(V)$, let $\psi \in \text{Prop}(V)$, let σ be a W -substitution, and let τ be a U -substitution. If $\Phi\sigma \vdash_{PL} \psi\sigma$ then $\Phi\tau \cup \Psi \vdash_{PL} \psi\tau$, where $\Psi = \{\tau(a) \leftrightarrow \tau(b) \mid a, b \in V, \sigma(a) = \sigma(b)\}$.

Proof Let \sim be the equivalence relation on V induced by σ , i.e. $a \sim b$ iff $\sigma(a) = \sigma(b)$. For every $a \in V$, fix a representative $\nu(a)$ of the equivalence class of a under \sim . For $a \in V$, put $\tau'(a) = \tau(\nu(a))$. By Lemma 23, $\Phi\sigma \vdash_{PL} \psi\sigma$ implies $\Phi\tau' \vdash_{PL} \psi\tau'$. The proof is finished by noting that by replacement of equivalents, $\Phi\tau \cup \Psi$ propositionally entails $\Phi\tau'$, and $\Psi \cup \{\psi\tau'\}$ propositionally entails $\psi\tau$. \square

Lemma 42 Let V_{Λ} be the set $\{a_{\phi} \mid \phi \in \mathcal{F}(\Lambda)\}$ of propositional variables, let $\Phi \subseteq \text{Prop}(\Lambda(V_{\Lambda}))$, let σ_{Λ} be the $\mathcal{F}(\Lambda)$ -substitution taking a_{ϕ} to ϕ , and let $\hat{\sigma}_{\Lambda}$ be the $\mathcal{P}(C_{\mathcal{L}})$ -valuation given by $\hat{\sigma}_{\Lambda}(a_{\phi}) = \hat{\phi}$. If $\Phi\sigma_{\Lambda}$ is consistent, then $\Phi\hat{\sigma}_{\Lambda}$ is one-step consistent.

Proof We may assume, purely for ease of notation, that $\mathcal{L} = (\Lambda, \mathbb{R})$ is given in terms of a set \mathbb{R} of one-step rules. Proceeding by contraposition, assume that $\Phi\hat{\sigma}_{\Lambda} \not\vdash_{\mathcal{L}} \perp$. By Proposition 24, we can assume that this one-step derivation uses only \mathfrak{A} -instances of rules, where $\mathfrak{A} = \{\hat{\phi} \mid \phi \in \mathcal{F}(\Lambda)\}$ (this set is already closed under Boolean operations); by suitable renaming of variables we may even assume that all

rules appearing in the one-step derivation use the set V_Λ of propositional variables and are instantiated using the valuation $\hat{\sigma}_\Lambda$. This means that $(\Phi \cup \Theta)\hat{\sigma}_\Lambda \vdash_{PL} \perp$, where Θ is the set of all conclusions χ of rules $\phi/\chi \in \mathbf{R}$ over V_Λ such that $C_\mathcal{L} \models \phi\hat{\sigma}_\Lambda$. By Lemma 40, the latter implies $\vdash_{\mathcal{L}} \phi\sigma_\Lambda$; hence all formulas in $\Theta\sigma_\Lambda$ are \mathcal{L} -derivable. By Lemma 41, we have moreover

$$(\Phi \cup \Theta)\sigma_\Lambda \cup \Psi \vdash_{PL} \perp, \quad (*)$$

where

$$\Psi = \{L(\rho_1, \dots, \rho_n) \leftrightarrow L(\rho'_1, \dots, \rho'_n) \mid \hat{\rho}_1 = \hat{\rho}'_1, \dots, \hat{\rho}_n = \hat{\rho}'_n\}.$$

Since by Lemma 40, $\hat{\rho} = \hat{\rho}'$ implies $\vdash_{\mathcal{L}} \rho \leftrightarrow \rho'$, all formulas in Ψ are derivable by the congruence rule. Thus, (*) implies that $\Phi\sigma_\Lambda \vdash_{\mathcal{L}} \perp$ as required. \square

Proof (Existence lemma) We obtain $\zeta(\Phi)$ with the required property by the one-step Lindenbaum lemma (Lemma 39) and by the definition of the interpretation of modal operators in $\mathcal{M}_\mathcal{L}$ once we show that the set

$$\{L(\hat{\phi}_1, \dots, \hat{\phi}_n) \mid L(\phi_1, \dots, \phi_n) \in \Phi\} \cup \{\neg L(\hat{\phi}_1, \dots, \hat{\phi}_n) \mid \neg L(\phi_1, \dots, \phi_n) \in \Phi\}$$

is one-step consistent. This follows from consistency of Φ and Lemma 42. \square

Lemma 43 (Truth lemma) In a canonical model $(C_\mathcal{L}, \zeta)$, $\Phi \models_{(C_\mathcal{L}, \zeta)} \phi$ iff $\phi \in \Phi$.

Proof Induction on ϕ ; coherence is precisely the induction step for modal operator application. \square

Theorem 44 (Strong completeness) The logic \mathcal{L} is strongly complete for $\mathcal{M}_\mathcal{L}$.

Proof By the existence lemma, there exists a canonical model $(C_\mathcal{L}, \zeta)$. Now suppose $\Phi \not\vdash_{\mathcal{L}} \psi$. Then $\Phi \cup \{\neg\psi\}$ is consistent, hence contained in a maximally consistent set Ψ by the Lindenbaum lemma (Lemma 38). By the truth lemma, we have $\Psi \models_{(C_\mathcal{L}, \zeta)} \neg\psi$ for $(C_\mathcal{L}, \zeta)$ canonical, thus $\Phi \not\models_{\mathcal{M}_\mathcal{L}} \psi$. \square

Remark 45 Alternatively, existence of canonical models and strong completeness may be established using the coalgebraic Jónsson-Tarski theorem as proved in [18]. The crucial prerequisite for application of this theorem is to define a family of maps, not necessarily natural,

$$h_A : SL_\mathcal{L}A \rightarrow SL_\mathcal{L}\bar{Q}SA$$

for all Boolean algebras A , where $L_\mathcal{L}$ is the functorial presentation of \mathcal{L} (Definition 28) and S and \bar{Q} constitute the dual adjunction between sets and Boolean algebras as explained above. The h_A are subject to a requirement that in this concrete case translates into the condition $h_A(u) \supseteq \{L_\mathcal{L}j_A(a) \mid a \in u\} =: v_0$ for all $u \in SL_\mathcal{L}A$, where $j_A : A \rightarrow \bar{Q}SA$, the unit of the mentioned adjunction, maps $a \in A$ to $\{u \in SA \mid a \in u\}$. To show that $h_A(u)$ exists as required, one has to show that v_0 has the finite intersection property, i.e. is consistent. As j_A is injective, this amounts to proving that $L_\mathcal{L}$ preserves injective Boolean homomorphisms; the latter follows by essentially the same arguments as employed in the proof of Proposition 24.

By virtue of h_A , every $L_\mathcal{L}$ -algebra $\alpha : L_\mathcal{L}A \rightarrow A$ gives rise to an $M_\mathcal{L}$ -coalgebra $h_A \circ S\alpha : SA \rightarrow SL_\mathcal{L}\bar{Q}SA = M_\mathcal{L}A$ (applying this construction to the initial \mathcal{L} -algebra

yields a canonical model), whose complex algebra is the *canonical extension* of A , an $L_{\mathcal{L}}$ -algebra on \bar{QSA} . The coalgebraic Jónsson-Tarski theorem then guarantees that j_A is an $L_{\mathcal{L}}$ -algebra homomorphism from α into its canonical extension; strong completeness follows as a direct corollary. In the development above, we have preferred a more explicit treatment that works entirely on the coalgebraic side, i.e. in terms of neighbourhood frames.

3 Finite Branching

We proceed to consider finitely branching structures, with a view to obtaining a Hennessy-Milner type result. The resulting finitely branching canonical structures will turn out to be closely related to the natural semantics of many modal logics. Finite branching generally implies non-compactness and hence the failure of strong completeness; we shall however prove that weak completeness continues to hold.

Due to naturality of predicate liftings, satisfaction of Λ -formulas is invariant under morphisms of coalgebras and hence under behavioural equivalence [27]. Conversely, by results of [33], Λ has the *Hennessy-Milner property* for a Λ -structure \mathcal{M} based on T , i.e. states that satisfy the same Λ -formulas are behaviourally equivalent, if T is ω -accessible and Λ is *separating* in the sense that $t \in TX$ is uniquely determined by the set

$$\{(L, (A_1, \dots, A_n)) \mid L \in \Lambda \text{ } n\text{-ary}, A_1, \dots, A_n \in \mathcal{P}(X), t \in \llbracket L \rrbracket(A_1, \dots, A_n)\},$$

equivalently if t is uniquely determined by its one-step theory. Recall here that a functor is ω -accessible if it preserves directed colimits; the following alternative characterisations will be useful below:

Lemma 46 ([1], Proposition 5.2) For a functor $T : \mathbf{Set} \rightarrow \mathbf{Set}$, the following are equivalent:

- (i) T is ω -accessible
- (ii) T preserves directed unions
- (iii) For every set X , $TX = \bigcup_{Y \subseteq X \text{ finite}} TY$ (recall Assumption 6).

For those unfamiliar with the notion of directed colimit, it will suffice to think of ω -accessibility as defined by condition (iii) for purposes of the further development. We say that a Λ -structure is ω -accessible if its underlying functor is ω -accessible.

Example 47 The covariant powerset functor \mathcal{P} fails to be accessible, as infinite subsets of an infinite set X fail to belong to any of the subsets $\mathcal{P}(Y) \subseteq \mathcal{P}(X)$ for $Y \subseteq X$ finite. Contrastingly, the finite powerset functor \mathcal{P}^{fin} is obviously ω -accessible.

The functor $M_{\mathcal{L}}$ fails to be ω -accessible for obvious cardinality reasons. Intuitively, $M_{\mathcal{L}}$ -models have unbounded branching, while the Hennessy-Milner property can only be expected for finitely branching systems (as is the case already for standard Kripke models). We thus consider a subfunctor $M_{\mathcal{L}}^{fin}$ of $M_{\mathcal{L}}$ that captures precisely the finitely branching models.

In order to construct $M_{\mathcal{L}}^{fin}$, we can rely on the following general mechanism. We define the ω -accessible part T^{fin} of a set functor T by

$$T^{fin}X = \bigcup_{Y \subseteq X \text{ finite}} TY \subseteq TX$$

(recall Assumption 6). It is easy to see that T^{fin} is a subfunctor of T . By Lemma 46, T^{fin} is ω -accessible. Moreover, T^{fin} agrees with T on finite sets. Given any structure \mathcal{M} based on T , we denote by \mathcal{M}^{fin} the substructure of \mathcal{M} induced by T^{fin} (Section 1). Then, we define the *canonical finitely branching \mathcal{L} -structure* to be the structure $\mathcal{M}_{\mathcal{L}}^{fin}$. One-step soundness of \mathcal{L} for $\mathcal{M}_{\mathcal{L}}^{fin}$ is immediate, as $\mathcal{M}_{\mathcal{L}}^{fin}$ is a substructure of $\mathcal{M}_{\mathcal{L}}$. We then obtain

Theorem 48 The logic \mathcal{L} is weakly complete and has the Hennessy-Milner property for $\mathcal{M}_{\mathcal{L}}^{fin}$.

The proof of weak completeness requires a notion of weak completeness at the single-step level, called one-step completeness, which we recall from earlier work. In the definition of one-step completeness, there is some latitude in the exact design of the one-step logic it concerns. In particular, one may either choose $\text{Prop}(\Lambda(\mathcal{P}(X)))$ as the set of one-step formulas, as we have done so far, or one may opt for a more syntactic treatment where one considers one-step formulas over sets V of variables, i.e. elements of $\text{Prop}(\Lambda(\text{Prop}(V)))$, instead of over $\mathcal{P}(X)$ [36, 34]. For the sake of clarification, we now prove the equivalence of the arising notions of one-step completeness. We begin by making the more syntactic version of one-step derivation explicit; the main point to note here is that giving up the immediate evaluation of the inner propositional layer necessitates introducing the congruence rule also at the one-step level.

Definition 49 (One-step derivation with variables) Let X be a set, let V be a set of propositional variables, let $\psi \in \text{Prop}(\Lambda(\text{Prop}(V)))$ be a one-step formula over V , and let τ be a $\mathcal{P}(X)$ -valuation for V . We say that ψ is *one-step derivable over (X, τ)* , and write $\vdash_{\mathcal{L}}^{X, \tau} \psi$, if ψ is propositionally entailed by conclusions of $\text{Prop}(V)$ -instances of \mathbf{R} and the congruence rule whose premises hold over (X, τ) ; formally: if $\Theta \vdash_{PL} \psi$, where

$$\Theta = \{\chi\sigma \mid \phi/\chi \in \mathbf{R} \cup \mathbf{C}; \sigma \text{ a } \text{Prop}(V)\text{-substitution}; X \models \phi\sigma\tau\}$$

and \mathbf{C} denotes the set of *congruence rules*

$$\frac{a_1 \leftrightarrow b_1; \dots; a_n \leftrightarrow b_n}{L(a_1, \dots, a_n) \leftrightarrow L(b_1, \dots, b_n)} \quad (L \in \Lambda \text{ } n\text{-ary}).$$

Proposition and Definition 50 (One-step completeness) We say that \mathcal{L} is *one-step complete* for an \mathcal{L} -structure \mathcal{M} based on T if for each set X , the following equivalent conditions are satisfied:

- (i) Whenever $TX \models \psi\tau$ for a set X , a one-step formula $\psi \in \text{Prop}(\Lambda(\text{Prop}(V)))$ over V , and a $\mathcal{P}(X)$ -valuation τ for V , then ψ is one-step derivable over (X, τ) (Definition 49).
- (ii) Whenever $TX \models \psi$ for a set X and $\psi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$, then ψ is one-step derivable over X (Definition 19).

(iii) Every one-step consistent $\psi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ is one-step satisfiable in TX .

Proof The equivalence of (ii) and (iii) is just the usual dualisation argument. The implication (i) \implies (ii) is trivial — essentially, (ii) is just a special case of (i) where τ assigns distinct values to all variables, and instances of the congruence rules become propositional tautologies when propositional formulas are replaced by their extensions in X .

We prove the implication (ii) \implies (i): by (ii), $\psi\tau$ as in the statement of (i) is one-step derivable over X . By Proposition 24, we can assume that the derivation uses only $\text{Prop}(\mathfrak{A})$ -instances of rules, where $\mathfrak{A} = \{\tau(a) \mid a \in V\}$; moreover, by suitable renaming of variables we may assume that all rule instances needed in the derivation use a common set W of variables, disjoint from V , and a common $\text{Prop}(\mathfrak{A})$ -valuation for W , which we extend to a $\text{Prop}(\mathfrak{A})$ -valuation $\bar{\tau}$ for the disjoint union $W + V$ using the given \mathfrak{A} -valuation τ for V . Thus, $\Theta\bar{\tau} \vdash_{PL} \psi\bar{\tau}$ for $\Theta = \{\chi \mid \phi/\chi \in \mathbf{R}; X \models \phi\bar{\tau}\}$. For each $a \in W$, we pick $\sigma(a) \in \text{Prop}(V)$ such that $\sigma(a)\tau = \bar{\tau}(a) \in \text{Prop}(\mathfrak{A})$; moreover, we put $\sigma(a) = a$ for $a \in V$, thus defining a $\text{Prop}(V)$ -substitution σ on $W + V$. Then $\Theta\sigma\tau \vdash_{PL} \psi\tau$. By Lemma 41, $\Theta\sigma \cup \Psi \vdash_{PL} \psi$, where Ψ states the equivalence of those atoms in $\Lambda(\text{Prop}(V))$ that are equalised under the valuation τ . These are precisely the conclusions of congruence rules whose premises hold over (X, τ) (Definition 49), so that $\vdash_{\mathcal{L}}^{X, \tau} \psi$ as required. \square

By standard results, one-step completeness implies weak completeness for the logic at large:

Theorem 51 [26, 34] If \mathcal{L} is one-step complete for \mathcal{M} , then \mathcal{L} is weakly complete (Definition 11) for \mathcal{M} .

Lemma 52 The logic \mathcal{L} is one-step complete for the \mathcal{L} -structure \mathcal{M} iff the equivalent conditions of Proposition 50 hold for *finite* sets X .

Proof This is part of [36], Proposition 3.10. \square

Corollary 53 If \mathcal{L} is one-step complete for \mathcal{M} , then \mathcal{L} is one-step complete for \mathcal{M}^{fn} . \square

We now have all the technical machinery in place to establish weak completeness and the Hennessy-Milner property for the finitary part of the canonical structure.

Proof (Theorem 48) It is easy to see that \mathcal{L} is separating w.r.t. $\mathcal{M}_{\mathcal{L}}$ and hence w.r.t. $\mathcal{M}_{\mathcal{L}}^{fn}$, so that the Hennessy-Milner property follows from the fact that $\mathcal{M}_{\mathcal{L}}^{fn}$ is ω -accessible.

To prove weak completeness, we prove one-step completeness. By Corollary 53, it suffices to prove one-step completeness of \mathcal{L} for $\mathcal{M}_{\mathcal{L}}$. Thus, we have to show that every one-step consistent formula $\phi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ is one-step satisfiable in $\mathcal{M}_{\mathcal{L}}(X)$. This is immediate by the one-step Lindenbaum lemma (Lemma 39) and the one-step truth lemma (Lemma 32). \square

Example 54 Referring back to Remark 34, we note that when $\mathcal{L} = (\Lambda, \mathbf{R})$ is one-step complete for an \mathcal{L} -structure \mathcal{M} based on T and \mathbf{R}' is a set of one-step formulas, then

$(\Lambda, \mathbf{R} \cup \mathbf{R}')$ need not be one-step complete for the substructure of \mathcal{M} induced by $T_{\mathbf{R}'}$. Consider the constant functor $TX = D$, where D is a fixed infinite set. Let Λ consist of nullary modal operators d , $d \in D$, interpreted over T by $\llbracket d \rrbracket_X(A) = \{d\}$. Then the set $\mathbf{R} = \{\neg d_1 \wedge d_2 \mid d_1, d_2 \in D, d_1 \neq d_2\}$ is easily seen to be one-step complete. Let $\mathbf{R}' = \{\neg d \mid d \in D\}$. Then $T_{\mathbf{R}'}X = \emptyset$ for all X , and hence the one-step consistent formula \top fails to be satisfiable in TX , i.e. $(\Lambda, \mathbf{R} \cup \mathbf{R}')$ fails to be one-step complete for the substructure induced by $T_{\mathbf{R}'}$.

N.B.: By the above results, the same accident does not happen in the canonical structure $\mathcal{M}_{\mathcal{L}}$ for $\mathcal{L} = (\Lambda, \mathbf{R})$ as above. Here, $M_{\mathcal{L}}$ is isomorphic to the constant functor for the set $D \cup \{*\}$, where $*$ corresponds to (the unique maximally one-step consistent extension of) the one-step consistent set $\{\neg d \mid d \in D\}$, and consequently $(M_{\mathcal{L}})_{\mathbf{R}'}$ is the constant functor $(M_{\mathcal{L}})_{\mathbf{R}'}(X) = \{*\}$.

Remark 55 Note that $\mathcal{M}_{\mathcal{L}}$ supports a strong version of one-step completeness: \mathcal{L} is *strongly one-step complete* for $\mathcal{M}_{\mathcal{L}}$, i.e. every one-step consistent subset of $\text{Prop}(\Lambda(\mathcal{P}(X)))$ (rather than just every consistent formula in $\text{Prop}(\Lambda(\mathcal{P}(X)))$) is satisfiable in $M_{\mathcal{L}}(X)$. This was the key to the proof of the crucial existence lemma above, and indeed the above arguments show that strong one-step completeness implies strong completeness. However, there are only few interesting examples of structures other than $\mathcal{M}_{\mathcal{L}}$ for which a logic \mathcal{L} is strongly one-step complete: every such structure \mathcal{M} based on T has a surjective natural transformation $\theta : T \rightarrow M_{\mathcal{L}}$ which assigns to each $t \in TX$ its one-step theory, and when Λ is separating for \mathcal{M} , then θ is a natural isomorphism. The single example we are aware of where strong one-step completeness does apply to a non-canonical structure (for which separation then necessarily fails) is Pauly's coalition logic [29], which has been provided with a coalgebraic model in [36]: in our notation, Theorem 3.2 of [29] can be read as saying that θ is surjective for coalition logic. A more widely applicable strong completeness criterion for coalgebraic modal logic, in generalisation of results of [20] for functors preserving finite sets, is forthcoming [35].

Example 56 We give explicit descriptions (up to natural isomorphism) of $M_{\mathcal{L}}^{\text{fin}}$ for some of the logics of Example 12.

1. For the standard modal logic K , we have already seen that M_K is the filter functor (Example 35.2). Its finitely branching sibling M_K^{fin} is the finite powerset functor \mathcal{P}^{fin} : since both functors are ω -accessible and \mathbf{Set} is generated by taking directed colimits of finite sets, it suffices to show that the two functors coincide on finite sets. For X finite, $M_K^{\text{fin}}(X) = M_K(X) \cong \mathbf{F}(X)$, and filters on finite sets are in natural bijection with subsets.

2. For graded modal logic GML (Example 12.3), $M_{\text{GML}}^{\text{fin}}$ is a modification \mathbf{B}_{∞} of the finite multiset functor where elements of multisets may have infinite multiplicity ∞ . More precisely, $\mathbf{B}_{\infty}(X)$ is the set of maps $X \rightarrow \mathbb{N} \cup \{\infty\}$ with finite support, and the interpretation of modal operators is as over \mathbf{B} , where ∞ has the expected behaviour w.r.t. \geq and sums. As \mathbf{B}_{∞} is ω -accessible, it suffices, again, to prove that \mathbf{B}_{∞} and $M_{\text{GML}}^{\text{fin}}$ coincide on finite sets X . We define an isomorphism $\mu : M_{\text{GML}}^{\text{fin}} \rightarrow \mathbf{B}_{\infty}$ as follows: for $\Phi \in M_{\text{GML}}^{\text{fin}}(X)$ and $x \in X$, put $\mu_X(\Phi)(x) = \sup\{k + 1 \mid \diamond_k\{x\} \in \Phi\}$, under the convention $\sup \emptyset = 0$. It is easy to see that μ is natural; moreover, the μ_X are injective, as the information contained in $\mu_X(\Phi)$ determines, by the axiomatisation

of GML, the set $\{\diamond_k\{x\} \mid k \in \mathbb{N}, x \in X, \diamond_k\{x\} \in \Phi\}$, and hence, again by the axioms, all of Φ . To prove that μ_X is surjective, it suffices to show that the set $\Psi = \{\diamond_k\{x\} \mid x \in X, k < B(x)\}$ is one-step consistent for every $B \in \mathbf{B}_\infty(X)$. As the one-step derivation system is finitary, it suffices to show that every finite subset of Ψ is one-step consistent; but every such finite subset is evidently one-step satisfiable.

3. For probabilistic modal logic PML (Example 12.4), M_{PML}^{fin} is a modification of the finite distribution functor where events A are assigned generalised probabilities PA which are downclosed subsets of the rational interval $[0, 1] \cap \mathbb{Q}$. These are either open intervals $[0, r)$, with $r \in [0, 1]$, or closed intervals $[0, q]$, with $q \in [0, 1] \cap \mathbb{Q}$; thus, the space of generalised probabilities essentially consists of the interval $[0, 1]$ and a second copy of $[0, 1] \cap \mathbb{Q}$ which is infinitesimally greater than the first. In such structures, the modality L_p is interpreted by the predicate lifting taking a set A to the set $\{P \mid PA \geq p\}$, with \geq the ordering just discussed. The distributions $P \in M_{PML}^{fin}(X)$ are required to obey the axiomatization of PML [36] w.r.t. the canonical semantics; it is presently unclear whether this requirement can be replaced by a simpler condition. (In particular, it does *not* suffice to assign generalised probabilities only to singletons, as generalised probability measures P are additive only up to identification of $[0, p)$ and $[0, p]$. For instance, for A, B disjoint, $PA = [0, p]$, and $PB = [0, q]$, we may have either $P(A \cup B) = [0, p + q]$ or $P(A \cup B) = [0, p + q)$.)

4. The ω -accessible part of the canonical structure \mathcal{M}_{CK} for the conditional logic CK (see Example 12.5) is isomorphic to the functor T^{fin} (see Example 35), where the elements of $T^{fin}(X)$ are the functions $f : \mathcal{Q}(X) \rightarrow \mathcal{P}^{fin}(X)$ that are *finitely based*, i.e. there exists a finite subset $Y \subseteq X$ such that $f(A) = f(A \cap Y)$ for all $A \subseteq X$.

Note how $M_{\mathcal{L}}$ works in essence as a compactification (and indeed is constructed in a way which is strongly reminiscent of the topological Stone-Ćech-compactification), while $M_{\mathcal{L}}^{fin}$ removes only those failures of compactness which do not have to do with finite branching (see [34] for examples).

4 An Adjunction between Syntax and Semantics

We now set up an adjoint correspondence between rank-1 logics and set functors as their semantic counterparts. This establishes the canonical structure of a rank-1 logic as indeed canonical in a precise sense, i.e. as a universal model capturing all other ones. This situation is analogous to similar correspondences in equational logics and type theory: e.g. to a single-sorted equational theory, interpreted over cartesian categories (i.e. categories with finite products) with a distinguished object, one associates a Lawvere theory, which is again a cartesian category with a distinguished object and may simultaneously be regarded as an initial model and as a semantic representation of the given theory. The situation is dual for modal logics: the canonical structure serves as a *final* model of the given rank-1 logic, into which all other models may be mapped.

We make the categorical setting precise by collecting all rank 1 modal logics in a category \mathbf{ModL} ; for ease of presentation, we assume w.l.o.g. that logics are given in terms of axioms only. A *morphism* $(\Lambda_1, \mathbf{R}_1) \rightarrow (\Lambda_2, \mathbf{R}_2)$ in \mathbf{ModL} is a map $h : \Lambda_1 \rightarrow \Lambda_2$ such that the induced translation of formulas takes axioms in \mathbf{R}_1 to derivable formulas in $(\Lambda_2, \mathbf{R}_2)$ (in the sense of Definition 49). The category of semantic objects is the

category $\mathbf{Fn} = [\mathbf{Set}, \mathbf{Set}]$ of set functors and natural transformations. We have a functor $\mathbf{Th} : \mathbf{Fn}^{\text{op}} \rightarrow \mathbf{ModL}$ which takes a functor T to the logic $(\Lambda_T, \mathbf{R}_T)$, where Λ_T is the set of all finitary predicate liftings for T , and \mathbf{R}_T is the set of all one-step Λ_T -formulas which are one-step sound for T (where each predicate lifting in Λ_T is interpreted by itself). Given a natural transformation $\mu : T \rightarrow S$, $\mathbf{Th}(\mu) : \mathbf{Th}(S) \rightarrow \mathbf{Th}(T)$ is the morphism taking an n -ary predicate lifting $\lambda : \mathcal{Q}^n \rightarrow \mathcal{Q} \circ S^{\text{op}}$ for S to the predicate lifting $\mathcal{Q}\mu \circ \lambda$ for T ; it is easy to see that $\mathbf{Th}(\mu)$ indeed preserves axioms. Note that, in this terminology, an \mathcal{L} -structure is just a morphism of the form $h : \mathcal{L} \rightarrow \mathbf{Th}(T)$. In particular, the canonical \mathcal{L} -structure can be cast as a morphism

$$\eta_{\mathcal{L}} : \mathcal{L} \rightarrow \mathbf{Th}(M_{\mathcal{L}}).$$

The arrows $\eta_{\mathcal{L}}$ are the unit of the announced adjunction:

Theorem 57 The canonical \mathcal{L} -structure $\eta_{\mathcal{L}}$ is universal; i.e. for each \mathcal{L} -structure $h : \mathcal{L} \rightarrow \mathbf{Th}(T)$, there exists a unique natural transformation $h^{\#} : T \rightarrow M_{\mathcal{L}}$ such that $\mathbf{Th}(h^{\#})\eta_{\mathcal{L}} = h$.

$$\begin{array}{ccccc} \mathcal{L} & \xrightarrow{\eta_{\mathcal{L}}} & \mathbf{Th}(M_{\mathcal{L}}) & & M_{\mathcal{L}} \\ & \searrow h & \swarrow \mathbf{Th}(h^{\#}) & \dashrightarrow h^{\#} & \\ & & \mathbf{Th}(T) & & T \end{array}$$

Proof The map $h_X^{\#} : TX \rightarrow M_{\mathcal{L}}(X)$ takes $t \in TX$ to its one-step theory (Definition 15) in the structure represented by h . It is clear that this set is maximally consistent. Naturality of $h^{\#}$ is immediate from naturality of predicate liftings. Commutation of the above diagram and uniqueness of $h^{\#}$ are established by straightforward unfolding of the definitions. \square

In other words,

the canonical structure is a terminal object in the category of \mathcal{L} -structures.

Phrased differently, for every \mathcal{L} -structure based on T there is a unique natural transformation $T \rightarrow M_{\mathcal{L}}$ and vice versa, resulting in an isomorphism of categories. Recall for the following that for a functor $F : \mathbf{A} \rightarrow \mathbf{B}$ and a \mathbf{B} -object B , one has *comma categories* $F \downarrow B$ and $B \downarrow F$. Objects of $F \downarrow B$ are *objects over B*, i.e. pairs (A, f) consisting of an \mathbf{A} -object A and a morphism $f : FA \rightarrow B$; morphisms $(A, f) \rightarrow (C, g)$ in $F \downarrow B$ are \mathbf{A} -morphisms $a : A \rightarrow C$ such that $g \circ Fa = f$. Dually, objects of $B \downarrow F$ are *objects under B*, i.e. pairs (f, A) consisting of an \mathbf{A} -object A and a morphism $f : B \rightarrow FA$; morphisms $(f, A) \rightarrow (g, C)$ in $B \downarrow F$ are \mathbf{A} -morphisms $a : A \rightarrow C$ such that $Fa \circ f = g$. E.g. the discussion at the beginning of the section has shown that the category of \mathcal{L} -structures is isomorphic to the comma category $\mathcal{L} \downarrow \mathbf{Th}$ of objects under \mathcal{L} , and by Theorem 57, the latter is isomorphic to the comma category $\mathbf{Fn} \downarrow M_{\mathcal{L}}$ of objects over $M_{\mathcal{L}}$, i.e. natural transformations into $M_{\mathcal{L}}$ (where we abuse \mathbf{Fn} to denote the identity functor $\mathbf{Fn} \rightarrow \mathbf{Fn}$). Hence, we have

Corollary 58 The category of \mathcal{L} -structures is isomorphic to the comma category $\mathbf{Fn} \downarrow M_{\mathcal{L}}$ of objects over $M_{\mathcal{L}}$.

In other words, a coalgebraic semantics, i.e. a structure for \mathcal{L} based on a functor T , may be seen as a map associating to each $t \in TX$ a one-step theory. Returning to the characterisation of $M_{\mathcal{L}}$ given in Proposition 31, we can phrase this equivalently in terms of distributive laws, a concept used in [3, 17]: a *distributive law* for \mathcal{L} is a natural transformation $\delta : L_{\mathcal{L}}\bar{Q} \rightarrow \bar{Q}T^{\text{op}}$, where $L_{\mathcal{L}}$ is the functorial presentation of $\mathcal{L} = (\Lambda, \mathbf{R})$ (Definition 28) and $\bar{Q} : \text{Set}^{\text{op}} \rightarrow \mathbf{BA}$ is part of the dual adjunction between sets and Boolean algebras (Section 2):

$$L_{\mathcal{L}} \left(\text{BA} \right) \begin{array}{c} \xleftarrow{\bar{Q}} \\ \downarrow \delta \\ \xrightarrow{\bar{Q}} \end{array} \text{Set}^{\text{op}} \left(T^{\text{op}} \right)$$

Notice that a distributive law may be seen as an object of the comma category $L_{\mathcal{L}}\bar{Q} \downarrow (\bar{Q} \circ -^{\text{op}})$ of objects over $L_{\mathcal{L}}\bar{Q}$, where $(\bar{Q} \circ -^{\text{op}}) : [\text{Set}, \text{Set}]^{\text{op}} \rightarrow [\text{Set}^{\text{op}}, \mathbf{BA}]$ maps T to $\bar{Q} \circ T^{\text{op}}$. Morphisms from δ as above to $\theta : L_{\mathcal{L}}\bar{Q} \rightarrow \bar{Q}S^{\text{op}}$ in $L_{\mathcal{L}}\bar{Q} \downarrow (\bar{Q} \circ -^{\text{op}})$ are described explicitly as natural transformations $\mu : S \rightarrow T$ such that $\bar{Q}\mu \circ \delta = \theta$:

$$\begin{array}{ccc} L_{\mathcal{L}}\bar{Q} & \xrightarrow{\delta} & \bar{Q} \circ T^{\text{op}} \\ & \searrow \theta & \swarrow \bar{Q}\mu \\ & \bar{Q} \circ S^{\text{op}} & \end{array} \quad \begin{array}{c} \\ \\ S \end{array} \xrightarrow{\mu} T$$

Recall that $L_{\mathcal{L}}A$ is essentially the set $\text{Prop}(\Lambda(A))$ modulo provable equivalence in \mathcal{L} ; hence, a distributive law δ intuitively assigns an extension in TX to every one-step formula in $\text{Prop}(\Lambda(\mathcal{P}(X)))$, respecting the axiomatisation of \mathcal{L} . Indeed, δ induces an \mathcal{L} -structure by stipulating $\llbracket L \rrbracket_X(A_1, \dots, A_n) = \delta_X(\llbracket L(A_1, \dots, A_n) \rrbracket)$. This gives rise to a second characterisation of the category of \mathcal{L} -structures:

Proposition 59 Let $\mathcal{L} = (\Lambda, \mathbf{R})$ be a rank-1 logic. Then the category of \mathcal{L} -structures is isomorphic to the comma category $L_{\mathcal{L}}\bar{Q} \downarrow (\bar{Q} \circ -^{\text{op}})$ of objects over $L_{\mathcal{L}}\bar{Q}$, i.e. to the category of distributive laws for \mathcal{L} .

Proof Recall from Proposition 31 that $M_{\mathcal{L}} \cong S \circ L_{\mathcal{L}} \circ \bar{Q}$ where S maps a Boolean algebra A to the set of its ultrafilters. Given a morphism of structures $\eta : T \rightarrow M_{\mathcal{L}} \cong SL_{\mathcal{L}}\bar{Q}$ we define $\delta : L_{\mathcal{L}}\bar{Q} \rightarrow \bar{Q}T^{\text{op}}$ by

$$\delta_X(e) = \{t \in TX \mid e \in \eta_X(t)\}.$$

The fact that δ_X is a morphism of Boolean algebras follows from $\eta_X(t)$ being an ultrafilter for all $t \in TX$. Conversely, given $\delta : L_{\mathcal{L}}\bar{Q} \rightarrow \bar{Q}T^{\text{op}}$, define $\eta : T \rightarrow M_{\mathcal{L}}$ by

$$\eta_X(t) = \{e \in L_{\mathcal{L}}\bar{Q}X \mid t \in \delta(e)\}.$$

It is easy to see that $\eta_X(t)$ is indeed an ultrafilter (it corresponds to the one-step theory of t). The simple verification that both constructions are functorial and mutually inverse is left to the reader. \square

Similar results hold for the canonical finitely branching \mathcal{L} -structure $M_{\mathcal{L}}^{\text{fin}}$, which now becomes a morphism

$$\eta_{\mathcal{L}}^{\text{fin}} : \mathcal{L} \rightarrow \text{Th}(M_{\mathcal{L}}^{\text{fin}}).$$

Theorem 60 The \mathcal{L} -structure $\eta_{\mathcal{L}}^{fin}$ is universal among the finitely branching \mathcal{L} -structures; i.e. for each \mathcal{L} -structure $h : \mathcal{L} \rightarrow \text{Th}(T)$ with T ω -accessible, there exists a unique natural transformation $h^\# : T \rightarrow M_{\mathcal{L}}^{fin}$ such that $\text{Th}(h^\#)\eta_{\mathcal{L}}^{fin} = h$.

The proof requires the following explicit description of $M_{\mathcal{L}}^{fin}$, which is obtained immediately from the definition of $M_{\mathcal{L}}^{fin}$ and the action of $M_{\mathcal{L}}$ on subset inclusions:

Lemma 61 A maximally consistent set $\Phi \in M_{\mathcal{L}}(X)$ is contained in $M_{\mathcal{L}}^{fin}(X)$ iff Φ has a *finite support*, i.e. a finite subset $Y \subseteq X$ such that for all $\phi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$, $\phi \in \Phi$ iff $\phi\sigma_Y \in \Phi$, where σ_Y is the $\mathcal{P}(Y)$ -valuation for $\mathcal{P}(X)$ defined by $\sigma_Y(A) = A \cap Y$. \square

Proof (Theorem 60) It suffices to show that for T ω -accessible, the theory $h^\#(t)$ of $t \in TX$ as defined in the proof of Theorem 57 is in $M_{\mathcal{L}}^{fin}(X)$. By accessibility, we have $Y \subseteq X$ finite such that $t \in TY$. By naturality of predicate liftings, Y is a finite support of $h^\#(t)$, and hence $h^\#(t) \in M_{\mathcal{L}}^{fin}(X)$ by Lemma 61. \square

As indicated above, Theorems 57 and 60 allow us to replace rank-1 logics by functors in the definition of the coalgebraic semantics: an \mathcal{L} -structure based on a functor T may equivalently be regarded as a natural transformation $T \rightarrow M_{\mathcal{L}}$; analogously, an \mathcal{L} -structure based on an ω -accessible functor T may be regarded as a natural transformation $T \rightarrow M_{\mathcal{L}}^{fin}$. One may then attempt to phrase properties of \mathcal{L} -structures in terms of natural transformations. E.g. we have

Proposition 62 An \mathcal{L} -structure \mathcal{M} based on T is separating iff the associated natural transformation $T \rightarrow M_{\mathcal{L}}$ is injective. \square

Thus, we have the following classification result.

Theorem 63 Up to natural isomorphism, the ω -accessible \mathcal{L} -structures for which \mathcal{L} is separating are precisely the substructures of the canonical finitely branching \mathcal{L} -structure $M_{\mathcal{L}}^{fin}$.

(It is almost but not quite true that one can replace separation by the Hennessy-Milner property in the above theorem, as in some corner cases, logics may have the Hennessy-Milner property without being separating [33]. One may however replace separation by the Hennessy-Milner property for a logic that extends \mathcal{L} with propositional atoms, as laid out in more detail in Remark 66 below.)

Proof All that remains to be shown is that subfunctors of ω -accessible set functors are again ω -accessible; this follows easily from [1], Proposition 5.2. \square

With a little additional infrastructure, we can also capture weak completeness, i.e. one-step completeness, at the level of natural transformations. For a subset Γ of Λ , let \mathcal{L}/Γ denote the modal logic with similarity type Γ whose axioms are all one-step Γ -formulas which are derivable in \mathcal{L} (in the sense of Definition 49). We have natural transformations

$$\begin{aligned} \pi^\Gamma : M_{\mathcal{L}} &\rightarrow M_{\mathcal{L}/\Gamma} \\ \Phi &\mapsto \Phi \cap \text{Prop}(\Gamma(\mathcal{P}(X))). \end{aligned}$$

Then one-step completeness is characterised as follows.

Proposition 64 The logic \mathcal{L} is one-step complete for an \mathcal{L} -structure \mathcal{M} based on T with associated natural transformation $\mu : T \rightarrow M_{\mathcal{L}}$ iff $\pi_X^{\Gamma} \mu_X$ is surjective for all finite subsets Γ of Λ and all finite sets X .

Proof ‘*Only if*’: Let $\Gamma \subseteq \Lambda$ and X be finite, and let $\Phi \in M_{\mathcal{L}/\Gamma}(X)$. Since $\Gamma(\mathcal{P}(X))$ is finite, Φ has a finite set Φ' of representatives modulo propositional equivalence. The conjunction $\bigwedge \Phi' \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ is one-step \mathcal{L} -consistent. By one-step completeness, there exists $t \in TX$ such that $t \models^X \bigwedge \Phi'$, and hence $\pi_X^{\Gamma}(\mu_X(t)) = \Phi$.

‘*If*’: Let $\psi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ be one-step \mathcal{L} -consistent. We have to prove that ψ is one-step satisfiable. By Lemma 52, we can assume that X is finite. Let Γ be the (finite) set of modal operators occurring in ψ . By the one-step Lindenbaum lemma (Lemma 39), ψ is contained in a maximally one-step consistent set $\Phi \subseteq \text{Prop}(\Gamma(\mathcal{P}(X)))$. By assumption, we have $t \in TX$ such that $\pi_X^{\Gamma}(\mu_X(t)) = \Phi$; then $t \models^X \psi$. \square

Corollary 65 If Λ is finite, then \mathcal{L} is one-step complete for an \mathcal{L} -structure \mathcal{M} based on T with associated natural transformation $\mu : T \rightarrow M_{\mathcal{L}}$ iff μ is surjective. In particular, $\mathcal{M}_{\mathcal{L}}^{\text{fin}}$ is, in this case, the only ω -accessible \mathcal{L} -structure for which \mathcal{L} is one-step complete and separating.

Proof Immediate by Proposition 64, as the π^{Γ} are surjective, and $\pi^{\Lambda} = id_{M_{\mathcal{L}}}$ is one of the π^{Γ} . \square

Remark 66 Given $\mathcal{L} = (\Lambda, \mathbf{R})$ and a countably infinite set U of propositional atoms, i.e. nullary modalities, let $\mathcal{L}[U]$ denote the logic $(\Lambda \cup U, \mathbf{R})$. Thus, $\mathcal{L}[U]$ is just the standardly expected extension of \mathcal{L} by propositional atoms. Let \mathcal{M} be a structure for \mathcal{L} with underlying functor T , and define the functor T_U by $T_U(X) = TX \times \mathcal{P}(U)$. Coalgebras for T_U are essentially T -models, i.e. T -coalgebras equipped with valuations for the propositional atoms. Then \mathcal{M} is extended to a structure $\mathcal{M}[U]$ for $\mathcal{L}[U]$ with underlying functor T_U by putting $\llbracket L \rrbracket_X^{\mathcal{M}[U]}(A_1, \dots, A_n) = \llbracket L \rrbracket_X^{\mathcal{M}}(A_1, \dots, A_n) \times \mathcal{P}(U)$ for $L \in \Lambda$ n -ary, and

$$\llbracket u \rrbracket_X^{\mathcal{M}[U]} = \{(t, B) \in TX \times \mathcal{P}(U) \mid u \in B\}$$

for $u \in U$. Essentially, this yields the expected interpretation of $\mathcal{L}[U]$ over T -models. It is easy to see that \mathcal{L} is separating for \mathcal{M} iff $\mathcal{L}[U]$ is separating for $\mathcal{M}[U]$, and that \mathcal{L} is one-step complete for \mathcal{M} iff $\mathcal{L}[U]$ is one-step complete for $\mathcal{M}[U]$ (this follows also from general modularity results [6, 37], as $\mathcal{M}[U]$ is a modular combination of \mathcal{M} and the obvious structure for U). Moreover, by [36], Proposition 5.3, $\mathcal{L}[U]$ is one-step complete for $\mathcal{M}[U]$ iff $\mathcal{L}[U]$ is weakly complete for $\mathcal{M}[U]$. Finally, if T is ω -accessible then T_U is also ω -accessible and hence has a final coalgebra C [2], and since each subset of U is realised as the set of valid atomic propositions in some state of C , C is infinite. By the results of [33], it follows that $\mathcal{L}[U]$ is separating for $\mathcal{M}[U]$ iff $\mathcal{L}[U]$ has the Hennessy-Milner property for $\mathcal{M}[U]$.

Summarising the above, we may replace the conditions that \mathcal{L} is separating or one-step complete for \mathcal{M} by requiring that $\mathcal{L}[U]$ has the Hennessy-Milner property or is weakly complete, respectively, for $\mathcal{M}[U]$ in the above results, thus obtaining classification theorems involving natural conditions on the logic at large rather than properties of the one-step logic. In particular, this allows us to turn Corollary 65 into the following statement:

If \mathcal{L} is finite and U is an infinite set of propositional atoms, then $\mathcal{M}_{\mathcal{L}}^{fin}$ is the only ω -accessible structure \mathcal{M} such that $\mathcal{L}[U]$ is weakly complete and has the Hennessy-Milner property for $\mathcal{M}[U]$.

Example 67 Let K_{ω} denote the multi-agent version of K with ω agents. In straightforward generalisation of Example 56.1, $M_{K_{\omega}}^{fin}$ is the countable power $(\mathcal{P}^{fin})^{\omega}$ of the finite powerset functor \mathcal{P}^{fin} . By the above classification results, the ω -accessible structures for which K_{ω} is one-step complete and separating are precisely the substructures of $M_{K_{\omega}}^{fin}$ induced by subfunctors T of $(\mathcal{P}^{fin})^{\omega}$ such that for each finite set X , each finite subset $I \subseteq \omega$ and each I -indexed family (A_i) of subsets $A_i \subseteq X$, there exists $(C_j)_{j < \omega} \in TX$ such that $C_i = A_i$ for all $i \in I$.

This illustrates that for infinite similarity types, a crisper classification result than Proposition 64 is unlikely. In particular, there is no smallest structure (w.r.t. subfunctor inclusion) among the class \mathfrak{S} of structures described above: the substructure induced by the subfunctor T of $(\mathcal{P}^{fin})^{\omega}$ where TX contains precisely the families $(C_j)_{j < \omega}$ such that $C_j = \emptyset$ for all but finitely many j is minimal in \mathfrak{S} , but not contained in the substructure $\mathcal{N} \in \mathfrak{S}$ induced by the subfunctor S of $(\mathcal{P}^{fin})^{\omega}$ where SX consists of all families $(C_j)_{j < \omega}$ such that $C_j \neq \emptyset$ for all but finitely many j .

Contrastingly, the structure $\mathcal{M}_{\mathbf{B}}$ based on the finite multiset functor \mathbf{B} is the smallest structure for which GML is one-step complete and separating: it is known that GML is one-step complete [36] and separating [33] for $\mathcal{M}_{\mathbf{B}}$. Moreover, since for X finite, every finite multiset over X can be uniquely described in $M_{\text{GML}}^{fin}(X) = \mathbf{B}_{\infty}(X)$ (see Example 56.2) by a (finite) one-step formula in GML, every further structure with these properties must contain $\mathcal{M}_{\mathbf{B}}$ as a substructure. Thus, the ω -accessible structures for which GML is one-step complete and separating (i.e. has the Hennessy-Milner property and is weakly complete when propositional atoms are included) are precisely the substructures of $\mathcal{M}_{\text{GML}}^{fin}$ induced by functors between \mathbf{B} and \mathbf{B}_{∞} .

5 Decidability

A benefit of the coalgebraic semantics constructed above is that we can now apply results on coalgebraic modal logic to arbitrary rank-1 modal logics, even when the latter lack a formal model-theoretic semantics. This includes in particular the generic decidability and complexity results of [34, 36], of which we now obtain purely syntactic versions. Throughout this section, we fix a rank-1 modal logic $\mathcal{L} = (\Lambda, \mathbf{R})$.

In [34], a generic finite model construction was given which yields criteria for decidability and upper complexity bounds for coalgebraic modal logics. The generic complexity bounds generally do not match known bounds in particular examples, typically *PSPACE*. This is remedied in [36], where a generic *PSPACE* decision procedure for coalgebraic modal logics based on a shallow model construction is given, at the price of stronger assumptions on the logic.

A crucial role in the algorithmic methods of [34] is played by the following localised version of the satisfiability problem:

Definition 68 (One-step satisfiability problem) The *one-step satisfiability problem* for a Λ -structure \mathcal{M} is to decide, given a finite set X and a conjunctive clause ψ over $\Lambda(\mathcal{P}(X))$, whether ψ is one-step satisfiable over \mathcal{M} .

By the results of [34], the satisfiability problem of a coalgebraic modal logic is

- decidable if its one-step satisfiability problem is decidable
- in *NEXPTIME* if one-step satisfiability is in *NP*
- in *EXPTIME* if one-step satisfiability is in *P*.

This instantiates to the canonical structure as follows.

Lemma 69 A formula $\psi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ is one-step satisfiable over $\mathcal{M}_{\mathcal{L}}$ iff ψ is one-step \mathcal{L} -consistent.

Proof Immediate by the one-step truth lemma (Lemma 32) and the one-step Lindenbaum lemma (Lemma 39). \square

Definition 70 (One-step consistency problem) The *one-step consistency problem* for \mathcal{L} is to decide, given a finite set X and a conjunctive clause ψ over $\Lambda(\mathcal{P}(X))$, whether ψ is one-step \mathcal{L} -consistent.

Corollary 71 The consistency problem of \mathcal{L} (i.e. whether an Λ -formula ϕ is \mathcal{L} -consistent) is

- decidable if the one-step consistency problem is decidable
- in *NEXPTIME* if one-step consistency is in *NP*
- in *EXPTIME* if one-step consistency is in *P*.

In order to turn the above corollary into a more directly applicable decidability criterion, we recall the notion of *rule contraction* from [36]:

Definition 72 (Closure under contraction) An instance $\phi\sigma/\psi\sigma$ of a rule ϕ/ψ over V is *contracted* if the clause $\psi\sigma$ over $\Lambda(V)$ is *contracted*, i.e. does not contain duplicate literals (over $\Lambda(V)$). We say that a set R of rules is *closed under contraction* if for every V -instance $\phi\sigma/\psi\sigma$ of a rule ϕ/ψ over V in R , there exists a contracted V -instance $\phi'\sigma'/\psi'\sigma'$ of a rule $\phi'/\psi' \in R$ such that $\psi'\sigma'$ propositionally entails $\psi\sigma$ and $\phi\sigma$ propositionally entails $\phi'\sigma'$.

The key feature of contraction closed rule sets is that indeed one only needs contracted instances of rules; we make this formal for one-step derivations:

Lemma 73 Let R be closed under contraction. Then a one-step formula $\phi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ over a set X is one-step derivable iff ϕ is propositionally entailed by the set

$$\{\psi\tau \mid \phi/\psi \in R, \tau \text{ a } \mathcal{P}(X)\text{-valuation, } X \models \phi\tau, \psi\tau \text{ contracted}\}.$$

\square

Sets of one-step rules can easily be closed under contraction: just add a rule $\phi\sigma/\psi'$ for every rule ϕ/ψ over V in R and every V -substitution σ , where ψ' is obtained from $\psi\sigma$ by removing duplicate literals (typically, the premise $\phi\sigma$ will be replaced by some suitable propositional equivalent). It is clear that the new rules are derivable from

the original ones, and moreover that every finite rule set has a finite closure under contraction. E.g. the modal logic K is represented in terms of one-step rules as

$$(N) \frac{a}{\Box a} \quad (RR) \frac{a \wedge b \rightarrow c}{\Box a \wedge \Box b \rightarrow \Box c},$$

and closing under contraction additionally yields (besides the trivial rule \top/\top) the monotonicity rule

$$(M) \frac{a \rightarrow b}{\Box a \rightarrow \Box b}.$$

Corollary 74 The consistency problem of $\mathcal{L} = (\Lambda, R)$ is decidable if Λ is finite and R is a set of one-step rules which is closed under contraction and recursive (i.e. it is decidable whether a one-step rule is contained in R up to propositional equivalence of premises).

Proof By Corollary 71, it suffices to prove that one-step consistency is decidable. Since propositional entailment from finite sets of formulas is decidable, this reduces by Lemma 73 to showing that given a finite set X , the set

$$\Psi = \{\psi\tau \mid \phi/\psi \in R, \tau \text{ a } \mathcal{P}(X)\text{-valuation, } X \models \phi\tau, \psi\tau \text{ contracted}\} \quad (*)$$

is computable, which due to finiteness of the set of *contracted* clauses over $\Lambda(\mathcal{P}(X))$ amounts to showing that it is decidable whether a given contracted clause χ belongs to Ψ . But this is clear: recall that by Remark 3, we can assume that all variables in a rule ϕ/ψ over V occur in ψ (in particular that V is finite). In the notation of (*), ψ and τ can hence be read off directly from χ , choosing a suitable standard naming scheme for propositional variables (e.g. a_i for the i -th variable in ψ). One then has to check for all possible choices of $\phi \in \text{Prop}(V)$ whether $X \models \phi\tau$ and whether $\phi/\psi \in R$; by the assumptions made, it suffices to check some finite set of representatives of $\text{Prop}(V)$ up to propositional equivalence. \square

Corollary 75 The consistency problem of $\mathcal{L} = (\Lambda, R)$ is decidable if Λ and R are finite.

Proof W.l.o.g., R consists of one-step rules. By the above, R has a finite (hence recursive) closure under contraction. \square

Remark 76 The generic *PSPACE* decision procedure of [36] is based on rule sets which are closed under contraction and additionally under rule resolution. While one can use the results presented in the present work to turn this originally semantics-based result into a purely syntactic criterion [38], it has transpired that there is also a direct syntactic proof of this criterion (see the extended version of [36]).

The results of [36] moreover imply a decidability criterion that complements Corollary 74 in that the finiteness assumption for Λ is removed, but the stronger assumption is made that R is resolution closed; as closure under resolution is generally a more complex process than closure under contraction, decidability is typically easier to establish for the contraction closure of a given recursive set of rules than for the resolution closure.

Corollary 75 partially reproves a result by Lewis [23]; the latter applies more generally to *non-iterative* modal logics, i.e. logics whose axioms do not nest modal

operators (but unlike one-step formulas may contain top-level propositional variables, as in the T -axiom $\Box a \rightarrow a$). It is the subject of future research to extend the present results, and in fact the entire framework of coalgebraic modal logic, to cover also non-iterative modal logics by means of copointed functors. On the other hand, Corollary 74 improves on Lewis' result by only requiring the axiomatisation to be recursive (rather than finite). However, we are not presently aware of a realistic example of an infinitely axiomatised modal logic with finitely many modal operators.

6 Example: Deontic Logic

A typical application area for the above results are modal logics that come from a philosophical background, such as epistemic and deontic logics, which are often defined either without any reference to semantics at all or with a neighbourhood semantics essentially equivalent to the canonical semantics described above. Deontic logics [14, 40], which have received much recent interest in computer science as logics for obligations of agents, are moreover often axiomatised in rank 1.

Standard deontic logic [5] is just the modal logic KD . This has been criticized on the grounds that it entails the *deontic explosion*: if O is the modal obligation operator 'it ought to be the case that', then the K -axiom, formulated as $(Oa \wedge Ob) \leftrightarrow O(a \wedge b)$, implies that in the presence of a single deontic dilemma, everything is obligatory, i.e. $Oa \wedge O\neg a \rightarrow Ob$. Some approaches to this problem are summarized in [11], where it is advocated to eliminate the deontic explosion by restricting at least one direction of K to the case that $a \wedge b$ is *permitted*, i.e. to $P(a \wedge b)$, where P is the dual $\neg O\neg$ of O . This leads to the axioms

$$\begin{aligned} \text{(PM)} \quad & O(a \wedge b) \wedge P(a \wedge b) \rightarrow Oa \\ \text{(PAND)} \quad & Oa \wedge Ob \wedge P(a \wedge b) \rightarrow O(a \wedge b) \end{aligned}$$

(in [11], (PM) is formulated as a rule (RPM)). Two systems are proposed (both including the congruence rule): given the further axioms (N) OT , (P) $\neg O\perp$, and

$$\text{(ADD)} \quad (Oa \wedge Ob) \rightarrow O(a \wedge b),$$

DPM.1 is determined by (PM), (N), and (ADD), while DPM.2 is given by (PM), (PAND), (N), and (P). A further system PA, consisting of (PAND), (P), (N), and the standard monotonicity axiom is rejected, as it still leads to a form of deontic explosion where everything permitted is obligatory in the presence of a dilemma.

It is shown in [11] that DPM.1 and DPM.2 are sound and *weakly* complete w.r.t. the obvious classes of neighbourhood frames, and that both logics are decidable; the proofs are rather involved. In our framework, the situation presents itself as follows. The neighbourhood semantics of [11] is easily seen to be precisely the canonical semantics; the new insight here is that the semantics is coalgebraic. The rest is for free: by Theorem 44, both DPM.1 and DPM.2 are even *strongly* complete (the reason that the strong completeness proof fails in [11] is that an explicit construction of a canonical neighbourhood model is attempted). Decidability is immediate by Corollary 75; the finite model property (proved in [11] using filtrations) follows from the results of [34]. (The same holds for PA, and in fact for rather arbitrary variations of the axiom system.) A challenge that remains is to establish that DPM.1 and DPM.2 are in $PSPACE$ by the methods of [36] or [39].

7 Conclusion

We have established that every modal logic \mathcal{L} of rank 1 has a canonical coalgebraic semantics $\mathcal{M}_{\mathcal{L}}$, for which \mathcal{L} is sound and strongly complete. Moreover, \mathcal{L} has a canonical finitely branching coalgebraic semantics $\mathcal{M}_{\mathcal{L}}^{fn}$, for which \mathcal{L} is sound and weakly complete and has the Hennessy-Milner property. All finitely branching semantics for which \mathcal{L} has the Hennessy-Milner property are obtained as substructures of $\mathcal{M}_{\mathcal{L}}^{fn}$. Interestingly, if the similarity type of \mathcal{L} is finite, then $\mathcal{M}_{\mathcal{L}}^{fn}$ is uniquely determined as a finitely branching semantics for which \mathcal{L} is weakly complete and has the Hennessy-Milner property.

These results provide a converse to the previous insight that every coalgebraic modal logic can be axiomatized in rank 1 [34]. They have allowed us to formulate purely syntactic versions of semantics-based generic decidability and complexity criteria for coalgebraic modal logic [34], including e.g. the result that every rank-1 logic with finitely many modal operators whose contraction closure is recursive is decidable. This result is related to a result by Lewis [23], which applies to a slightly more general class of logics called *non-iterative* logics (where axioms still cannot nest modal operators, but may contain top-level propositional variables), but which makes the stronger assumption that the axiomatisation is finite. Extending our results to non-iterative logics is the subject of future work.

We have applied our framework to recently defined versions of deontic logic which accommodate deontic dilemmas [11]. In particular, we have obtained decidability and strong completeness for these logics as immediate consequences of our generic results, while the original work has rather involved proofs and moreover establishes only decidability and weak completeness. Application of the generic *PSPACE* upper bound [36] to these logics remains an open problem.

We emphasise that the restriction to rank 1 is not an inherent limitation of the coalgebraic approach — the fact that coalgebraic modal logics are of rank 1 is due to the interpretation of these logics over the whole class of coalgebras for the relevant functor (in analogy to the standard modal logic K), and logics outside rank 1 may be modelled by passing to suitable subclasses (covarieties) of coalgebras, in generalisation of e.g. the interpretation of $K4$ over transitive Kripke frames (i.e. \mathcal{P} -coalgebras); general completeness and decidability results for such logics are the subject of ongoing [28] and future work. A further point of interest is to obtain completeness and decidability results for coalgebraic modal logics with iteration, i.e. the coalgebraic counterpart of CTL. Finally, results on generic strong completeness criteria for coalgebraic modal logics (other than canonical structures) are forthcoming [35].

Acknowledgements

The authors wish to thank Alexander Kurz for inspiring discussions, the anonymous referees of this work and [38] for their helpful suggestions, and Erwin R. Catesbeiana for interjections regarding maximally consistent sets.

References

- [1] J. Adámek and H.-E. Porst. On tree coalgebras and coalgebra presentations. *Theoret. Comput. Sci.*, 311:257–283, 2004.
- [2] M. Barr. Terminal coalgebras in well-founded set theory. *Theoret. Comput. Sci.*, 114:299–315, 1993.
- [3] M. Bonsangue and A. Kurz. Duality for logics of transition systems. In V. Sassone, editor, *Foundations of Software Science And Computations Structures, FOSSACS 2005*, volume 3441 of *Lect. Notes Comput. Sci.*, pages 455–469. Springer, 2005.
- [4] G. Boole. *An Investigation of the Laws of Thought on Which are Founded the Mathematical Theories of Logic and Probabilities*. Macmillan, 1854. Reprinted with corrections, Dover Publications, New York, 1958.
- [5] B. Chellas. *Modal Logic*. Cambridge, 1980.
- [6] C. Cirstea and D. Pattinson. Modular construction of modal logics. In *Concurrency Theory*, volume 3170 of *Lect. Notes Comput. Sci.*, pages 258–275. Springer, 2004.
- [7] G. D’Agostino and A. Visser. Finality regained: A coalgebraic study of Scott-sets and multisets. *Arch. Math. Logic*, 41:267–298, 2002.
- [8] S. Demri and D. Lugiez. Presburger modal logic is only PSPACE-complete. In *Automated Reasoning, IJCAR 06*, volume 4130 of *Lect. Notes Artificial Intelligence*, pages 541–556. Springer, 2006.
- [9] R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41:340–367, 1994.
- [10] K. Fine. In so many possible worlds. *Notre Dame J. Formal Logic*, 13:516–520, 1972.
- [11] L. Goble. A proposal for dealing with deontic dilemmas. In A. Lomuscio and D. Nute, editors, *Deontic Logic in Computer Science, DEON 2004*, volume 3065 of *Lect. Notes Artificial Intelligence*, pages 74–113. Springer, 2004.
- [12] H.-P. Gumm. Functors for coalgebras. *Algebra Universalis*, 45:135–147, 2001.
- [13] A. Heifetz and P. Mongin. Probabilistic logic for type spaces. *Games and Economic Behavior*, 35:31–53, 2001.
- [14] R. Hilpinen. Deontic logic. In L. Goble, editor, *The Blackwell Guide to Philosophical Logic*. Blackwell, 2001.
- [15] B. Jacobs. Towards a duality result in the modal logic of coalgebras. In H. Reichel, editor, *Coalgebraic Methods in Computer Science, CMCS 2000*, volume 33 of *Electron. Notes Theoret. Comput. Sci.* Elsevier, 2000.

- [16] B. Jacobs. Many-sorted coalgebraic modal logic: a model-theoretic study. *Theoret. Informatics and Applications*, 35(1):31–59, 2001.
- [17] C. Kupke, A. Kurz, and D. Pattinson. Algebraic semantics for coalgebraic logics. In J. Adámek and S. Milius, editors, *Coalgebraic Methods in Computer Science*, volume 106 of *Electron. Notes Theoret. Comput. Sci.*, pages 219–241. Elsevier, 2004.
- [18] C. Kupke, A. Kurz, and D. Pattinson. Ultrafilter extensions for coalgebras. In J. L. Fiadeiro, N. Harman, M. Roggenbach, and J. Rutten, editors, *Algebra and Coalgebra in Computer Science*, volume 3629 of *Lect. Notes Comput. Sci.*, pages 263–277. Springer, 2005.
- [19] A. Kurz. Specifying coalgebras with modal logic. *Theoret. Comput. Sci.*, 260:119–138, 2001.
- [20] A. Kurz and J. Rosický. Strongly complete logics for coalgebras. Draft, July 2006.
- [21] K. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inform. Comput.*, 94:1–28, 1991.
- [22] V. S. Levchenkov. Solution of equations in Boolean algebra. *Comput. Math. Model.*, 11:154–163, 2000. Translated from *Nelineinaya Dinamika i Upravlenie*, pp. 119–132, 1999.
- [23] D. Lewis. Intensional logics without iterative axioms. *J. Philos. Logic*, 3:457–466, 1975.
- [24] L. Moss. Coalgebraic logic. *Ann. Pure Appl. Logic*, 96:277–317, 1999.
- [25] D. Pattinson. Semantical principles in the modal logic of coalgebras. In A. Ferreira and H. Reichel, editors, *Theoretical Aspects of Computer Science, STACS 2001*, volume 2010 of *Lect. Notes Comput. Sci.*, pages 514–526. Springer, 2001.
- [26] D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.*, 309:177–193, 2003.
- [27] D. Pattinson. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Logic*, 45:19–33, 2004.
- [28] D. Pattinson and L. Schröder. Beyond rank 1: Algebraic semantics and finite models for coalgebraic logics. In R. Amadio, editor, *Foundations of Software Science and Computation Structures, FOSSACS 2008*, volume 4962 of *Lect. Notes Comput. Sci.*, pages 66–80, 2008.
- [29] M. Pauly. A modal logic for coalitional power in games. *J. Logic Comput.*, 12:149–166, 2002.
- [30] M. Röbiger. Coalgebras and modal logic. In B. Jacobs, L. Moss, H. Reichel, and J. Rutten, editors, *Coalgebraic Methods in Computer Science*, volume 33 of *Electron. Notes Theoret. Comput. Sci.* Elsevier, 2000.

- [31] M. Rößiger. From Modal Logic to Terminal Coalgebras. *Theoret. Comput. Sci.*, 260:209–228, 2001.
- [32] J. Rutten. Universal coalgebra: A theory of systems. *Theoret. Comput. Sci.*, 249:3–80, 2000.
- [33] L. Schröder. Expressivity of coalgebraic modal logic: the limits and beyond. *Theoret. Comput. Sci.* In press. Earlier version in V. Sassone, editor, Foundations of Software Science And Computation Structures, FOSSACS 2005, vol. 3441 of *Lect. Notes Comput. Sci.*, Springer, 2005, pp. 440–454.
- [34] L. Schröder. A finite model construction for coalgebraic modal logic. *J. Logic Algebraic Programming*, 73:97–110, 2007.
- [35] L. Schröder and D. Pattinson. Generic strong completeness results in modal logic. In preparation.
- [36] L. Schröder and D. Pattinson. PSPACE reasoning for rank-1 modal logics. In R. Alur, editor, *Logic in Computer Science, LICS 2006*, pages 231–240. IEEE, 2006. Extended version to appear in *ACM Trans. Comput. Logic*; preprint available as e-print arXiv:0706.4044.
- [37] L. Schröder and D. Pattinson. Modular algorithms for heterogeneous modal logics. In L. Age, A. Tarlecki, and C. Cachin, editors, *Automata, Languages and Programming, ICALP 07*, volume 4596 of *Lect. Notes Comput. Sci.*, pages 459–471. Springer, 2007.
- [38] L. Schröder and D. Pattinson. Rank-1 modal logics are coalgebraic. In W. Thomas and P. Weil, editors, *Theoretical Aspects of Computer Science, STACS 2007*, volume 4393 of *Lect. Notes Comput. Sci.*, pages 573–585. Springer, 2007.
- [39] L. Schröder and D. Pattinson. Shallow models for non-iterative modal logics. Available as e-print arXiv:0802.0116, 2008.
- [40] G. H. von Wright. Deontic logic. *Mind*, 60:1–15, 1951.