


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Annals of Pure and Applied Logic

journal homepage: www.elsevier.com/locate/apal

Full Length Article

Arboreal categories and equi-resource homomorphism preservation theorems [☆]

Samson Abramsky, Luca Reggio ^{*}

Department of Computer Science, University College London, 66–72 Gower Street, London, WC1E 6EA, United Kingdom

ARTICLE INFO

Article history:

Received 28 January 2023

Received in revised form 30 January 2024

Accepted 12 February 2024

Available online 15 February 2024

MSC:

03B10

03B70

03C13

18C10

68Q19

68Q55

Keywords:

Homomorphism preservation theorems

Logical resources

Game comonads

First-order logic

Modal logic

Guarded logics

ABSTRACT

The classical homomorphism preservation theorem, due to Łoś, Lyndon and Tarski, states that a first-order sentence φ is preserved under homomorphisms between structures if, and only if, it is equivalent to an existential positive sentence ψ . Given a notion of (syntactic) complexity of sentences, an “equi-resource” homomorphism preservation theorem improves on the classical result by ensuring that ψ can be chosen so that its complexity does not exceed that of φ .

We describe an axiomatic approach to equi-resource homomorphism preservation theorems based on the notion of arboreal category. This framework is then employed to establish novel homomorphism preservation results, and improve on known ones, for various logic fragments, including first-order, guarded and modal logics.

© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

1.1. Summary

In this paper, we take steps towards an *axiomatic model theory* which is resource-sensitive, and well adapted for finite and algorithmic model theory and descriptive complexity. This builds on previous work

[☆] Research supported by the EPSRC grant EP/V040944/1, and by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 837724.

^{*} Corresponding author.

E-mail addresses: s.abramsky@ucl.ac.uk (S. Abramsky), l.reggio@ucl.ac.uk (L. Reggio).

on game comonads [3,8,17,4,19,5,28], which describe model comparison games in terms of resource-indexed comonads on the category of relational structures. This allows the syntax-free description of a wide range of logics, indexed by resource parameters such as quantifier rank and number of variables. As explained in [2], this opens up a landscape in which properties of the comonad are related to the degree of tractability of the logic.

Game comonads and the associated (comonadic) adjunctions have been axiomatised at a general categorical level in terms of *arboreal categories* [6,7], in which the objects have an intrinsic tree structure. This structure can be used to define notions of bisimulation, game, etc. at an abstract level, and to transfer these process notions via an adjunction to an *extensional category*—the main concrete example in applications being the category of relational structures.

In this paper, we show that this axiomatic framework can be used to prove a substantial resource-sensitive model-theoretic result, Rossman’s equirank homomorphism preservation theorem, as well as a number of extensions and variations, e.g. to modal and guarded logics, and relativisations to a wide range of subclasses of structures.

Some striking features of this axiomatic approach are its generality, its resource sensitivity, and the fact that it does not use compactness. It is interesting to compare this with the use of accessible categories as an axiomatic framework for the study of abstract elementary classes [14]. This is aimed at extending first-order model theory into the infinite, replacing compactness by λ -accessibility. By contrast, motivated by the needs of finite model theory and descriptive complexity, we are interested in capturing fine structure “down below”, in fragments of first-order logic.¹

1.2. Extended discussion

Homomorphism preservation theorems relate the syntactic shape of a sentence with the semantic property of being preserved under homomorphisms between structures. Recall that a first-order sentence φ in a vocabulary τ is said to be *preserved under homomorphisms* if, whenever there is a homomorphism of τ -structures $A \rightarrow B$ and $A \models \varphi$, then also $B \models \varphi$. Further, an *existential positive sentence* is a first-order sentence that uses only the connectives \vee, \wedge and the quantifier \exists . The following classical result, known as the *homomorphism preservation theorem*, is due to Łoś, Lyndon and Tarski [33,34,42] and applies to arbitrary (first-order) vocabularies.

Theorem 1.1. *A first-order sentence is preserved under homomorphisms if, and only if, it is equivalent to an existential positive sentence.*

The homomorphism preservation theorem is a fairly straightforward consequence of the compactness theorem, see e.g. [43, Lemma 3.1.2]. However, applying the compactness theorem means that we lose control over the syntactic shape of an existential positive sentence ψ that is equivalent to a sentence φ preserved under homomorphisms. In particular, it is an ineffective approach if we want to determine to which extent the passage from φ to ψ increases the “complexity”.

One way to measure the complexity of a formula φ is in terms of its *quantifier rank*, i.e. the maximum number of nested quantifiers appearing in φ . Rossman’s *equirank homomorphism preservation theorem* [41], which applies to relational vocabularies (i.e. vocabularies that contain no constant or function symbols), shows that it is possible to find a ψ whose quantifier rank is less than or equal to the quantifier rank of φ :

Theorem 1.2. *A first-order sentence of quantifier rank at most k is preserved under homomorphisms if, and only if, it is equivalent to an existential positive sentence of quantifier rank at most k .*

¹ Monadic second-order logic is also of interest and within scope, see [28].

This is a considerable improvement on the classical homomorphism preservation theorem and was proved by Rossman on the way to his celebrated *finite homomorphism preservation theorem*, stating that Theorem 1.1 admits a relativisation to finite structures.² In the proof of the equirank homomorphism preservation theorem, the application of the compactness theorem is replaced with a model construction which is similar in spirit to the construction of a saturated elementary extension of a given structure.

The main contribution of this paper consists in laying out a categorical framework in which “equi-resource” homomorphism preservation theorems can be proved in an axiomatic fashion. In [3,8], *game comonads* were introduced to capture in a structural way a number of logic fragments and corresponding combinatorial and game-theoretic notions, both at the level of finite and infinite structures. For a recent survey article, see [2]. The template of game comonads was axiomatised in [6], see also [7], by means of the notion of *arboreal category*. Our proof strategy consists in establishing abstract homomorphism preservation theorems at the general level of arboreal categories, and then instantiating these results for specific choices of game comonads.

We thus obtain novel equi-resource homomorphism preservation theorems for modal and guarded logics (Theorems 5.6 and 5.8, respectively), along with relativisations to appropriate subclasses of structures—e.g. the class of finite structures (Theorems 5.20 and 5.21, respectively). Further, we derive a relativisation result (Theorem 6.12) which refines Rossman’s equirank homomorphism preservation theorem.

This paper is organised as follows. In Section 2 we provide a brief introduction to game comonads, and in Section 3 we recall the necessary definitions and facts concerning arboreal categories. Homomorphism preservation theorems are recast into categorical statements and proved at the level of arboreal categories in Sections 4, 5 and 6. Section 7 contains the proof of our main technical result, namely Theorem 6.9. Finally, in Section 8 we present a simple class of relativisation results for classes of well-behaved structures; these are discussed separately as they are not *equi-resource* homomorphism preservation theorems.

Throughout this article, we shall assume the reader is familiar with the basic notions of category theory; standard references include e.g. [10,35].

Let us conclude this introduction with some observations to guide the reader through the article:

From games to arboreal categories: This paper builds on our prior work [7] on arboreal categories. Familiarity with the latter work would be beneficial, but we do recall the necessary material in Section 3 in order to make the article self-contained. Let us briefly overview the different levels of abstractions leading from model-comparison games to arboreal categories.

Preservation of many fragments of first-order logic, and variations thereof, can be described in terms of the existence of a winning strategy in a game. Game comonads arise from the insight that the construction associating to a structure A its collection GA of “explorations” in the game induces a comonad G , which is an enrichment of the (finite-)list comonad. The coalgebras for the comonad G capture compatible forest covers of structures, which can be thought of as unfoldings of games, and winning strategies in the game correspond to various types of morphisms between coalgebras.

Arboreal categories capture the main features of the categories of coalgebras for game comonads, and the minimal properties required to fruitfully imitate the notions of back-and-forth equivalence and bisimulation, and to show that they coincide, gave rise to the axioms of arboreal categories. The filtering of coalgebra covers corresponding to the stratification induced by the resource parameter in the logic (e.g. the quantifier rank of formulas) begets resource-indexed arboreal adjunctions.

A dividing line: Analysing homomorphism preservation theorems from a categorical perspective, we identify the *bisimilar companion property* as a key dividing line, and Section 5 is divided accordingly. Suppose

² The finite homomorphism preservation theorem is a major result in finite model theory, as well as a surprising one given that most preservation theorems fail when restricted to finite structures. Note that the finite homomorphism preservation theorem and the classical one are incomparable results.

that G is a game comonad on a category of structures corresponding to a logic fragment \mathcal{L} . Typically, any structure A is equivalent to GA in the *existential positive* fragment of \mathcal{L} , but there could be a sentence in \mathcal{L} that is valid in precisely one of the two structures. The bisimilar companion property asserts that this is not the case: A and GA are indistinguishable in the logic fragment \mathcal{L} .

The bisimilar companion property is satisfied e.g. by the comonads corresponding to modal and guarded logics. For these logics, we obtain simple proofs of equi-resource homomorphism preservation theorems.

On the other hand, the bisimilar companion property fails for first-order logic with bounded quantifier rank, or with a bounded number of variables (the latter falls outside of the scope of this paper, see Remark 6.1), and showing that this property can be “forced” requires a considerable amount of work. Our approach is inspired by that of Rossman [41], and shows that the methods and ideas put forward in [41] can be lifted to the axiomatic setting. There is a difference though: whereas Rossman uses *n-cores*, approximating the notion of *core* [24], we use cofree coalgebras. The former are (minimal) retracts of structures, while the latter can be thought of as “coverings”.

Equality: The framework of game comonads typically encodes the preservation of *equality-free* fragments of logics. Generally, the difference is invisible when dealing with existential positive fragments, or extensions with counting quantifiers, as these often admit equality elimination (cf. [19, §VI] for the case of logics with counting quantifiers). The difference is also immaterial for modal logic, because its standard translation is contained in the equality-free fragment of first-order logic. Nevertheless, this gap needs to be addressed when considering e.g. first-order logic with bounded quantifier rank.

This was handled in [3,9] and subsequent works by considering a fresh binary relation I and assuming that (i) I is interpreted as the identity relation on the class of structures we are interested in, and (ii) all homomorphisms preserve I . Note however that, if A is a structure in which I coincides with the identity relation, the structure GA obtained by applying a game comonad G to A need not have the same property.

In the present paper, we adopt the approach outlined in [7], which is better suited for an axiomatic treatment. This consists in composing the comonadic adjunction induced by a game comonad with an adjunction that “introduces and eliminates I ” (see Example 3.18 for more details). This is why, throughout, we need to work with possibly non-comonadic adjunctions.

2. Logic fragments and game comonads

We shall mainly deal with two types of vocabularies:

- *Relational vocabularies*, i.e. first-order vocabularies that contain no function or constant symbols.
- *Multi-modal vocabularies*, i.e. relational vocabularies in which every relation symbol has arity at most 2.

Multi-modal vocabularies will be referred to simply as *modal vocabularies*. If σ is a modal vocabulary, we can assign to each unary relation symbol $P \in \sigma$ a propositional variable p , and to each binary relation symbol $R \in \sigma$ modalities \diamond_R and \square_R . We refer to σ -structures as *Kripke structures*. For any Kripke structure A , the interpretation of P in A , denoted by P^A , corresponds to the valuation of the propositional variable p , and the binary relation R^A to the accessibility relation for the modalities \diamond_R and \square_R .

For our running examples and intended applications, we will be interested in the following resource-bounded fragments of first-order logic (always including the equality symbol) and modal logic, for a relational vocabulary σ and positive integer k :

- FO_k and $\exists^+\text{FO}_k$: these denote, respectively, the set of sentences of quantifier rank $\leq k$ in the vocabulary σ , and its existential positive fragment.

- ML_k and \exists^+ML_k : if σ is a modal vocabulary, ML_k is the set of modal formulas of modal depth $\leq k$ in the vocabulary σ (recall that the *modal depth* is the maximum number of nested modalities in a formula). Moreover, \exists^+ML_k denotes the existential positive fragment of ML_k , i.e. the set of formulas in ML_k that use only the connectives \vee, \wedge and the diamond modalities.
- $ML_k(\#)$: this is the extension of ML_k with *graded modalities*. Recall that graded modalities have the form \diamond_R^n and \square_R^n , with $n \geq 0$, and $A, a \models \diamond_R^n \varphi$ if there are at least n R -successors of a satisfying φ ; $\square_R^n \varphi$ is logically equivalent to $\neg \diamond_R^n \neg \varphi$.

Further logics, namely *guarded logics*, will be considered in Section 5.1.

Any logic fragment³ \mathcal{L} , i.e. any set of first-order sentences in a vocabulary τ , induces an equivalence relation $\equiv^{\mathcal{L}}$ on τ -structures defined as follows: for all τ -structures A, B ,

$$A \equiv^{\mathcal{L}} B \iff \forall \varphi \in \mathcal{L}. (A \models \varphi \iff B \models \varphi).$$

If \mathcal{L} consists of all first-order sentences, $\equiv^{\mathcal{L}}$ coincides with elementary equivalence.

Syntax-free characterisations of the equivalence relations $\equiv^{\mathcal{L}}$ play an important role in model theory. For example, the Keisler-Shelah theorem states that two τ -structures are elementarily equivalent if, and only if, they have isomorphic ultrapowers. A different approach is through *model comparison games*. These have a wide range of applications in model theory, see e.g. [26, §3], and are central to finite model theory where tools such as the compactness theorem and ultraproducts are not available. Model comparison games lead to a perspective which may be described as “model theory without compactness”.

Game comonads arise from the insight that model comparison games can be seen as semantic constructions in their own right. Although we shall not employ games as a tool, we recall two examples of games to motivate the framework of game comonads.

Henceforth we shall work with a relational vocabulary σ . Let A, B be σ -structures. Both types of game are two-player games played between Spoiler and Duplicator. Whereas Spoiler aims to show that A and B are different, Duplicator aims to show that they are similar. Each game is played in a number of rounds:

Ehrenfeucht-Fraïssé game: In the i th round, Spoiler chooses an element in one of the structures and Duplicator responds by choosing an element in the other structure. After k rounds have been played, we have sequences $[a_1, \dots, a_k]$ and $[b_1, \dots, b_k]$ of elements from A and B respectively. Duplicator wins this play if the ensuing relation $r := \{(a_i, b_i) \mid 1 \leq i \leq k\}$ is a partial isomorphism. Duplicator wins the k -round game if they have a strategy which is winning after i rounds, for all $1 \leq i \leq k$.

Bisimulation game: Suppose σ is a modal vocabulary. The game is played between pointed Kripke structures (A, a) and (B, b) , where $a \in A$ and $b \in B$. The initial position is $(a_0, b_0) = (a, b)$. In the i th round, with current position (a_{i-1}, b_{i-1}) , Spoiler chooses one of the two structures, say A , a binary relation symbol R in σ , and an element $a_i \in A$ such that $(a_{i-1}, a_i) \in R^A$. Duplicator responds by choosing an element in the other structure, say $b_i \in B$, such that $(b_{i-1}, b_i) \in R^B$. If Duplicator has no such response, they lose. Duplicator wins the k -round game if, for all unary relation symbols P in σ , we have $P^A(a_i) \iff P^B(b_i)$ for all $0 \leq i \leq k$.

Assume the vocabulary σ is finite. The classical Ehrenfeucht-Fraïssé theorem [21,22] states that Duplicator has a winning strategy in the k -round Ehrenfeucht-Fraïssé game played between A and B if, and only if, A and B satisfy the same first-order sentences of quantifier rank at most k . Similarly, Duplicator has a

³ We employ the nomenclature “logic fragment”, rather than the more customary term “theory”, as we are mainly interested in the situation where \mathcal{L} is defined by constraining the syntactic shape of sentences.

winning strategy in the k -round bisimulation game played between pointed Kripke structures (A, a) and (B, b) if, and only if, (A, a) and (B, b) satisfy the same modal formulas of modal depth at most k [25].

Remark 2.1. In both Ehrenfeucht-Fraïssé and bisimulation games, the resource parameter is the number of rounds. This need not be the case in general. For instance, the resource parameter for pebble games [13,27], which correspond to finite-variable fragments of first-order logic, is the number of pebbles available to the players.

Next, we introduce the comonads corresponding to Ehrenfeucht-Fraïssé and bisimulation games, respectively. For each σ -structure A , denote by $\mathbb{E}_k(A)$ the set of all non-empty sequences of length at most k of elements from A . In other words, $\mathbb{E}_k(A)$ is the set of all plays in A in the k -round Ehrenfeucht-Fraïssé game. The interpretations of the relation symbols can be lifted from A to $\mathbb{E}_k(A)$ as follows. Let $\varepsilon_A: \mathbb{E}_k(A) \rightarrow A$ be the function sending a sequence to its last element. For each relation symbol $R \in \sigma$ of arity j , we define $R^{\mathbb{E}_k(A)}$ to be the set of all tuples $(s_1, \dots, s_j) \in \mathbb{E}_k(A)^j$ such that:

- (i) The sequences s_1, \dots, s_j are pairwise comparable in the prefix order.
- (ii) $(\varepsilon_A(s_1), \dots, \varepsilon_A(s_j)) \in R^A$.

For every homomorphism $f: A \rightarrow B$, let

$$\mathbb{E}_k(f): \mathbb{E}_k(A) \rightarrow \mathbb{E}_k(B), \quad [a_1, \dots, a_l] \mapsto [f(a_1), \dots, f(a_l)].$$

This yields a comonad (in fact, a *family* of comonads indexed by $k > 0$) on the category $\mathbf{Struct}(\sigma)$ of σ -structures and their homomorphisms, known as *Ehrenfeucht-Fraïssé comonad* [8]. The underlying functor of this comonad is $\mathbb{E}_k: \mathbf{Struct}(\sigma) \rightarrow \mathbf{Struct}(\sigma)$, the counit is ε , and the comultiplication at A is the homomorphism

$$\mathbb{E}_k(A) \rightarrow \mathbb{E}_k \mathbb{E}_k(A), \quad [a_1, \dots, a_l] \mapsto [[a_1], [a_1, a_2], \dots, [a_1, \dots, a_l]].$$

A similar construction applies to k -round bisimulation games. Suppose σ is a modal vocabulary and let (A, a) be a pointed Kripke structure. We define a Kripke structure $\mathbb{M}_k(A, a)$ whose carrier is the set of all paths p of length at most k starting from a :

$$a \xrightarrow{R_1} a_1 \xrightarrow{R_2} a_2 \rightarrow \dots \xrightarrow{R_n} a_n$$

where R_1, \dots, R_n are binary relation symbols in σ . If $P \in \sigma$ is unary then $P^{\mathbb{M}_k(A, a)}$ is the set of paths p whose last element a_n belongs to P^A . For a binary relation symbol R in σ , $R^{\mathbb{M}_k(A, a)}$ is the set of pairs of paths (p, p') such that p' is obtained by extending p by one step along R . The distinguished element of the Kripke structure $\mathbb{M}_k(A, a)$ is the trivial path $\langle a \rangle$ of length 0, and the function $\varepsilon_{(A, a)}: (\mathbb{M}_k(A, a), \langle a \rangle) \rightarrow (A, a)$ sending a path to its last element is a morphism of pointed Kripke structures. By similar reasoning to the one above, the assignment $(A, a) \mapsto (\mathbb{M}_k(A, a), \langle a \rangle)$ can be upgraded to a comonad on the category $\mathbf{Struct}_\bullet(\sigma)$ of pointed Kripke structures and their homomorphisms, the *modal comonad* [8].

In addition to the examples mentioned above, the framework of game comonads covers a number of model comparison games, cf. e.g. [3–5,17,28,38]. In each case, they yield structural (syntax-free) characterisations of equivalence in the corresponding logic fragments. This will be illustrated from an axiomatic perspective in Section 3.4.

3. Arboreal categories

In this section, we recall from [7] the basic definitions and facts concerning arboreal categories. All categories under consideration are assumed to be locally small and *well-powered*, i.e. every object has a *small* set of subobjects (as opposed to a proper class).

3.1. Proper factorisation systems

Given arrows e and m in a category \mathcal{C} , we say that e has the *left lifting property* with respect to m , or that m has the *right lifting property* with respect to e , if for every commutative square as on the left-hand side below

$$\begin{array}{ccc}
 \cdot & \longrightarrow & \cdot \\
 e \downarrow & & \downarrow m \\
 \cdot & \longrightarrow & \cdot
 \end{array}
 \qquad
 \begin{array}{ccc}
 \cdot & \longrightarrow & \cdot \\
 e \downarrow & \nearrow d & \downarrow m \\
 \cdot & \longrightarrow & \cdot
 \end{array}$$

there is an arrow d such that the right-hand diagram above commutes. For any class \mathcal{H} of morphisms in \mathcal{C} , let ${}^{\text{rl}}\mathcal{H}$ (respectively \mathcal{H}^{rl}) be the class of morphisms having the left (respectively right) lifting property with respect to every morphism in \mathcal{H} .

Definition 3.1. A pair of classes of morphisms $(\mathcal{Q}, \mathcal{M})$ in a category \mathcal{C} is a *weak factorisation system* provided it satisfies the following conditions:

- (i) Every morphism f in \mathcal{C} can be written as $f = m \circ e$ with $e \in \mathcal{Q}$ and $m \in \mathcal{M}$.
- (ii) $\mathcal{Q} = {}^{\text{rl}}\mathcal{M}$ and $\mathcal{M} = \mathcal{Q}^{\text{rl}}$.

A *proper factorisation system* is a weak factorisation system $(\mathcal{Q}, \mathcal{M})$ such that each arrow in \mathcal{Q} is epic and each arrow in \mathcal{M} is monic. We refer to \mathcal{M} -morphisms as *embeddings* and denote them by \hookrightarrow . \mathcal{Q} -morphisms will be referred to as *quotients* and denoted by \twoheadrightarrow .

Next, we state some well known properties of proper factorisation systems (cf. e.g. [23] or [40]) which will be used throughout the paper without further reference.

Lemma 3.2. *Let $(\mathcal{Q}, \mathcal{M})$ be a proper factorisation system in \mathcal{C} . The following hold:*

- (a) \mathcal{Q} and \mathcal{M} are closed under compositions.
- (b) \mathcal{Q} contains all retractions, \mathcal{M} contains all sections, and $\mathcal{Q} \cap \mathcal{M} = \{\text{isomorphisms}\}$.
- (c) The pullback of an \mathcal{M} -morphism along any morphism, if it exists, is in \mathcal{M} .
- (d) $g \circ f \in \mathcal{Q}$ implies $g \in \mathcal{Q}$.
- (e) $g \circ f \in \mathcal{M}$ implies $f \in \mathcal{M}$.

Assume \mathcal{C} is a category admitting a proper factorisation system $(\mathcal{Q}, \mathcal{M})$. In the same way that one usually defines the poset of subobjects of a given object $X \in \mathcal{C}$, we can define the poset $\mathbb{S}X$ of \mathcal{M} -subobjects of X . Given embeddings $m: S \hookrightarrow X$ and $n: T \hookrightarrow X$, let us say that $m \leq n$ provided there is a morphism $i: S \rightarrow T$ such that $m = n \circ i$ (if it exists, i is necessarily an embedding).

$$\begin{array}{ccc}
 S & \xrightarrow{m} & X \\
 \downarrow i & \nearrow n & \\
 T & &
 \end{array}$$

This yields a preorder on the class of all embeddings with codomain X . The symmetrisation \sim of \leq is given by $m \sim n$ if, and only if, there exists an isomorphism $i: S \rightarrow T$ such that $m = n \circ i$. Let $\mathbb{S}X$ be the class of \sim -equivalence classes of embeddings with codomain X , equipped with the natural partial order \leq induced by \leq . We shall systematically represent a \sim -equivalence class by any of its representatives. As \mathcal{C} is well-powered and every embedding is a monomorphism, $\mathbb{S}X$ is a set.

For any morphism $f: X \rightarrow Y$ and embedding $m: S \rightarrow X$, consider the (quotient, embedding) factorisation of $f \circ m$:

$$S \twoheadrightarrow \exists_f S \xrightarrow{\exists_f m} Y.$$

This yields a monotone map $\exists_f: \mathbb{S}X \rightarrow \mathbb{S}Y$ sending m to $\exists_f m$. Note that the map \exists_f is well-defined because factorisations are unique up to isomorphism. If f is an embedding (or, more generally, $f \circ m$ is an embedding), $\exists_f m$ can be identified with $f \circ m$. For the following observation, cf. e.g. [7, Lemma 2.7(a)].

Lemma 3.3. *Let \mathcal{C} be a category equipped with a proper factorisation system and let $f: X \rightarrow Y$ be an embedding in \mathcal{C} . Then $\exists_f: \mathbb{S}X \rightarrow \mathbb{S}Y$ is an order-embedding.*

3.2. Paths and arboreal categories

Let \mathcal{C} be a category endowed with a proper factorisation system $(\mathcal{Q}, \mathcal{M})$.

Definition 3.4. An object X of \mathcal{C} is called a *path* provided the poset $\mathbb{S}X$ is a finite linear order. Paths will be denoted by P, Q and variations thereof.

The collection of paths is closed under embeddings and quotients. That is, given an arrow $f: X \rightarrow Y$, if f is an embedding and Y is a path then X is a path, and if f is a quotient and X is a path then Y is a path [7, Lemma 3.5].

A *path embedding* is an embedding $P \rightarrow X$ whose domain is a path. We let $\mathbb{P}X$ denote the sub-poset of $\mathbb{S}X$ consisting of the path embeddings. Because paths are closed under quotients, for any arrow $f: X \rightarrow Y$ the monotone map $\exists_f: \mathbb{S}X \rightarrow \mathbb{S}Y$ restricts to a monotone map

$$\mathbb{P}f: \mathbb{P}X \rightarrow \mathbb{P}Y, \quad (m: P \rightarrow X) \mapsto (\exists_f m: \exists_f P \rightarrow Y). \quad (1)$$

Henceforth, we shall assume that \mathcal{C} contains, up to isomorphism, only a set of paths. Thus, whenever convenient, we shall work with a small skeleton of the full subcategory of paths. For any object X of \mathcal{C} , we have a diagram with vertex X consisting of all path embeddings with codomain X ; the morphisms between paths are those which make the obvious triangles commute (note that they are necessarily embeddings):

$$\begin{array}{ccc} & X & \\ \curvearrowright & & \curvearrowleft \\ P & \xrightarrow{\quad} & Q \end{array}$$

This yields a cocone over the small diagram $\mathbb{P}X$. We say that X is *path-generated* provided this is a colimit cocone in \mathcal{C} .

Suppose for a moment that coproducts of sets of paths exist in \mathcal{C} . An object X of \mathcal{C} is *connected* if, for all non-empty sets of paths $\{P_i \mid i \in I\}$ in \mathcal{C} , any morphism

$$X \rightarrow \coprod_{i \in I} P_i$$

factors through some coproduct injection $P_j \rightarrow \coprod_{i \in I} P_i$.

In order to state the definition of arboreal category, let us say that a proper factorisation system is *stable* if, for any quotient e and embedding m with common codomain, the pullback of e along m exists and is a quotient.

Definition 3.5. An *arboreal category* is a category \mathcal{C} , equipped with a stable proper factorisation system, that satisfies the following conditions:

- (i) \mathcal{C} has all coproducts of sets of paths.
- (ii) For any paths P, Q, Q' in \mathcal{C} , if a composite $P \rightarrow Q \rightarrow Q'$ is a quotient then so is $P \rightarrow Q$.
- (iii) Every object of \mathcal{C} is path-generated.
- (iv) Every path in \mathcal{C} is connected.

Remark 3.6. Item (ii) in the previous definition is equivalent to the following *2-out-of-3 condition*: For any paths P, Q, Q' and morphisms

$$P \xrightarrow{f} Q \xrightarrow{g} Q',$$

if any two of f, g , and $g \circ f$ are quotients, then so is the third. See [7, Remark 3.9]. Moreover, item (iii) is equivalent to saying that the inclusion functor $\mathcal{C}_p \hookrightarrow \mathcal{C}$ is dense, where \mathcal{C}_p is the full subcategory of \mathcal{C} defined by the paths [7, Lemma 5.1].

Finally, note that any arboreal category admits an initial object, obtained as the coproduct of the empty set, and any initial object is a path because its poset of \mathcal{M} -subobjects has a single element—namely, the equivalence class of the identity.

If (P, \leq) is a poset, then $C \subseteq P$ is a *chain* if it is linearly ordered. (P, \leq) is a *forest* if, for all $x \in P$, the set $\downarrow x := \{y \in P \mid y \leq x\}$ is a finite chain. The *height* of a forest is the supremum of the cardinalities of its chains. The *covering relation* \prec associated with a partial order \leq is defined by $u \prec v$ if and only if $u < v$ and there is no w such that $u < w < v$. The *roots* of a forest are the minimal elements, and a *tree* is a forest with at most one root. Morphisms of forests are maps that preserve roots and the covering relation. The category of forests is denoted by \mathcal{F} , and the full subcategory of trees by \mathcal{T} .

Example 3.7. Let σ be a relational vocabulary. A *forest-ordered σ -structure* (A, \leq) is a σ -structure A equipped with a forest order \leq . A morphism of forest-ordered σ -structures $f: (A, \leq) \rightarrow (B, \leq')$ is a σ -homomorphism $f: A \rightarrow B$ that is also a forest morphism. This determines a category $\mathcal{R}(\sigma)$. We equip $\mathcal{R}(\sigma)$ with the factorisation system given by (surjective morphisms, embeddings), where an embedding is a morphism which is an embedding *qua* σ -homomorphism. Let $\mathcal{R}^E(\sigma)$ be the full subcategory of $\mathcal{R}(\sigma)$ determined by those objects (A, \leq) satisfying the following condition:

- (E) If $a, b \in A$ are distinct elements that appear in a tuple of related elements $(a_1, \dots, a_l) \in R^A$ for some $R \in \sigma$, then either $a < b$ or $b < a$.⁴

For each $k > 0$, let $\mathcal{R}_k^E(\sigma)$ be the full subcategory of $\mathcal{R}^E(\sigma)$ of forest-ordered structures of height $\leq k$. In [9, Theorem 9.1], it is shown that $\mathcal{R}_k^E(\sigma)$ is isomorphic to the category of coalgebras for the Ehrenfeucht-Fraïssé comonad \mathbb{E}_k on $\mathbf{Struct}(\sigma)$. The objects (A, \leq) of $\mathcal{R}_k^E(\sigma)$ are forest covers of A witnessing that its *tree-depth* is at most k [36].

⁴ I.e., if a and b are adjacent in the *Gaifman graph* of A , then they are comparable in the forest order.

The category $\mathcal{R}^E(\sigma)$ is arboreal when equipped with the restriction of the factorisation system on $\mathcal{R}(\sigma)$. The paths in $\mathcal{R}^E(\sigma)$ are those objects in which the order is a finite chain. Similarly, $\mathcal{R}_k^E(\sigma)$ is an arboreal category for all $k > 0$, when equipped with the restriction of the factorisation system on $\mathcal{R}(\sigma)$. See [7, Examples 5.4].

Example 3.8. Assume that σ is a modal vocabulary. Let $\mathcal{R}^M(\sigma)$ be the full subcategory of $\mathcal{R}(\sigma)$ consisting of the tree-ordered σ -structures (A, \leq) satisfying:

(M) For $a, b \in A$, $a < b$ if and only if $(a, b) \in R^A$ for some unique binary relation R .

For each $k > 0$, the full subcategory $\mathcal{R}_k^M(\sigma)$ of $\mathcal{R}^M(\sigma)$ consisting of the tree-ordered structures of height at most k is isomorphic to the category of coalgebras for the modal comonad \mathbb{M}_k on $\mathbf{Struct}_\bullet(\sigma)$ [9, Theorem 9.5]. When equipped with the restriction of the factorisation system on $\mathcal{R}(\sigma)$, the category $\mathcal{R}^M(\sigma)$ is arboreal and its paths are those objects in which the order is a finite chain. Likewise for $\mathcal{R}_k^M(\sigma)$.

It follows from the definition of path that, for any object X of an arboreal category, the poset $\mathbb{P} X$ is a tree; in fact, a non-empty tree. Crucially, this assignment extends to a functor into the category of trees (for a proof, see [7, Theorem 3.11]):

Theorem 3.9. Let \mathcal{C} be an arboreal category. The assignment $f \mapsto \mathbb{P} f$ in equation (1) induces a functor $\mathbb{P}: \mathcal{C} \rightarrow \mathcal{T}$ into the category of trees.

Finally, recall from [7, Lemma 3.15, Proposition 5.6 and Remark 5.7] the following properties of paths and posets of embeddings.

Lemma 3.10. The following statements hold in any arboreal category \mathcal{C} :

- (a) Between any two paths there is at most one embedding.
- (b) For all objects X of \mathcal{C} , the poset $\mathbb{S} X$ of its \mathcal{M} -subobjects is a complete lattice.⁵
- (c) Let X be an object of \mathcal{C} and let $\mathcal{U} \subseteq \mathbb{P} X$ be a non-empty subset. A path embedding $m \in \mathbb{P} X$ is below $\bigvee \mathcal{U} \in \mathbb{S} X$ if, and only if, it is below some element of \mathcal{U} .

If it exists, the unique embedding between paths P, Q in an arboreal category is denoted by

$$!_{P,Q}: P \rightarrow Q.$$

If no confusion arises, we simply write $!: P \rightarrow Q$.

3.3. Bisimilarity and back-and-forth systems

Throughout this section, we fix an arbitrary arboreal category \mathcal{C} . A morphism $f: X \rightarrow Y$ in \mathcal{C} is said to be *open* if it satisfies the following path-lifting property: Given any commutative square

$$\begin{array}{ccc} P & \xrightarrow{\quad} & X \\ \downarrow & \nearrow \text{---} & \downarrow f \\ Q & \xrightarrow{\quad} & Y \end{array}$$

⁵ In fact, $\mathbb{S} X$ is a *perfect* lattice, cf. [39] or [18].

with P, Q paths, there is an arrow $Q \rightarrow X$ making the two triangles commute. (If such an arrow exists, it is automatically an embedding.) Further, f is a *pathwise embedding* if, for all path embeddings $m: P \rightarrow X$, the composite $f \circ m: P \rightarrow Y$ is a path embedding.

Combining these notions, we can define a bisimilarity relation between objects of \mathcal{C} :

Definition 3.11. Two objects X, Y of \mathcal{C} are said to be *bisimilar* if there exist an object Z of \mathcal{C} and a span of open pathwise embeddings $X \leftarrow Z \rightarrow Y$.

Remark 3.12. The definition of open morphism given above is a refinement of the one introduced in [30] (cf. [7, §4.1] for a discussion of the relation between these notions), which is a special case of the (axiomatic) concept of open map in toposes [29].

As we shall see next, if \mathcal{C} has binary products the bisimilarity relation can be characterised in terms of back-and-forth systems. Let X, Y be objects of \mathcal{C} . Given $m \in \mathbb{P} X$ and $n \in \mathbb{P} Y$, we write $\llbracket m, n \rrbracket$ to indicate that $\text{dom}(m) \cong \text{dom}(n)$. Intuitively, the pair $\llbracket m, n \rrbracket$ encodes a partial isomorphisms between X and Y “of shape P ”, with P a path.

Definition 3.13. A *back-and-forth system* between objects X and Y of \mathcal{C} is a set

$$\mathcal{B} = \{\llbracket m_i, n_i \rrbracket \mid m_i \in \mathbb{P} X, n_i \in \mathbb{P} Y, i \in I\}$$

satisfying the following conditions:

- (i) $\llbracket \perp_X, \perp_Y \rrbracket \in \mathcal{B}$, where \perp_X, \perp_Y are the roots of $\mathbb{P} X$ and $\mathbb{P} Y$, respectively.
- (ii) If $\llbracket m, n \rrbracket \in \mathcal{B}$ and $m' \in \mathbb{P} X$ are such that $m \prec m'$, there exists $n' \in \mathbb{P} Y$ satisfying $n \prec n'$ and $\llbracket m', n' \rrbracket \in \mathcal{B}$.
- (iii) If $\llbracket m, n \rrbracket \in \mathcal{B}$ and $n' \in \mathbb{P} Y$ are such that $n \prec n'$, there exists $m' \in \mathbb{P} X$ satisfying $m \prec m'$ and $\llbracket m', n' \rrbracket \in \mathcal{B}$.

Two objects X and Y of \mathcal{C} are said to be *back-and-forth equivalent* if there exists a back-and-forth system between them.

For a proof of the following result, see [7, Theorem 6.4].

Theorem 3.14. *Let X, Y be objects of an arboreal category admitting a product. Then X and Y are bisimilar if, and only if, they are back-and-forth equivalent.*

The existence of a back-and-forth system between X and Y can be equivalently described in terms of the existence of a Duplicator winning strategy in a two-player game played between $\mathbb{P} X$ and $\mathbb{P} Y$ [7, §6.2]. Since winning strategies can be composed to yield again a winning strategy, in any arboreal category with binary products the bisimilarity relation is transitive, hence an equivalence relation.

3.4. Resource-indexed arboreal adjunctions

Let \mathcal{C} be an arboreal category, with full subcategory of paths \mathcal{C}_p . We say that \mathcal{C} is *resource-indexed* if for all positive integers k there is a full subcategory \mathcal{C}_p^k of \mathcal{C}_p closed under embeddings⁶ with

⁶ That is, for any embedding $P \rightarrow Q$ in \mathcal{C} with P, Q paths, if $Q \in \mathcal{C}_p^k$ then also $P \in \mathcal{C}_p^k$. We shall further assume that each category \mathcal{C}_p^k contains the initial object of \mathcal{C} .

$$\mathcal{C}_p^1 \hookrightarrow \mathcal{C}_p^2 \hookrightarrow \mathcal{C}_p^3 \hookrightarrow \dots$$

This induces a corresponding tower of full subcategories \mathcal{C}_k of \mathcal{C} , with the objects of \mathcal{C}_k those whose cocone of path embeddings with domain in \mathcal{C}_p^k is a colimit cocone in \mathcal{C} . It turns out that each category \mathcal{C}_k is arboreal. Furthermore, the paths in \mathcal{C}_k are precisely the objects of \mathcal{C}_p^k , i.e. $(\mathcal{C}_k)_p = \mathcal{C}_p^k$. Cf. [7, Proposition 7.7] and its proof.

Example 3.15. Consider the arboreal category $\mathcal{R}^E(\sigma)$ from Example 3.7. This can be regarded as a resource-indexed arboreal category by taking as \mathcal{C}_p^k the full subcategory of $\mathcal{R}^E(\sigma)$ consisting of the objects in which the order is a finite chain of cardinality $\leq k$. The generated subcategory \mathcal{C}_k then coincides with $\mathcal{R}_k^E(\sigma)$.

Similar reasoning shows that the arboreal category $\mathcal{R}^M(\sigma)$ from Example 3.8 can also be regarded as a resource-indexed arboreal category.

Definition 3.16. Let $\{\mathcal{C}_k\}$ be a resource-indexed arboreal category and let \mathcal{E} be a category. A *resource-indexed arboreal adjunction* between \mathcal{E} and \mathcal{C} is an indexed family of adjunctions

$$\mathcal{C}_k \begin{array}{c} \xrightarrow{L_k} \\ \perp \\ \xleftarrow{R_k} \end{array} \mathcal{E}.$$

A *resource-indexed arboreal cover* of \mathcal{E} by \mathcal{C} is a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} such that all adjunctions $L_k \dashv R_k$ are comonadic, i.e. for all $k > 0$ the comparison functor from \mathcal{C}_k to the category of Eilenberg-Moore coalgebras for the comonad $G_k := L_k R_k$ is an isomorphism.

Example 3.17. Let σ be a relational vocabulary and let $\mathcal{E} = \mathbf{Struct}(\sigma)$. Consider the resource-indexed arboreal category $\mathcal{R}^E(\sigma)$ in Example 3.15. For each $k > 0$, there is a forgetful functor

$$L_k^E : \mathcal{R}_k^E(\sigma) \rightarrow \mathbf{Struct}(\sigma)$$

which forgets the forest order. This functor is comonadic. The right adjoint R_k^E sends a σ -structure A to $\mathbb{E}_k(A)$ equipped with the prefix order, and the comonad induced by this adjunction coincides with the k -round Ehrenfeucht-Fraïssé comonad \mathbb{E}_k . This gives rise to a resource-indexed arboreal cover of $\mathbf{Struct}(\sigma)$ by $\mathcal{R}^E(\sigma)$.

Similarly, if σ is a modal vocabulary, there is a resource-indexed arboreal cover of $\mathbf{Struct}_\bullet(\sigma)$ by $\mathcal{R}^M(\sigma)$ such that each adjunction $L_k^M \dashv R_k^M$ induces the k -round modal comonad \mathbb{M}_k .

Example 3.18. To deal with the equality symbol in the logic, we consider resource-indexed arboreal adjunctions constructed as follows. Let

$$\sigma^I := \sigma \cup \{I\}$$

be the vocabulary obtained by adding a fresh binary relation symbol I to σ . Any σ -structure can be expanded to a σ^I -structure by interpreting I as the identity relation. This yields a fully faithful functor $J : \mathbf{Struct}(\sigma) \rightarrow \mathbf{Struct}(\sigma^I)$. The functor J has a left adjoint $H : \mathbf{Struct}(\sigma^I) \rightarrow \mathbf{Struct}(\sigma)$ which sends a σ^I -structure A to the quotient of the σ -reduct of A by the equivalence relation generated by I^A [19, Lemma 25]. We can then compose the adjunction $H \dashv J$ with e.g. the Ehrenfeucht-Fraïssé resource-indexed arboreal cover of $\mathbf{Struct}(\sigma^I)$ by $\mathcal{R}^E(\sigma^I)$ from Example 3.17.

$$\begin{array}{ccc}
 & \xrightarrow{L_k^E} & \\
 \mathcal{R}_k^E(\sigma^I) & \perp & \mathbf{Struct}(\sigma^I) \\
 & \xleftarrow{R_k^E} & \\
 & \xrightarrow{H} & \\
 & \perp & \mathbf{Struct}(\sigma) \\
 & \xleftarrow{J} &
 \end{array}$$

The composite adjunctions $HL_k^E \dashv R_k^E J$, which are not comonadic, yield the *Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction* between $\mathbf{Struct}(\sigma)$ and $\mathcal{R}^E(\sigma^I)$.

Crucially, a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} can be used to define several resource-indexed relations between objects of \mathcal{E} :

Definition 3.19. Consider a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} , with adjunctions $L_k \dashv R_k$, and any two objects a, b of \mathcal{E} . For all $k > 0$, we define:

- $a \rightarrow_k^{\mathcal{C}} b$ if there exists a morphism $R_k a \rightarrow R_k b$ in \mathcal{C}_k .
- $a \leftrightarrow_k^{\mathcal{C}} b$ if $R_k a$ and $R_k b$ are bisimilar in \mathcal{C}_k .
- $a \cong_k^{\mathcal{C}} b$ if $R_k a$ and $R_k b$ are isomorphic in \mathcal{C}_k .

Further, we write $\rightleftarrows_k^{\mathcal{C}}$ for the symmetrisation of the preorder $\rightarrow_k^{\mathcal{C}}$. There are inclusions

$$\cong_k^{\mathcal{C}} \subseteq \leftrightarrow_k^{\mathcal{C}} \subseteq \rightleftarrows_k^{\mathcal{C}}.$$

The first inclusion is trivial, the second follows from [7, Lemma 6.20]. For a proof of the following easy observation, see [7, Lemma 7.20].

Lemma 3.20. Consider a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} , with adjunctions $L_k \dashv R_k$. The following hold for all $a, b \in \mathcal{E}$ and all $k > 0$:

- (a) If there exists a morphism $a \rightarrow b$ in \mathcal{E} then $a \rightarrow_k^{\mathcal{C}} b$.
- (b) $a \rightleftarrows_k^{\mathcal{C}} L_k R_k a$.

To conclude, we recall how the relations in Definition 3.19 capture, in our running examples, preservation of the logics introduced at the beginning of Section 2. Given a set of sentences (or modal formulas) \mathcal{L} , let $\Rightarrow^{\mathcal{L}}$ be the preorder on (pointed) σ -structures given by

$$A \Rightarrow^{\mathcal{L}} B \iff \forall \varphi \in \mathcal{L}. (A \models \varphi \Rightarrow B \models \varphi).$$

The equivalence relation $\equiv^{\mathcal{L}}$ is the symmetrisation of $\Rightarrow^{\mathcal{L}}$.

Example 3.21. Let σ be a finite relational vocabulary. Consider the Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction between $\mathbf{Struct}(\sigma)$ and $\mathcal{R}^E(\sigma^I)$ in Example 3.18, and write \rightarrow_k^E and \leftrightarrow_k^E for the relations on $\mathbf{Struct}(\sigma)$ induced according to Definition 3.19. For all σ -structures A, B and all $k > 0$, we have

$$A \rightarrow_k^E B \iff A \Rightarrow^{\exists^+ \text{FO}_k} B$$

and

$$A \leftrightarrow_k^E B \iff A \equiv^{\text{FO}_k} B.$$

Cf. [9, Theorems 3.2 and 5.1] and [9, Theorem 10.5], respectively. We also mention that \cong_k^E coincides with equivalence in the extension of FO_k with *counting quantifiers* [9, Theorem 5.3(2)], although we shall not need this fact.

Example 3.22. Suppose σ is a finite modal vocabulary and consider the relations \rightarrow_k^M and \cong_k^M on $\mathbf{Struct}_\bullet(\sigma)$ induced by the modal resource-indexed arboreal cover of $\mathbf{Struct}_\bullet(\sigma)$ by $\mathcal{R}_k^M(\sigma)$ in Example 3.17. For all pointed Kripke structures $(A, a), (B, b)$ and all $k > 0$, we have

$$(A, a) \rightarrow_k^M (B, b) \iff (A, a) \Rrightarrow^{\exists^+ \text{ML}_k} (B, b),$$

see [8, Theorem 9]. Furthermore,

$$(A, a) \cong_k^M (B, b) \implies (A, a) \equiv^{\text{ML}_k(\#)} (B, b),$$

cf. [8, Proposition 15] and [20, Proposition 3.6]. We mention in passing that the relation \leftrightarrow_k^M coincides with equivalence in ML_k [9, Theorem 10.13].

4. Homomorphism preservation theorems

In this section, we recast the statement of a generic equi-resource homomorphism preservation theorem into a property (HP)—and its strengthening (HP[#])—that a resource-indexed arboreal adjunction may or may not satisfy. Sufficient conditions under which properties (HP) and (HP[#]) hold are provided in Section 5.

4.1. (HP) and (HP[#])

Given a first-order sentence φ in a relational vocabulary σ , its “model class” $\mathbf{Mod}(\varphi)$ is the full subcategory of $\mathbf{Struct}(\sigma)$ defined by the σ -structures A such that $A \models \varphi$. To motivate the formulation of properties (HP) and (HP[#]), we recall a well-known characterisation of model classes of sentences in FO_k , i.e. first-order sentences of quantifier rank at most k , and in its existential positive fragment $\exists^+ \text{FO}_k$. Since a sentence can only contain finitely many relation symbols, for the purpose of investigating homomorphism preservation theorems we can safely assume that σ is finite.

For a full subcategory \mathcal{D} of a category \mathcal{A} , and a relation ∇ on the class of objects of \mathcal{A} , we say that \mathcal{D} is *upwards closed (in \mathcal{A}) with respect to ∇* if

$$\forall a, b \in \mathcal{A}, \text{ if } a \in \mathcal{D} \text{ and } a \nabla b \text{ then } b \in \mathcal{D}.$$

If ∇ is an equivalence relation and the latter condition is satisfied, we say that \mathcal{D} is *saturated under ∇* . The following lemma follows from the fact that, for all $k \geq 0$, there are finitely many sentences in FO_k up to logical equivalence. Cf. e.g. [32, Lemma 3.13].

Lemma 4.1. *The following hold for all $k \geq 0$ and all full subcategories \mathcal{D} of $\mathbf{Struct}(\sigma)$:*

- (a) $\mathcal{D} = \mathbf{Mod}(\varphi)$ for some $\varphi \in \text{FO}_k$ if, and only if, \mathcal{D} is saturated under \equiv^{FO_k} .
- (b) $\mathcal{D} = \mathbf{Mod}(\psi)$ for some $\psi \in \exists^+ \text{FO}_k$ if, and only if, \mathcal{D} is upwards closed with respect to $\Rrightarrow^{\exists^+ \text{FO}_k}$.

Remark 4.2. The previous lemma remains true if FO_k is replaced with any fragment of first-order logic that is closed under Boolean connectives and contains, up to logical equivalence, finitely many sentences.

Now, fix a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} , with adjunctions

$$\mathcal{C}_k \begin{array}{c} \xrightarrow{L_k} \\ \perp \\ \xleftarrow{R_k} \end{array} \mathcal{E}.$$

Let us say that a full subcategory \mathcal{D} of \mathcal{E} is *closed (in \mathcal{E}) under morphisms* if, whenever there is an arrow $a \rightarrow b$ in \mathcal{E} with $a \in \mathcal{D}$, also $b \in \mathcal{D}$. Note that, when $\mathcal{E} = \mathbf{Struct}(\sigma)$ and $\mathcal{D} = \mathbf{Mod}(\varphi)$ is the model class of some sentence φ , the category \mathcal{D} is closed under morphisms precisely when φ is preserved under homomorphisms.

Consider the following statement, where $\rightarrow_k^{\mathcal{C}}$ and $\leftrightarrow_k^{\mathcal{C}}$ are the relations on the objects of \mathcal{E} induced by the resource-indexed arboreal adjunction as in Definition 3.19:

(HP) For any full subcategory \mathcal{D} of \mathcal{E} saturated under $\leftrightarrow_k^{\mathcal{C}}$, \mathcal{D} is closed under morphisms precisely when it is upwards closed with respect to $\rightarrow_k^{\mathcal{C}}$.

Replacing the relation $\leftrightarrow_k^{\mathcal{C}}$ with $\cong_k^{\mathcal{C}}$, we obtain a strengthening of (HP), namely:

(HP#) For any full subcategory \mathcal{D} of \mathcal{E} saturated under $\cong_k^{\mathcal{C}}$, \mathcal{D} is closed under morphisms precisely when it is upwards closed with respect to $\rightarrow_k^{\mathcal{C}}$.

Just recall that $\cong_k^{\mathcal{C}} \subseteq \leftrightarrow_k^{\mathcal{C}}$, and so any full subcategory \mathcal{D} saturated under $\leftrightarrow_k^{\mathcal{C}}$ is also saturated under $\cong_k^{\mathcal{C}}$. Thus, (HP#) entails (HP).

Remark 4.3. By Lemma 3.20(a), any full subcategory of \mathcal{E} that is upwards closed with respect to $\rightarrow_k^{\mathcal{C}}$ is closed under morphisms. Hence, the right-to-left implications in (HP) and (HP#) are always satisfied.

In view of Example 3.21 and Lemma 4.1, for the Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction between $\mathbf{Struct}(\sigma)$ and $\mathcal{R}^E(\sigma^I)$, Property (HP) coincides with Rossman’s equirank homomorphism preservation theorem (Theorem 1.2).

In Section 6 we will prove that (HP) holds for any resource-indexed arboreal adjunction satisfying appropriate properties (see Corollary 6.10), which are satisfied in particular by the Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction.

5. Exploring the landscape: tame and not-so-tame

We shall identify, in Section 5.1, a class of “tame” resource-indexed arboreal adjunctions, namely those satisfying the *bisimilar companion property*, for which (HP) always holds. As an application, we derive equi-resource homomorphism preservation theorems for graded modal logic and guarded first-order logics.

In the absence of the bisimilar companion property, one may try to “force” it; this leads us, in Section 5.2, to the notion of *extendability*, inspired by the work of Rossman [41]. Finally, in Section 5.3, we provide simple sufficient conditions under which properties (HP) and (HP#) admit a relativisation to a full subcategory.

5.1. Tame: bisimilar companion property and idempotency

For all $k > 0$, write $G_k := L_k R_k$ for the comonad on \mathcal{E} induced by the adjunction $L_k \dashv R_k: \mathcal{E} \rightarrow \mathcal{C}_k$.

Definition 5.1. A resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} , with induced comonads G_k , has the *bisimilar companion property* if $a \leftrightarrow_k^{\mathcal{C}} G_k a$ for all $a \in \mathcal{E}$ and $k > 0$.

Proposition 5.2. (HP) holds for any resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} satisfying the bisimilar companion property.

Proof. For the left-to-right implication in (HP), let \mathcal{D} be a full subcategory of \mathcal{E} closed under morphisms and saturated under $\leftrightarrow_k^{\mathcal{C}}$. Suppose that $a \rightarrow_k^{\mathcal{C}} b$ for objects a, b of \mathcal{E} . By definition, this means that there is an arrow $R_k a \rightarrow R_k b$ and so, as L_k is left adjoint to R_k , there is an arrow $G_k a \rightarrow b$. Using the bisimilar companion property, we get

$$a \leftrightarrow_k^{\mathcal{C}} G_k a \rightarrow b.$$

Therefore, if $a \in \mathcal{D}$ then also $b \in \mathcal{D}$. That is, \mathcal{D} is upwards closed with respect to $\rightarrow_k^{\mathcal{C}}$.

The converse direction follows from Remark 4.3. \square

In order to establish a similar result for Property (HP[#]), recall that a comonad G is *idempotent* if its comultiplication $G \Rightarrow GG$ is a natural isomorphism.

Definition 5.3. A resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} is *idempotent* if so are the induced comonads G_k , for all $k > 0$.

Proposition 5.4. (HP[#]) holds for any idempotent resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} .

Proof. Recall that G_k is idempotent if, and only if, ηR_k is a natural isomorphism, where η is the unit of the adjunction $L_k \dashv R_k$. In particular, for any $a \in \mathcal{E}$, the component of ηR_k at a yields an isomorphism $R_k a \cong R_k G_k a$ in \mathcal{C} . Hence, $a \cong_k^{\mathcal{C}} G_k a$ for all $a \in \mathcal{E}$. Reasoning as in the proof of Proposition 5.2, it is easy to see that (HP[#]) holds. \square

Remark 5.5. Consider an idempotent resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} with induced comonads G_k on \mathcal{E} . The previous proof shows that, for all $a \in \mathcal{E}$ and $k > 0$, we have $a \cong_k^{\mathcal{C}} G_k a$. A fortiori, $a \leftrightarrow_k^{\mathcal{C}} G_k a$. Therefore, any idempotent resource-indexed arboreal adjunction satisfies the bisimilar companion property.

Next, we show how Propositions 5.2 and 5.4 can be exploited to obtain equi-resource homomorphism preservation theorems for (graded) modal logic and guarded first-order logics. Relativisations of these results to subclasses of structures, e.g. to the class of all finite structures, are discussed in Section 5.3.

Graded modal logic. Let σ be a finite modal vocabulary. As observed in [9, §9.3], the modal comonads \mathbb{M}_k on $\mathbf{Struct}_{\bullet}(\sigma)$ are idempotent, hence so is the modal resource-indexed arboreal cover of $\mathbf{Struct}_{\bullet}(\sigma)$ by $\mathcal{R}_k^M(\sigma)$. This corresponds to the fact that a tree-ordered Kripke structure is isomorphic to its tree unravelling. Thus, Proposition 5.4 entails the following *equidepth homomorphism preservation theorem* for graded modal formulas (i.e., modal formulas that possibly contain graded modalities):

Theorem 5.6. The following statements are equivalent for any graded modal formula φ of modal depth at most k in a modal vocabulary:

- (1) φ is preserved under homomorphisms between pointed Kripke structures.
- (2) φ is logically equivalent to an existential positive modal formula of modal depth at most k .

Proof. Fix a graded modal formula $\varphi \in \text{ML}_k(\#)$. Since a single modal formula contains only finitely many modalities and propositional variables, we can assume without loss of generality that φ is a formula in a finite modal vocabulary. By Example 3.22 we have

$$\rightarrow_k^M = \Rightarrow^{\exists^+ \text{ML}_k} \quad \text{and} \quad \cong_k^M \subseteq \equiv^{\text{ML}_k(\#)}.$$

In particular, the latter inclusion entails that the full subcategory $\mathbf{Mod}(\varphi)$ of $\mathbf{Struct}_\bullet(\sigma)$ is saturated under \cong_k^M . As the modal resource-indexed arboreal cover is idempotent, Proposition 5.4 implies that $\mathbf{Mod}(\varphi)$ is closed under morphisms if, and only if, it is upwards closed with respect to \rightarrow_k^M . Note that $\mathbf{Mod}(\varphi)$ is closed under morphisms precisely when φ is preserved under homomorphisms between pointed Kripke structures. On the other hand, the equality $\rightarrow_k^M = \Rightarrow^{\exists^+ \text{ML}_k}$ implies that $\mathbf{Mod}(\varphi)$ is upwards closed with respect to \rightarrow_k^M if, and only if, $\mathbf{Mod}(\varphi) = \mathbf{Mod}(\psi)$ for some $\psi \in \exists^+ \text{ML}_k$ (this is akin to Lemma 4.1(b) and hinges on the fact that $\exists^+ \text{ML}_k$ contains finitely many formulas up to logical equivalence). Thus the statement follows. \square

Remark 5.7. Forgetting about both graded modalities and modal depth, Theorem 5.6 implies that a modal formula is preserved under homomorphisms if, and only if, it is equivalent to an existential positive modal formula. This improves the well known result that a modal formula is preserved under simulations precisely when it is equivalent to an existential positive modal formula (see e.g. [16, Theorem 2.78]).

Guarded fragments of first-order logic. The study of guarded fragments of first-order logic was initiated by Andréka, van Benthem and Némethi in [12] to analyse, and extend to the first-order setting, the good algorithmic and model-theoretic properties of modal logic. *Guarded formulas* (over a relational vocabulary σ) are defined by structural induction, starting from atomic formulas and applying Boolean connectives and the following restricted forms of quantification: if $\varphi(\bar{x}, \bar{y})$ is a guarded formula, then so are

$$\exists \bar{x}. G(\bar{x}, \bar{y}) \wedge \varphi(\bar{x}, \bar{y}) \quad \text{and} \quad \forall \bar{x}. G(\bar{x}, \bar{y}) \rightarrow \varphi(\bar{x}, \bar{y})$$

where G is a so-called *guard*. The (syntactic) conditions imposed on guards determine different guarded fragments of first-order logic. We shall consider the following two:

- *Atom guarded:* $G(\bar{x}, \bar{y})$ is an atomic formula in which all variables in \bar{x}, \bar{y} occur.
- *Loosely guarded:* $G(\bar{x}, \bar{y})$ is a conjunction of atomic formulas such that each pair of variables, one in \bar{x} and the other in \bar{y} , occurs in one of the conjuncts.

The atom guarded fragment of first-order logic was introduced in [12] under the name of F2 (“Fragment 2”), whereas the loosely guarded fragment was defined by van Benthem in [15]. The atom guarded fragment can be regarded as an extension of modal logic, in the sense that the standard translation of the latter is contained in the former. In turn, the loosely guarded fragment extends the atomic one and can express e.g. (the translation of) the *Until* modality in temporal logic, cf. [15, p. 9].

For each notion of guarding \mathfrak{g} (atomic or loose), denote by

$$\mathfrak{gFO}^n \quad \text{and} \quad \exists^+ \mathfrak{gFO}^n,$$

respectively, the n -variable \mathfrak{g} -guarded fragment of first-order logic and its existential positive fragment. In [4], *guarded comonads* $\mathbb{G}_n^{\mathfrak{g}}$ on $\mathbf{Struct}(\sigma)$ are defined for all $n > 0$. The associated categories of Eilenberg-Moore coalgebras are arboreal and induce the \mathfrak{g} -*guarded resource-indexed arboreal cover* of $\mathbf{Struct}(\sigma)$ with resource parameter n . For an explicit description of the resource-indexed arboreal category in question, cf. [4, §IV].

Assume the vocabulary σ is finite and let $\rightarrow_n^{\mathfrak{g}}$ and $\leftrightarrow_n^{\mathfrak{g}}$ be the resource-indexed relations on $\mathbf{Struct}(\sigma)$ induced by the \mathfrak{g} -guarded resource-indexed arboreal cover. It follows from [4, Theorems III.4 and V.2] that, for all σ -structures A, B and all $n > 0$,

$$A \rightarrow_n^{\mathfrak{g}} B \iff A \Rightarrow^{\exists^+ \mathfrak{gFO}^n} B$$

and

$$A \leftrightarrow_n^{\mathfrak{g}} B \iff A \equiv^{\mathfrak{gFO}^n} B.$$

To obtain an analogue of Lemma 4.1, we consider finite fragments of \mathfrak{gFO}^n by stratifying in terms of *guarded-quantifier rank* (cf. Remark 4.2). Note that, as guarded quantifiers bound *tuples* of variables, rather than single variables, the guarded-quantifier rank of a guarded formula is typically lower than its ordinary quantifier rank. Nevertheless, for all $k \geq 0$, the fragment \mathfrak{gFO}_k^n of \mathfrak{gFO}^n consisting of those sentences with guarded-quantifier rank at most k contains finitely many sentences up to logical equivalence.

This stratification can be modelled in terms of comonads $\mathbb{G}_{n,k}^{\mathfrak{g}}$ on $\mathbf{Struct}(\sigma)$, for all $n, k > 0$, as explained in [4, §VII]. Fixing n and letting k vary, we obtain an n -variable \mathfrak{g} -guarded resource-indexed arboreal cover of $\mathbf{Struct}(\sigma)$, with resource parameter k . The induced relations $\rightarrow_{n,k}^{\mathfrak{g}}$ and $\leftrightarrow_{n,k}^{\mathfrak{g}}$ on $\mathbf{Struct}(\sigma)$ coincide, respectively, with preservation of $\exists^+ \mathfrak{gFO}_k^n$ and equivalence in \mathfrak{gFO}_k^n . Thus, for any full subcategory \mathcal{D} of $\mathbf{Struct}(\sigma)$:

- $\mathcal{D} = \mathbf{Mod}(\varphi)$ for some $\varphi \in \mathfrak{gFO}_k^n$ if, and only if, \mathcal{D} is saturated under $\leftrightarrow_{n,k}^{\mathfrak{g}}$.
- $\mathcal{D} = \mathbf{Mod}(\psi)$ for some $\psi \in \exists^+ \mathfrak{gFO}_k^n$ if, and only if, \mathcal{D} is upwards closed with respect to $\rightarrow_{n,k}^{\mathfrak{g}}$.

As observed in [5, §6.1], the \mathfrak{g} -guarded resource-indexed arboreal cover of $\mathbf{Struct}(\sigma)$ satisfies the bisimilar companion property, and so does the n -variable \mathfrak{g} -guarded resource-indexed arboreal cover for all $n > 0$. Therefore, Proposition 5.2 implies the following *equivrank-variable homomorphism preservation theorem* for guarded logics:

Theorem 5.8. *Let \mathfrak{g} be a notion of guarding (either atom or loose). The following statements are equivalent for any \mathfrak{g} -guarded sentence φ in n variables of guarded-quantifier rank at most k in a relational vocabulary:*

- (1) φ is preserved under homomorphisms.
- (2) φ is logically equivalent to an existential positive \mathfrak{g} -guarded sentence in n variables of guarded-quantifier rank at most k .

5.2. Not-so-tame: extendability

A resource-indexed arboreal adjunction may fail to satisfy the bisimilar companion property, in which case Proposition 5.2 does not apply. This is the case e.g. for the Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction:


Example 5.9. The Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction between $\mathbf{Struct}(\sigma)$ and $\mathcal{R}^E(\sigma^I)$ does not have the bisimilar companion property. Suppose that $\sigma = \{R\}$ consists of a single binary relation symbol and let A be the σ -structure with underlying set $\{a, b\}$ satisfying $R^A = \{(a, b), (b, a)\}$. In view of Examples 3.18 and 3.21, it suffices to find $k > 0$ and a first-order sentence φ of quantifier rank $\leq k$ such that

$$A \models \varphi \quad \text{and} \quad H\mathbb{E}_k JA \not\models \varphi.$$

Let φ be the sentence $\forall x \forall y (x \neq y \Rightarrow xRy)$ of quantifier rank 2 stating that any two distinct elements are R -related. Then φ is satisfied by A but not by $H\mathbb{E}_k JA$, because the sequences $[a]$ and $[b]$ are not R -related in $H\mathbb{E}_k JA$. This shows that the bisimilar companion property fails for all $k \geq 2$.

When the bisimilar companion property fails, i.e. $a \not\leftrightarrow_k^{\mathcal{C}} G_k a$ for some $a \in \mathcal{C}$ and $k > 0$, we may attempt to force it by finding appropriate extensions a^* and $(G_k a)^*$ of a and $G_k a$, respectively, satisfying $a^* \leftrightarrow_k^{\mathcal{C}} (G_k a)^*$. This motivates the notion of k -*extendability* (see Definition 5.14 below), inspired by the work of Rossman [41] and its categorical interpretation in [1].

To start with, we introduce the following notations. Given objects a, b of a category \mathcal{A} , we write $a \rightarrow b$ to denote the existence of an arrow from a to b . Further, we write $a \rightleftarrows b$ to indicate that $a \rightarrow b$ and $b \rightarrow a$, i.e. a and b are *homomorphically equivalent*. This applies in particular to coslice categories. Recall that, for any $c \in \mathcal{A}$, the *coslice category* c/\mathcal{A} (also known as *under category*) has as objects the arrows in \mathcal{A} with domain c . For any two objects $m: c \rightarrow a$ and $n: c \rightarrow b$ of c/\mathcal{A} , an arrow $f: m \rightarrow n$ in c/\mathcal{A} is a morphism $f: a \rightarrow b$ in \mathcal{A} such that $f \circ m = n$. Hence, $m \rightleftarrows n$ in c/\mathcal{A} precisely when there are arrows $f: a \rightarrow b$ and $g: b \rightarrow a$ in \mathcal{A} satisfying $f \circ m = n$ and $g \circ n = m$. We shall represent this situation by means of the following diagram:



Remark 5.10. Note that $m \rightleftarrows n$ in c/\mathcal{A} whenever there is a section $f: a \rightarrow b$ in \mathcal{A} satisfying $f \circ m = n$. Just observe that the left inverse f^{-1} of f satisfies

$$f^{-1} \circ n = f^{-1} \circ f \circ m = m$$

and so $n \rightarrow m$. Further, $f \circ m = n$ entails $m \rightarrow n$. Hence, $m \rightleftarrows n$.

Now, let a be an object of an arboreal category \mathcal{C} and let $m: P \rightarrow a$ be a path embedding. As $\mathbb{S} a$ is a complete lattice by Lemma 3.10(b), the supremum $\bigvee \uparrow m$ in $\mathbb{S} a$ of all path embeddings above m exists and we shall denote it by

$$i_m: \mathcal{S}_m \rightarrow a.$$

Clearly, $m \leq i_m$ in $\mathbb{S} a$, and so there is a path embedding

$$m_\star: P \rightarrow \mathcal{S}_m$$

satisfying $i_m \circ m_\star = m$. (Note that i_m is well defined only up to isomorphism in the coslice category \mathcal{C}/a , but as usual we work with representatives for isomorphism classes.)

Remark 5.11. To provide an intuition for the previous definition, let us say for the sake of this remark that a path embedding $m: P \rightarrow a$ is “dense in a ” if all elements of $\mathbb{P} a$ are comparable with m . Then Lemma 5.12(a) below implies that $i_m: \mathcal{S}_m \rightarrow a$ is the largest \mathcal{M} -subobject of a in which m is dense.

Lemma 5.12. *The following statements hold for all path embeddings $m: P \rightarrow a$:*

- (a) $\mathbb{P} \mathcal{S}_m$ is isomorphic to the subtree of $\mathbb{P} a$ consisting of the elements that are comparable with m .
- (b) For all path embeddings $n: P \rightarrow b$ and arrows $f: a \rightarrow b$ such that $f \circ m = n$, there is a unique $g: \mathcal{S}_m \rightarrow \mathcal{S}_n$ making the following diagram commute.

$$\begin{array}{ccc}
 \mathcal{S}_m & \overset{g}{\dashrightarrow} & \mathcal{S}_n \\
 \downarrow i_m & & \downarrow i_n \\
 a & \xrightarrow{f} & b
 \end{array}$$

(c) For all path embeddings $n: P \rightarrow b$, if $m \rightarrow n$ then $m_\star \rightarrow n_\star$.

Proof. (a) The map $i_m \circ -: \mathbb{S} \mathcal{S}_m \rightarrow \mathbb{S} a$ is an order-embedding by Lemma 3.3, and so its restriction $\mathbb{P} \mathcal{S}_m \rightarrow \mathbb{P} a$ is an injective forest morphism. Hence, $\mathbb{P} \mathcal{S}_m$ is isomorphic to the subtree of $\mathbb{P} a$ consisting of those elements that factor through i_m , i.e. that are below $\bigvee \uparrow m$. By Lemma 3.10(c), an element of $\mathbb{P} a$ is below $\bigvee \uparrow m$ precisely when it is below some element of $\uparrow m$. In turn, the latter is equivalent to being comparable with m .

(b) Since \mathcal{S}_m is path-generated, it is the colimit of the canonical cocone C of path embeddings over the small diagram $\mathbb{P} \mathcal{S}_m$ which, by item (a), can be identified with the subdiagram of $\mathbb{P} a$ consisting of those elements comparable with m . As $\mathbb{P}(f \circ i_m)$ is monotone, it sends path embeddings comparable with m to path embeddings comparable with n , and so the cocone $\{f \circ i_m \circ p \mid p \in C\}$ factors through $i_n: \mathcal{S}_n \rightarrow b$. Hence, there is $g: \mathcal{S}_m \rightarrow \mathcal{S}_n$ such that $i_n \circ g = f \circ i_m$. Finally, note that if $g': \mathcal{S}_m \rightarrow \mathcal{S}_n$ satisfies $i_n \circ g' = f \circ i_m$ then we have $i_n \circ g = i_n \circ g'$, and so $g = g'$ because i_n is monic.


(c) Suppose there exists $f: a \rightarrow b$ such that $f \circ m = n$. By item (b), there is $g: \mathcal{S}_m \rightarrow \mathcal{S}_n$ such that $i_n \circ g = f \circ i_m$. Therefore,

$$i_n \circ g \circ m_\star = f \circ i_m \circ m_\star = f \circ m = n = i_n \circ n_\star$$

and so $g \circ m_\star = n_\star$ because i_n is a monomorphism. \square

Remark 5.13. Lemma 5.12(c) entails that $m_\star \rightarrow n_\star$ whenever $m_\star \rightarrow n$. Just observe that $(m_\star)_\star$ can be identified with m_\star .

Definition 5.14. Consider a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} , with adjunctions $L_k \dashv R_k$. An object a of \mathcal{E} is k -extendable if it satisfies the following property for all $e \in \mathcal{E}$: For all path embeddings $m: P \rightarrow R_k a$ and $n: P \rightarrow R_k e$ such that $m_\star \rightleftharpoons n_\star$ in the coslice category P/\mathcal{C}_k (see the leftmost diagram below),



if $n': Q \rightarrow R_k e$ is a path embedding satisfying $n \leq n'$ in $\mathbb{P}(R_k e)$, there is a path embedding $m': Q \rightarrow R_k a$ such that $m \leq m'$ in $\mathbb{P}(R_k a)$ and $m'_\star \rightleftharpoons n'_\star$ in Q/\mathcal{C}_k (as displayed in the rightmost diagram above).

We shall see in Proposition 5.16 below that, under appropriate assumptions, the k -extendability property allows us to upgrade the relation $\rightleftharpoons_k^{\mathcal{C}}$ to the finer relation $\leftrightarrow_k^{\mathcal{C}}$. For the next lemma, recall that a category is *locally finite* if there are finitely many arrows between any two of its objects.

Lemma 5.15. Let \mathcal{C} be an arboreal category whose full subcategory \mathcal{C}_p consisting of the paths is locally finite. If $f: P \rightarrow Q$ and $g: Q \rightarrow P$ are quotients between paths, then f and g are inverse to each other.

Proof. The set M of quotients $P \rightarrow P$ is a finite monoid with respect to composition, and it satisfies the right-cancellation law because every quotient is an epimorphism. Hence M is a group, and so $g \circ f \in M$ has

an inverse. It follows that $g \circ f$ is an embedding. Because there is at most one embedding between any two paths by Lemma 3.10(a), $g \circ f = \text{id}_P$. By symmetry, also $f \circ g = \text{id}_Q$. \square

Proposition 5.16. *Consider a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} such that \mathcal{C}_p^k is locally finite for all $k > 0$. For all k -extendable objects a, b of \mathcal{E} admitting a product, we have $a \rightleftarrows_k^{\mathcal{C}} b$ if and only if $a \leftrightarrow_k^{\mathcal{C}} b$.*

Proof. Fix an arbitrary $k > 0$ and recall that \mathcal{C}_k is an arboreal category. The “if” part of the statement follows from the inclusion $\leftrightarrow_k^{\mathcal{C}} \subseteq \rightleftarrows_k^{\mathcal{C}}$.

For the “only if” part suppose that $a \rightleftarrows_k^{\mathcal{C}} b$, i.e. $R_k a$ and $R_k b$ are homomorphically equivalent in \mathcal{C}_k . To improve readability, let $X := R_k a$ and $Y := R_k b$. We must prove that X and Y are bisimilar. As X and Y admit a product in \mathcal{C}_k , namely the image under R_k of the product of a and b in \mathcal{E} , by Theorem 3.14 it suffices to show that X and Y are back-and-forth equivalent.

Fix arbitrary morphisms $f: X \rightarrow Y$ and $g: Y \rightarrow X$, and let m and n denote generic elements of $\mathbb{P} X$ and $\mathbb{P} Y$, respectively. We claim that

$$\mathcal{B} := \{ \llbracket m, n \rrbracket \mid \exists s: \mathcal{S}_m \rightarrow \mathcal{S}_n, t: \mathcal{S}_n \rightarrow \mathcal{S}_m \text{ s.t. } \mathbb{P} s(m_\star) = n_\star \text{ and } \mathbb{P} t(n_\star) = m_\star \}$$

is a back-and-forth system between X and Y , i.e. it satisfies items (i)–(iii) in Definition 3.13. For item (i), let \perp_X, \perp_Y be the roots of $\mathbb{P} X$ and $\mathbb{P} Y$, respectively. Note that \mathcal{S}_{\perp_X} and \mathcal{S}_{\perp_Y} can be identified, respectively, with X and Y . As $\mathbb{P} f$ and $\mathbb{P} g$ are forest morphisms, $\mathbb{P} f(\perp_X) = \perp_Y$ and $\mathbb{P} g(\perp_Y) = \perp_X$. So, $\llbracket \perp_X, \perp_Y \rrbracket \in \mathcal{B}$.

For item (ii), suppose $\llbracket m, n \rrbracket \in \mathcal{B}$ and let $m' \in \mathbb{P} X$ satisfy $m \prec m'$. We seek $n' \in \mathbb{P} Y$ such that $n \prec n'$ and $\llbracket m', n' \rrbracket \in \mathcal{B}$. By assumption, there are arrows $s: \mathcal{S}_m \rightarrow \mathcal{S}_n$ and $t: \mathcal{S}_n \rightarrow \mathcal{S}_m$ such that $\mathbb{P} s(m_\star) = n_\star$ and $\mathbb{P} t(n_\star) = m_\star$. Writing $P := \text{dom}(m)$ and $P' := \text{dom}(n)$, we have the following diagrams

$$\begin{array}{ccc} P & \xrightarrow{e} \twoheadrightarrow \cdot & \xrightarrow{\varphi} P' \\ \downarrow m_\star & \Downarrow \mathbb{P} s(m_\star) & \downarrow n_\star \\ \mathcal{S}_m & \xrightarrow{s} \mathcal{S}_n & \xleftarrow{t} \mathcal{S}_m \end{array} \quad \begin{array}{ccc} P' & \xrightarrow{e'} \twoheadrightarrow \cdot & \xrightarrow{\psi} P \\ \downarrow n_\star & \Downarrow \mathbb{P} t(n_\star) & \downarrow m_\star \\ \mathcal{S}_n & \xrightarrow{t} \mathcal{S}_m & \xleftarrow{s} \mathcal{S}_n \end{array}$$

where φ and ψ are isomorphisms. By Lemma 5.15, $\varphi \circ e$ and $\psi \circ e'$ are inverse to each other, thus the left-hand diagram below commutes.

$$\begin{array}{ccc} & P & \\ m_\star \swarrow & & \searrow n_\star \circ \varphi \circ e \\ \mathcal{S}_m & \xrightarrow{s} & \mathcal{S}_n \\ & \xleftarrow{t} & \end{array} \quad \begin{array}{ccc} & Q & \\ m'_\star \swarrow & & \searrow n'_\star \\ \mathcal{S}_{m'} & \xrightarrow{s'} & \mathcal{S}_{n'} \\ & \xleftarrow{t'} & \end{array}$$

Let $Q := \text{dom}(m')$. Since b is k -extendable, there exist a path embedding $n': Q \rightarrow Y$ such that $n \leq n'$ in $\mathbb{P} Y$, and arrows $s': \mathcal{S}_{m'} \rightarrow \mathcal{S}_{n'}$ and $t': \mathcal{S}_{n'} \rightarrow \mathcal{S}_{m'}$ making the right-hand diagram above commute. It follows that $\llbracket m', n' \rrbracket \in \mathcal{B}$; just observe that $\mathbb{P} s'(m'_\star) = n'_\star$ because the composite $s' \circ m'_\star$ is an embedding, and similarly $\mathbb{P} t'(n'_\star) = m'_\star$.

It remains to show that $n \prec n'$. For any element x of a tree, denote by $\text{ht}(x)$ its height. As $n \leq n'$ in $\mathbb{P} Y$, it is enough to show that $\text{ht}(n') = \text{ht}(n) + 1$. Using the fact that forest morphisms preserve the height of points and $\mathbb{P} s'(m'_\star) = n'_\star$, we get

$$\text{ht}(n') = \text{ht}(n'_\star) = \text{ht}(m'_\star) = \text{ht}(m') = \text{ht}(m) + 1.$$

Item (iii) is proved in a similar way using the fact that a is k -extendable. \square

Definition 5.17. Consider a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} , and an object $a \in \mathcal{E}$. For all $k > 0$, a k -extendable cover of a is a section $a \rightarrow a^*$ in \mathcal{E} such that a^* is k -extendable.

Proposition 5.18. (HP) holds for any resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} satisfying the following properties for all $k > 0$:

- (i) \mathcal{C}_p^k is locally finite.
- (ii) \mathcal{E} has binary products and each of its objects admits a k -extendable cover.

Proof. Fix a full subcategory \mathcal{D} of \mathcal{E} saturated under $\leftrightarrow_k^{\mathcal{C}}$. For the non-trivial implication in (HP), assume \mathcal{D} is closed under morphisms and let $a, b \in \mathcal{E}$ satisfy $a \rightarrow_k^{\mathcal{C}} b$ and $a \in \mathcal{D}$.

If $G_k := L_k R_k$, then $a \rightarrow_k^{\mathcal{C}} b$ implies $G_k a \rightarrow b$. Let $s: a \rightarrow a^*$ and $t: G_k a \rightarrow (G_k a)^*$ be sections with a^* and $(G_k a)^*$ k -extendable objects. It follows from Lemma 3.20(a) that $a \rightleftarrows_k^{\mathcal{C}} a^*$ and $G_k a \rightleftarrows_k^{\mathcal{C}} (G_k a)^*$. By Lemma 3.20(b) we have $a \rightleftarrows_k^{\mathcal{C}} G_k a$ and so, by transitivity, $a^* \rightleftarrows_k^{\mathcal{C}} (G_k a)^*$. An application of Proposition 5.16 yields $a^* \leftrightarrow_k^{\mathcal{C}} (G_k a)^*$. We thus have the following diagram, where t^{-1} denotes the left inverse of t .

$$\begin{array}{ccc} a^* & \leftrightarrow_k^{\mathcal{C}} & (G_k a)^* \\ \uparrow s & & \downarrow t^{-1} \\ a & \rightleftarrows_k^{\mathcal{C}} & G_k a \longrightarrow b \end{array}$$

Since $a \in \mathcal{D}$, and \mathcal{D} is saturated under $\leftrightarrow_k^{\mathcal{C}}$ and closed under morphisms, all the objects in the diagram above sit in \mathcal{D} . In particular, $b \in \mathcal{D}$ and thus (HP) holds. \square

5.3. Relativising to full subcategories

Let \mathcal{E}' be a full subcategory of \mathcal{E} . We say that (HP) holds *relative to* \mathcal{E}' if the following condition is satisfied: For any full subcategory \mathcal{D} of \mathcal{E} saturated under $\leftrightarrow_k^{\mathcal{C}}$, $\mathcal{D} \cap \mathcal{E}'$ is closed under morphisms in \mathcal{E}' precisely when it is upwards closed in \mathcal{E}' with respect to $\rightarrow_k^{\mathcal{C}}$. Likewise for (HP#).

For the next proposition, observe that in order for a comonad G on \mathcal{E} to restrict to a full subcategory \mathcal{E}' of \mathcal{E} it is necessary and sufficient that $Ga \in \mathcal{E}'$ for all $a \in \mathcal{E}'$.

Proposition 5.19. Consider a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} , and let \mathcal{E}' be a full subcategory of \mathcal{E} such that the induced comonads $G_k := L_k R_k$ restrict to \mathcal{E}' . If the resource-indexed arboreal adjunction has the bisimilar companion property then (HP) holds relative to \mathcal{E}' . If it is idempotent then (HP#) holds relative to \mathcal{E}' .

Proof. The same, mutatis mutandis, as for Propositions 5.2 and 5.4, respectively. \square

Since the modal comonads \mathbb{M}_k restrict to finite pointed Kripke structures, the previous result yields a variant of Theorem 5.6 for finite structures:

Theorem 5.20. The following statements are equivalent for any graded modal formula φ of modal depth at most k in a modal vocabulary:

- (1) φ is preserved under homomorphisms between finite pointed Kripke structures.
- (2) φ is logically equivalent over finite pointed Kripke structures to an existential positive modal formula of modal depth at most k .

Similarly, since the guarded comonads $\mathbb{G}_{n,k}^{\mathfrak{g}}$ restrict to finite structures, we obtain the following variant of Theorem 5.8 for finite structures:

Theorem 5.21. *Let \mathfrak{g} be a notion of guarding (either atom or loose). The following statements are equivalent for any \mathfrak{g} -guarded sentence φ in n variables of guarded-quantifier rank at most k in a relational vocabulary:*

- (1) φ is preserved under homomorphisms between finite structures.
- (2) φ is logically equivalent over finite structures to an existential positive \mathfrak{g} -guarded sentence in n variables of guarded-quantifier rank at most k .

When the counits of the induced comonads $G_k := L_k R_k$ are componentwise surjective, the next observation, combined with Proposition 5.19, gives a criterion to relativise equi-resource homomorphism preservation theorems to subclasses of structures. Recall that a *negative formula* is one obtained from negated atomic formulas and $\vee, \wedge, \exists, \forall$.

Lemma 5.22. *Let G be a comonad on $\mathbf{Struct}(\sigma)$ whose counit is componentwise surjective and let T be a set of negative sentences. Then G restricts to $\mathbf{Mod}(T)$.*

Proof. If $\psi \in T$, its negation is logically equivalent to a positive sentence χ . Positive sentences are preserved under surjective homomorphisms and so, for all $A \in \mathbf{Struct}(\sigma)$, considering the component of the counit $GA \rightarrow A$ we obtain

$$GA \models \chi \implies A \models \chi.$$

I.e., G restricts to $\mathbf{Mod}(\psi)$. As $\mathbf{Mod}(T) = \bigcap_{\psi \in T} \mathbf{Mod}(\psi)$, the statement follows. \square

The counits of the guarded comonads $\mathbb{G}_{n,k}^{\mathfrak{g}}$ are componentwise surjective, thus the equirank-variable homomorphism preservation theorem for guarded logics and its finite variant (Theorems 5.8 and 5.21, respectively) admit a relativisation to any full subcategory of the form $\mathbf{Mod}(T)$ where T is a set of negative \mathfrak{g} -guarded sentences.

Relativisations to subclasses of structures can be obtained even in the absence of the bisimilar companion property; in that case, we need to ensure that k -extendable covers can be constructed within the subclass. We defer the statement of this result to Section 6.2 (see Corollary 6.11).

6. An axiomatic approach

In this section, we consider resource-indexed arboreal adjunctions between \mathcal{E} and \mathcal{C} that need not satisfy the bisimilar companion property. Using the notion of k -extendable cover introduced in Section 5.2, we identify sufficient conditions ensuring that Property (HP) is satisfied (see Corollary 6.10). We introduce first conditions (E1)–(E2) on \mathcal{E} in Section 6.1 and then, in Section 6.2, conditions (A1)–(A4) on the adjunctions. In Section 6.3, we derive a slight generalisation of the equirank homomorphism preservation theorem by showing that these conditions are satisfied by the Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction.

Remark 6.1. Let us point out that we cannot derive from this axiomatic approach an *equivariable* homomorphism preservation theorem, whereby the number of variables in a sentence is preserved, let alone an *equirank-variable* one. In fact, the corresponding (k -round) n -pebble comonads do not satisfy Property (A3) below. On the other hand, under the additional assumption that $k \leq n + 2$, where k is the quantifier rank and n is the number of variables in a sentence, an equirank-variable homomorphism preservation theorem

was proved by Paine in [38]. Also, our approach does not readily apply to *hybrid logic* because the hybrid comonads \mathbb{H}_k in [5] do not satisfy the path restriction Property (A4) (cf. Definition 6.8).

6.1. Axioms for the extensional category

We require that the category \mathcal{E} have the following properties:

- (E1) \mathcal{E} has all finite limits and small colimits.
 (E2) \mathcal{E} is equipped with a proper factorisation system such that:

- Embeddings are stable under pushouts along embeddings.
- Pushout squares of embeddings are also pullbacks.
- Pushout squares of embeddings are stable under pullbacks along embeddings.

Remark 6.2. Note that Property (E2) only involves one half of the factorisation system, namely the embeddings. In fact, it could be weakened to the requirement that \mathcal{E} admit a class of monomorphisms \mathcal{N} satisfying appropriate properties. When \mathcal{N} is the class of all monomorphisms, these are akin to the conditions for an *adhesive category*, cf. [31].

Example 6.3. If σ is a relational vocabulary, $\mathbf{Struct}(\sigma)$ satisfies (E1)–(E2). In fact, it is well known that $\mathbf{Struct}(\sigma)$ is complete and cocomplete, hence it satisfies (E1). For (E2), consider the proper factorisation system given by surjective homomorphisms and embeddings. Up to isomorphism, embeddings can be identified with inclusions of induced substructures. The pushout of a span of embeddings in $\mathbf{Struct}(\sigma)$ can be identified with a union of structures, and so embeddings are stable under pushouts along embeddings. The remaining two conditions in (E2) hold because they are satisfied in \mathbf{Set} by the class of monomorphisms, and the forgetful functor $\mathbf{Struct}(\sigma) \rightarrow \mathbf{Set}$ preserves and reflects pullback and pushout diagrams consisting of embeddings.

We note in passing that Properties (E1)–(E2) are stable under taking coslices:

Lemma 6.4. *If a category \mathcal{E} satisfies (E1)–(E2), then so does e/\mathcal{E} for all $e \in \mathcal{E}$.*

Proof. Fix an arbitrary object $e \in \mathcal{E}$. It is well known that limits and colimits in e/\mathcal{E} are inherited from \mathcal{E} , so e/\mathcal{E} satisfies (E1) because \mathcal{E} does.

By assumption, \mathcal{E} admits a proper factorisation system satisfying (E2). Let \mathcal{Q} and \mathcal{M} be the classes of arrows in e/\mathcal{E} whose underlying morphisms in \mathcal{E} are quotients and embeddings, respectively. It is folklore that $(\mathcal{Q}, \mathcal{M})$ is a weak factorisation system in e/\mathcal{E} . Moreover, this factorisation system is proper because the codomain functor $\text{cod}: e/\mathcal{E} \rightarrow \mathcal{E}$ is faithful. Recall that $\text{cod}: e/\mathcal{E} \rightarrow \mathcal{E}$ preserves pushouts, so embeddings in e/\mathcal{E} are stable under pushouts along embeddings because the corresponding property is satisfied in \mathcal{E} . The remaining two properties in (E2) follow by similar reasoning, using the fact that $\text{cod}: e/\mathcal{E} \rightarrow \mathcal{E}$ preserves and reflects limits and pushouts. \square

Example 6.5. It follows from Example 6.3 and Lemma 6.4 that, for all relational vocabularies σ , the category $\mathbf{Struct}_\bullet(\sigma)$ of pointed σ -structures satisfies (E1)–(E2).

6.2. Axioms for the resource-indexed adjunctions

We now assume that the extensional category \mathcal{E} satisfies (E1)–(E2), and proceed to introduce conditions on the resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} . To start with, consider an arbitrary adjunction $L \dashv R: \mathcal{E} \rightarrow \mathcal{C}$. As with any adjunction, there are hom-set bijections

$$\mathcal{E}(Lc, e) \longrightarrow \mathcal{C}(c, Ra), \quad f \mapsto f^b$$

natural in $c \in \mathcal{C}$ and $e \in \mathcal{E}$. Explicitly, f^b is defined as the composite

$$c \xrightarrow{\eta_c} RLc \xrightarrow{Rf} Re$$

where η is the unit of the adjunction $L \dashv R$. The inverse of the function $f \mapsto f^b$ sends $g \in \mathcal{C}(c, Re)$ to the morphism $g^\#$ given by the composition

$$Lc \xrightarrow{Lg} LRe \xrightarrow{\varepsilon_e} e,$$

where ε is the counit of the adjunction. Naturality of these bijections means that

$$(f_1 \circ f_2)^b = Rf_1 \circ f_2^b$$

for all morphisms $f_1: e \rightarrow e'$ and $f_2: Lc \rightarrow e$ in \mathcal{E} , and

$$(g_1 \circ g_2)^\# = g_1^\# \circ Lg_2$$

for all morphisms $g_1: c \rightarrow Re$ and $g_2: c' \rightarrow c$ in \mathcal{C} .

Next, we introduce the path restriction property for resource-indexed arboreal adjunctions. In a nutshell, this states that whenever $a \in \mathcal{E}$ embeds into the image under L_k of a path, a itself can be equipped with a path structure. Furthermore, these path structures can be chosen in a coherent fashion. We start with an auxiliary definition:

Definition 6.6. Consider a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} . A path $P \in \mathcal{C}_p^k$ is *smooth* if there exist $e \in \mathcal{E}$ and an embedding $P \mapsto R_k e$.

Remark 6.7. The motivation for considering smooth paths arises from the fact that, when considering a fresh binary relation symbol I modelling equality in the logic (cf. Example 3.18), the interpretation of I in these paths is always an equivalence relation.

Definition 6.8. A resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} has the *path restriction property* if, for all smooth paths $Q \in \mathcal{C}_p^k$ and embeddings $j: a \mapsto L_k Q$, there is a path $Q_a \in \mathcal{C}_p^k$ such that $L_k Q_a \cong a$ and the following conditions are satisfied:

- (i) For all path embeddings $!_{P,Q}: P \mapsto Q$ in \mathcal{C}_k and all commutative diagrams

$$\begin{array}{ccc} & \xrightarrow{L_k(!_{P,Q})} & \\ \downarrow & & \downarrow \\ L_k P & \xrightarrow{f} a \xrightarrow{j} & L_k Q \end{array}$$

there is an arrow $\ell: P \rightarrow Q_a$ such that $L_k \ell = f$.

- (ii) For all path embeddings $!_{P,Q}: P \rightarrow Q$ such that $L_k(!_{P,Q})$ is an embedding, the pullback of $L_k(!_{P,Q})$ along j is of the form $L_k\ell$ for some $\ell: Q_b \rightarrow Q_a$.

$$\begin{array}{ccc} b & \xrightarrow{\quad} & L_k P \\ \downarrow L_k \ell & \lrcorner & \downarrow L_k(!_{P,Q}) \\ a & \xrightarrow{j} & L_k Q \end{array}$$

Finally, recall that an object a of a category \mathcal{A} is *finitely presentable* if the associated hom-functor $\mathcal{A}(a, -): \mathcal{A} \rightarrow \mathbf{Set}$ preserves directed colimits.

With regards to the resource-indexed arboreal adjunction, we assume that the following properties are satisfied for all $k > 0$ and all paths $P \in \mathcal{C}_p^k$:

- (A1) The category \mathcal{C}_p^k is locally finite and has finitely many objects up to isomorphism.
(A2) $L_k P$ is finitely presentable in \mathcal{E} .
(A3) For all arrows $m: P \rightarrow R_k a$ in \mathcal{C}_k , if m is an embedding then so is $m^\# : L_k P \rightarrow a$. The converse holds whenever P is smooth.
(A4) The path restriction property is satisfied.

Theorem 6.9. *Consider a resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} satisfying (E1)–(E2) and (A1)–(A4). For all $a \in \mathcal{E}$ and all $k > 0$, there exists a k -extendable cover of a .*

The proof of the previous key fact is deferred to Section 7. Let us point out the following immediate consequence:

Corollary 6.10. *(HP) holds for all resource-indexed arboreal adjunctions satisfying (E1)–(E2) and (A1)–(A4).*

Proof. By Proposition 5.18 and Theorem 6.9. \square

We can also deduce the following relativisation result. Let us say that a full subcategory \mathcal{D} of a category \mathcal{A} is *closed (in \mathcal{A}) under co-retracts* provided that, whenever $A \in \mathcal{D}$ and there is a section $A \rightarrow B$ in \mathcal{A} , also $B \in \mathcal{D}$.

Corollary 6.11. *Let σ be a relational vocabulary and consider a resource-indexed arboreal adjunction between $\mathbf{Struct}(\sigma)$ and \mathcal{C} satisfying (A1)–(A4). The following hold:*

- (1) *If \mathcal{D} is a full subcategory of $\mathbf{Struct}(\sigma)$ closed under co-retracts such that each induced comonad $G_k := L_k R_k$ restricts to \mathcal{D} , then (HP) holds relative to \mathcal{D} .*
(2) *If the counits of the comonads G_k are componentwise surjective and T is a set of negative sentences in the vocabulary σ , then (HP) holds relative to $\mathbf{Mod}(T)$.*

Proof. The proof of item 1 is the same, mutatis mutandis, as for Proposition 5.18, using Theorem 6.9 and the fact that if \mathcal{D} is closed under co-retracts then k -extendable covers can be constructed within \mathcal{D} .

Item 2 is an immediate consequence of item 1. Just observe that the comonads G_k restrict to $\mathbf{Mod}(T)$ by Lemma 5.22, and $\mathbf{Mod}(T)$ is closed under co-retracts (cf. the proof of the aforementioned lemma). \square

6.3. The equirank homomorphism preservation theorem

As observed in Section 4.1, Rossman's equirank homomorphism preservation theorem is equivalent to Property (HP) for the Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction between $\mathbf{Struct}(\sigma)$ and

$\mathcal{R}^E(\sigma^I)$. In turn, by Corollary 6.10, to establish (HP) it suffices to show that the latter resource-indexed arboreal adjunction satisfies (E1)–(E2) and (A1)–(A4). By Example 6.3, the category $\mathbf{Struct}(\sigma)$ satisfies (E1)–(E2) when equipped with the (surjective homomorphisms, embeddings) factorisation system, so it remains to show that (A1)–(A4) hold. Before doing so, note that Corollary 6.11 yields the following slight generalisation of the equirank homomorphism preservation theorem (just observe that the counits of the induced comonads on $\mathbf{Struct}(\sigma)$ are componentwise surjective).

Theorem 6.12. *Let σ be a relational vocabulary and let \mathcal{D} be a full subcategory of $\mathbf{Struct}(\sigma)$ closed under co-retracts such that the comonads on $\mathbf{Struct}(\sigma)$ induced by the Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction restrict to \mathcal{D} . Then the equirank homomorphism preservation theorem holds relative to \mathcal{D} .*

In particular, the equirank homomorphism preservation theorem holds relative to $\mathbf{Mod}(T)$ whenever T is a set of negative sentences in the vocabulary σ .

Remark 6.13. A consequence of the first part of Theorem 6.12 is that the equirank homomorphism preservation theorem admits a relativisation to any class of structures that is *co-homomorphism closed*, i.e. downwards closed with respect to the homomorphism preorder on $\mathbf{Struct}(\sigma)$, a fact already pointed out by Rossman in [41, §7.1.2].

We proceed to verify conditions (A1)–(A4) for the adjunctions $L_k \dashv R_k$, where

$$L_k := HL_k^E \quad \text{and} \quad R_k := R_k^E J$$

with the notation of Example 3.18. Note that, since a first-order sentence contains only finitely many relation symbols, in order to deduce the equirank homomorphism preservation theorem, as well as Theorem 6.12 above, we can assume without loss of generality that the relational vocabulary σ is finite.

(A1) Recall from Example 3.7 that, for all $k > 0$, the paths in $\mathcal{R}_k^E(\sigma^I)$ are those forest-ordered σ^I -structures (A, \leq) such that the order is a chain of cardinality at most k . Thus, A has cardinality at most k . It follows at once that there are finitely many paths in $\mathcal{R}_k^E(\sigma^I)$ up to isomorphism, and at most one arrow between any two paths.

(A2) For any path $P = (A, \leq)$ in $\mathcal{R}_k^E(\sigma^I)$, $L_k P$ is the quotient of the σ -reduct of A with respect to the equivalence relation generated by I^A . As A is finite, so is $L_k P$. The finitely presentable objects in $\mathbf{Struct}(\sigma)$ are precisely the finite σ -structures (see e.g. [11, §5.1]), hence $L_k P$ is finitely presentable.

(A3) Consider an arrow $m: P \rightarrow R_k B$ in $\mathcal{R}_k^E(\sigma^I)$, with $P = (A, \leq)$ a path and B a σ -structure. Let $B' := J(B)$ be the σ^I -structure obtained from B by interpreting I as the identity relation. Then $R_k B$ is obtained by equipping $\mathbb{E}_k(B')$ with the prefix order. Consider the σ -homomorphism

$$L_k m = Hm: H(A) \rightarrow H\mathbb{E}_k(B').$$

For convenience of notation, given an element $x \in A$ we write $[x]$ for the corresponding element of $H(A)$, and likewise for elements of $\mathbb{E}_k(B')$. Then $m^\#: L_k P \rightarrow B$ is the composite of Hm with the homomorphism $H\mathbb{E}_k(B') \rightarrow B$ sending the equivalence class of an element of $\mathbb{E}_k(B')$ to the last element of any of its representatives. This map is well defined because, for any pair of sequences in $I^{\mathbb{E}_k(B')}$, their last elements coincide.

Suppose m is an embedding. If $m^\#([x]) = m^\#([y])$ then $(m(x), m(y))$ belongs to the equivalence relation generated by $I^{\mathbb{E}_k(B')}$, and so (x, y) belongs to the equivalence relation generated by I^A . It follows that $[x] = [y]$ in $H(A)$, and so $m^\#$ is injective. The same argument, mutatis mutandis, shows that $m^\#$ reflects the interpretation of the relation symbols, hence is an embedding.

Conversely, suppose that $m^\#$ is an embedding and P is smooth. Consider an embedding $n: P \hookrightarrow R_k C$ with $C \in \mathbf{Struct}(\sigma)$. Note that the restriction of $I^{\mathbb{E}_k(C')}$ to the image of n is an equivalence relation, hence

I^A is an equivalence relation. As any forest morphism whose domain is linearly ordered is injective, m is injective. So, it remains to show that m reflects the interpretation of the relation symbols. For all $x, y \in A$, if

$$(m(x), m(y)) \in I^{\mathbb{E}_k(B')}$$

then $m^\#([x]) = m^\#([y])$ and so $[x] = [y]$ because $m^\#$ is injective. That is, $(x, y) \in I^A$, showing that m reflects the interpretation of the relation I . Suppose now that S is a relation symbol different from I . For convenience of notation we shall assume that S has arity 2; the general case is a straightforward adaptation. For all $x, y \in A$, if

$$(m(x), m(y)) \in S^{\mathbb{E}_k(B')}$$

then $(m^\#([x]), m^\#([y])) \in S^B$ and so $([x], [y]) \in S^{H(A)}$ because $m^\#$ is an embedding. That is, there are $x', y' \in A$ such that $(x, x'), (y, y') \in I^A$ and $(x', y') \in S^A$. We claim that the following property holds, from which it follows that m is an embedding:

$$(x, x'), (y, y') \in I^A \quad \text{and} \quad (x', y') \in S^A \implies (x, y) \in S^A. \tag{*}$$

In turn, this is a consequence of the fact that n is an embedding and, in $\mathbb{E}_k(C')$,

$$(n(x), n(x')), (n(y), n(y')) \in I^{\mathbb{E}_k(C')} \quad \text{and} \quad (n(x'), n(y')) \in S^{\mathbb{E}_k(C')}$$

imply $(n(x), n(y)) \in S^{\mathbb{E}_k(C')}$.

(A4) Finally, we show that the path restriction property is satisfied. Let $Q = (B, \leq)$ be a smooth path in $\mathcal{R}_k^E(\sigma^I)$, and let $j: A \rightarrow H(B)$ be an embedding in $\mathbf{Struct}(\sigma)$. Without loss of generality, we can identify A with a substructure of $H(B)$, and j with the inclusion map. As observed above, since Q is smooth, I^B is an equivalence relation and property (*) is satisfied (with B in place of A). If $q_B: B \rightarrow H(B)$ is the canonical quotient map, let Q_A denote the substructure of B whose underlying set is

$$\{x \in B \mid q_B(x) \in A\}.$$

Then Q_A is a path in $\mathcal{R}_k^E(\sigma^I)$ when equipped with the restriction of the order on B and, using (*), we get $H(Q_A) \cong A$.

It follows from the definition of Q_A that item (i) in Definition 6.8 is satisfied. Just observe that any substructure of Q_A that is downwards closed in Q is also downwards closed in Q_A . With regards to item (ii), consider a path embedding $!_{P,Q}: P \rightarrow Q$ with $P = (C, \leq)$ and form the following pullback square in $\mathbf{Struct}(\sigma)$.

$$\begin{array}{ccc} D & \hookrightarrow & H(C) \\ \downarrow & \lrcorner & \downarrow L_k(!_{P,Q}) \\ A & \xrightarrow{j} & H(B) \end{array}$$

Identifying $H(C)$ with a substructure of $H(B)$, we can assume D is the intersection of A and $H(C)$. Because C is a substructure of B , it follows that Q_D is a substructure of Q_A . Moreover, because C is downwards closed in B , Q_D is downwards closed in Q_A . That is, there is an inclusion $Q_D \hookrightarrow Q_A$ whose image under L_k coincides with the pullback of $L_k(!_{P,Q})$ along j . Hence the path restriction property holds.

7. Proof of Theorem 6.9

For the remainder of this section, we fix an arbitrary resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} , with adjunctions $L_k \dashv R_k: \mathcal{E} \rightarrow \mathcal{C}_k$, satisfying (E1)–(E2) and (A1)–(A4).

7.1. Relative extendability

For all $k > 0$, we denote by $\overline{\mathcal{C}}_k$ the full subcategory of \mathcal{C}_k whose objects are colimits of finite diagrams of embeddings in \mathcal{C}_p^k . Further, we write $L_k[\overline{\mathcal{C}}_k]$ for the full subcategory of \mathcal{E} defined by the objects of the form $L_k c$ for $c \in \overline{\mathcal{C}}_k$.

Remark 7.1. It follows from (A1) that $\overline{\mathcal{C}}_k$ is equivalent to a finite category. Therefore $L_k[\overline{\mathcal{C}}_k]$ contains, up to isomorphism, only finitely many objects.

As we shall see in the following lemma, every path embedding $P \mapsto R_k a$ is homomorphically equivalent to one of the form $P \mapsto R_k \tilde{a}$ with $\tilde{a} \in L_k[\overline{\mathcal{C}}_k]$. Consequently, in the definition of k -extendable object (see Definition 5.14) we can assume without loss of generality that $e \in L_k[\overline{\mathcal{C}}_k]$. This observation, combined with Remark 7.1, will allow us to control the size of the diagrams featuring in the proof of Theorem 6.9.

Lemma 7.2. *For all path embeddings $m: P \mapsto R_k a$, there are $\tilde{a} \in L_k[\overline{\mathcal{C}}_k]$ and a path embedding $\tilde{m}: P \mapsto R_k \tilde{a}$ such that $m \rightleftharpoons \tilde{m}$ in P/\mathcal{C}_k .*

Proof. Fix an arbitrary path embedding $m: P \mapsto R_k a$. By (A1), there is a finite set of paths $\mathcal{P} = \{P_1, \dots, P_j\} \subseteq \mathcal{C}_p^k$ such that each path in \mathcal{C}_k is isomorphic to exactly one member of \mathcal{P} . We can assume without loss of generality that $P \in \mathcal{P}$.

For each path embedding $p \in \mathbb{P}(R_k a)$, denote by T_p the tree obtained by first considering the tree $\uparrow p \subseteq \mathbb{P}(R_k a)$ and then replacing each node q (which is an isomorphism class of a path embedding) with the unique path $P_i \in \mathcal{P}$ such that $P_i \cong \text{dom}(q)$. We assume that T_p is *reduced*, i.e. given any two nodes x and y of T_p that cover the same node, if the trees $\uparrow x$ and $\uparrow y$ are equal then $x = y$. (If T_p is not reduced, we can remove branches in the obvious manner to obtain a maximal reduced subtree T'_p .) We refer to T_p as the *type* of p ; note that this is a finite tree. In particular, if \perp is the root of $\mathbb{P}(R_k a)$, we get a finite tree T_\perp .

Now, for each node x of T_\perp , we shall define a path embedding m_x into $R_k a$ whose domain belongs to \mathcal{P} . The definition of m_x is by induction on the height of x . Suppose x has height 0, i.e. x is the root of T_\perp . Then $x = P_i$ for a unique $i \in \{1, \dots, j\}$. Define m_x as the restriction of m to P_i , i.e. the composition of $m: P \mapsto R_k a$ with the unique embedding $P_i \mapsto P$. Next, suppose m_z has been defined for all nodes z of height at most l , and let x be a node of height $l + 1$ labelled by some P_j . We distinguish two cases:

- If there is a node $y \geq x$ such that T_m coincides with the tree $\uparrow y \subseteq T_\perp$, then we let m_x be the restriction of m to P_j . Note that, in this case, the type of m_x coincides with the tree $\uparrow x \subseteq T_\perp$. Moreover, if z is the predecessor of x then m_z will also be an appropriate restriction of m , and thus m_x extends m_z .
- Otherwise, we let $m_x: P_j \mapsto R_k a$ be any path embedding such that:

- (i) The type of m_x coincides with the tree $\uparrow x \subseteq T_\perp$.
- (ii) m_x extends m_z , where z is the predecessor of x .

Note that such an embedding m_x exists because $x \in \uparrow z \subseteq T_\perp$ and, by inductive hypothesis, $\uparrow z$ coincides with the type of m_z .

The set

$$V := \{m_x \mid x \in T_\perp\}$$

is finite and contains m . We regard V as a cocone over a finite diagram D of paths and embeddings between them. Let $\tilde{a} := L_k(\text{colim } D)$ and note that $\tilde{a} \in L_k[\overline{\mathcal{C}}_k]$. The functor L_k preserves colimits because it is left adjoint, hence \tilde{a} is the colimit in \mathcal{E} of the diagram $L_k D$. The cocone $\{n^\# \mid n \in V\}$ with vertex a over $L_k D$ then factors through a unique morphism $f: \tilde{a} \rightarrow a$. By construction, $m: P \rightarrow R_k a$ factors through $R_k f$, and so there is $\tilde{m}: P \rightarrow \tilde{a}$ such that $m = R_k f \circ \tilde{m}$. Hence, $\tilde{m} \rightarrow m$.

Next, with the aim of showing that $m \rightarrow \tilde{m}$, we shall define a morphism $R_k a \rightarrow R_k \tilde{a}$. As $R_k a$ is path generated, it suffices to define a cocone

$$W = \{\varphi_p \mid p \in \mathbb{P}(R_k a)\}$$

with vertex $R_k \tilde{a}$ over the diagram of path embeddings into $R_k a$. Suppose $p \in \mathbb{P}(R_k a)$. We define the corresponding arrow φ_p by induction on the height of p :

- (i) If p is the root of $\mathbb{P}(R_k a)$, then it factors through $R_k f: R_k \tilde{a} \rightarrow R_k a$, and so it yields an embedding $\varphi_p: \text{dom}(p) \rightarrow R_k \tilde{a}$.
- (ii) Suppose that p has height $l+1$ and φ_q has been defined whenever q has height at most l . We distinguish two cases: if p factors through $R_k f: R_k \tilde{a} \rightarrow R_k a$, i.e. $p = R_k f \circ s_p$ for some embedding s_p , then we set $\varphi_p := s_p$. This is the case, in particular, when $p \leq m$ in $\mathbb{P}(R_k a)$. Clearly, if p extends q then φ_p extends φ_q .


Otherwise, let q be such that $p \succ q$. By inductive hypothesis, we can suppose that $R_k f \circ \varphi_q$ coincides with an embedding $m_x: P_j \rightarrow R_k a$ in V (up to an isomorphism $\text{dom}(q) \cong P_j$) whose type coincides with the tree $\uparrow x \subseteq T_\perp$. As q corresponds to a node y covering x labelled by some $P_h \cong \text{dom}(p)$, by definition of V there is an embedding $m_y: P_h \rightarrow R_k a$ in V such that m_y extends m_x , and the type m_y coincides with $\uparrow y$. Since m_y factors through $R_k f$, precomposing with the isomorphism $\text{dom}(p) \cong P_h$ we get an embedding $\varphi_p: \text{dom}(p) \rightarrow R_k \tilde{a}$. Observe that φ_p extends φ_q .

The compatibility condition for the cocone W states that φ_p extends φ_q whenever p extends q , which is ensured by the definition above. Thus, W induces a morphism $g: R_k a \rightarrow R_k \tilde{a}$ and, by construction, $g \circ m = \tilde{m}$. Hence, $m \rightarrow \tilde{m}$. \square

The construction of k -extendable covers is akin to that of ω -saturated elementary extensions in model theory, where one starts with a first-order structure M and constructs an elementary extension M_1 of M in which all types over (finite subsets of) M are realised, then an elementary extension M_2 of M_1 in which all types over M_1 are realised, and so forth. The union of the induced elementary chain of models yields the desired ω -saturated elementary extension of M .

In the same spirit, we introduce a notion of k -extendability relative to a homomorphism, which models the one-step construction just outlined.

Definition 7.3. Let $h: a \rightarrow b$ be an arrow in \mathcal{E} . We say that b is k -extendable relative to h if the following property is satisfied for all $e \in L_k[\overline{\mathcal{C}}_k]$: For all path embeddings $m: P \rightarrow R_k a$ and $n: P \rightarrow R_k e$ such that $m_\star \rightleftarrows n_\star$,



if $n' : Q \rightarrow R_k e$ is a path embedding such that $n \leq n'$ in $\mathbb{P}(R_k e)$, there is a path embedding $m' : Q \rightarrow R_k b$ such that the leftmost diagram below commutes and $m'_\star \rightleftharpoons n'_\star$.

$$\begin{array}{ccc}
 P & \xrightarrow{!} & Q \\
 m \downarrow & & \downarrow m' \\
 R_k a & \xrightarrow{R_k h} & R_k b
 \end{array}
 \quad
 \begin{array}{ccc}
 & Q & \\
 m'_\star \swarrow & & \searrow n'_\star \\
 \mathcal{S}_{m'} & \xrightarrow{\quad \quad \quad} & \mathcal{S}_{n'}
 \end{array}$$

Suppose that, given an object $a \in \mathcal{E}$, we are able to construct a section $s : a \rightarrow b$ such that b is k -extendable relative to s . Iterating this process countably many times, we obtain a k -extendable cover $a \rightarrow a^*$, thus settling Theorem 6.9. The main hurdle consists in establishing the following proposition; a proof is offered in Section 7.2.

Proposition 7.4. *For all $a \in \mathcal{E}$ and all $k > 0$ there is a section $s : a \rightarrow b$ such that b is k -extendable relative to s .*

We can finally prove Theorem 6.9:

Proof of Theorem 6.9. Let $a \in \mathcal{E}$. By applying Proposition 7.4 repeatedly, we obtain a chain of sections

$$a \xrightarrow{s_1} b_1 \xrightarrow{s_2} b_2 \xrightarrow{s_3} b_3 \xrightarrow{s_4} \dots$$

such that b_i is k -extendable relative to s_i , for all $i \geq 1$. Denote the previous diagram by D and let a^* be the colimit of D in \mathcal{E} , which exists by (E1). Let $h_i : b_i \rightarrow a^*$ be the colimit map with domain b_i , and $s : a \rightarrow a^*$ the one with domain a . As all the arrows in D are sections, so are the colimit maps; in particular, s is a section.

We claim that a^* is k -extendable. Suppose $m : P \rightarrow R_k(a^*)$ and $n : P \rightarrow R_k e$ are path embeddings such that $m_\star \rightleftharpoons n_\star$. By Lemma 7.2, we can assume without loss of generality that $e \in L_k[\overline{\mathcal{E}}_k]$.

By (A2), $L_k P$ is finitely presentable in \mathcal{E} and so $m^\# : L_k P \rightarrow a^*$ factors through one of the colimit maps. Assume without loss of generality that $m^\#$ factors through $h_j : b_j \rightarrow a^*$ for some $j \geq 1$, so there is an arrow $r : L_k P \rightarrow b_j$ satisfying $m^\# = h_j \circ r$. If $m_j := r^\flat$, it follows that $m = R_k h_j \circ m_j$. In particular, m_j is an embedding because so is m . Since $R_k h_j$ is a section, Remark 5.10 entails $m_j \rightleftharpoons m$, and so $(m_j)_\star \rightleftharpoons m_\star$ by Lemma 5.12(c). Because $m_\star \rightleftharpoons n_\star$, also $(m_j)_\star \rightleftharpoons n_\star$.

Now, let $n' : Q \rightarrow R_k e$ be any path embedding such that $n \leq n'$ in $\mathbb{P}(R_k e)$. Since b_{j+1} is k -extendable relative to s_{j+1} , there is a path embedding $m' : Q \rightarrow R_k b_{j+1}$ such that $m'_\star \rightleftharpoons n'_\star$ and the following diagram commutes.

$$\begin{array}{ccc}
 P & \xrightarrow{!} & Q \\
 m_j \downarrow & & \downarrow m' \\
 R_k b_j & \xrightarrow{R_k s_{j+1}} & R_k b_{j+1}
 \end{array}$$

It follows that

$$m = R_k h_j \circ m_j = R_k h_{j+1} \circ R_k s_{j+1} \circ m_j = R_k h_{j+1} \circ m' \circ !$$

and so $m'' := R_k h_{j+1} \circ m' : Q \rightarrow R_k(a^*)$ satisfies $m \leq m''$ in $\mathbb{P}(R_k(a^*))$. Again by Remark 5.10 and Lemma 5.12(c) we get $m''_\star \rightleftharpoons m'_\star$, and thus $m''_\star \rightleftharpoons n'_\star$. This shows that a^* is k -extendable. \square

7.2. Proof of Proposition 7.4

Fix an object a of \mathcal{E} and a positive integer k . To improve readability, we drop the subscript from L_k and R_k , and simply write L and R (but continue to denote by \mathcal{C}_k the arboreal category). We must find a section $s: a \rightarrow b$ such that b is k -extendable relative to s . Consider all pairs of path embeddings

$$(u: P \rightarrow Ra, v: P \rightarrow Re)$$

in \mathcal{C}_k such that $e \in L[\overline{\mathcal{C}}_k]$ and $Lv_\star \rightarrow u^\#$ in LP/\mathcal{E} .

Remark 7.5. Note that $Lv_\star \rightarrow u^\#$ entails that Lv_\star is an embedding. Just observe that $u^\#$ is an embedding by the first part of (A3).

Each such pair (u, v) induces a pushout square in \mathcal{E} as follows.

$$\begin{array}{ccc} LP & \xrightarrow{Lv_\star} & L\mathcal{S}_v \\ \downarrow u^\# & \lrcorner & \downarrow \lambda_{(u,v)} \\ a & \xrightarrow{\iota_{(u,v)}} & a +_{LP} L\mathcal{S}_v \end{array}$$

This pushout square exists by (E1) and consists entirely of embeddings by virtue of (E2).

Lemma 7.6. $\iota_{(u,v)}$ is a section.

Proof. Just observe that, since $Lv_\star \rightarrow u^\#$, there is $g: L\mathcal{S}_v \rightarrow a$ such that $g \circ Lv_\star = u^\#$. By the universal property of the pushout, there is an arrow $h: a +_{LP} L\mathcal{S}_v \rightarrow a$ such that $h \circ \iota_{(u,v)}$ is the identity of a . \square

We let D be the diagram in \mathcal{E} consisting of all the morphisms

$$\iota_{(u,v)}: a \rightarrow a +_{LP} L\mathcal{S}_v$$

as above. Because \mathcal{C}_k is locally finite and e varies among the objects of $L[\overline{\mathcal{C}}_k]$, choosing representatives for isomorphism classes in an appropriate way we can assume by Remark 7.1 that D is a small diagram.

By (E1), D admits a colimit $b := \text{colim } D$. In other words, b is obtained as a wide pushout in \mathcal{E} . Denote by $s: a \rightarrow b$ the colimit map with domain a , and by

$$t_{(u,v)}: a +_{LP} L\mathcal{S}_v \rightarrow b$$

the colimit map corresponding to the arrow $\iota_{(u,v)}$. As all arrows in D are sections by Lemma 7.6, so are the colimit maps. In particular, $s: a \rightarrow b$ is a section.

We claim that b is k -extendable relative to s , thus settling Proposition 7.4. Assume we are given path embeddings $m: P \rightarrow Ra$ and $n: P \rightarrow Re$, with $e \in L[\overline{\mathcal{C}}_k]$, such that $m_\star \rightleftharpoons n_\star$ as displayed in the leftmost diagram below.

$$\begin{array}{ccc} & P & \\ m_\star \swarrow & & \searrow n_\star \\ \mathcal{S}_m & \xrightarrow{f} & \mathcal{S}_n \\ & \xleftarrow{g} & \end{array} \qquad \begin{array}{ccc} & LP & \\ m^\# \swarrow & & \searrow Ln_\star \\ a & \xleftarrow{i_m^\# \circ Lg} & L\mathcal{S}_n \end{array}$$

If $i_m: \mathcal{S}_m \rightarrow Ra$ is the canonical embedding, we get a commutative triangle as on the right-hand side above. Just observe that

$$i_m^\# \circ Lg \circ Ln_\star = i_m^\# \circ Lm_\star = (i_m \circ m_\star)^\# = m^\#.$$

Hence $Ln_\star \rightarrow m^\#$. Let

$$\iota_{(m,n)}: a \rightarrow a +_{LP} L\mathcal{S}_n$$

be the corresponding arrow in the diagram D . To improve readability we shall write, respectively, ι , λ and t instead of $\iota_{(m,n)}$, $\lambda_{(m,n)}$ and $t_{(m,n)}$.

Let $n': Q \rightarrow Re$ be a path embedding with $n \leq n'$ in $\mathbb{P}(Re)$. We must exhibit a path embedding $m': Q \rightarrow Rb$ such that the leftmost square below commutes and $m'_\star \rightleftharpoons n'_\star$.

$$\begin{array}{ccc}
 P & \xrightarrow{!} & Q \\
 \downarrow m & & \downarrow m' \\
 Ra & \xrightarrow{Rs} & Rb
 \end{array}
 \qquad
 \begin{array}{ccc}
 & Q & \\
 m'_\star \swarrow & & \nwarrow n'_\star \\
 \mathcal{S}_{m'} & \xleftarrow{\quad} & \mathcal{S}_{n'}
 \end{array}
 \tag{2}$$

Note that, because $n \leq n'$, we get $\mathcal{S}_{n'} \leq \mathcal{S}_n$ in $\mathbb{S}Re$. Thus, n'_\star can be identified with a path embedding into \mathcal{S}_n . Consider the arrow

$$\xi := (\lambda \circ Ln'_\star)^\flat: Q \rightarrow R(a +_{LP} L\mathcal{S}_n).$$

Lemma 7.7. ξ is an embedding.

Proof. By the first part of (A3), $(n')^\#$ is an embedding. Then Ln' is an embedding since $(n')^\# = \varepsilon_e \circ Ln'$, and so is Ln'_\star . It follows that $\lambda \circ Ln'_\star$ is an embedding because it is a composition of embeddings, and ξ is an embedding by the second part of (A3). \square

Lemma 7.7, combined with the fact that Rt is a section (hence an embedding), entails that the composite $m' := Rt \circ \xi: Q \rightarrow Rb$ is an embedding. Moreover

$$\begin{aligned}
 m' \circ ! &= ((m' \circ !)^\#)^\flat = ((m')^\# \circ L!)^\flat = (t \circ \lambda \circ Ln'_\star \circ L!)^\flat \\
 &= (t \circ \lambda \circ Ln_\star)^\flat = (t \circ \iota \circ m^\#)^\flat = (s \circ m^\#)^\flat = Rs \circ m,
 \end{aligned}$$

showing that the leftmost diagram in equation (2) commutes. Since Rt is a section, we have $\xi \rightleftharpoons m'$ by Remark 5.10, and so $\xi_\star \rightleftharpoons m'_\star$ by Lemma 5.12(c). Therefore, in order to show that $m'_\star \rightleftharpoons n'_\star$ it suffices to prove that $\xi_\star \rightleftharpoons n'_\star$. We have

$$\lambda^\flat \circ n'_\star = ((\lambda^\flat \circ n'_\star)^\#)^\flat = (\lambda \circ Ln'_\star)^\flat = \xi$$

and thus $n'_\star \rightarrow \xi$. It follows from Remark 5.13 that $n'_\star \rightarrow \xi_\star$.

It remains to show that $\xi_\star \rightarrow n'_\star$; the proof of this fact will occupy us for the rest of this section. As \mathcal{C}_k is an arboreal category, \mathcal{S}_ξ is the colimit of its path embeddings. Thus, in order to define a morphism $\mathcal{S}_\xi \rightarrow \mathcal{S}_{n'}$ it suffices to define a compatible cocone with vertex $\mathcal{S}_{n'}$ over the diagram of path embeddings into \mathcal{S}_ξ . By Lemma 5.12(a), the path embeddings into \mathcal{S}_ξ can be identified with the path embeddings into $R(a +_{LP} L\mathcal{S}_n)$ that are comparable with ξ . For each such path embedding $q: Q' \rightarrow R(a +_{LP} L\mathcal{S}_n)$, we shall define an arrow $\zeta_q: Q' \rightarrow Re$ and prove that these form a compatible cocone. We will then deduce, using the induced mediating morphism $\mathcal{S}_\xi \rightarrow Re$, that $\xi_\star \rightarrow n'_\star$.

Fix an arbitrary path embedding $q: Q' \rightarrow R(a +_{LP} L\mathcal{S}_n)$ above ξ and consider the following diagram in \mathcal{E} , where the four vertical faces are pullbacks.

$$\begin{array}{ccccc}
 \overline{LP} & \xrightarrow{\nu} & \overline{L\mathcal{S}_n} & & \\
 \downarrow \mu_2 & \searrow \mu_1 & \downarrow \tau_1 & & \\
 LP & \xrightarrow{\sigma_1} & LQ' & & \\
 \downarrow \sigma_2 & \searrow \sigma_2 & \downarrow \tau_2 & & \\
 a & \xrightarrow{Ln_*} & L\mathcal{S}_n & & \\
 \downarrow m^\# & \searrow m^\# & \downarrow \lambda & & \\
 a & \xrightarrow{\iota} & a +_{LP} L\mathcal{S}_n & & \\
 & & \downarrow q^\# & &
 \end{array} \tag{3}$$

Note that the previous diagram consists entirely of embeddings because $q^\#$ is an embedding by the first part of (A3), and the pullback in \mathcal{E} of an embedding exists by (E1) and is again an embedding.

Because q is above ξ , there is an embedding $Q \rightarrow Q'$, and thus also an embedding $P \rightarrow Q'$. By the universal property of pullbacks, there are unique arrows

$$\vartheta: LP \rightarrow \bar{a} \quad \text{and} \quad \Delta: LP \rightarrow \overline{LP}$$

making the following diagrams commute.

$$\begin{array}{ccc}
 \begin{array}{ccc}
 \overbrace{LP \xrightarrow{\vartheta} \bar{a} \xrightarrow{\sigma_1} LQ'}^{L!} & & \\
 \downarrow m^\# & \searrow \sigma_2 & \downarrow q^\# \\
 a & \xrightarrow{\iota} & a +_{LP} L\mathcal{S}_n
 \end{array} & &
 \begin{array}{ccc}
 \overbrace{LP \xrightarrow{\Delta} \overline{LP} \xrightarrow{\mu_1} \bar{a}}^{\vartheta} & & \\
 \downarrow \text{id}_{LP} & \searrow \mu_2 & \downarrow \sigma_2 \\
 LP & \xrightarrow{m^\#} & a
 \end{array}
 \end{array}$$

Note in particular that μ_2 is a retraction whose right inverse is Δ . As μ_2 is also an embedding, it must be an isomorphism with (two-sided) inverse Δ .

By (A4) (more precisely, by item (i) in Definition 6.8) there is an arrow $w: P \rightarrow Q_{\bar{a}}$ between paths such that $Lw = \vartheta$. Hence, we can consider $\sigma_2^b: Q_{\bar{a}} \rightarrow Ra$. Note that

$$(\sigma_2^b \circ w)^\# = \sigma_2 \circ Lw = \sigma_2 \circ \vartheta = m^\#,$$

and so $\sigma_2^b \circ w = m$. In particular, w is an embedding. As $Q_{\bar{a}}$ is a path, $i_w: \mathcal{S}_w \rightarrow Q_{\bar{a}}$ can be identified with the identity $Q_{\bar{a}} \rightarrow Q_{\bar{a}}$. By Lemma 5.12(b), there is a unique arrow $\psi_q: Q_{\bar{a}} \rightarrow \mathcal{S}_m$ making the following diagram commute.

$$\begin{array}{ccc}
 Q_{\bar{a}} & \xrightarrow{\psi_q} & \mathcal{S}_m \\
 \downarrow \sigma_2^b & & \downarrow i_m \\
 & Ra &
 \end{array}$$

Lemma 7.8. *The following diagram commutes.*

$$\begin{array}{ccc}
 & \overline{LP} & \\
 \mu_1 \swarrow & & \searrow Ln_* \circ \mu_2 \\
 \bar{a} & \xrightarrow{L(f \circ \psi_q)} & L\mathcal{S}_n
 \end{array}$$

Proof. Note that

$$i_m \circ \psi_q \circ w = \sigma_2^b \circ w = m = i_m \circ m_*$$

and so $\psi_q \circ w = m_*$ since i_m is a monomorphism. Applying the functor L to the outer commutative diagram on the left-hand side below, we obtain the commutative diagram on the right-hand side.

$$\begin{array}{ccc}
 & P & \\
 w \swarrow & \downarrow m_* & \searrow n_* \\
 Q_{\bar{a}} & \mathcal{S}_m & \mathcal{S}_n \\
 \psi_q \searrow & \downarrow f & \\
 & \mathcal{S}_m & \mathcal{S}_n
 \end{array}
 \qquad
 \begin{array}{ccc}
 & LP & \\
 \vartheta \swarrow & & \searrow Ln_* \\
 \bar{a} & \xrightarrow{L(f \circ \psi_q)} & L\mathcal{S}_n
 \end{array}$$

Hence, precomposing with μ_2 we get

$$Ln_* \circ \mu_2 = L(f \circ \psi_q) \circ \vartheta \circ \mu_2 = L(f \circ \psi_q) \circ \mu_1. \quad \square$$

For convenience of notation, let us write

$$\tilde{\gamma}_q := L(f \circ \psi_q): \bar{a} \rightarrow L\mathcal{S}_n \quad \text{and} \quad \gamma_q := \varepsilon_e \circ Li_n \circ \tilde{\gamma}_q: \bar{a} \rightarrow e.$$

With this notation we have

$$\begin{aligned}
 \gamma_q \circ \mu_1 &= \varepsilon_e \circ Li_n \circ \tilde{\gamma}_q \circ \mu_1 \\
 &= \varepsilon_e \circ Li_n \circ Ln_* \circ \mu_2 && \text{Lemma 7.8} \\
 &= \varepsilon_e \circ Li_n \circ \tau_2 \circ \nu
 \end{aligned}$$

and so the leftmost diagram below commutes.

$$\begin{array}{ccc}
 \overline{LP} & \xrightarrow{\nu} & \overline{L\mathcal{S}_n} \\
 \mu_1 \downarrow & & \downarrow \varepsilon_e \circ Li_n \circ \tau_2 \\
 \bar{a} & \xrightarrow{\gamma_q} & e
 \end{array}
 \qquad
 \begin{array}{ccc}
 \overline{LP} & \xrightarrow{\nu} & \overline{L\mathcal{S}_n} \\
 \mu_1 \downarrow & & \downarrow \tau_1 \\
 \bar{a} & \xrightarrow{\sigma_1} & LQ'
 \end{array}$$

Now, note that by (E2) the top face of diagram (3), displayed in the rightmost diagram above, is a pushout in \mathcal{E} . By the universal property of pushouts, there is a unique $\delta_q: LQ' \rightarrow e$ satisfying

$$\delta_q \circ \sigma_1 = \gamma_q \quad \text{and} \quad \delta_q \circ \tau_1 = \varepsilon_e \circ Li_n \circ \tau_2.$$

Define $\zeta_q := (\delta_q)^b: Q' \rightarrow Re$ for all path embeddings $q: Q' \rightarrow R(a +_{LP} L\mathcal{S}_n)$ above ξ . Further, if q is below ξ , we let ζ_q be the obvious restriction of ζ_ξ .

Lemma 7.9. *The following family of arrows forms a compatible cocone over the diagram of path embeddings into \mathcal{S}_ξ :*

$$\{\zeta_q \mid q: Q' \rightarrow R(a +_{LP} L\mathcal{S}_n) \text{ is a path embedding comparable with } \xi\}.$$


Proof. Fix arbitrary path embeddings

$$q: Q' \rightarrow R(a +_{LP} L\mathcal{S}_n) \quad \text{and} \quad q': Q'' \rightarrow R(a +_{LP} L\mathcal{S}_n)$$

comparable with ξ . The compatibility condition for the cocone states that ζ_q extends $\zeta_{q'}$ whenever $q \geq q'$. It suffices to settle the case where $\xi \leq q \leq q'$. Also, it is enough to show that δ_q extends $\delta_{q'}$, i.e. $\delta_q \circ L! = \delta_{q'}$ where $!: Q'' \rightarrow Q'$ is the unique embedding. Just observe that $\delta_q \circ L! = \delta_{q'}$ entails

$$\zeta_q \circ ! = \delta_q^b \circ ! = ((\delta_q^b \circ !)^{\#})^b = (\delta_q \circ L!)^b = \delta_{q'}^b = \zeta_{q'}.$$

Consider the following diagram all whose vertical faces are pullbacks and note that by (A4), and more precisely by item (ii) in Definition 6.8, the pullback of $L!$ along σ_1 is of the form $L\ell$ for some arrow $\ell: Q_{\bar{a}} \rightarrow Q_{\bar{a}}$.



In view of the definition of $\delta_{q'}$ in terms of the universal property of pushouts, it suffices to show that $\delta_q \circ L!$ satisfies

$$(\delta_q \circ L!) \circ \bar{\sigma}_1 = \gamma_{q'} \quad \text{and} \quad (\delta_q \circ L!) \circ \bar{\tau}_1 = \varepsilon_e \circ Li_n \circ \tau_2 \circ \bar{\tau}_2.$$

The latter equation follows at once from the identity $\delta_q \circ \tau_1 = \varepsilon_e \circ Li_n \circ \tau_2$. With regards to the former, we have

$$(\delta_q \circ L!) \circ \bar{\sigma}_1 = \delta_q \circ \sigma_1 \circ L\ell = \gamma_q \circ L\ell.$$

Thus it suffices to show that $\gamma_q \circ L\ell = \gamma_{q'}$, and this clearly follows if we prove that $\tilde{\gamma}_q \circ L\ell = \tilde{\gamma}_{q'}$. Recall that $\psi_{q'}$ is the unique morphism such that the composite

$$Q_{\bar{a}} \xrightarrow{\psi_{q'}} \mathcal{S}_m \xrightarrow{i_m} Ra$$

coincides with $(\sigma_2 \circ L\ell)^b$. But

$$(\sigma_2 \circ L\ell)^b = ((\sigma_2^b \circ \ell)^\#)^b = \sigma_2^b \circ \ell,$$

so $i_m \circ \psi_{q'} \circ \ell = \sigma_2^b \circ \ell$ entails $\psi_{q'} = \psi_q \circ \ell$. Therefore,

$$\tilde{\gamma}_{q'} = L(f \circ \psi_{q'}) = L(f \circ \psi_q \circ \ell) = \tilde{\gamma}_q \circ L\ell. \quad \square$$

The previous lemma entails the existence of a unique morphism $h: \mathcal{S}_\xi \rightarrow Re$ satisfying $h \circ q = \zeta_q$ for all path embeddings q into $R(a +_{LP} L\mathcal{S}_n)$ that are comparable with ξ . In order to conclude that $\xi_\star \rightarrow n'_\star$ as desired, we prove the following useful property of the cocone consisting of the morphisms ζ_q .

Lemma 7.10. *Let $G := LR$ and consider the composite morphism*

$$RL\mathcal{S}_n \xrightarrow{RLi_n} RGe \xrightarrow{R\varepsilon_e} Re.$$

If $q = R\lambda \circ \alpha$ for some arrow $\alpha: Q' \rightarrow RL\mathcal{S}_n$, then $\zeta_q = R\varepsilon_e \circ RLi_n \circ \alpha$.

Proof. Suppose that $q = R\lambda \circ \alpha$ for some $\alpha: Q' \rightarrow RL\mathcal{S}_n$. We have

$$(R\varepsilon_e \circ RLi_n \circ \alpha)^\# = ((\varepsilon_e \circ Li_n \circ \alpha^\#)^\flat)^\# = \varepsilon_e \circ Li_n \circ \alpha^\#.$$

By the universal property of $\delta_q = \zeta_q^\#$, $\zeta_q = R\varepsilon_e \circ RLi_n \circ \alpha$ if, and only if,

$$(\varepsilon_e \circ Li_n \circ \alpha^\#) \circ \tau_1 = \varepsilon_e \circ Li_n \circ \tau_2 \tag{4}$$

and

$$(\varepsilon_e \circ Li_n \circ \alpha^\#) \circ \sigma_1 = \varepsilon_e \circ Li_n \circ \tilde{\gamma}_q. \tag{5}$$

Observe that $\alpha^\# \circ \tau_1 = \tau_2$ because

$$\lambda \circ \alpha^\# \circ \tau_1 = q^\# \circ \tau_1 = \lambda \circ \tau_2$$

and λ is a monomorphism. Thus, equation (4) holds. Further, note that

$$q^\# = (R\lambda \circ \alpha)^\# = ((\lambda \circ \alpha^\#)^\flat)^\# = \lambda \circ \alpha^\#$$

and so τ_1 in diagram (3) is an isomorphism. As pushout squares of embeddings in \mathcal{E} are also pullbacks by (E2), μ_1 in diagram (3) is also an isomorphism. Therefore,

$$\begin{aligned} \lambda \circ \alpha^\# \circ \sigma_1 &= q^\# \circ \sigma_1 = q^\# \circ \sigma_1 \circ \mu_1 \circ \mu_1^{-1} = \lambda \circ Ln_\star \circ \mu_2 \circ \mu_1^{-1} \\ &= \lambda \circ \tilde{\gamma}_q \circ \mu_1 \circ \mu_1^{-1} = \lambda \circ \tilde{\gamma}_q \end{aligned}$$

and so $\alpha^\# \circ \sigma_1 = \tilde{\gamma}_q$. Equation (5) then follows at once. \square

Since $\xi = R\lambda \circ (Ln'_\star)^\flat$, recalling that we identify n'_\star with a path embedding into \mathcal{S}_n , an application of the previous lemma with $\alpha := (Ln'_\star)^\flat$ yields

$$\zeta_\xi = R\varepsilon_e \circ RLi_n \circ (Ln'_\star)^\flat = (\varepsilon_e \circ Li_n \circ Ln'_\star)^\flat = (\varepsilon_e \circ Ln')^\flat = ((n')^\#)^\flat = n'.$$

In other words, $h \circ \xi_\star = n'$ and so $\xi_\star \rightarrow n'$. It follows from Remark 5.13 that $\xi_\star \rightarrow n'_\star$, thus concluding the proof of Proposition 7.4.

8. Relativisations to classes of well-behaved structures

In this section, we discuss a family of simple relativisation results which are, in a sense, oblivious to the “tame” versus “not-so-tame” divide of the landscape. Recall from Section 5 that, typically, the essence of a homomorphism preservation theorem is captured by the following configuration:

$$a \xleftrightarrow[k]{\mathcal{E}} G_k a \longrightarrow b.$$

The aim consists in proving that, if a belongs to a full subcategory \mathcal{D} of the extensional category \mathcal{E} , then so does b . The category \mathcal{D} is assumed to be closed in \mathcal{E} under morphisms, and saturated under the relation $\leftrightarrow_k^{\mathcal{E}}$.

The results presented in this paper are based on *upgrading arguments* (see [37] for an overview of this method): the relation $\xleftrightarrow[k]{\mathcal{E}}$ between a and $G_k a$ is upgraded to the finer relation $\leftrightarrow_k^{\mathcal{E}}$, possibly between

appropriate extensions a^* and $(G_k a)^*$. However, if a admits an arrow to $G_k a$, and thus also to b , then no upgrading is needed; in this case, the assumption that \mathcal{D} be saturated under $\leftrightarrow_k^{\mathcal{C}}$ is superfluous. The existence of a morphism $a \rightarrow G_k a$ is a strong requirement but, as we shall see below, it is satisfied in several cases of interest.

Fix a finite relational vocabulary σ and recall that every finite structure in $\mathbf{Struct}(\sigma)$ admits a *core*, denoted by $\text{core}(A)$, which is unique up to isomorphism. This can be characterised as the minimal substructure of A (with respect to set-theoretic inclusion) onto which A retracts, and any two homomorphically equivalent finite σ -structures have isomorphic cores; see e.g. [24]. For the next lemma, consider any comonadic adjunction

$$L \dashv R: \mathbf{Struct}(\sigma) \rightarrow \mathcal{C}.$$

We let $G := LR$ be the induced comonad, and denote by $L[\mathcal{C}]$ the full subcategory of $\mathbf{Struct}(\sigma)$ on those objects that are in the image of the functor L .

Lemma 8.1. *If $L[\mathcal{C}]$ is closed in $\mathbf{Struct}(\sigma)$ under induced substructures, then the following statements are equivalent for all $A \in \mathbf{Struct}(\sigma)$:*

- (1) *There exists a homomorphism $A \rightarrow GA$.*
- (2) *A is homomorphically equivalent to a structure $A' \in L[\mathcal{C}]$.*

Moreover, if A is finite, the previous conditions are equivalent to the following:

- (3) $\text{core}(A) \in L[\mathcal{C}]$.

Proof. $1 \Rightarrow 2$. Suppose there is a homomorphism $A \rightarrow GA$, and decompose it as

$$A \twoheadrightarrow A' \hookrightarrow GA$$

where A' is an induced substructure of GA . Composing the inclusion $A' \hookrightarrow GA$ with the counit $\varepsilon: GA \rightarrow A$, we get a homomorphism $A' \rightarrow A$. It follows that A and A' are homomorphically equivalent. Note that A' belongs to $L[\mathcal{C}]$ because it is an induced substructure of GA , which belongs to $L[\mathcal{C}]$.

$2 \Rightarrow 1$. If A is homomorphically equivalent to some $A' \in L[\mathcal{C}]$ then, by functoriality, GA is homomorphically equivalent to GA' . Since the adjunction $L \dashv R$ is comonadic and $A' \in L[\mathcal{C}]$, there is a coalgebra map $A' \rightarrow GA'$. We thus obtain a diagram

$$\begin{array}{ccc} A & \xrightleftharpoons{\quad} & A' \\ & & \downarrow \\ GA & \xrightleftharpoons{\quad} & GA' \end{array}$$

and so $A \rightarrow GA$.

$2 \Rightarrow 3$. Suppose that A is homomorphically equivalent to some $A' \in L[\mathcal{C}]$. Let B be the image of A under any homomorphism $A \rightarrow A'$ and note that B is finite, because so is A , and it is homomorphically equivalent to A . Further, B belongs to $L[\mathcal{C}]$ because the latter contains A' and is closed under induced substructures. Thus,

$$\text{core}(A) \cong \text{core}(B) \hookrightarrow B \in L[\mathcal{C}]$$

entails $\text{core}(A) \in L[\mathcal{C}]$.

3 \Rightarrow 2. Take $A' := \text{core}(A)$. \square

Note that the equivalence between items 1 and 2 in Lemma 8.1 holds, more generally, if we replace the category of σ -structures with any category \mathcal{E} equipped with a class \mathcal{M} of embeddings, assuming that $L[\mathcal{E}]$ is closed under taking \mathcal{M} -subobjects. For simplicity, we opted to state this lemma, as well as the following proposition, in the particular case where $\mathcal{E} = \mathbf{Struct}(\sigma)$.

Consider any resource-indexed arboreal cover of $\mathbf{Struct}(\sigma)$ by \mathcal{C} , with adjunctions $L_k \dashv R_k$, such that each category $L_k[\mathcal{C}_k]$ is closed in $\mathbf{Struct}(\sigma)$ under induced substructures. We obtain the following abstract homomorphism preservation theorem relative to a class of structures \mathcal{E}' :

Proposition 8.2. *Let \mathcal{E}' be a full subcategory of $\mathbf{Struct}(\sigma)$, and suppose there is k such that every object of \mathcal{E}' is homomorphically equivalent to an object of $L_k[\mathcal{C}_k]$. For any full subcategory \mathcal{D} of $\mathbf{Struct}(\sigma)$, $\mathcal{D} \cap \mathcal{E}'$ is closed under morphisms in \mathcal{E}' precisely when it is upwards closed in \mathcal{E}' with respect to $\rightarrow_k^{\mathcal{C}}$.*

Proof. For the non-trivial direction, assume $\mathcal{D} \cap \mathcal{E}'$ is closed under morphisms in \mathcal{E}' . We must show that, for all $A, B \in \mathcal{E}'$ such that $A \rightarrow_k^{\mathcal{C}} B$, if $A \in \mathcal{D}$ then also $B \in \mathcal{D}$. By Lemma 8.1, there is a homomorphism $A \rightarrow G_k A$, where $G_k := L_k R_k$. Further, there is a homomorphism $G_k A \rightarrow B$ because $A \rightarrow_k^{\mathcal{C}} B$. The composite yields a homomorphism $A \rightarrow B$, and so $B \in \mathcal{D}$. \square

Let us say that a class of σ -structures has *tree-depth at most k* if each of its members has tree-depth at most k . Consider the Ehrenfeucht-Fraïssé resource-indexed arboreal cover of $\mathbf{Struct}(\sigma)$ by $\mathcal{R}^E(\sigma)$ from Example 3.17. In view of [9, Theorem 6.2], the category $L_k[\mathcal{R}_k^E(\sigma)]$ can be identified with the class of σ -structures having tree-depth at most k . If A is a σ -structure of tree-depth at most k , then any induced substructure of A has also tree-depth at most k ; hence, $L_k[\mathcal{R}_k^E(\sigma)]$ is closed in $\mathbf{Struct}(\sigma)$ under induced substructures. Proposition 8.2 then implies that, on any class of σ -structures with tree-depth at most k , any first-order sentence that is preserved under homomorphisms is equivalent to an existential positive sentence of quantifier rank at most k .

Remark 8.3. It is enough to consider the Ehrenfeucht-Fraïssé resource-indexed arboreal cover of $\mathbf{Struct}(\sigma)$ by $\mathcal{R}^E(\sigma)$, rather than the resource-indexed arboreal adjunction between $\mathbf{Struct}(\sigma)$ and $\mathcal{R}^E(\sigma^I)$ in Example 3.21, because every existential positive sentence (with quantifier rank at most k) is equivalent to one (with quantifier rank at most k) that does not use the equality symbol.

In a similar fashion, we deduce the following result for classes of finite σ -structures:

Corollary 8.4. *On any class of finite σ -structures whose cores have tree-depth at most k , every first-order sentence that is preserved under homomorphisms is also definable by an existential positive sentence of quantifier rank at most k .*

Note that these are not equi-resource homomorphism preservation theorems, as the quantifier rank of the ensuing existential positive sentence is bounded by the tree-depth of the class of (cores of) structures under consideration, which may be larger than the quantifier rank of the first-order sentence we started with.

On the other hand, since we did not assume in Proposition 8.2 that \mathcal{D} is saturated under $\leftrightarrow_k^{\mathcal{C}}$, Corollary 8.4 holds, more generally, for all (not necessarily first-order definable) Boolean queries; see e.g. [32, §2.1] for the latter notion. Generalisations to non-Boolean queries are also available, by considering extensions of the Ehrenfeucht-Fraïssé resource-indexed arboreal cover to categories of n -pointed σ -structures.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgement

We are grateful to the anonymous referees for the numerous comments and suggestions that helped us to improve the presentation of our results.

References

- [1] S. Abramsky, Whither semantics?, *Theor. Comput. Sci.* 807 (2020) 3–14.
- [2] S. Abramsky, Structure and power: an emerging landscape, in: *Special Issue in Honor of the Boris Trakhtenbrot Centenary*, *Fundam. Inform.* 186 (1–4) (2022) 1–26.
- [3] S. Abramsky, A. Dawar, P. Wang, The pebbling comonad in finite model theory, in: *32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2017, pp. 1–12.
- [4] S. Abramsky, D. Marsden, Comonadic semantics for guarded fragments, in: *Proceedings of the 36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS*, 2021.
- [5] S. Abramsky, D. Marsden, Comonadic semantics for hybrid logic, in: *47th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, in: *Leibniz International Proceedings in Informatics*, vol. 241, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022, 7.
- [6] S. Abramsky, L. Reggio, Arboreal categories and resources, in: *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, *Leibniz International Proceedings in Informatics*, vol. 198, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, 115.
- [7] S. Abramsky, L. Reggio, Arboreal categories: an axiomatic theory of resources, *Log. Methods Comput. Sci.* 19 (2023) 14.
- [8] S. Abramsky, N. Shah, Relating structure and power: comonadic semantics for computational resources, in: *27th EACSL Annual Conference on Computer Science Logic (CSL)*, 2018, 2.
- [9] S. Abramsky, N. Shah, Relating structure and power: comonadic semantics for computational resources, *J. Log. Comput.* 31 (6) (2021) 1390–1428.
- [10] J. Adámek, H. Herrlich, G.E. Strecker, *Abstract and Concrete Categories. The Joy of Cats*, online edition, 2004.
- [11] J. Adámek, J. Rosický, *Locally Presentable and Accessible Categories*, *London Mathematical Society Lecture Note Series*, vol. 189, Cambridge University Press, Cambridge, 1994.
- [12] H. Andr eka, J. van Benthem, I. N emeti, Modal languages and bounded fragments of predicate logic, *J. Philos. Log.* 27 (3) (1998) 217–274.
- [13] J. Barwise, On Moschovakis closure ordinals, *J. Symb. Log.* 42 (2) (1977) 292–296.
- [14] T. Beke, J. Rosický, Abstract elementary classes and accessible categories, *Ann. Pure Appl. Log.* 163 (12) (2012) 2008–2017.
- [15] J. van Benthem, *Dynamic bits and pieces*, ILLC research report, University of Amsterdam, 1997.
- [16] P. Blackburn, M. De Rijke, Y. Venema, *Modal Logic*, *Cambridge Tracts in Theoretical Computer Science*, vol. 53, Cambridge University Press, 2002.
- [17] A.   Conghaile, A. Dawar, Game comonads & generalised quantifiers, in: *29th EACSL Annual Conference on Computer Science Logic, CSL*, in: *LIPICs*, vol. 183, Schloss Dagstuhl - Leibniz-Zentrum f ur Informatik, 2021, 16.
- [18] B.A. Davey, H.A. Priestley, *Introduction to Lattices and Order*, second ed., Cambridge University Press, New York, 2002.
- [19] A. Dawar, T. Jakl, L. Reggio, Lov asz-type theorems and game comonads, in: *Proceedings of the 36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS*, 2021.
- [20] M. de Rijke, A note on graded modal logic, *Stud. Log.* 64 (2) (2000) 271–283.
- [21] A. Ehrenfeucht, An application of games to the completeness problem for formalized theories, *Fundam. Math.* 49 (1960/1961) 129–141.
- [22] R. Fraiss e, Sur quelques classifications des syst emes de relations, *Publ. Sci. Univ. Alger., S er. A* 1 (1954) 35–182.
- [23] P.J. Freyd, G.M. Kelly, Categories of continuous functors, I, *J. Pure Appl. Algebra* 2 (3) (1972) 169–191.
- [24] P. Hell, J. Nešetřil, The core of a graph, *Discrete Math.* 109 (1) (1992) 117–126.
- [25] M. Hennessy, R. Milner, On observing nondeterminism and concurrency, in: *Automata, Languages and Programming*, Springer Berlin Heidelberg, Berlin, Heidelberg, 1980, pp. 299–309.
- [26] W. Hodges, *Model Theory*, *Encyclopedia of Mathematics and Its Applications*, vol. 42, Cambridge University Press, 1993.
- [27] N. Immerman, Upper and lower bounds for first order expressibility, *J. Comput. Syst. Sci.* 25 (1) (1982) 76–98.
- [28] T. Jakl, D. Marsden, N. Shah, A game comonadic account of Courcelle and Feferman-Vaught-Mostowski theorems, Preprint, available at <https://arxiv.org/abs/2205.05387>, 2022.

- [29] A. Joyal, I. Moerdijk, A completeness theorem for open maps, *Ann. Pure Appl. Log.* 70 (1) (1994) 51–86.
- [30] A. Joyal, M. Nielsen, G. Winskel, Bisimulation and open maps, in: *Proceedings of 8th Annual IEEE Symposium on Logic in Computer Science*, 1993, pp. 418–427.
- [31] S. Lack, P. Sobociński, Adhesive categories, in: *Foundations of Software Science and Computation Structures*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp. 273–288.
- [32] L. Libkin, *Elements of Finite Model Theory*, Texts in Theoretical Computer Science. An EATCS Series, Springer, 2004.
- [33] J. Łoś, On the extending of models. I, *Fundam. Math.* 42 (1955) 38–54.
- [34] R.C. Lyndon, Properties preserved under homomorphism, *Pac. J. Math.* 9 (1959) 143–154.
- [35] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., GTM, vol. 5, Springer-Verlag, New York, 1998.
- [36] J. Nešetřil, P. Ossona de Mendez, Tree-depth, subgraph coloring and homomorphism bounds, *Eur. J. Comb.* 27 (6) (2006) 1022–1041.
- [37] M. Otto, Model theoretic methods for fragments of FO and special classes of (finite) structures, in: J. Esparza, C. Michaux, C. Steinhorn (Eds.), *Finite and Algorithmic Model Theory*, in: *London Mathematical Society Lecture Note Series*, vol. 379, Cambridge University Press, 2011, pp. 271–341.
- [38] T. Paine, A pebbling comonad for finite rank and variable logic, and an application to the equirank-variable homomorphism preservation theorem, *Electron. Notes Theor. Comput. Sci.* 352 (2020) 191–209.
- [39] G.N. Raney, Completely distributive complete lattices, *Proc. Am. Math. Soc.* 3 (1952) 677–680.
- [40] E. Riehl, *Factorization systems*, Notes available at <https://emilyriehl.github.io/files/factorization.pdf>.
- [41] B. Rossman, Homomorphism preservation theorems, *J. ACM* 55 (2008) 15:1–15:53.
- [42] A. Tarski, Contributions to the theory of models. III, *Ned. Akad. Wet. Proc., Ser. A* 58 (1955) 56–64; *Indag. Math.* 17 (1955) 56–64.
- [43] K. Tent, M. Ziegler, *A Course in Model Theory*, Lecture Notes in Logic, vol. 40, Cambridge University Press, 2012.