Relating Structure and Power: Comonadic Semantics for Computational Resources

Samson Abramsky[1]

Oxford University Department of Computer Science
Wolfson Building, Parks Road, Oxford OX1 3QD, U.K.
samson.abramsky@cs.ox.ac.uk
https://orcid.org/0000-0003-3921-6637

Nihil Shah

Oxford University Department of Computer Science
Wolfson Building, Parks Road, Oxford OX1 3QD, U.K.
nihil@berkeley.edu
https://orcid.org/0000-0003-2844-0828

—— **Abstract** ——

Combinatorial games are widely used in finite model theory, constraint satisfaction, modal logic and concurrency theory to characterize logical equivalences between structures. In particular, Ehrenfeucht-Fraïssé games, pebble games, and bisimulation games play a central role. We show how each of these types of games can be described in terms of an indexed family of comonads on the category of relational structures and homomorphisms. The index k is a resource parameter which bounds the degree of access to the underlying structure. The coKleisli categories for these comonads can be used to give syntax-free characterizations of a wide range of important logical equivalences. Moreover, the coalgebras for these indexed comonads can be used to characterize key combinatorial parameters: tree-depth for the Ehrenfeucht-Fraïssé comonad, tree-width for the pebbling comonad, and synchronization-tree depth for the modal unfolding comonad. These results pave the way for systematic connections between two major branches of the field of logic in computer science which hitherto have been almost disjoint: categorical semantics, and finite and algorithmic model theory.

1 Introduction

There is a remarkable divide in the field of logic in Computer Science, between two distinct strands: one focussing on semantics and compositionality ("Structure"), the other on expressiveness and complexity ("Power"). It is remarkable because these two fundamental aspects of our field are studied using almost disjoint technical languages and methods, by almost disjoint research communities. We believe that bridging this divide is a major issue in Computer Science, and may hold the key to fundamental advances in the field.

In this paper, we develop a novel approach to relating categorical semantics, which exemplifies the first strand, to finite model theory, which exemplifies the second. It builds on the ideas introduced in [2], but goes much further, showing clearly that there is a strong and robust connection, which can serve as a basis for many further developments.

The setting

Relational structures and the homomorphisms between them play a fundamental rôle in finite model theory, constraint satisfaction and database theory. The existence of a homomorphism $A \to B$ is an equivalent formulation of constraint satisfaction, and also equivalent to the preservation of existential positive sentences [7]. This setting also generalizes what has become a central perspective in graph theory [15].

Model theory and deception

In a sense, the purpose of model theory is "deception". It allows us to see structures not "as they really are", *i.e.* up to isomorphism, but only up to *definable properties*, where definability is relative to a logical language $\mathcal{L}$. The key notion is *logical equivalence* $\equiv^{\mathcal{L}}$. Given structures $\mathcal{A}$, $\mathcal{B}$ over the same vocabulary:

$$\mathcal{A} \equiv^{\mathcal{L}} \mathcal{B} \quad \stackrel{\Delta}{\iff} \quad \forall \varphi \in \mathcal{L}. \, \mathcal{A} \models \varphi \iff \mathcal{B} \models \varphi.$$

If a class of structures $\mathcal{K}$ is definable in $\mathcal{L}$, then it must be saturated under $\equiv^{\mathcal{L}}$. Moreover, for a wide class of cases of interest in finite model theory, the converse holds [20].

The idea of syntax-independent characterizations of logical equivalence is quite a classical one in model theory, exemplified by the Keisler-Shelah theorem [30]. It acquires additional significance in finite model theory, where model comparison games such as Ehrenfeucht-Fraïssé games, pebble games and bisimulation games play a central role [21].

We offer a new perspective on these ideas. We shall study these games, not as external artefacts, but as semantic constructions in their own right. Each model-theoretic comparison game encodes "deception" in terms of limited access to the structure. These limitations are indexed by a parameter which quantifies the resources which control this access. For Ehrenfeucht-Fraïssé games and bisimulation games, this is the number of rounds; for pebble games, the number of pebbles.

Main Results

We now give a conceptual overview of our main results. Technical details will be provided in the following sections.

We shall consider three forms of model comparison game: Ehrenfeucht-Fraïssé games, pebble games and bisimulation games [21]. For each of these notions of game G, and value of the resource parameter k, we shall define a corresponding *comonad* $\mathbb{C}_k$ on the category of relational structures and homomorphisms over some relational vocabulary. For each structure $\mathcal{A}$, $\mathbb{C}_k\mathcal{A}$ is another structure over the same vocabulary, which encodes the limited access to $\mathcal{A}$ afforded by playing the game on $\mathcal{A}$ with k resources. There is always an associated homomorphism $\varepsilon_{\mathcal{A}} : \mathbb{C}_k\mathcal{A} \to \mathcal{A}$ (the *counit* of the comonad), so that $\mathbb{C}_k\mathcal{A}$ "covers" $\mathcal{A}$. Moreover, given a homomorphism $h : \mathbb{C}_k\mathcal{A} \to \mathcal{B}$, there is a *Kleisli coextension* homomorphism $h^* : \mathbb{C}_k\mathcal{A} \to \mathbb{C}_k\mathcal{B}$. This allows us to form the *coKleisli category* $\mathsf{Kl}(\mathbb{C}_k)$ for the comonad. The objects are relational structures, while the morphisms from $\mathcal{A}$ to $\mathcal{B}$ in $\mathsf{Kl}(\mathbb{C}_k)$ are exactly the homomorphisms of the form $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$. Composition of these morphisms uses the Kleisli coextension. The connection between this construction and the corresponding form of game G is expressed by the following result:

▶ **Theorem 1.** *The following are equivalent:*
1. *There is a coKleisli morphism $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$*
2. *Duplicator has a winning strategy for the existential G-game with k resources, played from $\mathcal{A}$ to $\mathcal{B}$.*

The existential form of the game has only a "forth" aspect, without the "back". This means that Spoiler can only play in $\mathcal{A}$, while Duplicator only plays in $\mathcal{B}$. This corresponds to the asymmetric form of the coKleisli morphisms $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$. Intuitively, Spoiler plays in $\mathbb{C}_k\mathcal{A}$, which gives them limited access to $\mathcal{A}$, while Duplicator plays in $\mathcal{B}$. The Kleisli coextension guarantees that Duplicator's strategies can always be lifted to $\mathbb{C}_k\mathcal{B}$; while we can always compose a strategy $\mathbb{C}_k\mathcal{A} \to \mathbb{C}_k\mathcal{B}$ with the counit on $\mathcal{B}$ to obtain a coKleisli morphism.

This asymmetric form may seem to limit the scope of this approach, but in fact this is not the case. For each of these comonads $\mathbb{C}_k$, we have the following equivalences:
- $\mathcal{A} \rightleftarrows_k \mathcal{B}$ iff there are coKleisli morphisms $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$ and $\mathbb{C}_k\mathcal{B} \to \mathcal{A}$. Note that there need be no relationship between these morphisms.
- $\mathcal{A} \cong_{\mathsf{Kl}(\mathbb{C}_k)} \mathcal{B}$ iff $\mathcal{A}$ and $\mathcal{B}$ are isomorphic in the coKleisli category $\mathsf{Kl}(\mathbb{C}_k)$. This means that there are morphisms $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$ and $\mathbb{C}_k\mathcal{B} \to \mathcal{A}$ which are inverses of each other in $\mathsf{Kl}(\mathbb{C}_k)$.

Clearly, $\cong_{\mathsf{Kl}(\mathbb{C}_k)}$ strictly implies $\rightleftarrows_k$. We can also define an intermediate "back-and-forth" equivalence $\leftrightarrow_k$, parameterized by a winning condition $\mathsf{W}_{\mathcal{A},\mathcal{B}} \subseteq \mathbb{C}_k\mathcal{A} \times \mathbb{C}_k\mathcal{B}$.

For each of our three types of game, there are corresponding fragments $\mathcal{L}_k$ of first-order logic:
- For Ehrenfeucht-Fraïssé games, $\mathcal{L}_k$ is the fragment of quantifier-rank $\leq k$.
- For pebble games, $\mathcal{L}_k$ is the k-variable fragment.
- For bismulation games over relational vocabularies with symbols of arity at most 2, $\mathcal{L}_k$ is the modal fragment [4] with modal depth $\leq k$.

In each case, we write $\exists\mathcal{L}_k$ for the existential positive fragment of $\mathcal{L}_k$, and $\mathcal{L}_k(\#)$ for the extension of $\mathcal{L}_k$ with counting quantifiers [21].

We can now state our first main result, in a suitably generic form.

▶ **Theorem 2.** *For finite structures $\mathcal{A}$ and $\mathcal{B}$:*

(1) $\quad \mathcal{A} \equiv^{\exists\mathcal{L}_k} \mathcal{B} \quad\Longleftrightarrow\quad \mathcal{A} \rightleftarrows_k \mathcal{B}.$
(2) $\quad \mathcal{A} \equiv^{\mathcal{L}_k} \mathcal{B} \quad\Longleftrightarrow\quad \mathcal{A} \leftrightarrow_k \mathcal{B}.$
(3) $\quad \mathcal{A} \equiv^{\mathcal{L}_k(\#)} \mathcal{B} \quad\Longleftrightarrow\quad \mathcal{A} \cong_{\mathsf{Kl}(\mathbb{C}_k)} \mathcal{B}.$

Note that this is really a family of three theorems, one for each type of game G. Thus in each case, we capture the salient logical equivalences in syntax-free, categorical form.

We now turn to the significance of indexing by the resource parameter k. When $k \leq l$, we have a natural inclusion morphism $\mathbb{C}_k\mathcal{A} \to \mathbb{C}_l\mathcal{A}$, since playing with k resources is a special case of playing with $l \geq k$ resources. This tells us that the smaller k is, the easier it is to find a morphism $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$. Intuitively, the more we restrict Spoiler's abilities to access the structure of $\mathcal{A}$, the easier it is for Duplicator to win the game.

The contrary analysis applies to morphisms $A \to \mathbb{C}_k\mathcal{B}$. The smaller k is, the *harder* it is find such a morphism. Note, however, that if $\mathcal{A}$ is a finite structure of cardinality k, then $\mathcal{A} \rightleftarrows_k \mathbb{C}_k\mathcal{A}$. In this case, with k resources we can access the whole of $\mathcal{A}$. What can we say when k is strictly smaller than the cardinality of $\mathcal{A}$?

It turns out that there is a beautiful connection between these indexed comonads and combinatorial invariants of structures. This is mediated by the notion of *coalgebra*, another fundamental (and completely general) aspect of comonads. A coalgebra for a comonad $\mathbb{C}_k$ on a structure $\mathcal{A}$ is a morphism $\mathcal{A} \to \mathbb{C}_k\mathcal{A}$ satisfying certain properties. We define the *coalgebra number* of a structure $\mathcal{A}$, with respect to the indexed family of comonads $\mathbb{C}_k$, to be the least k such that there is a $\mathbb{C}_k$-coalgebra on $\mathcal{A}$.

We now come to our second main result.

▶ **Theorem 3.**

▬ *For the pebbling comonad, the coalgebra number of $\mathcal{A}$ corresponds precisely to the* tree-width *of $\mathcal{A}$.*

▬ *For the Ehrenfeucht-Fraïssé comonad, the coalgebra number of $\mathcal{A}$ corresponds precisely to the* tree-depth *of $\mathcal{A}$ [27].*

▬ *For the modal comonad, the coalgebra number of $\mathcal{A}$ corresponds precisely to the* modal unfolding depth *of $\mathcal{A}$.*

The main idea behind these results is that coalgebras on $\mathcal{A}$ are in bijective correspondence with decompositions of $\mathcal{A}$ of the appropriate form. We thus obtain categorical characterizations of these key combinatorial parameters.

2 Game Comonads

In this section we will define the comonads corresponding to each of the forms of model comparison game we consider.

Firstly, a few notational preliminaries. A relational vocabulary σ is a set of relation symbols R, each with a specified positive integer arity. A σ-structure $\mathcal{A}$ is given by a set A, the universe of the structure, and for each R in σ with arity k, a relation $R^{\mathcal{A}} \subseteq A^k$. A homomorphism $h : \mathcal{A} \to \mathcal{B}$ is a function $h : A \to B$ such that, for each relation symbol R of arity k in σ, for all $a_1, \ldots, a_k$ in A: $R^{\mathcal{A}}(a_1, \ldots, a_k) \Rightarrow R^{\mathcal{B}}(h(a_1), \ldots, h(a_k))$. We write $\mathcal{R}(\sigma)$ for the category of σ-structures and homomorphisms.

We shall write $A^{\leq k}$ for the set of non-empty sequences of length $\leq k$ on a set A. We use list notation $[a_1, \ldots, a_j]$ for such sequences, and indicate concatenation by juxtaposition. We write $s \sqsubseteq t$ for the prefix ordering on sequences. If $s \sqsubseteq t$, there is a unique s' such that $ss' = t$, which we refer to as the suffix of s in t. For each positive integer n, we define $\mathbf{n} := \{1, \ldots, n\}$.

We shall need a few notions on posets. The comparability relation on a poset $(P, \leq)$ is $x \uparrow y$ iff $x \leq y$ or $y \leq x$. A chain in a poset $(P, \leq)$ is a subset $C \subseteq P$ such that, for all $x, y \in C$, $x \uparrow y$. A *forest* is a poset $(F, \leq)$ such that, for all $x \in F$, the set of predecessors $\downarrow(x) := \{y \in F \mid y \leq x\}$ is a finite chain. The height $\mathsf{ht}(F)$ of a forest F is $\max_C |C|$, where C ranges over chains in F.

We recall that a comonad (G, ε, δ) on a category $\mathcal{C}$ is given by a functor $G : \mathcal{C} \to \mathcal{C}$, and natural transformations $\varepsilon : G \Rightarrow I$ (the counit), and $\delta : G \Rightarrow G^2$ (the comultiplication), subject to the conditions that the following diagrams commute, for all objects A of $\mathcal{C}$:

$$
\begin{array}{ccc}
GA & \xrightarrow{\delta_A} & GGA \\
{\scriptstyle \delta_A}\downarrow & & \downarrow{\scriptstyle G\delta_A} \\
GGA & \xrightarrow{\delta_{GA}} & GGGA
\end{array}
\qquad\qquad
\begin{array}{ccc}
GA & \xrightarrow{\delta_A} & GGA \\
{\scriptstyle \delta_A}\downarrow & & \downarrow{\scriptstyle G\varepsilon_A} \\
GGA & \xrightarrow{\varepsilon_{GA}} & GA
\end{array}
$$

An equivalent formulation is *comonad in Kleisli form* [23]. This is given by an object map G, arrows $\varepsilon_A : GA \to A$ for every object A of $\mathcal{C}$, and a Kleisli coextension operation which takes $f : GA \to B$ to $f^* : GA \to GB$. These must satisfy the following equations:

$$\varepsilon_A^* = \mathsf{id}_{GA}, \qquad \varepsilon \circ f^* = f, \qquad (g \circ f^*)^* = g^* \circ f^*.$$

We can then extend G to a functor by $Gf = (f \circ \varepsilon)^*$; and if we define the comultiplication $\delta : G \Rightarrow G^2$ by $\delta_A = \mathsf{id}_{GA}^*$, then (G, ε, δ) is a comonad in the standard sense. Conversely, given a comonad (G, ε, δ), we can define the coextension by $f^* = Gf \circ \delta_A$. This allows us to define the coKleisli category $\mathsf{Kl}(G)$, with objects the same as those of $\mathcal{C}$, and morphisms from A to B given by the morphisms in $\mathcal{C}$ of the form $GA \to B$. Kleisli composition of $f : GA \to B$ with $g : GB \to C$ is given by $g \bullet f := g \circ f^*$.

2.1 The Ehrenfeucht-Fraïssé Comonad

We shall define a comonad $\mathbb{E}_k$ on $\mathcal{R}(\sigma)$ for each positive integer k. It will be convenient to define $\mathbb{E}_k$ in Kleisli form. For each structure $\mathcal{A}$, we define a new structure $\mathbb{E}_k\mathcal{A}$, with universe $\mathbb{E}_k A := A^{\leq k}$. We define the map $\varepsilon_{\mathcal{A}} : \mathbb{E}_k A \to A$ by $\varepsilon_{\mathcal{A}}[a_1, \ldots, a_j] = a_j$. For each relation symbol R of arity n, we define $R^{\mathbb{E}_k\mathcal{A}}$ to be the set of n-tuples $(s_1, \ldots, s_n)$ of sequences which are pairwise comparable in the prefix ordering, and such that $R^{\mathcal{A}}(\varepsilon_{\mathcal{A}} s_1, \ldots, \varepsilon_{\mathcal{A}} s_n)$. Finally, we define the coextension. Given a homomorphism $f : \mathbb{E}_k\mathcal{A} \to \mathcal{B}$, we define $f^* : A^{\leq k} \to B^{\leq k}$ by $f^*[a_1, \ldots, a_j] = [b_1, \ldots, b_j]$, where $b_i = f[a_1, \ldots, a_i]$, $1 \leq i \leq j$.

▶ **Proposition 4.** *The triple* $(\mathbb{E}_k, \varepsilon, (\cdot)^*)$ *is a comonad in Kleisli form.*

   Intuitively, an element of $A^{\leq k}$ represents a play in $\mathcal{A}$ of length $\leq k$. A coKleisli morphism $\mathbb{E}_k\mathcal{A} \to \mathcal{B}$ represents a Duplicator strategy for the existential Ehrenfeucht-Fraïssé game with k rounds, where Spoiler plays only in $\mathcal{A}$, and $b_i = f[a_1, \ldots, a_i]$ represents Duplicator's response in $\mathcal{B}$ to the i'th move by Spoiler. The winning condition for Duplicator in this game is that, after k rounds have been played, the induced relation $\{(a_i, b_i) \mid 1 \leq i \leq k\}$ is a partial homomorphism from $\mathcal{A}$ to $\mathcal{B}$.

   These intuitions are confirmed by the following result.

▶ **Theorem 5.** *The following are equivalent:*
1. *There is a homomorphism* $\mathbb{E}_k\mathcal{A} \to \mathcal{B}$.
2. *Duplicator has a winning strategy for the existential Ehrenfeucht-Fraïssé game with k rounds, played from $\mathcal{A}$ to $\mathcal{B}$.*

2.2 The Pebbling Comonad

We now turn to the case of pebble games. The following construction appeared in [2]. Given a structure $\mathcal{A}$, we define $\mathbb{P}_k\mathcal{A}$, which will represent plays of the k-pebble game on $\mathcal{A}$.[2] The universe is $(\mathbf{k} \times A)^+$, the set of finite non-empty sequences of moves (p, a), where $p \in \mathbf{k}$ is a pebble index, and $a \in A$. We shall use the notation $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ for these sequences, which may be of arbitrary length. Thus the universe of $\mathbb{P}_k\mathcal{A}$ is always infinite, even if $\mathcal{A}$ is a finite structure. This is unavoidable, by [2, Theorem 7]. We define $\varepsilon_{\mathcal{A}} : \mathbb{P}_k A \to A$ to send a play $[(p_1, a_1), \ldots, (p_n, a_n)]$ to a_n, the A-component of its last move.

   Given an n-ary relation $R \in \sigma$, we define $R^{\mathbb{P}_k\mathcal{A}}(s_1, \ldots, s_n)$ iff (1) the s_i are pairwise comparable in the prefix ordering; (2) the pebble index of the last move in each s_i does not appear in the suffix of s_i in s_j for any $s_j \sqsupseteq s_i$; and (3) $R^{\mathcal{A}}(\varepsilon_{\mathcal{A}}(s_1), \ldots, \varepsilon_{\mathcal{A}}(s_n))$.

   Finally, given a homomorphism $f : \mathbb{P}_k\mathcal{A} \to \mathcal{B}$, we define $f^* : \mathbb{P}_k A \to \mathbb{P}_k B$ by
$f^*[(p_1, a_1), \ldots, (p_j, a_j)] = [(p_1, b_1), \ldots, (p_j, b_j)]$, where $b_i = f[(p_1, a_1), \ldots, (p_i, a_i)]$, $1 \leq i \leq j$.

▶ **Proposition 6.** *The triple* $(\mathbb{P}_k, \varepsilon, (\cdot)^*)$ *is a comonad in Kleisli form.*

   The following is [2, Theorem 13].

▶ **Theorem 7.** *The following are equivalent:*
1. *There is a homomorphism* $\mathbb{P}_k\mathcal{A} \to \mathcal{B}$.
2. *There is a winning strategy for Duplicator in the existential k-pebble game from $\mathcal{A}$ to $\mathcal{B}$.*

---

[2] In [2] we used the notation $\mathbb{T}_k$ for this comonad.

2.3   The Modal Comonad

For the modal case, we assume that the relational vocabulary σ contains only symbols of arity at most 2. We can thus regard a σ-structure as a Kripke structure for a multi-modal logic, where the universe is thought of as a set of worlds, each binary relation symbol R_α gives the accessibility relation for one of the modalities, and each unary relation symbol P give the valuation for a corresponding propositional variable. If there are no unary symbols, such structures are exactly the labelled transition systems widely studied in concurrency [25].

Modal logic localizes its notion of satisfaction in a structure to a world. We shall reflect this by using the category of *pointed relational structures* $\mathcal{R}_\star(\sigma)$. Objects of this category are pairs $(\mathcal{A}, a)$ where $\mathcal{A}$ is a σ-structure and $a \in A$. Morphisms $h : (\mathcal{A}, a) \to (\mathcal{B}, b)$ are homomorphisms $h : \mathcal{A} \to \mathcal{B}$ such that $h(a) = b$. Of course, the same effect could be achieved by expanding the vocabulary σ with a constant, but pointed categories appear in many mathematical contexts.

For each $k > 0$, we shall define a comonad $\mathbb{M}_k$, where $\mathbb{M}_k(\mathcal{A}, a)$ corresponds to unravelling the structure $\mathcal{A}$, starting from a, to depth k. The universe of $\mathbb{M}_k(\mathcal{A}, a)$ comprises the unit sequence $[a]$, which is the distinguished element, together with all sequences of the form $[a_0, \alpha_1, a_1, \ldots, \alpha_j, a_j]$, where $a = a_0$, $1 \leq j \leq k$, and $R_{\alpha_i}^{\mathcal{A}}(a_i, a_{i+1})$, $0 \leq i < j$. The map $\varepsilon_{\mathcal{A}} : \mathbb{M}_k(A, a) \to (A, a)$ sends a sequence to its last element. Unary relation symbols P are interpreted by $P^{\mathbb{M}_k(\mathcal{A}, a)}(s)$ iff $P^{\mathcal{A}}(\varepsilon_{\mathcal{A}} s)$. For binary relations R_α, the interpretation is $R_\alpha^{\mathbb{M}_k(\mathcal{A}, a)}(s, t)$ iff for some $a' \in A$, $t = s[\alpha, a']$. Given a morphism $f : \mathbb{M}_k(\mathcal{A}, a) \to (\mathcal{B}, b)$, we define $f^* : \mathbb{M}_k(\mathcal{A}, a) \to \mathbb{M}_k(\mathcal{B}, b)$ recursively by $f^*[a] = [b]$, $f^*(s[\alpha, a']) = f^*(s)[\alpha, b']$ where $b' = f(s[\alpha, a'])$. This is well-defined since f is a morphism by assumption.

▶ **Proposition 8.** *The triple* $(\mathbb{M}_k, \varepsilon, (\cdot)^*)$ *is a comonad in Kleisli form on* $\mathcal{R}_\star(\sigma)$.

We recall the notion of *simulation* between Kripke structures [5]. Given structures $\mathcal{A}$, $\mathcal{B}$, we define relations $\preceq_k \subseteq A \times B$, $k \geq 0$, by induction on k: $\preceq_0 = A \times B$, and $a \preceq_{k+1} b$ iff (1) for all unary P, $P^{\mathcal{A}}(a)$ implies $P^{\mathcal{B}}(b)$, and (2) for all R_α, if $R_\alpha^{\mathcal{A}}(a, a')$, then for some b', $R_\alpha^{\mathcal{B}}(b, b')$ and $a' \preceq_k b'$. It is standard that these relations are equivalently formulated in terms of a modified existential Ehrenfeucht-Fraïssé game [5, 14].

▶ **Theorem 9.** *Let* $\mathcal{A}$, $\mathcal{B}$ *be Kripke structures, with* $a \in A$ *and* $b \in B$, *and* $k > 0$. *The following are equivalent:*
1. *There is a homomorphism* $f : \mathbb{M}_k(\mathcal{A}, a) \to (\mathcal{B}, b)$.
2. $a \preceq_k b$.
3. *There is a winning strategy for Duplicator in the* k-*round simulation game from* $(\mathcal{A}, a)$ *to* $(\mathcal{B}, b)$.

3   Logical Equivalences

We now show how our game comonads can be used to give syntax-free characterizations of a range of logical equivalences, which play a central rôle in finite model theory and modal logic.

We shall be considering logics $\mathcal{L}$ which arise as fragments of $\mathcal{L}_{\infty, \omega}$, the extension of first-order logic with infinitary conjunctions and disjunctions, but where formulas contain only finitely many variables. In particular, we will consider the fragments $\mathcal{L}_k$, of formulas with quantifier rank $\leq k$, and $\mathcal{L}^k$, the k-variable fragment. These play a fundamental rôle in finite model theory.

We shall also consider two variants for each of these fragments $\mathcal{L}$. One is the existential positive fragment $\exists\mathcal{L}$, which contains only those formulas of $\mathcal{L}$ built using existential quantifiers, conjunction and disjunction. The other is $\mathcal{L}(\#)$, the extension of $\mathcal{L}$ with counting

quantifiers. These have the form $\exists_{\leq n}, \exists_{\geq n}$, where the semantics of $\mathcal{A} \models \exists_{\geq n} x. \psi$ is that there exist at least n distinct elements of A satisying ψ.

Each of these logics $\mathcal{L}$ induces an equivalence on structures in $\mathcal{R}(\sigma)$:

$$\mathcal{A} \equiv^{\mathcal{L}} \mathcal{B} \quad \overset{\Delta}{\Longleftrightarrow} \quad \forall \varphi \in \mathcal{L}. \, \mathcal{A} \models \varphi \iff \mathcal{B} \models \varphi.$$

Our aim is to characterize these equivalences in terms of our game comonads, and more specifically, to use morphisms in the coKleisli categories as witnesses for these equivalences.

Two equivalences can be defined uniformly for any indexed family of comonads $\mathbb{C}_k$:

- $\mathcal{A} \rightleftarrows_k^{\mathbb{C}} \mathcal{B}$ iff there are coKleisli morphisms $\mathbb{C}_k \mathcal{A} \to \mathcal{B}$ and $\mathbb{C}_k \mathcal{B} \to \mathcal{A}$. Note that there need be no relationship between these morphisms. This is simply the equivalence induced by the preorder collapse of the coKleisli category.
- $\mathcal{A} \cong_k^{\mathbb{C}} \mathcal{B}$ iff $\mathcal{A}$ and $\mathcal{B}$ are isomorphic in the coKleisli category $\mathsf{Kl}(\mathbb{C}_k)$. This means that there are morphisms $\mathbb{C}_k \mathcal{A} \to \mathcal{B}$ and $\mathbb{C}_k \mathcal{B} \to \mathcal{A}$ which are inverses of each other in $\mathsf{Kl}(\mathbb{C}_k)$.

Clearly, $\cong_k^{\mathbb{C}}$ strictly implies $\rightleftarrows_k^{\mathbb{C}}$.

We shall also define an intermediate, "back-and-forth" equivalence $\leftrightarrow_k^{\mathbb{C}}$. This will be more specific to "game comonads" defined on a concrete category such as $\mathcal{R}(\sigma)$, but it will still be defined and shown to have the appropriate properties in considerable generality. We assume that for each structure $\mathcal{A}$, the universe $\mathbb{C}_k A$ has a forest order $\sqsubseteq$, as seen in our concrete constructions using the prefix ordering on sequences. We add a root $\bot$ for convenience. We write the covering relation for this order as $\prec$; thus $s \prec t$ iff $s \sqsubseteq t$, $s \neq t$, and for all u, $s \sqsubseteq u \sqsubseteq t$ implies $u = s$ or $u = t$. We shall also assume that, for any coKleisli morphism $f : \mathbb{C}_k \mathcal{A} \to \mathcal{B}$, the Kleisli coextension preserves the covering relation: $s \prec s'$ implies $f^*(s) \prec f^*(s')$.

The definition will be parameterized on a set $\mathsf{W}_{\mathcal{A},\mathcal{B}} \subseteq \mathbb{C}_k A \times \mathbb{C}_k B$ of "winning positions" for each pair of structures $\mathcal{A}, \mathcal{B}$. We assume that a function $f : \mathbb{C}_k A \to B$ such that, for all $s \in \mathbb{C}_k A$, $(s, f^*(s)) \in \mathsf{W}_{\mathcal{A},\mathcal{B}}$, is a coKleisli morphism.

We define the back-and-forth $\mathbb{C}_k$ game between $\mathcal{A}$ and $\mathcal{B}$ as follows.

At the start of each round of the game, the position is specified by $(s, t) \in \mathbb{C}_k A \times \mathbb{C}_k B$. The initial position is $(\bot, \bot)$. The round proceeds as follows. Either Spoiler chooses some $s' \succ s$, and Duplicator responds with $t' \succ t$, resulting in a new position (s', t'); or Spoiler chooses some $t'' \succ t$ and Duplicator responds with $s'' \succ s$, resulting in (s'', t''). Duplicator wins the round if the new position is in $\mathsf{W}_{\mathcal{A},\mathcal{B}}$.

We can then define $\mathcal{S}(\mathcal{A}, \mathcal{B})$ to be the set of all functions $f : \mathbb{C}_k A \to B$ such that, for all $s \in \mathbb{C}_k A$, $(s, f^*(s)) \in \mathsf{W}_{\mathcal{A},\mathcal{B}}$.

We define a *locally invertible pair* (F, G) from $\mathcal{A}$ to $\mathcal{B}$ to be a pair of sets $F \subseteq \mathcal{S}(\mathcal{A}, \mathcal{B})$, $G \subseteq \mathcal{S}(\mathcal{B}, \mathcal{A})$, satisfying the following conditions:

1. For all $f \in F$, $s \in \mathbb{C}_k A$, for some $g \in G$, $g^* f^*(s) = s$.
2. For all $g \in G$, $t \in \mathbb{C}_k B$, for some $f \in F$, $f^* g^*(t) = t$.

We define $\mathcal{A} \leftrightarrow_k^{\mathbb{C}} \mathcal{B}$ iff there is a non-empty locally invertible pair from $\mathcal{A}$ to $\mathcal{B}$.

▶ **Proposition 10.** *The following are equivalent:*

1. $\mathcal{A} \leftrightarrow_k^{\mathbb{C}} \mathcal{B}$.
2. *There is a winning strategy for Duplicator in the $\mathbb{C}_k$ game between $\mathcal{A}$ and $\mathcal{B}$.*

Proof. Assuming (1), with a locally invertible pair (F, G), we define a strategy for Duplicator inductively, such that after each round, the play is within the set

$$\{(s, f^*(s)) \mid s \in \mathbb{C}_k A, f \in F\} \;=\; \{(g^*(t), t) \mid t \in \mathbb{C}_k B, g \in G\}.$$

Assume (s,t) has been played. If Spoiler now plays $s' \succ s$ in $\mathbb{C}_k A$, then there is $f \in F$ such that $f^*(s) = t$, and we respond with $t' = f^*(s') \succ f^*(s)$. Since $f \in \mathcal{S}(\mathcal{A}, \mathcal{B})$, $(s', t') \in \mathsf{W}_{\mathcal{A}, \mathcal{B}}$. The case when Spoiler plays in $\mathbb{C}_k B$ is symmetric.

Assuming (2), let Φ be the set of all plays (s,t) following the Duplicator strategy. Define

$$F := \{f : \mathbb{C}_k A \to B \mid \forall s \in \mathbb{C}_k A . (s, f^*(s)) \in \Phi\},$$
$$G := \{g : \mathbb{C}_k B \to A \mid \forall t \in \mathbb{C}_k B . (g^*(t), t) \in \Phi\}.$$

Since the strategy is winning, $\Phi \subseteq \mathsf{W}_{\mathcal{A}, \mathcal{B}}$, and $F \subseteq \mathcal{S}(\mathcal{A}, \mathcal{B})$, $G \subseteq \mathcal{S}(\mathcal{B}, \mathcal{A})$. We claim that for all $(s,t) \in \Phi$: (A) $\exists f \in F . f^*(s) = t$, and (B) $\exists g \in G . g^*(t) = s$. (A) follows by extending (s,t) to a morphism $f : \mathbb{C}_k \mathcal{A} \to \mathcal{B}$. For any $s' \sqsubseteq s$, we assign the corresponding predecessor of t. For any s' which is not a predecessor of s, let $s_1 = s \sqcap s'$, the meet of s and s'. We write t_1 for the corresponding predecessor of t. We define f on s' by assigning t_1 in response to s_1, and then following Duplicator's responses as Spoiler plays according to s' in $\mathbb{C}_k A$. (B) follows by a symmetric argument.

Now for any $f \in F$ and $s \in \mathbb{C}_k A$, $(s, f^*(s)) \in \Phi$, and hence by (B) we can find $g \in G$ to witness local invertibility; the case for $g \in G$ and $t \in \mathbb{C}_k B$ is symmetric. ◀

The local invertibility condition on a pair of sets (F, G) has a fixpoint characterization, which may be of some interest. We define set functions $\Gamma : \mathcal{P}(\mathcal{S}(\mathcal{A}, \mathcal{B})) \to \mathcal{P}(\mathcal{S}(\mathcal{B}, \mathcal{A}))$, $\Delta : \mathcal{P}(\mathcal{S}(\mathcal{B}, \mathcal{A})) \to \mathcal{P}(\mathcal{S}(\mathcal{A}, \mathcal{B}))$:

$$\Gamma(F) = \{g \in T \mid \forall t \in \mathbb{C}_k B . \exists f \in F . f^* g^* t = t\},$$
$$\Delta(G) = \{f \in S \mid \forall s \in \mathbb{C}_k A . \exists g \in G . g^* f^* s = s\}.$$

These functions are monotone. Moreover, a pair of sets (F, G) is locally invertible iff $F \subseteq \Delta(G)$ and $G \subseteq \Gamma(F)$. These conditions in turn imply that $F \subseteq \Delta\Gamma(F)$, and if this holds, then we can set $G := \Gamma(F)$ to obtain a locally invertible pair (F, G). Thus existence of a locally invertible pair is equivalent to the existence of non-empty F such that $F \subseteq \Theta(F)$, where $\Theta = \Delta\Gamma$. Since Θ is monotone, by Knaster-Tarski this is equivalent to the greatest fixpoint of Θ being non-empty. (Note that $\Theta(\varnothing) = \varnothing$).

If $\mathcal{A}$ and $\mathcal{B}$ are finite, so is S, and we can construct the greatest fixpoint by a finite descending sequence $S \supseteq \Theta(S) \supseteq \Theta^2(S) \supseteq \cdots$. This fixpoint is non-empty iff $\mathcal{A} \leftrightarrow_k^{\mathbb{E}} \mathcal{B}$.

We shall now turn to a detailed study of each of our comonads in turn.

3.1 The Ehrenfecht-Fraïssé comonad

A coKelisli morphism $f : \mathbb{E}_k \mathcal{A} \to \mathcal{B}$ is an I-morphism if $s \sqsubseteq t$ and $\varepsilon_{\mathcal{A}}(s) = \varepsilon_{\mathcal{A}}(t)$ implies that $f(s) = f(t)$. An equivalent statement is that, if we add a binary relation symbol I to the vocabulary, and set $I^{\mathcal{A}}$ to be the identity relation on A, and $I^{\mathcal{B}}$ to be the identity relation on B, then f is also a homomorphism with respect to I. The significance of this condition is that, if $f : \mathbb{E}_k \mathcal{A} \to \mathcal{B}$ and $g : \mathbb{E}_k \mathcal{B} \to \mathcal{A}$ are I-morphisms, then $f^*(s) = t$, $g^*(t) = s$ imply that (s,t) defines a partial isomorphism from $\mathcal{A}$ to $\mathcal{B}$. We refine the definition of the equivalence $\cong_k^{\mathbb{E}}$ accordingly. We say that $\mathcal{A} \cong_k^{\mathbb{E}} \mathcal{B}$ iff there are I-morphisms $f : \mathbb{E}_k \mathcal{A} \to \mathcal{B}$ and $g : \mathbb{E}_k \mathcal{B} \to \mathcal{A}$ with $f^{*-1} = g^*$.

Note that, for any coKleisli morphism $f : \mathbb{E}_k \mathcal{A} \to \mathcal{B}$, there is an I-morphism $f_I : \mathbb{E}_k \mathcal{A} \to \mathcal{B}$, obtained by firstly restricting f to non-repeating sequences, then extending it by applying the I-morphism condition for repetitions. It is easy to verify that f_I is a homomorphism. Thus there is no need to modify the equivalence $\rightleftharpoons_k^{\mathbb{E}}$.

We define $\mathsf{W}^{\mathbb{E}_k}_{\mathcal{A},\mathcal{B}}$ to be the set of pairs $(s,t) \in \mathbb{E}_k\mathcal{A} \times \mathbb{E}_k\mathcal{B}$ such that $s = [a_1, \ldots, a_j]$, $t = [b_1, \ldots, b_j]$, and $\{(a_i, b_i) \mid 1 \le i \le j\}$ defines a partial isomorphism from $\mathcal{A}$ to $\mathcal{B}$. This specifies the back-and-forth equivalence $\leftrightarrow^{\mathbb{E}}_k$.

We now recall the *bijection game* [16]. In this variant of the Ehrenfeucht-Fraïssé game, Spoiler wins if the two structures have different cardinality. Otherwise, at round i, Duplicator chooses a bijection ψ_i between A and B, and Spoiler chooses an element a_i of A. This determines the choice by Duplicator of $b_i = \psi_i(a_i)$. Duplicator wins after k rounds if the relation $\{(a_i, b_i) \mid 1 \le i \le k\}$ is a partial isomorphism.

▶ **Proposition 11.** *The following are equivalent, for finite structures $\mathcal{A}$ and $\mathcal{B}$:*
1. *$\mathcal{A} \cong^{\mathbb{E}}_k \mathcal{B}$.*
2. *There is a winning strategy for Duplicator in the k-round bijection game.*

Proof. Assuming (1), we have I-morphisms $f : \mathbb{E}_k\mathcal{A} \to \mathcal{B}$ and $g : \mathbb{E}_k\mathcal{B} \to \mathcal{A}$ with $g^* = f^{*-1}$. For each $s \in \{[]\} \cup A^{<k}$, we can define a map $\psi_s : A \to B$, by $\psi_s(a) = f(s[a])$. This is a bijection, with inverse defined similarly from g. These bijections provide a strategy for Duplicator. Since each $(s, f^*(s))$ is a partial isomorphism, this is a winning strategy.

Conversely, a winning strategy provides bjections ψ_s, which we can use to define f by $f(s[a]) = \psi_s(a)$. The winning conditions imply that this is an I-isomorphism in the coKleisli category. ◀

We can now state our main result on logical equivalences for the Ehrenfeucht-Fraïssé comonad.

▶ **Theorem 12.**
1. *For all structures $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} \equiv^{\exists\mathcal{L}_k} \mathcal{B} \iff \mathcal{A} \rightleftarrows^{\mathbb{E}}_k \mathcal{B}$.*
2. *For all structures $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} \equiv^{\mathcal{L}_k} \mathcal{B} \iff \mathcal{A} \leftrightarrow^{\mathbb{E}}_k \mathcal{B}$.*
3. *For all finite structures $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} \equiv^{\mathcal{L}_k(\#)} \mathcal{B} \iff \mathcal{A} \cong^{\mathbb{E}}_k \mathcal{B}$.*

Proof. (1) follows from Theorem 5 and standard results [19]. (2) follows from Proposition 10 and the Ehrenfeucht-Fraïssé theorem [11]. (3) follows from Proposition 11 and results originating in [16] and expounded in [21]. ◀

If we modify $\mathsf{W}^{\mathbb{E}_k}_{\mathcal{A},\mathcal{B}}$, and hence $\leftrightarrow^{\mathbb{E}}_k$, by asking for partial correspondences rather than partial isomorphisms, we obtain a characterization of elementary equivalence for equality-free logic [6].

3.2 The Pebbling Comonad

A similar notion of I-morphism applies to the pebbling comonad as we saw previously with the Ehrenfeucht-Fraïssé comonad [2].

Given $s = [(p_1, a_1), \ldots, (p_n, a_n)] \in \mathbb{P}_k\mathcal{A}$ and $t = [(p_1, b_1), \ldots, (p_n, b_n)] \in \mathbb{P}_k\mathcal{B}$, we define $\phi_{s,t} = \{(a_p, b_p) \mid p \in \mathbf{k}, \ p \text{ occurs in } s\}$, where the last occurrence of p in s is on a_p, and the corresponding last occurrence in t is on b_p. We define $\mathsf{W}^{\mathbb{P}_k}_{\mathcal{A},\mathcal{B}}$ to be the set of all such (s,t) for which $\phi_{s,t}$ is a partial isomorphism. This specifies the back-and-forth equivalence $\leftrightarrow^{\mathbb{P}}_k$.

We now state the following result, characterizing the equivalences induced by finite-variable logics $\mathcal{L}^k$.

▶ **Theorem 13.**
1. *For all structures $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} \equiv^{\exists\mathcal{L}^k} \mathcal{B} \iff \mathcal{A} \rightleftarrows^{\mathbb{P}}_k \mathcal{B}$.*
2. *For all finite structures $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} \equiv^{\mathcal{L}^k} \mathcal{B} \iff \mathcal{A} \leftrightarrow^{\mathbb{P}}_k \mathcal{B}$.*
3. *For all finite structures $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} \equiv^{\mathcal{L}^k(\#)} \mathcal{B} \iff \mathcal{A} \cong^{\mathbb{P}}_k \mathcal{B}$.*

Proof. This follows from Theorems 14, 18 and 20 of [2]. ◀

3.3 The Modal Comonad

The key notion of equivalence in modal logic is bisimulation [5, 29]. We shall define the finite approximants to bisimulation [17].[3] Given Kripke structures $\mathcal{A}$ and $\mathcal{B}$, we define a family of relations $\sim_k \subseteq A \times B$: $\sim_0 = A \times B$; $a \sim_{k+1} b$ iff (1) for all unary P, $P^{\mathcal{A}}(a)$ iff $P^{\mathcal{B}}(b)$; and (2) for all binary R_α, $R_\alpha^{\mathcal{A}}(a, a')$ implies for some b', $R_\alpha^{\mathcal{B}}(b, b')$ and $a' \sim_k b'$, and $R_\alpha^{\mathcal{B}}(b, b')$ implies for some a', $R_\alpha^{\mathcal{A}}(a, a')$ and $a' \sim_k b'$.

We define $\mathsf{W}_{\mathcal{A}, \mathcal{B}}^{\mathbb{M}_k}$ to be the set of all $(s, t) \in \mathbb{M}_k(\mathcal{A}, a) \times \mathbb{M}_k(\mathcal{B}, b)$ such that $s = [a_0, \alpha_1, a_1, \ldots, \alpha_j, a_j]$, $t = [b_0, \alpha_1, b_1, \ldots, \alpha_j, b_j]$, and for all i and all unary P, $P^{\mathcal{A}}(a_i)$ iff $P^{\mathcal{B}}(b_i)$. This specifies the back-and-forth equivalence $\leftrightarrow_k^{\mathbb{M}}$.

▶ **Theorem 14.** *For pointed Kripke structures* $(\mathcal{A}, a)$ *and* $(\mathcal{B}, b)$: $a \sim_k b$ *iff* $(\mathcal{A}, a) \leftrightarrow_k^{\mathbb{M}} (\mathcal{B}, b)$.

Turning to logic, we will consider $\mathcal{M}_k$, the *modal fragment* of modal depth $\leq k$. This arises from the standard translation of (multi)modal logic into $\mathcal{L}_{\infty, \omega}$ [5]. Let us fix a relational vocabulary σ with symbols of arity ≤ 2. For each unary symbol P, there will be a corresponding propositional variable p. Formulas are built from these propositional variables by propositional connectives, and modalities $\Box_\alpha, \Diamond_\alpha$ corresponding to the binary relation symbols R_α in σ. Modal formulas φ then admit a translation into formulas $[\![\varphi]\!] = \psi(x)$ in one free variable. The translation sends propositional variables p to $P(x)$, commutes with the propositional connectives, and sends $\Diamond_\alpha \varphi$ to $\exists y.\, R_\alpha(x, y) \wedge \psi(y)$, where $\psi(x) = [\![\varphi]\!]$. This translation is semantics-preserving: given a σ-structure $\mathcal{A}$ and $a \in A$, then $\mathcal{A}, a \models \varphi$ in the sense of Kripke semantics iff $\mathcal{A} \models \psi(a)$ in the standard model-theoretic sense, where $\psi(x) = [\![\varphi]\!]$.

We define the modal depth of a modal formula φ as the maximum nesting depth of modalities occurring in φ. $\mathcal{M}_k$ is then the image of the translation of modal formulas of modal depth $\leq k$. The existential positive fragment $\exists \mathcal{M}_k$ arises from the modal sublanguage in which formulas are built from propositional variables using only conjunction, disjunction and the diamond modalities $\Diamond_\alpha$.

Extensions of the modal language with counting capabilities have been studied in the form of *graded modalities* [10]. These have the form $\Diamond_\alpha^n, \Box_\alpha^n$, where $\mathcal{A}, a \models \Diamond_\alpha^n \varphi$ if there are at least n R_α-successors of a which satisfy φ. We define $\mathcal{M}_k(\#)$ to be the extension of the modal fragment with graded modalities.

A corresponding notion of graded bisimulation is given in [10]. This is in turn related to *resource bismulation* [8], which has been introduced in the concurrency setting. The two notions are shown to coincide for image-finite Kripke structures in [3], who also show that they can be presented in a simplified form. We recall that a Kripke structure $\mathcal{A}$ is image-finite if for all $a \in A$ and R_α, $R_\alpha(a) := \{a' \mid R^{\mathcal{A}}(a, a')\}$ is finite.

Adapting the results in [3], we define approximants $\sim_k^{\mathsf{g}}$ for graded bisimulation: $\sim_0^{\mathsf{g}} = A \times B$, and $a \sim_{k+1}^{\mathsf{g}} b$ if for all P, $P^{\mathcal{A}}(a)$ iff $P^{\mathcal{B}}(b)$, and for all R_α, there is a bijection $\theta : R^{\mathcal{A}}(a) \cong R^{\mathcal{B}}(b)$ such that, for all $a' \in R^{\mathcal{A}}(a)$, $a' \sim_k^{\mathsf{g}} \theta(a')$.

We can also define a corresponding graded bisimulation game between $(\mathcal{A}, a)$ and $(\mathcal{B}, b)$. At round 0, the elements $a_0 = a$ and $b_0 = b$ are chosen. Duplicator wins if for all P, $P^{\mathcal{A}}(a)$ iff $P^{\mathcal{B}}(b)$, otherwise Spoiler wins. At round $i + 1$, Spoiler chooses some R_α, and Duplicator chooses a bijection $\theta_i : R_\alpha^{\mathcal{A}}(a_i) \cong R_\alpha^{\mathcal{B}}(b_i)$. If there is no such bijection, Spoiler wins. Otherwise, Spoiler then chooses $a_{i+1} \in R^{\mathcal{A}}(a_i)$, and $b_{i+1} := \theta_i(a_{i+1})$. Duplicator wins this round if for all P, $P^{\mathcal{A}}(a_{i+1})$ iff $P^{\mathcal{B}}(b_{i+1})$, otherwise Spoiler wins.

---

[3] Our focus on finite approximants in this paper is for uniformity, and because they are relevant in resource terms. We can extend the comonadic semantics beyond the finite levels. We shall return to this point in the final section.

This game is evidently analogous to the bijection game we encountered previously.

▶ **Proposition 15.** *The following are equivalent:*

1. *There is a winning strategy for Duplicator in the k-round graded bisimulation game between $(\mathcal{A}, a)$ and $(\mathcal{B}, b)$.*
2. $a \sim_k^{\mathbf{g}} b$.
3. $(\mathcal{A}, a) \cong_k^{\mathbb{M}} (\mathcal{B}, b)$.

▶ **Theorem 16.**

1. *For all Kripke structures $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} \equiv^{\exists \mathcal{M}_k} \mathcal{B} \iff \mathcal{A} \rightleftarrows_k^{\mathbb{M}} \mathcal{B}$.*
2. *For all Kripke structures $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} \equiv^{\mathcal{M}_k} \mathcal{B} \iff \mathcal{A} \leftrightarrow_k^{\mathbb{M}} \mathcal{B}$.*
3. *For all image-finite Kripke structures $\mathcal{A}$ and $\mathcal{B}$: $\mathcal{A} \equiv^{\mathcal{M}_k(\#)} \mathcal{B} \iff \mathcal{A} \cong_k^{\mathbb{M}} \mathcal{B}$.*

Proof. (1) follows from Proposition 9 and standard results on preservation of existential positive modal formulas by simulations [5]. (2) follows from Theorem 14 and the Hennesy-Milner Theorem [17, 5]. (3) follows from Proposition 15 and the results in [10, 3]. ◀

4 Coalgebras and combinatorial parameters

Another fundamental aspect of comonads is that they have an associated notion of *coalgebra*. A coalgebra for a comonad (G, ε, δ) is a morphism $\alpha : A \to GA$ such that the following diagrams commute:

$$
\begin{array}{ccc}
A & \xrightarrow{\alpha} & GA \\
\alpha \downarrow & & \downarrow \delta_A \\
GA & \xrightarrow{G\alpha} & G^2 A
\end{array}
\qquad\qquad
\begin{array}{ccc}
A & \xrightarrow{\alpha} & GA \\
& \mathrm{id}_A \searrow & \downarrow \varepsilon_A \\
& & A
\end{array}
$$

Our use of indexed comonads $\mathbb{C}_k$ opens up a new kind of question for coalgebras. Given a structure $\mathcal{A}$, we can ask: what is the least value of k such that a $\mathbb{C}_k$-coalgebra exists on $\mathcal{A}$? We call this the *coalgebra number* of $\mathcal{A}$. We shall find that for each of our comonads, the coalgebra number is a significant combinatorial parameter of the structure.

4.1 The Ehrenfeucht-Fraïssé comonad and tree-depth

A graph is $G = (V, \frown)$, where V is the set of vertices, and $\frown$ is the adjacency relation, which is symmetric and irreflexive. A forest cover for G is a forest $(F, \leq)$ such that $V \subseteq F$, and if $v \frown v'$, then $v \uparrow v'$. The tree-depth $\mathsf{td}(G)$ is defined to be $\min_F \mathsf{ht}(F)$, where F ranges over forest covers of G.[4] It is clear that we can restrict to forest covers of the form $(V, \leq)$, since given a forest cover $(F, \leq)$ of $G = (V, \frown)$, $(V, \leq \cap V^2)$ is also a forest cover of G, and $\mathsf{ht}(V) \leq \mathsf{ht}(F)$. Henceforth, by forest covers of G we shall mean those with universe V.

Given a σ-structure $\mathcal{A}$, the Gaifman graph $\mathcal{G}(\mathcal{A})$ is $(A, \frown)$, where $a \frown a'$ iff for some relation $R \in \sigma$, for some $(a_1, \ldots, a_n) \in R^{\mathcal{A}}$, $a = a_i$, $a' = a_j$, $a \neq a'$. The tree-depth of $\mathcal{A}$ is $\mathsf{td}(\mathcal{G}(\mathcal{A}))$.

▶ **Theorem 17.** *Let $\mathcal{A}$ be a finite σ-structure, and $k > 0$. There is a bijective correspondence between*

1. $\mathbb{E}_k$-*coalgebras $\alpha : \mathcal{A} \to \mathbb{E}_k \mathcal{A}$.*
2. *Forest covers of $\mathcal{G}(\mathcal{A})$ of height $\leq k$.*

---

[4] We formulate this notion in order-theoretic rather than graph-theoretic language, but it is equivalent to the definition in [27].

Proof. Suppose that $\alpha : \mathcal{A} \rightarrow \mathbb{E}_k \mathcal{A}$ is a coalgebra. For $a \in A$, let $\alpha(a) = [a_1, \ldots, a_j]$. The first coalgebra equation says that $\alpha(a_i) = [a_1, \ldots, a_i]$, $1 \leq i \leq j$. The second says that $a_j = a$. Thus $\alpha : A \rightarrow A^{\leq k}$ is an injective map whose image is a prefix-closed subset of $A^{\leq k}$. Defining $a \leq a'$ iff $\alpha(a) \sqsubseteq \alpha(a')$ yields a forest order on A, of height $\leq k$. If $a \frown a'$ in $\mathcal{G}(\mathcal{A})$, for some $a_1, \ldots, a_n$ with $a = a_i$, $a' = a_j$, we have $R^{\mathcal{A}}(a_1, \ldots, a_n)$. Since α is a homomorphism, we must have $R^{\mathbb{E}_k \mathcal{A}}(\alpha(a_1), \ldots, \alpha(a_n))$, hence $\alpha(a_i) \uparrow \alpha(a_j)$, and so $a_i \uparrow a_j$. Thus $(A, \leq)$ is a forest cover of $\mathcal{A}$, of height $\leq k$.

Conversely, given such a forest cover $(A, \leq)$, for each $a \in A$, its predecessors form a chain $a_1 < \cdots < a_j$, with $a_j = a$, and $j \leq k$. We define $\alpha(a) = [a_1, \ldots, a_j]$, which yields a map $\alpha : A \rightarrow A^{\leq k}$, which evidently satisfies the coalgebra equations. If $R^{\mathcal{A}}(a_1, \ldots, a_n)$, then since $(A, \leq)$ is a forest cover, we must have $a_i \uparrow a_j$ for all i, j, and hence $\alpha(a_i) \uparrow \alpha(a_j)$. Thus α is a homomorphism.   ◄

We write $\kappa^{\mathbb{E}}(\mathcal{A})$ for the coalgebra number of $\mathcal{A}$ with respect to the the Ehrenfeucht-Fraïssé comonad.

▶ **Theorem 18.** *For all finite structures $\mathcal{A}$: $\mathsf{td}(\mathcal{A}) = \kappa^{\mathbb{E}}(\mathcal{A})$.*

4.2   The pebbling comonad and tree-width

We review the notions of tree decompositions and tree-width. A tree $(T, \leq)$ is a forest with a least element (the root). A tree is easily seen to be a meet-semilattice: every pair of elements x, x' has a greatest lower bound $x \wedge x'$ (the greatest common ancestor). The path from x to x' is the set $\mathsf{path}(x, x') := [x \wedge x', x] \cup [x \wedge x', x']$, where we use interval notation: $[y, y'] := \{z \in T \mid y \leq z \leq y'\}$.

A tree-decomposition of a graph $G = (V, \frown)$ is a tree $(T, \leq)$ together with a labelling function $\lambda : T \rightarrow \mathcal{P}(V)$ satisfying the following conditions:

- (TD1) for all $v \in V$, for some $x \in T$, $v \in \lambda(x)$;
- (TD2) if $v \frown v'$, then for some $x \in T$, $\{v, v'\} \subseteq \lambda(x)$;
- (TD3) if $v \in \lambda(x) \cap \lambda(x')$, then for all $y \in \mathsf{path}(x, x')$, $v \in \lambda(y)$.

The width of a tree decomposition is given by $\max_{x \in T} |\lambda(x)| - 1$. We define the tree-width $\mathsf{tw}(G)$ of a graph G as $\min_T \mathsf{width}(T)$, where T ranges over tree decompositions of G.

We shall now give an alternative formulation of tree-width which will provide a useful bridge to the coalgebraic characterization. It is also interesting in its own right: it clarifies the relationship between tree-width and tree-depth, and shows how pebbling arises naturally in connection with tree-width.

A k-pebble forest cover for a graph $G = (V, \frown)$ is a forest cover $(V, \leq)$ together with a pebbling function $p : V \rightarrow \mathbf{k}$ such that, if $v \frown v'$ with $v \leq v'$, then for all $w \in (v, v']$, $p(v) \neq p(w)$.

The following result is implicit in [2], but it seems worthwhile to set it out more clearly.

▶ **Theorem 19.** *Let G be a finite graph. The following are equivalent:*
1. *G has a tree decomposition of width $< k$.*
2. *G has a k-pebble forest cover.*

Proof. $(1) \Rightarrow (2)$. Assume that $G = (V, \frown)$ has a tree decomposition $(T, \leq, \lambda)$ of width $< k$. We say that a tree decomposition is *orderly* if it has the following property: for all $x \in T$, there is at most one $v \in \lambda(x)$ such that for all $x' < x$, $v \notin \lambda(x')$. Nice tree decompositions are orderly [18]; hence by standard results, without loss of generality we can assume that the given tree decomposition is orderly.

For any $v \in V$, the set of $x \in T$ such that $v \in \lambda(x)$ is non-empty by (TD1), and closed under meets by (TD3). Since T is a tree, this implies that this set has a least element $\tau(v)$. This defines a function $\tau : V \to T$. The fact that tree decomposition is orderly implies that τ is injective. We can define an order on V by $v \leq v'$ iff $\tau(v) \leq \tau(v')$. This is isomorphic to a sub-poset of T, and hence is a forest order.

We define $p : V \to \mathbf{k}$ by induction on this order. Assuming $p(v')$ is defined for all $v' < v$, we consider $\tau(v)$. Since the tree decomposition is orderly, this means in particular that $p(v')$ is defined for all $v' \in S := \lambda(\tau(v)) \setminus \{v\}$. Since the decomposition is of width $< k$, we must have $|S| < k$. We set $p(v) := \min(\mathbf{k} \setminus \{p(v') \mid v' \in S\})$.

To verify that $(V, \leq)$ is a forest cover, suppose that $v \frown v'$. By (TD2), for some $x \in T$, $\{v, v'\} \subseteq \lambda(x)$. We have $\tau(v) \leq x \geq \tau(v')$, and since T is a tree, we must have $\tau(v) \uparrow \tau(v')$, whence $v \uparrow v'$.

Finally, we must verify the condition on the pebbling function p. Suppose that $v \frown v'$, and $v < w \leq v'$. Since $v \frown v'$, for some x, $\{v, v'\} \subseteq \lambda(x)$. But then $\tau(v) < \tau(w) \leq \tau(v') \leq x$. Since $v \in \lambda(\tau(v)) \cap \lambda(x)$, by (TD3), $v \in \lambda(\tau(w))$. By construction of the pebbling function, this implies $p(v) \neq p(w)$.

$(2) \Rightarrow (1)$. Suppose that $(V, \leq, p)$ is a k-pebble forest cover of G. We define a tree $T = V_\perp$ by adjoining a least element $\perp$ to V. We say that v is an active predecessor of v' if $v \leq v'$, and for all $w \in (v, v']$, $p(v) \neq p(w)$. We define the labelling function by setting $\lambda(v)$ to be the set of active predecessors of v; $\lambda(\perp) := \varnothing$. Since $p|_{\lambda(v)}$ is injective, $|\lambda(v)| \leq k$.

We verify the tree decomposition conditions. (TD1) holds, since $v \in \lambda(v)$. (TD2) If $v \frown v'$, then $v \uparrow v'$. Suppose $v \leq v'$. Then v is an active predecessor of v', and $\{v, v'\} \subseteq \lambda(v')$. (TD3) Suppose $v \in \lambda(v_1) \cap \lambda(v_2)$. Then v is an active predecessor of both v_1 and v_2. This implies that for all $w \in \mathsf{path}(v_1, v_2)$, v is an active predecessor of w, and hence $v \in \lambda(w)$. ◀

▶ **Theorem 20.** *Let $\mathcal{A}$ be a finite σ-structure. There is a bijective correspondence between:*
1. *$\mathbb{P}_k$-coalgebras $\alpha : \mathcal{A} \to \mathbb{P}_k \mathcal{A}$*
2. *k-pebble forest covers of $\mathcal{G}(\mathcal{A})$.*

Proof. See [2, Theorem 6]. ◀

We write $\kappa^{\mathbb{P}}(\mathcal{A})$ for the coalgebra number of $\mathcal{A}$ with respect to the the pebbling comonad.

▶ **Theorem 21.** *For all finite structures $\mathcal{A}$: $\mathsf{tw}(\mathcal{A}) = \kappa^{\mathbb{P}}(\mathcal{A}) - 1$.*

4.3 The modal comonad and synchronization tree depth

Let $\mathcal{A}$ be a Kripke structure. It will be convenient to write labelled transitions $a \xrightarrow{\alpha} a'$ for $R_\alpha(a, a')$. Given $a \in A$, the submodel generated by a is obtained by restricting the universe to the set of a' such that there is a path $a \xrightarrow{\alpha_1} \cdots \xrightarrow{\alpha_k} a'$. This submodel forms a synchronization tree [24] if for all a', there is a unique such path. The height of such a tree is the maximum length of any path from the root a.

▶ **Proposition 22.** *Let $\mathcal{A}$ be a Kripke structure, with $a \in A$. The following are equivalent:*
1. *There is a coalgebra $\alpha : (\mathcal{A}, a) \to \mathbb{M}_k(\mathcal{A}, a)$.*
2. *The submodel generated by a is a synchronization tree of height $\leq k$.*
We define the modal depth $\mathsf{md}(\mathcal{A}, a) = k$ if the submodel generated by a is a synchronization tree of height k.

▶ **Theorem 23.** *Let $\mathcal{A}$ be a Kripke structure, and $a \in A$ be such that the submodel generated by a is a synchronization tree of finite height. Then $\mathsf{md}(\mathcal{A}, a) = \kappa^{\mathbb{M}}(\mathcal{A}, a)$.*

Note the conditional nature of this result, which contrasts with those for the other comonads. The modal comonad is defined in such a way that the universe $\mathbb{M}_k(A, a)$ reflects information about the possible transitions. Thus having a coalgebra at all, regardless of the value of the resource parameter, is a strong constraint on the structure of the transition system.

5   Further Directions

From the categorical perspective, there is considerable additional structure which we have not needed for the results in this paper, but which may be useful for further investigations.

Coequaliser requirements.   In Moggi's work on computational monads, there is an "equaliser requirement" [26]. The dual version for a comonad (G, ε, δ) is that for every object A, the following diagram is a coequaliser:

$$G^2 A \; \underset{\varepsilon_{GA}}{\overset{G\varepsilon_A}{\rightrightarrows}} \; GA \; \xrightarrow{\varepsilon_A} \; A$$

This says in particular that the counit is a regular epi, and hence GA "covers" A in a strong sense.

This coequaliser requirement holds for all our comonads. For $\mathbb{E}_k$, this is basically the observation that, given a sequence of sequences $[s_1, \ldots, s_j]$, we have $\varepsilon[\varepsilon s_1, \ldots, \varepsilon s_j] = \varepsilon s_j$. The other cases are similar.

Indexed and graded structure.   Our comonads $\mathbb{E}_k, \mathbb{P}_k, \mathbb{M}_k$ are not merely discretely indexed by the resource parameter. In each case, there is a functor $(\mathbb{Z}^+, \leq) \to \mathsf{Comon}(\mathcal{R}(\sigma))$ from the poset category of the positive integers to the category of comonads on $\mathcal{R}(\sigma)$. Thus if $k \leq l$ there is a natural transformation with components $i_A^{k,l} : \mathbb{E}_k \mathcal{A} \to \mathbb{E}_l \mathcal{A}$, which preserves the counit and comultiplication; and similarly for the other comonads. Concretely, this is just including the plays of up to k rounds in the plays of up to l rounds, $k \leq l$.

Another way of parameterizing comonads by resource information is grading [12]. Recall that comonads on $\mathcal{C}$ are exactly the comonoids in the strict monoidal category $([\mathcal{C}, \mathcal{C}], \circ, I)$ of endofunctors on $\mathcal{C}$ [22]. Generalizing this description, a graded comonad is an oplax monoidal functor $G : (M, \cdot, 1) \to ([\mathcal{C}, \mathcal{C}], \circ, I)$ from a monoid of grades into this endofunctor category. This means that for each $m \in M$, there is an endofunctor G_m, there is a graded counit natural transformation $\varepsilon : G_1 \Rightarrow I$, and for all $m, m' \in M$, there is a graded comultiplication $\delta^{m,m'} : G_{m \cdot m'} \Rightarrow G_m G_{m'}$.

The two notions can obviously be combined. We can see our comonads as (trivially) graded, by viewing them as oplax monoidal functors $(\mathbb{Z}^+, \leq, \min, 1) \to ([\mathcal{C}, \mathcal{C}], \circ, I)$. Given $k \leq l$, we have e.g. $\mathbb{E}_k \Rightarrow \mathbb{E}_k \mathbb{E}_k \Rightarrow \mathbb{E}_k \mathbb{E}_l$. The question is whether there are more interesting graded structures which arise naturally in considering richer logical and computational settings.

Colimits and infinite behaviour.   In this paper, we have dealt exclusively with finite resource levels. However, there is an elegant means of passing to infinite levels. We shall illustrate this with the modal comonad. Using the inclusion morphisms described in the previous discussion of indexed structure, for each structure $\mathcal{A}$ we have a diagram

$$\mathbb{M}_1 \mathcal{A} \to \mathbb{M}_2 \mathcal{A} \to \cdots \to \mathbb{M}_k \mathcal{A} \to \cdots$$

By taking the colimits of these diagrams, we obtain a comonad $\mathbb{M}_\omega$, which corresponds to the usual unfolding of a Kripke structure to all finite levels. This will correspond to the bisimulation approximant $\sim_\omega$, which coincides with bisimulation itself on image-finite structures [17]. Transfinite extensions are also possible. Similar constructions can be applied to the other comonads. This provides a basis for lifting the comonadic analysis to the level of infinite models.

Relations between fragments and parameters. We can define morphisms between the different comonads we have discussed, which yield proofs about the relationships between the logical fragments they characterize. This categorical perspective avoids the cumbersome syntactic translations in the standard proofs of these results. For illustration, there is a comonad morphism $t : \mathbb{E}_k \Rightarrow \mathbb{P}_k$ with components $t_A : \mathbb{E}_k\mathcal{A} \to \mathbb{P}_k\mathcal{A}$ given by $[a_1, \ldots, a_j] \mapsto [(1, a_1), \ldots, (j, a_j)]$. Together with theorems 13 and 12, this shows that $\exists \mathcal{L}_k \subseteq \exists \mathcal{L}^k$ and $\mathcal{L}_k(\#) \subseteq \mathcal{L}^k(\#)$. Moreover, composing t with a coalgebra $\mathcal{A} \to \mathbb{E}_k\mathcal{A}$ yields a coalgebra $A \to \mathbb{P}_k\mathcal{A}$, demonstrating that $\mathsf{tw}(\mathcal{A}) + 1 \leq \mathsf{td}(\mathcal{A})$. Another morphism $\mathbb{M}_\omega \Rightarrow \mathbb{P}_2$ shows that modal logic can be embedded into 2-variable logic.

Concluding remarks

Our comonadic constructions for the three major forms of model comparison games show a striking unity, on the one hand, but also some very interesting differences. For the latter, we note the different forms of logical "deception" associated with each comonad, the different forms of back-and-forth equivalences, and the different combinatorial parameters which arise in each case.

One clear direction for future work is to gain a deeper understanding of what makes these constructions work. Another is to understand how widely the comonadic analysis of resources can be applied. We are currently investigating the guarded fragment [4, 14]; other natural candidates include existential second-order logic, and branching quantifiers and dependence logic [32].

Since comonads arise naturally in type theory and functional programming [31, 28], can we connect the study of finite model theory made here with a suitable type theory? Can this lead, via the Curry-Howard correspondence, to the systematic derivation of some significant meta-algorithms, such as decision procedures for guarded logics based on the tree model property [13], or algorithmic metatheorems such as Courcelle's theorem [9]?

Another intriguing direction is to connect these ideas with the graded quantum monad studied in [1], which provides a basis for the study of quantum advantage in $\mathcal{R}(\sigma)$. This may lead to a form of quantum finite model theory.

References

1   Samson Abramsky, Rui Soares Barbosa, Nadish de Silva, and Octavio Zapata. The quantum monad on relational structures. *To appear in proceedings of MFCS 2017*, 2018.

2   Samson Abramsky, Anuj Dawar, and Pengming Wang. The pebbling comonad in finite model theory. In *Logic in Computer Science (LICS), 2017 32nd Annual ACM/IEEE Symposium on*, pages 1–12. IEEE, 2017.

3   Luca Aceto, Anna Ingolfsdottir, and Joshua Sack. Resource bisimilarity and graded bisimilarity coincide. *Information Processing Letters*, 111(2):68–76, 2010.

4   Hajnal Andréka, István Németi, and Johan van Benthem. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.

5   Patrick Blackburn, Maarten De Rijke, and Yde Venema. *Modal Logic*, volume 53. Cambridge University Press, 2002.

6   E. Casanovas, P. Dellunde, and R. Jansana. On Elementary Equivalence for Equality-free Logic. *Notre Dame Journal of Formal Logic*, 37(3):506–522, 1996. `doi:10.1305/ndjfl/1039886524`.

7   Ashok K Chandra and Philip M Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proceedings of the Ninth Annual ACM Symposium on Theory of Computing*, pages 77–90. ACM, 1977.

8   Flavio Corradini, Rocco De Nicola, and Anna Labella. Graded modalities and resource bisimulation. In *International Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 381–393. Springer, 1999.

9   Bruno Courcelle. The monadic second-order logic of graphs. i. recognizable sets of finite graphs. *Information and computation*, 85(1):12–75, 1990.

10   M. de Rijke. A Note on Graded Modal Logic. *Studia Logica*, 64(2):271–283, 2000.

11   Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite model theory*. Springer Science & Business Media, 2005.

12   Marco Gaboardi, Shin-ya Katsumata, Dominic Orchard, Flavien Breuvart, and Tarmo Uustalu. Combining effects and coeffects via grading. *ACM SIGPLAN Notices*, 51(9):476–489, 2016.

13   Erich Grädel. Decision procedures for guarded logics. In *International Conference on Automated Deduction*, pages 31–51. Springer, 1999.

14   Erich Grädel and Martin Otto. The freedoms of (guarded) bisimulation. In *Johan van Benthem on Logic and Information Dynamics*, pages 3–31. Springer, 2014.

15   Pavol Hell and Jaroslav Něsetřil. *Graphs and homomorphisms*. Oxford University Press, 2004.

16   L. Hella. Logical hierarchies in PTIME. *Information and Computation*, 121:1–19, 1996.

17   Matthew Hennessy and Robin Milner. On observing nondeterminism and concurrency. In *International Colloquium on Automata, Languages, and Programming*, pages 299–309. Springer, 1980.

18   Ton Kloks. *Treewidth: computations and approximations*, volume 842. Springer Science & Business Media, 1994.

19   Phokion G Kolaitis and Moshe Y Vardi. On the expressive power of Datalog: tools and a case study. In *Proceedings of the ninth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 61–71. ACM, 1990.

20   Phokion G Kolaitis and Moshe Y Vardi. Infinitary logics and 0–1 laws. *Information and Computation*, 98(2):258–294, 1992.

21   Leonid Libkin. *Elements of Finite Model Theory (Texts in Theoretical Computer Science. An EATCS Series)*. Springer, 2004.

22   Saunders Mac Lane. *Categories for the working mathematician*, volume 5. Springer Science & Business Media, 2013.

23   Ernest G Manes. *Algebraic Theories*, volume 26. Springer Science & Business Media, 2012.

24   Robin Milner. *A calculus of communicating systems*. Number 92 in Lecture Notes in Comput. Science. Springer-Verlag, 1980.

25   Robin Milner. *Communication and concurrency*, volume 84. Prentice Hall New York etc., 1989.

26   Eugenio Moggi. Notions of computation and monads. *Information and computation*, 93(1):55–92, 1991.

27   Jaroslav Nešetřil and Patrice Ossona De Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *European Journal of Combinatorics*, 27(6):1022–1041, 2006.

28 Dominic Orchard. Programming contextual computations. Technical Report UCAM-CL-TR-854, University of Cambridge, 2014.

29 Davide Sangiorgi. On the origins of bisimulation and coinduction. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 31(4):15, 2009.

30 Saharon Shelah. Every two elementarily equivalent models have isomorphic ultrapowers. *Israel Journal of Mathematics*, 10(2):224–233, 1971.

31 Tarmo Uustalu and Varmo Vene. Comonadic notions of computation. *Electronic Notes in Theoretical Computer Science*, 203(5):263–284, 2008.

32 Jouko Väänänen. *Dependence logic: A new approach to independence friendly logic*, volume 70. Cambridge University Press, 2007.