Arboreal Covers Over Relational Structures

Nihil Shah

Linacre College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy in Computer Science

Hilary 2024

Acknowledgements

I thank my supervisor, Prof. Samson Abramsky, for the many years he spent mentoring and teaching me, both as a MSc. and DPhil student. His insights, support, patience and collaboration helped foster my growth from student to early researcher. Undoubtedly this thesis would not have been possible without his supervision.

I am also appreciative of Prof. Bartek Klin for taking over as my formal Oxford supervisor in the latter years of my studentship. I thank him for not only signing the necessary paperwork, but also the helpful discussions around category theory and distributive laws.

I thank the Oxford-DeepMind scholarship and the University of Oxford without whose financial support the research documented in this thesis could not have been carried out.

I thank Prof. Anuj Dawar for strengthening my intuitions in finite model theory and descriptive complexity. I am also grateful for his patience while I finished this thesis during the start of my postdoc with him. I have also greatly benefited from the advice of Prof. Dan Marsden in navigating academia and approaching problems from a categorical perspective.

As a participant in the game comonads community and the Structure meets Power workshops, I have had collaborations and discussions working with some outstanding scientists and mathematicians including Yoàv Montacute, Tomáš Jakl, Amin Karmalou, Luca Reggio, Adam O'Conghaile, Thomas Paine, Rue Soares Barbosa.

In my undergraduate institution, University of California, Berkeley, I had the great privilege of learning philosophy of mathematics from Prof. Paolo Mancosu and rhetoric of science from Prof. Nasser Zakariya. Their perspectives on the philosophical roots and cultural impact of mathematical logic have greatly influenced my view on these topics.

During the COVID-19 pandemic lockdown, I was fortunate to have the company of my friends and housemates at Ursula Hicks House of Linacre College. Special thanks go to Mark Quinlan and Mariko Nakahara for our weekend support bubble dinners.

I would also like to thank my many supportive friends and colleagues, including but not limited to: Tony Chen, Praagya Singh, Waleed Ovase, Sivert Aasnaess, Vincent Wang, Nick Ormond, Nick Hu, Cole Comfort, and Swaraj Dash. For always being my biggest fans, I thank my friends Natalia Reyes, Rodrigo Ochigame, and Hunan Rostomyan.

For their encouragement even when they struggle to understand my work, I thank my parents, Anju and Kamal, and my sister, Nisha.

Abstract

Model comparison games are an important tool in both finite and unrestricted model theory to prove when two structures satisfy the same sentences in a logic. Game comonads encode particular model comparison games as comonads on the categories of structures. This started with the work of Abramsky, Dawar, and Wang in 2017 on encoding pebble games for finite variable logics. Since this initial work, several examples of comonads have been engineered to capture a wide-range of logics relevant across different sub-fields of theoretical computer science. Each game comonad provides a categorical characterisation of an equivalence in a logic and its variants. The categorical constructions that involve these comonads have proved to be a nice tool for organising tacit connections between syntactic resources in logic, hierarchical approximations to constraint satisfaction and isomorphism, and well-known combinatorial parameters such as treewidth and tree-depth. Determining the commonalities shared amongst game comonads that enable them to have these features lead to the axiomatic formulation of arboreal category and arboreal cover by Abramsky and Reggio in 2021. Arboreal categories axiomatise the notion of category which has 'tree-shaped' objects and provide a native setting for dynamic notions like simulation, bisimulation, and resource-indexing. Arboreal covers are comonadic adjunctions to any 'extensional' category that allow application of these dynamic notions to the 'static' objects of the extensional category. Game comonads all arise as arboreal covers over categories of relational structures. The bridge drawn between finite model theory and category theory via arboreal covers over relational structures provides benefits to both fields. For finite model theorists, arboreal covers provide a setting for resource-sensitive constructions in model theory and allow categorical constructions to be fruitfully applied–yielding new perspectives to concrete problems. For categorical semanticists, tailoring and investigating abstract categorical constructions to this setting poses new challenges for these constructions leading to interesting discoveries. In this thesis, we exhibit a few contributions that serve to highlight these benefits:

- Relative-liftings of game comonads are used to capture logical equivalence for 'enriched' logics that contain symbols which carry specific interpretations, e.g. equality and connectivity predicates. In this thesis, we extend the Ehrenfeucht-Fraïssé comonad to multi-sorted structures and use relative lifting of this comonad to capture equivalence in monadic second-order logics.

- We explore a strengthening of the axioms of an arboreal category to arrive at the notion of 'linear arboreal category'. Using this notion, we show how to derive a linear variant of a game comonad. For logical equivalences that are captured by game comonads, this generalises the passage from branching-time behavioural relations, e.g. bisimulation to linear-time behavioural relations, e.g. trace equivalence. For combinatorial parameters captured by the coalgebras of a game comonad, this generalises the passage from tree-like parameters, e.g. treewidth to their path-like variant, e.g. pathwidth.

- We show that coKleisli laws of functorial operations over game comonads recover Feferman-Vaught-Mostowski theorems. Feferman-Vaught-Mostowski theorems are a basic tool used to show when a operation that combines structures preserves logical equivalence.

- One of the initial questions posed in the game comonads program was on the relationship between the pebbling comonad and the distribution-like quantum monad, defined by Abramsky, Barabosa, de Silva and Zapata in 2017, used to capture the notion of quantum homomorphism. By tailoring the definition of the Ehrenfeucht-Fraïssé comonad and quantum monad to commeasurable structures, we are able to produce a distributive law suggesting a path towards 'quantum finite model theory'. As a side effect of this research, we also discovered a no-go theorem demonstrating that there is no comonad-monad distributive law between non-linear polynomial comonads and monads which have a notion of 'uniform sampling'. The proof method and statement of the no-go theorem cover diverse range of (co)monads relevant to categorical semantics in general.

Contents

Chapter 1

Introduction

Category theorists look at structures as they “really are”, that is up to isomorphism. Model theorists, on the other hand, look at structures through the deceptive lens of expressiblity in logic. That is, model theorists can only distinguish structures by properties which are definable in a logic $\mathcal{J}$. Formally, given two mathematical structures of the same signature $\mathcal{A}$ and $\mathcal{B}$, model theory is concerned with equivalence under the relation

$$\mathcal{A} \equiv_{\mathcal{J}} \mathcal{B} := \forall\phi \in \mathcal{J}, \mathcal{A} \vDash \phi \Leftrightarrow \mathcal{B} \vDash \phi.$$

The equivalence $\equiv_{\mathcal{J}}$ is a key step in showing that a property P of structures is not expressible in the logic $\mathcal{J}$. To demonstrate inexpressibilty of P in $\mathcal{J}$ using the equivalence relation $\equiv_{\mathcal{J}}$, it suffices to construct structures $\mathcal{A}, \mathcal{B}$ with $\mathcal{A} \equiv_{\mathcal{J}} \mathcal{B}$ such that P is true in $\mathcal{A}$, P is false in $\mathcal{B}$. This is because if P were definable in $\mathcal{J}$, then there would be a sentence $\phi_P \in \mathcal{J}$, such that $\mathcal{A} \vDash \phi_P$ and $\mathcal{B} \nvDash \phi_P$. However, this cannot be the case, since by $\mathcal{A} \equiv_{\mathcal{J}} \mathcal{B}$, $\mathcal{A}$ and $\mathcal{B}$ satisfy the exact same sentences in $\mathcal{J}$.

Model comparison games, also known as Ehrenfeucht-Fraïssé games or bisimulation games, are a tool to prove, given $\mathcal{A}$ and $\mathcal{B}$, that $\mathcal{A} \equiv^{\mathcal{J}} \mathcal{B}$. An instance of a model comparison game from $\mathcal{A}$ to $\mathcal{B}$ has two players, Spoiler and Duplicator. Spoiler attempts to disprove that $\mathcal{A}$ and $\mathcal{B}$ are equivalent in $\equiv^{\mathcal{J}}$ by choosing elements from these structures in order to demonstrate a local difference between them. Duplicator, knowing Spoiler’s limitations, must respond with matching elements that satisfy the same relations as Spoiler’s choices. Formally, Duplicator wins the game by maintaining a partial isomorphism, consisting of the chosen elements, between the two structures. In the case of non-isomorphic structures $\mathcal{A}$ and $\mathcal{B}$, Duplicator’s task is difficult as she must deceive Spoiler by responding with elements that evade the difference between $\mathcal{A}$ and $\mathcal{B}$. Typically, these model comparison games are graded by some (usually, finite) ordinal resource parameter which constrains the local windows that Spoiler is allowed to play. Since equivalence in first order logic is the same as isomorphism over finite structures, grading by equivalence by some resource parameter is in fact essential for proving inexpressiblity results for logics over finite structures.

1.1 Model comparison game comonads

One such example of a model comparison game is the k-pebble game capturing equivalence in infinitary logic with at most k-many variables. This game was introduced by Immerman in [47] to study **PTIME** properties of finite structures. The k-pebble game has two variants: the one-sided variant introduced by Kolaitis and Vardi in [58] which has a tight connection with the k-consistency test, and the bijection variant introduced in [48] which has a tight connection with the k-Weisfeiler-Leman graph isomorphism test. Abramsky, Dawar, and Wang [5] provided these games with a categorical interpretation in terms of morphisms involving the k-pebble comonad $\mathbb{P}_k$. The coalgebras of this pebbling comonad are equivalent to tree decompositions of width $< k$ providing a coalgebraic definition of treewidth. Consequently, the pebbling comonad formalises the previously folkloric, tacit connections between finite variable logics, local consistency testing, Weisfeiler-Leman isomorphism testing, and bounded treewidth classes.

In my masters thesis [81] and in a joint paper with Abramsky [12], we adapted the pebbling comonad construction to obtain comonads which capture the k-round Ehrenfeucht-Fraïssé game characterising equivalence in first-order logic up to rank k, and the k-round bisimulation game characterising equivalence in modal logic up to depth k. These constructions were further developed upon in the journal version [13] of [12] by giving a sufficiently categorical definition of back-and-forth systems (leading to the later axiomatic developments in [9]), a theorem linking combinatorial parameters with positive-existential fragments, and a family of parameterised Chandra-Merlin style correspondences. Since then similar game comonads have been engineered to capture equivalence in guarded logics [7], hybrid logic [8], description logic [20], logics with restricted conjunction [68], and finite variable logics extended with all n-ary generalised quantifiers [25].

Every game comonad $\mathbb{C}_k$ is a comonad over a category of σ-structures $\mathsf{Struct}(\sigma)$ or pointed σ-structures $\mathsf{Struct}_\star(\sigma)$ for a relational signature σ. As model comparison games for logics have no formal definition, these game comonads are not constructed systematically via a general definition that applies to all cases. However, each game comonad $\mathbb{C}_k$ is associated to a logic $\mathcal{J}$ bounded by a resource parameter k. Every such game comonad exhibits the same pattern of results which we review below.

Morphism power theorem Given a logic $\mathcal{J}_k$ graded by some syntactic resource, e.g. number of variables, quantifier rank, modal depth, there is an associated back-and-forth model comparison game between structures characterising the equivalence $\equiv_{\mathcal{J}_k}$. Every such game can be restricted to a one-sided or 'forth' variant, where Spoiler and Duplicator do not switch structures, and Duplicator only has to preserve the relations between the chosen elements (rather than preserve and reflect the relations). This game turns out to characterise the relation

$$\mathcal{A} \Rrightarrow^{\exists^+\mathcal{J}_k} \mathcal{B} := \forall \phi \in \exists^+\mathcal{J}, \mathcal{A} \vDash \phi \Rightarrow \mathcal{B} \vDash \phi$$

for the positive-existential fragment $\exists^+\mathcal{J}_k$ of $\mathcal{J}_k$. The existence of a Duplicator winning strategy in a one-sided game from $\mathcal{A}$ to $\mathcal{B}$ is equivalent to the existence of a morphism in the coKleisli category of $\mathbb{C}_k$, i.e. a σ-morphism of type $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$.

Isomorphism power theorem Every model comparison game we consider has a bijection variant where Duplicator, prior to seeing Spoiler's choice, chooses a bijection between the universes of the two structures. After Spoiler makes his choice, Duplicator makes her choice by applying the bijection to Spoiler's chosen element. This bijection game characterises equivalence $\equiv^{\#\mathcal{J}_k}$ where $\#\mathcal{J}_k$ is $\mathcal{J}_k$ extended with quantifiers of the form $\exists_{\leq n}x$ and $\exists_{\geq n}x$. These quantifiers state that there exists at most or at least n, respectively, witnesses to x which satisfy the bound formula. The existence of a Duplicator winning strategy in a bijection game is equivalent to the existence of an isomorphism in the coKleisli category of $\mathbb{C}_k$.

Coalgebra characterisation theorem The coalgebras $\mathcal{A} \to \mathbb{C}_k\mathcal{A}$ for each Spoiler-Duplicator game comonad provides the data of a forest cover or tree-like decomposition of the σ-structure $\mathcal{A}$. Resource-indexing these comonads give coalgebraic definitions of common combinatorial parameters, e.g. tree-depth, treewidth, pathwidth, hypertree-width, used in graph theory and fixed parameter complexity.

Bisimulation power theorems The early research on these game comonads suffered from a 'Goldilocks problem'. In the first paper on the pebbling comonad [5], the preservation in the positive-existential fragment of k-variable logic and equivalence in k-variable counting language had clean characterisations in terms of the coKleisli category of $\mathbb{P}_k$. However, the full (unaltered) k-variable fragment was characterised in terms of an operation on the positional forms induced by coKleisli morphisms of $\mathbb{P}_k$. This positional form was insufficiently categorical and thus, not easily adapted to the game comonads explored in [12]. Nevertheless, the paper [12] did give an algebraic characterisation of these back-and-forth equivalences in terms of a fixpoint operation on the set of coKleisli morphisms between two structures. Finally, due to the insights of my co-author Samson Abramsky, this Goldilocks problem was resolved, in a satisfactorily categorical way, in the journal extension [13] of [12]. There it was shown that for all the Ehrenfeucht-Fraïssé, pebbling, and modal comonads, the corresponding logical equivalence $\equiv^{\mathcal{J}_k}$ could be captured by using a modified notion of Joyal, Nielsen, and Winskel's open map bisimulation [53] in the category of coalgebras for $\mathbb{C}_k$. This laid the foundation for the axioms of an arboreal category which isolated the categorical properties that enable this notion of bisimulation to be meaningful and applicable to all game comonads.

1.2 Arboreal categories

An arboreal category $\mathcal{C}$, first formulated by Abramsky and Reggio in [9], axiomatises the common properties shared amongst each game comonad's Eilenberg-Moore category. Intuitively, an arboreal category captures the notion of category with "tree-shaped" objects. An arboreal cover of category $\mathcal{E}$ by $\mathcal{C}$ is a comonadic adjunction $L \dashv R$, where $L \colon \mathcal{C} \to \mathcal{E}$, $R \colon \mathcal{E} \to \mathcal{C}$, and $\mathcal{E}$ is any category whose objects we view as 'static'. The comonadic property of the adjunction then allows us to view the object $R(X)$ in $\mathcal{C}$ as a tree-shaped cover of X which encodes a process for building up (parts of) of the object X in $\mathcal{E}$. Via this adjunction objects in $\mathcal{E}$ can be compared using a notion of bisimulation that naturally arises from the arboreal structure of $\mathcal{C}$. Each example model comparison game comonad is the induced comonad of an arboreal cover on a category of (pointed)-relational structures by an arboreal category $\mathcal{C}$.

Many concepts and results in both unrestricted and finite model theory can be formulated more generically via arboreal covers. Applying these generic formulations to different instances of arboreal covers has led to new results.

Lovász-type theorem A classical result by Lovász in his seminal paper [61] is that for all finite relational structures $\mathcal{A}$ and $\mathcal{B}$, $\mathcal{A}$ is is isomorphic to $\mathcal{B}$ if and only if, for all finite relational structures (over the same signature) $\mathcal{C}$, the number of homomorphisms from $\mathcal{C}$ to $\mathcal{A}$ is the same as the number of homomorphisms from $\mathcal{C}$ to $\mathcal{B}$. Lovász's theorem is a strengthening of the full and faithfulness of the Yoneda embedding for the category of finite σ-structures $\mathcal{E} = \mathsf{Struct}^f(\sigma)$. Namely, a corollary of Yoneda's lemma is that isomorphism between representable presheaves $\mathcal{E}(-, \mathcal{A}) \cong \mathcal{E}(-, \mathcal{B})$ in $\mathcal{E}$ is equivalent to $\mathcal{A} \cong \mathcal{B}$. Lovász had showed that, in the category $\mathcal{E}$, this characterisation of isomorphism could be weakened to pointwise isomorphism of presheaves, i.e. for all

$\mathcal{C} \in \mathbf{Ob}(\mathcal{E})$, there is a bijection between the homsets $\mathcal{E}(\mathcal{C}, \mathcal{A}) \cong \mathcal{E}(\mathcal{C}, \mathcal{B})$. Following this result, weaker notions of equivalence were discovered to have analogous homomorphism counting characterisations via restricting the class of test structures $\mathcal{C}$. For instance, Dvořák in [33] showed that counting homomorphisms from finite structures of treewidth $< k$ yields equivalence in k-variable logic with counting quantifiers. Grohe in [43] showed that counting homomorphisms from finite structures of tree-depth $\leq k$ yields equivalence in counting logic with quantifier rank $\leq k$.

Dawar, Jakl and Reggio in [30] introduced a generalisation of Lovász-type theorems that applies to a class of categories which includes arboreal categories. Applying this result to specific arboreal categories recover Dvořák and Grohe's results. This result was applied in joint work [68] with Yoàv Montacute to obtain a new Lovász-type theorem demonstrating that counting homomorphisms from finite structures of pathwidth $< k$ yields equivalence in a restricted conjunction, walk counting fragment of k-variable counting logic.

Rossman-style equiresource preservation theorems Preservation theorems characterise the syntactic form of sentence by the relation on models that preserve the truth of the sentence. For instance, a classical result in (unrestricted) model theory attributed to Łoś, Lyndon, and Tarski is the homomorphism preservation theorem (HPT) [85, 62, 82] stating that a first-order sentence is preserved under homomorphisms if and only if, the sentence is equivalent to a positive-existential sentence. Benjamin Rossman proved that the HPT remains true when restricted to finite structures [79]. Rossman's result is in sharp contrast to many other preservation theorems in classical model theory which fail over finite structures. As a warm-up for this finite HPT, Rossman proved an equirank HPT demonstrating that a first-order sentence of quantifier rank $\leq k$ is preserved under homomorphisms if and only if, the sentence is equivalent to a positive-existential sentence of quantifier rank $\leq k$ [79]. Beginning with work from Thomas Paine in [73] and culminating in Samson Abramsky and Luca Reggio's paper [11], the proof of Rossman's equirank HPT was reformulated in terms of arboreal covers. Generalising this proof yields similar Rossman-style equiresource theorems covering a wider range of resource-bounded logics with preservation under different semantic relations, e.g. modal depth and simulation preservation, guarded depth and guarded simulation preservation. These Rossman-style equiresource theorems point towards arboreal covers providing a framework for resource-sensitive model theory without compactness. Given the inherent non-constructive nature of compactness, this is a necessary restriction to address the needs of finite and algorithmic model theory.

1.3 Contributions of this thesis

Game comonads and their arboreal cover formulation have opened up the possibility of applying categorical concepts to not only finite model theory, but other areas of logic in computer science. The results of this thesis demonstrate how a categorical notion, applied in the arboreal setting, unifies similar constructions or theorems within logic in computer science:

(C1) *Second-order logics via relative liftings of comonads.* Each structure $\mathcal{A}$ of a language in signature σ specifies an interpretation $R^{\mathcal{A}}$ for each relational symbol $R \in \sigma$. These interpretations of atomic symbols determine the base case in the semantics of the language. However, often in model theory it is useful to have structures where some of the atomic symbols are restricted to a fixed interpretation. That is, these interpreted symbols are defined "in the same way" across different structures. For instance, many logics contain an equality symbol which is interpreted, for every structure $\mathcal{A}$, as the equality of elements in the underlying universe A. Another example is first-order logic with a connectivity predicate **conn** as in [80] defined as the transitive closure of the edge relation in the Gaifman Graph $\mathcal{G}(\mathcal{A})$ of $\mathcal{A}$. As connectivity of $\mathcal{G}(\mathcal{A})$ is inexpressible in first-order logic over finite structures, this logic is genuinely more expressive than 'unenriched' first order logic. Through relative-liftings of game comonad, we can capture equivalence in these enriched logics. First, we illustrate the technique of using relative liftings on Ehrenfeucht-Fraïssé comonads in order to capture equivalence in logics with equality and connectivity predicates. Second, we adapt the definition of the Ehrenfeucht-Fraïssé comonad to multi-sorted structures. Finally, we use a relative lifting of this comonad to capture equivalence in monadic second-order logic.

(C2) *Linear time equivalences via an adjunction.* By strengthening the notion of arboreal cover to linear arboreal cover, we unify three disparate ways in which the concept of 'linearity' appears within logic for computer science. The first such thread is in the study of linear-time variants, e.g. trace equivalence and trace inclusion, of branching-time equivalences between processes, e.g. bisimulation and simulation [83, 84]. The second such thread is in the study of fragments of infinitary logic related to linear Datalog [29]. The final thread is in the study of linear variants of combinatorial parameters that are defined in terms of tree-like decompositions. For instance, pathwidth is the linear variant of treewidth. We begin by adding two axioms to the notion of arboreal category to arrive at the notion of linear arboreal category allowing us to define a linear arboreal cover. Next, we show that every arboreal category $\mathcal{C}$ whose path objects satisfy a strong connected condition has a linear arboreal subcategory $\mathcal{C}^L$ containing the "linear tree-shaped objects" and "leaf-preserving" morphisms. We also show how to construct, from an arboreal cover of $\mathcal{E}$ by $\mathcal{C}$, a linear arboreal cover of $\mathcal{E}$ by $\mathcal{C}^L$. This provides a general method for deriving linear variants of game comonads. Finally, we exhibit linear variants of the pebbling and modal comonads, demonstrating how their relationship with linear variants of combinatorial parameters and behavioural equivalences.

(C3) *Feferman-Vaught-Mostowski theorems via CoKleisli laws.* Though the equivalence relation $\equiv_{\mathcal{J}}$ for many logics $\mathcal{J}$ is characterised using a corresponding Spoiler-Duplicator game, in practice using these games directly can be cumbersome. The composition method is a tool that finite model theorists use to mitigate this difficulty by enabling modular reasoning [59]. This method involves building structures by inductively applying operations on simple component structures. Feferman-Vaught-Mostowski (FVM) theorems enable this method by demonstrating an operation H preserves logical equivalence $\equiv_{\mathcal{J}}$. More generally, we consider FVM theorems of the form:

$$\text{if } \forall i \in [n] \text{ and } \mathcal{A}_i, \mathcal{B}_i \in \mathbf{Ob}(\mathcal{C}_i), \mathcal{A}_i \equiv^{\mathcal{J}_i} \mathcal{B}_i, \text{ then } H(\mathcal{A}_1, \dots, \mathcal{A}_n) \equiv^{\mathcal{J}} H(\mathcal{B}_1, \dots, \mathcal{B}_n)$$

for logics $\mathcal{J}_1, \dots, \mathcal{J}_n, \mathcal{J}$ and a operation $H\colon \mathcal{C}_1 \times \cdots \times \mathcal{C}_n \to \mathcal{C}$ on categories of (pointed) σ-structures $\mathcal{C}_1, \dots, \mathcal{C}_n, \mathcal{C}$. First, we show how discrete transformations can be re-interpreted as witnessing FVM theorems for positive-existential fragments. Next, we show how coKleisli laws can be re-interpreted as FVM theorems for logics extended with counting quantifiers. We exhibit translations between logics and FVM theorems for clique-width operations as examples.

(C4) *Towards quantum finite model theory via a mixed distributive law.* Parallel to the discovery of the pebbling comonad $\mathbb{P}_k$ was the discovery of the quantum monad $\mathbf{Q}_d$ [4]. The quantum monad characterises a notion of quantum homomorphism [65], originally defined in terms of perfect strategies of non-local games, between relational structures. As alluded to in [4] and conjectured explicitly in [24], a mixed distributive law of type $\mathbb{P}_k\mathbf{Q}_d \to \mathbf{Q}_d\mathbb{P}_k$ relating these two constructions may potentially give rise to "quantum finite model theory". More generally, we could ask for what game comonads $\mathbb{C}_k$ does there exist a distributive law of type $\mathbb{C}_k\mathbf{Q}_d \to \mathbf{Q}_d\mathbb{C}_k$? We address this question by first moving to a category of 'commeasurable' relational structures $\mathsf{Struct}^{\odot}(\sigma)$. We then produce an Ehrenfeucht-Fraïssé game comonad $\mathbb{E}_k^{\odot}$ and quantum monad $\mathbf{Q}_d^{\odot}$ tailored to this setting. Finally, we produce a distributive law of type $\mathbb{E}_k^{\odot}\mathbf{Q}_d^{\odot} \to \mathbf{Q}_d^{\odot}\mathbb{E}_k^{\odot}$.

Another aspect of the study of arboreal categories has been to motivate research into concrete category theory. As part of the search for a distributive law of a game comonad over the quantum monad, we made the following contribution to the general study of mixed comonad-monad distributive laws in concrete category theory:

(C5) *No-go theorems for directed containers over uniform choice.* In a sense, the powerset monad and distribution monads are crude version of the quantum monad operating at the level of 'possibilities' or 'probabilities' rather than assigning measurement outcomes to these possibilities. Similarly, most of the game comonads coincide (or are minor variants of) the prefix list comonad over **Set**. Therefore, in order to address the question of finding a distributive law of $\mathbb{P}_k$ over $\mathbf{Q}_d$, we first explored the question of whether the prefix list comonad distributes over the powerset monad or distribution monad. Using a Plotkin-style argument, we show that there can be no mixed distributive law of the prefix list comonad over the powerset monad. We then generalise this result in two ways. First, we generalise the comonad and show that every directed container which is not a coreader comonad has no distributive law over the powerset monad. Second, we generalise the monad and show that this same class of comonads has no distributive law over any monad which has a meaningful notion of 'uniform distribution'. This class of uniform choice monads is axiomatised using semi-algebraic conditions. We then determine necessary and sufficient conditions on semirings $\mathcal{S}$ for the multiset $\mathbf{M}_{\mathcal{S}}$ and distribution monads $\mathbf{D}_{\mathcal{S}}$ on that semiring to be uniform choice. We also exhibit the filter monad as a uniform choice monad.

1.4 Outline and relationship to other work

In the following, I contributed underlying ideas, candidate constructions, proofs, and/or exposition to the portions of joint work that appear in this thesis.

Chapter 2 sets out background material necessary for understanding the rest of the thesis.

Chapter 3 is a re-presentation of the content in Samson Abramsky and Luca Reggio's papers [9, 10]. The axioms of arboreal categories were motivated by the concrete examples of game comonads from Samson Abramsky, Anuj Dawar, and Penming Wang's paper [5], and my joint work with Samson Abramsky in the paper [12] and its extended journal version [13]. These examples, presented as arboreal covers, are included in Sections 3.2.2, 3.2.3, and 3.2.4. Section 3.3 of Chapter 3 is about using relative liftings of game comonads to capture equivalence in logics with interpreted symbols, e.g. equality and connectivity predicates. This work generalises the use of relative comonads to formalise I-morphisms which appeared in [13]. The concept of I-morphisms first appeared without this formalisation in [5] and was used to capture finite variable logic *with equality.* Section 3.3.2 is on the contribution (C1) where we construct a multi-sorted Ehrenfeucht-Fraïssé comonad and use a relative-lifting of this comonad to capture equivalence in monadic second-order logic. This construction and its application to capturing equivalence in monadic-second order logic is joint work with Tomáš Jakl and Dan Marsden appearing in the preprint [50].

Chapter 4 includes the details of contribution (C2) about linear arboreal subcategories. The content of this chapter is from joint work with Samson Abramsky and Yoàv Montacute in the conference paper [69]. This was motivated by the example of 'linearising' the pebbling comonad to obtain the pebble-relation comonad. Thus, the example of the linear pebbling arboreal cover in Section 4.2.2 includes content from my joint work with Yoàv Montacute in the papers [68, 70] on the pebble-relation comonad. The example of the linear modal arboreal cover in Section 4.2.3 is also from the preprint [69].

Chapter 5 is on the contribution (C3) demonstrating how coKleisli laws over game comonads correspond to FVM theorems in logic. Section 5.2 explains how discrete transformations and coKleisli law can be interpreted as abstract FVM theorems for positive-existential and counting logics. Section 5.3 illustrates this interpretation by semantically recovering translations between logics using game comonads, and proving FVM theorems about clique-width operations. This content of this chapter derives from joint work with Tomáš Jakl and Dan Marsden in the paper [52].

Chapter 6 is on the contributions (C5) and (C4) pertaining to comonad-monad distributive laws. Section 6.2 is on the contribution (C5) proving that there is no comonad-monad distributive law for a class of comonads induced by directed containers and a class of monads which have a meaningful notion of 'uniform choice' formalised using semi-algebraic notions. The content of Section 6.2.2 is joint work with Amin Karamlou in the forthcoming paper [55]. Section 6.3.3 is on the contribution (C4) and where we construct a distributive law of the Ehrenfeucht-Fraïssé comonad $\mathbb{E}_k^{\odot}$ over the quantum monad $\mathbf{Q}_d^{\odot}$ tailored to 'commeasurable' relational structures. The content of Section 6.3.3 is joint work with Amin Karamlou and is ongoing [56].

Chapter 7 concludes this thesis with a summary and explanations of further research directions.

Chapter 2

Background

This chapter gives an introduction to the concepts and notations in category theory and finite model theory that I will employ in this thesis.

2.1 Set and order notations

For every $n \in \mathbb{N}$, we write $[n]$ to denote the set $\{1, \ldots, n\}$ and (n) to denote the set $\{0, \ldots, n\}$. Given a sequence $s = [x_1, \ldots, x_n]$, we write $s[i, j]$, where $i < j$, for the subsequence $[x_i, \ldots, x_j]$ and $s(i, j]$ for the subsequence $[x_{i+1}, \ldots, x_j]$. Given a sequence $s = [x_1, \ldots, x_n]$ and index $i \in [n]$, we will write $s[i]$ to denote the i-th element x_i. Given two sequences s, s', we write $s \sqsubseteq s'$ if s is a prefix of s' and ss' for the concatenation of the the two sequences. We write $\mathsf{suffix}(s, s')$ for the sequence t such that $st = s'$. If $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ is a sequence of pairs from a set $[k] \times A$, then we write $\mathsf{last}_p(s)$ to denote an element a_i, whenever $p = p_i \in [k]$ is the last occurrence of p in the sequence s. We will also denotes lists as $\vec{a} = (a_1, \ldots, a_n) \in A^n$, we will use $[\vec{a}]$ to denote the support set $\{a_1, \ldots, a_n\}$.

Given a partially ordered set $(X, \leq)$, for $x \in X$, the *down set of* x is

$$\downarrow x = \{y \in X \mid y \leq x\}; \text{ and}$$

the *up set of* x is

$$\uparrow x = \{y \in X \mid x \leq y\}.$$

The *strict relation* $<$ *generated by* $\leq$ is defined as

$$x < y \text{ if } x \leq y \text{ and } x \neq y; \text{ and}$$

the *covering relation* $\prec$ *generated by* $\leq$ is defined as

$$x \prec y \text{ if } x < y \text{ and there is no z such that } x < z < y.$$

A partially ordered set $(L, \leq)$ is *linearly ordered* if for all $x, y \in L$

$$x \leq y \text{ or } y \leq x.$$

Definition 2.1.1. A partially ordered set $(T, \leq)$ is a *forest* if for all $x \in T$, $\downarrow x$ is finite and linearly ordered by $\leq$. We say that $\leq$ *forest orders* T.

Given a forest $(T, \leq)$, if $x \in T$ is such that for all $y \leq x$, $y = x$, then x is a *minimal element of* T. The *height of an element* $x \in T$, denoted $\mathsf{ht}(x)$, in a forest $(T, \leq)$ is the cardinality of the set $\downarrow x \backslash \{x\}$. The *height of the forest* $(T, \leq)$, denoted $\mathsf{ht}(T)$, is $\max_{x \in T} \mathsf{ht}(x)$. The set of successors of an element $x \in T$ is $\mathsf{succ}(x) := \{x' \mid x' \succ x\}$.

Definition 2.1.2. A function $f \colon X \to Y$ between two forests $(X, \leq_X)$ and $(Y, \leq_Y)$ is a *forest morphism* $f \colon (X, \leq_X) \to (Y, \leq_Y)$ if for every minimal element $x \in X$, $f(x)$ is minimal and for every x, x', if $x \prec_X x'$, then $f(x) \prec_Y f(x')$.

Viewing forests as partially ordered sets, this definition of forest morphism is equivalent to height-preserving monotone map.

Definition 2.1.3. A forest $(T, \leq)$ is a *tree* if there exists at most one minimal element $\perp$ called the *root*, i.e $\perp \in T$ such that for all $x \in T$, $\perp \leq x$. We say that $\leq$ *tree orders* T. The tree with no root is the unique empty tree.

Definition 2.1.4. A $(T, \leq)$ is a forest is a *linear forest* if every up-set $\uparrow x$ is finite and linearly ordered.

If $(T, \leq)$ is a linear forest and a tree, then it follows that $\leq$ is finite and linearly ordered.

Definition 2.1.5. A tree $(T, \leq)$ with root $\perp$ is a *linear tree* if $(T \backslash \{\perp\}, \leq)$ is a linear forest.

For a set X, we will use $\mathcal{P}(X)$ to denote the powerset of X. For a set function f, $\mathcal{P}(f) \colon \mathcal{P}(X) \to \mathcal{P}(Y)$ is the direct-image function where

$$f(S) = \{f(x) \in Y \mid x \in S\}.$$

2.2 Model theory

Definition 2.2.1. A finite, relational *signature* σ is a finite set of relational symbols where every relational symbol $R \in \sigma$ has an associated positive integer arity.

Remark 2.2.2. In model theory, signatures may contain constant symbols, function symbols with associated arities, and moreover, are not necessarily finite sets. However, a common practice in the sub-field of finite model theory is to restrict to finite relational signatures. This restriction is not essential, but enables the theorems to be less cumbersome to state.

Definition 2.2.3. A *σ-structure* $\mathcal{A}$ is given by a set A, the *universe* of the structure $\mathcal{A}$, and interpretations $R^{\mathcal{A}} \subseteq A^r$ for every relational symbol $R \in \sigma$ with arity r.

Formally, model theorists define a *logic* $\mathcal{J}$ to be a function from (relational) signatures σ to a collection $\mathcal{J}(\sigma)$ of formulas $\phi(\vec{x})$ for some list of free variables $\vec{x}$. Each logic $\mathcal{J}$ is equipped with a *semantics relation* $\vDash_{\mathcal{J}(\sigma)}$ defining when σ-structure $\mathcal{A}$ paired with a choice of parameters $\vec{a} \in A^n$ *satisfy* a formula $\varphi(\vec{x})$. As the logic and signature will often be implicit, the assertion that $\mathcal{A}, \vec{a}$ satisfies a formula $\varphi(\vec{x})$ will be denoted as:

$$\mathcal{A}, \vec{a} \vDash \varphi(\vec{x}).$$

For any two σ-structures $\mathcal{A}$, $\mathcal{B}$, parameters $\vec{a} \in A^n, \vec{b} \in B^n$, and free variables $\vec{x} = (x_1, \dots, x_n)$ we write

$$(\mathcal{A}, \vec{a}) \Rrightarrow^{\mathcal{J}} (\mathcal{B}, \vec{b}) \text{ if } \forall \phi(\vec{x}) \in \mathcal{J}(\sigma), (\mathcal{A}, \vec{a}) \vDash \phi(\vec{x}) \Rightarrow (\mathcal{B}, \vec{b}) \vDash \phi(\vec{x}); \text{ and}$$
$$(\mathcal{A}, \vec{a}) \equiv^{\mathcal{J}} (\mathcal{B}, \vec{b}) \text{ if } \forall \phi(\vec{x}) \in \mathcal{J}(\sigma), (\mathcal{A}, \vec{a}) \vDash \phi(\vec{x}) \Leftrightarrow (\mathcal{B}, \vec{b}) \vDash \phi(\vec{x}).$$

A *sentence* in $\mathcal{J}(\sigma)$ is formula of $\mathcal{J}(\sigma)$ with no free variables.

2.2.1 First-order and infinitary logic

We assume the standard semantics and syntax of first-order logic, denoted **FO**, as defined in [59]. Infinitary logic, denoted $\mathcal{L}_\infty$, extends first-order logic with disjunctions/conjunction over arbitrary (possibly infinite) sets of formulas. For both **FO** and $\mathcal{L}_\infty$, we will assume these logics are equality-free, i.e. there is no symbol that denotes equality in the universe of the structure. From both first-order logic and infinitary logic, we will consider three variants:

- $\exists^+\mathcal{L}_\infty$ and $\exists^+$**FO** are the fragments of $\mathcal{L}_\infty$ and **FO**, respectively, restricted to positive-existential formulas, i.e. those without negations or universal quantifiers.
- $\exists\mathcal{L}_\infty$ and $\exists$**FO** are the fragments of $\mathcal{L}_\infty$ and **FO**, respectively, restricted to existential formulas, those without universal quantifiers and negations only applied to atomic subformulas.
- $\#\mathcal{L}_\infty$ and #**FO** are the extensions of $\mathcal{L}_\infty$ and **FO**, respectively, with formulas that involve counting quantifiers of the form $\exists_{\leq i}x$ and $\exists_{\geq i}x$ for each $i \in \mathbb{N}$. The formula $\exists_{\leq i}x\varphi(x, \vec{y})$ has the semantics that there exists at most i many witnesses to x

satisfying the formula $\varphi(x, \vec{y})$. Similarly, the formula $\exists_{\geq i} x \varphi(x, \vec{y})$ has the semantics that there exists at least i many witnesses to x satisfying the formula $\varphi(x, \vec{y})$.

We will grade **FO**, $\mathcal{L}_\infty$, and their variants, by two resource parameters:

- Variables. For every $k \in \mathbb{N}$ with $k > 0$ and any logic $\mathcal{J}$ mentioned in this section, $\mathcal{J}^k$ is the $\mathcal{J}$ restricted to formulas which only use the variables $\{x_1, \ldots, x_k\}$.
- Quantifier Rank. For every $k \in \mathbb{N}$ with $k > 0$ and any logic $\mathcal{J}$ mentioned in this section, $\mathcal{J}_k$ is the $\mathcal{J}$ restricted to formulas whose nesting depth of quantifiers is at most k.

For instance, $\mathbf{FO}^k$, $\mathcal{L}^k$, $\exists^+\mathcal{L}^k$, $\#\mathcal{L}^k$ denote the k-variable fragments of first-order logic, infinitary logic, positive-existential infinitary logic, and infinitary logic with counting quantifiers. Similarly, $\mathbf{FO}_k$, $\mathcal{L}_k$, $\exists^+\mathcal{L}_k$, $\#\mathcal{L}_k$ denote the fragments of first-order logic, infinitary logic, positive-existential infinitary logic, and infinitary logic with counting quantifiers restricted to formulas with quantifier rank $\leq k$.

When necessary, for a quantifier Q and a vector of variables $\vec{y} = (y_1, \ldots, y_m)$, we will use $Q\vec{y}$ as a shorthand for the quantifier block $Qy_1Qy_2 \ldots Qy_m$. A formula $\varphi(\vec{y}) \in \exists^+\mathcal{L}_\infty$ is *primitive* if $\varphi(\vec{y})$ only involves only atomics, conjunctions and existential quantifiers, i.e. does not involve disjunctions.

2.2.2 Modal logic

We say signature σ is a *modal signature* if for all symbols $R \in \sigma$, the arity of R is ≤ 2. We will assume each binary relation R_α is indexed by a transition alphabet $\alpha \in \mathsf{Act}$ and each unary relation P is associated to a propositional variable in PV. The reason for this terminology is that pointed σ-structures $(\mathcal{A}, a)$ of modal signatures σ provide a semantics for multi-modal logic $\mathcal{M}$. Namely, if σ is modal, for each unary relation symbol $P \in \sigma$ we assign a propositional variable p, and to each binary relational symbol $R_\alpha \in \sigma$, we have modalities $\Box_\alpha$ and $\Diamond_\alpha$ which are interpreted using $R_\alpha^\mathcal{A}$ as the accessibility relation between worlds $a \in A$ in a σ-structure $\mathcal{A}$. As with first-order logic and infinitary logic, we will investigate the analogous variants of $\mathcal{M}$:

- $\Diamond^+\mathcal{M}$ is the fragment of $\mathcal{M}$ restricted to positive, $\Diamond_\alpha$-formulas, i.e. those without negations or $\Box_\alpha$ modalities.
- $\Diamond\mathcal{M}$ is the fragment of $\mathcal{M}$ restricted to $\Diamond_\alpha$-formulas and negations only applied to propositions.
- $\#\mathcal{M}$ is the extension of $\mathcal{M}$ with graded modalities $\Diamond_\alpha^{\leq n}$ and $\Diamond_\alpha^{\geq n}$ for every $n \in \mathbb{N}$ as in [32]. The semantic relation $\vDash$ is extended to interpret these formulas as:

- $(\mathcal{A}, a) \vDash \Diamond_{\alpha}^{\leq n}\psi$ if there exists at least n-many worlds $a' \in A$ where $(a, a') \in R_{\alpha}^{\mathcal{A}}$ such that $(\mathcal{A}, a') \vDash \psi$.
- $(\mathcal{A}, a) \vDash \Diamond_{\alpha}^{\geq n}\psi$ if there exists at most n-many worlds $a' \in A$ where $(a, a') \in R_{\alpha}^{\mathcal{A}}$ such that $(\mathcal{A}, a') \vDash \psi$.

We will also grade $\mathcal{M}$ and its variants by the *modal depth*, i.e. the nesting depth of modalities. Thus, $\mathcal{M}_k, \Diamond^{+}\mathcal{M}_k, \Diamond\mathcal{M}_k, \#\mathcal{M}_k$ are the fragments of $\mathcal{M}, \Diamond^{+}\mathcal{M}, \Diamond\mathcal{M}, \#\mathcal{M}$, respectively, restricted to formulas with modal depth at most k.

2.3 Category Theory

We assume the reader is familiar with the standard definitions of category, functor, natural transformation, and (co)limits (see e.g. [14, 78] for details). For a category $\mathcal{C}$, the class of objects will be denoted $\mathbf{Ob}(\mathcal{C})$ and the class of morphisms will be denoted $\mathbf{Mor}(\mathcal{C})$. Given objects $X, Y \in \mathbf{Ob}(\mathcal{C})$, $\mathcal{C}(X, Y)$ denotes the class of morphisms of type $X \to Y$ in $\mathbf{Mor}(\mathcal{C})$. For a category $\mathcal{C}$, we will write $X \to_{\mathcal{C}} Y$ if there exists a morphism $X \to Y \in \mathbf{Mor}(\mathcal{C})$ and $X \cong_{\mathcal{C}} Y$ if there exists an isomorphism $X \to Y \in \mathbf{Mor}(\mathcal{C})$. As all the categories we consider are locally-small, $\mathcal{C}(X, Y)$ is always a set. Familiarity with the definition of comonad and adjunction is helpful, but not strictly necessary, as we will define generalisations of these concepts.

The category of locally small categories with functor will be denoted **Cat**. The category of sets and functions will be denote **Set**.

2.3.1 Relative comonads

The following definitions are the dual of the definitions given in [17] on relative monads.

Definition 2.3.1. Given a functor $J \colon \mathcal{J} \to \mathcal{C}$, *a relative comonad on* J is a triple $(\mathbb{T}, \varepsilon, (\cdot)^*)$ where $\mathbb{T} \colon \mathbf{Ob}(\mathcal{J}) \to \mathbf{Ob}(\mathcal{C})$ is an object mapping, for every $X \in \mathbf{Ob}(\mathcal{J})$, there is a morphism $\varepsilon_X \colon \mathbb{T}X \to JX \in \mathbf{Mor}(\mathcal{C})$, and for every morphism $f \colon \mathbb{T}X \to JY \in \mathbf{Mor}(\mathcal{C})$, there is a coextension morphism $f^* \colon \mathbb{T}X \to \mathbb{T}Y \in \mathbf{Mor}(\mathcal{C})$ satisfying the following equations:

$$(\varepsilon_X)^* = \mathsf{id}_{\mathbb{T}X} \quad \varepsilon \circ f^* = f \quad (g \circ f^*)^* = g^* \circ f^*$$

These equations allow us to extend the object mapping $\mathbb{T} \colon \mathbf{Ob}(\mathcal{J}) \to \mathbf{Ob}(\mathcal{C})$ to a functor $\mathbb{T} \colon \mathcal{J} \to \mathcal{C}$ where $\mathbb{T}(f) = (J(f) \circ \varepsilon_X)^*$.

Definition 2.3.2. Given a relative comonad $(\mathbb{T}, \varepsilon, (\cdot)^*)$ over $J \colon \mathcal{J} \to \mathcal{C}$, the *CoKleisli category of* $\mathbb{T}$, denoted $\mathbf{Kl}(\mathbb{T})$, is defined such that

- $\mathbf{Ob}(\mathbf{Kl}(\mathbb{T})) = \mathbf{Ob}(\mathcal{J})$; and
- $\mathbf{Kl}(\mathbb{T})(X, Y) = \mathcal{C}(\mathbb{T}X, JY)$.
- The composition $g \circ_{\mathbf{Kl}(\mathbb{T})} f \colon \mathbb{T}X \to JZ$ of two morphisms $f \colon \mathbb{T}X \to JY$ and $g \colon \mathbb{T}Y \to JZ$ is given by
$$\mathbb{T}X \xrightarrow{f^*} \mathbb{T}Y \xrightarrow{g} JZ.$$
- The identity morphisms are given by the counit
$$\varepsilon_X \colon \mathbb{T}X \to JX.$$

The ordinary notion of a *comonad in coKleisli form* [63] and the corresponding coKleisli category can be recovered when $\mathcal{J} = \mathcal{C}$ and $J = \mathsf{id}_{\mathcal{C}}$. Observe that given a comonad $(\mathbb{T}, \varepsilon, (\cdot)^*)$ and a functor $J \colon \mathcal{J} \to \mathcal{C}$, the functor $\mathbb{T}^J = \mathbb{T} \circ J$ can be made into a relative comonad $(\mathbb{T}^J, \varepsilon', (\cdot)')$ on J, where $\varepsilon'_X = \varepsilon_{JX}$ and the coextension mapping $(\cdot)'$ is defined for $f \colon \mathbb{T}^J X \to JY$ to be the same as f^*, i.e. $f' = f^* \colon \mathbb{T}^J \to \mathbb{T}^J Y$. We say that $\mathbb{T}^J$ is the *relative-lifting of* $\mathbb{T}$ *along* J.

Given an ordinary comonad in coKleisli form, we can define a comultiplication morphism $\delta_X \colon \mathbb{T}X \to \mathbb{T}\mathbb{T}X$, where $\delta_X = (\mathsf{id}_{\mathbb{T}X})^*$, which satisfies the following equations:

$$\mathbb{T}\delta_X \circ \delta_X = \delta_{\mathbb{T}X} \circ \delta_X; \qquad \mathbb{T}\varepsilon_X \circ \delta_X = \varepsilon_{\mathbb{T}X} \circ \delta_X = \mathsf{id}_{\mathbb{T}X}.$$

The triple $(\mathbb{T} \colon \mathcal{C} \to \mathcal{C}, \varepsilon, \delta)$, where $\mathbb{T}$ is a functor, ε is a natural transformation, and δ is a natural transformation, is a *comonad in standard form.* Its coKleisli form can be recovered by defining the coextension mapping $(\cdot)^*$ as $f^* = \mathbb{T}f \circ \delta$.

Definition 2.3.3. Given a comonad $(\mathbb{T} \colon \mathcal{C} \to \mathcal{C}, \varepsilon, \delta)$, a *coalgebra* over $\mathbb{T}$ is defined as a pair $(A, \alpha \colon A \to \mathbb{T}A)$, where $A \in \mathbf{Ob}(\mathcal{C})$ and $\alpha \in \mathbf{Mor}(\mathcal{C})$, satisfying the following diagrams:

$$\delta_A \circ \alpha = \mathbb{T}\alpha \circ \alpha \tag{2.1}$$

$$\varepsilon_A \circ \alpha = \mathsf{id}_A. \tag{2.2}$$

Definition 2.3.4. The *category of coalgebras*, or *Eilenberg-Moore category*, associated with a comonad $\mathbb{T} \colon \mathcal{C} \to \mathcal{C}$, denoted by $\mathbf{EM}(\mathbb{T})$, is defined as

- $\mathbf{Ob}(\mathbf{EM}(\mathbb{T}))$ consists of coalgebras $(A, \alpha \colon A \to \mathbb{T}A)$;
- A morphism of type $h \colon (A, \alpha) \to (B, \beta) \in \mathbf{Mor}(\mathbf{EM}(\mathbb{T}))$ is a morphism $h \colon A \to B \in \mathbf{Mor}(\mathcal{C})$ such that $\mathbb{T}h \circ \alpha = \beta \circ h$;
- Identity and composition are defined as in $\mathcal{C}$.

2.3.2 Relative adjunctions

Definition 2.3.5. A *relative adjunction* between a functor $J\colon \mathcal{J} \to \mathcal{D}$ and a category $\mathcal{C}$ is two functors $L\colon \mathcal{C} \to \mathcal{D}$, $R\colon \mathcal{J} \to \mathcal{C}$ and family of bijections:

$$\phi_{X,Y}\colon \mathcal{D}(L(X), J(Y)) \cong \mathcal{C}(X, R(Y))$$

which is natural in both arguments X, Y.

We recover the notion of ordinary adjunction when $\mathcal{J} = \mathcal{D}$ and $J = \mathsf{Id}_{\mathcal{D}}$. Given an ordinary adjunction $L \dashv R$, the unit of the adjunction has components $\eta_X = \phi_{X,L(X)}(\mathsf{id}_{L(X)})$ and the counit of the adjunction has components $\varepsilon_X = \phi^{-1}_{R(Y),Y}(\mathsf{id}_{R(Y)})$. Every ordinary adjunction $L\colon \mathcal{C} \to \mathcal{D} \dashv R\colon \mathcal{D} \to \mathcal{C}$ yields a comonad $(\mathbb{T}, \varepsilon, \delta)$ over $\mathcal{D}$, where $\mathbb{T} = LR$, ε is the counit of the adjunction, and the component $\delta_X\colon LR(X) \to LRLR(X)$ for $X \in \mathcal{D}$ is defined as $L(\eta_{R(X)})$. A *resolution* of $\mathbb{T}\colon \mathcal{D} \to \mathcal{D}$ is an adjunction $L\colon \mathcal{C} \to \mathcal{D} \dashv R\colon \mathcal{D} \to \mathcal{C}$ that gives rise to $\mathbb{T}$ as a comonad. There is a category of resolutions of $\mathbb{T}$, denoted by $\mathbf{Res}(\mathbb{T})$, defined as:

- The class $\mathbf{Ob}(\mathbf{Res}(\mathbb{T}))$ consists of adjunctions $(L \dashv R, \phi)$ with $L\colon \mathcal{C} \to \mathcal{D}$ and $R\colon \mathcal{D} \to \mathcal{C}$ such that $\mathbb{T} = LR$;
- A morphism of type $K\colon (L \dashv R, \phi) \to (L' \dashv R', \phi')$ in $\mathbf{Mor}(\mathbf{Res}(\mathbb{T}))$ where $L\colon \mathcal{C} \to \mathcal{D}$ and $L'\colon \mathcal{C}' \to \mathcal{D}$ is a functor $K\colon \mathcal{C} \to \mathcal{C}'$ where the diagram

$$\begin{tikzcd}
\mathcal{C} \arrow[dd, "K"'] \arrow[dr, shift left, "L"] & \\
& \mathcal{D} \arrow[ul, shift left, "R"] \arrow[dl, shift left, "R'"] \\
\mathcal{C}' \arrow[ur, shift left, "L'"] &
\end{tikzcd}$$

  commutes in $\mathbf{Cat}$;
- Identity and composition are defined as in $\mathbf{Cat}$.

The terminal object of $\mathbf{Res}(\mathbb{T})$ is the forgetful-cofree adjunction $U \dashv F$ where $U\colon \mathbf{EM}(\mathbb{T}) \to \mathcal{D}$ and $F\colon \mathcal{D} \to \mathbf{EM}(\mathbb{T})$. Thus for every resolution $(L \dashv R, \phi)$ of $\mathbb{T}$ with unit $\eta\colon \mathsf{Id}_{\mathcal{D}} \to RL$, there is a unique morphism $K_!\colon \mathcal{C} \to \mathbf{EM}(\mathbb{T})$ given by the functor such that

- For every $X \in \mathbf{Ob}(X)$ $K_!(X) = (L(X), L(\eta_X)\colon LX \to LRL(X))$. By $\mathbb{T} = LR$ and the properties of the resolution $(L \dashv R, \phi)$, $L(\eta_X)\colon L(X) \to \mathbb{T}(L(X))$ is indeed a coalgebra.

- For every morphism $f \in \mathcal{C}(X, Y)$, $K_!(f)$ is the commutative diagram:

$$\begin{tikzcd}
L(X) \arrow[r, "L(\eta_X)"] \arrow[d, "f"'] & \mathbb{T}(L(X)) \arrow[d, "L(\phi(f))"] \\
L(Y) \arrow[r, "L(\eta_Y)"] & \mathbb{T}(L(Y))
\end{tikzcd}$$

A resolution $L \dashv R$ is a *comonadic adjunction* if this unique morphism $K_! \colon \mathcal{C} \to \mathbf{EM}(\mathbb{T})$ in $\mathbf{Res}(\mathbb{T})$ is an equivalence.

2.3.3 Factorisation Systems

Definition 2.3.6 (lifting property)**.** Given a category $\mathcal{C}$ and morphisms e and m in $\mathcal{C}$, 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 diagram as the left diagram below, there exists a *diagonal filler* d such that the right diagram below commutes.

$$\begin{tikzcd}
\bullet \arrow[r, "e"] \arrow[d] & \bullet \arrow[d] \\
\bullet \arrow[r, "m"] & \bullet
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
\bullet \arrow[r, "e"] \arrow[d] & \bullet \arrow[d] \arrow[dl, "d"'] \\
\bullet \arrow[r, "m"] & \bullet
\end{tikzcd}$$

We denote this property by $e \pitchfork m$. For any class $\mathcal{M}$ of morphisms in $\mathcal{C}$, let ${}^{\pitchfork}\mathcal{M}$ denote the class of morphisms with the left lifting property with respect to every morphism in $\mathcal{M}$. The class $\mathcal{M}^{\pitchfork}$ is defined analogously.

Definition 2.3.7 (weak factorisation system)**.** Given two classes of morphisms $\mathcal{Q}$ and $\mathcal{M}$ in a category $\mathcal{C}$, the pair $(\mathcal{Q}, \mathcal{M})$ is a *weak factorisation system* if the following conditions are satisfied:

(1) For every morphism f in $\mathcal{C}$, $f = m \circ e$, where $e \in \mathcal{Q}$ and $m \in \mathcal{M}$;

(2) $\mathcal{Q} = {}^{\pitchfork}\mathcal{M}$ and $\mathcal{M} = \mathcal{Q}^{\pitchfork}$.

A weak factorisation system is *proper* if every $e \in \mathcal{Q}$ is an epimorphism and every $m \in \mathcal{M}$ is a monomorphism. A proper weak factorisation system is *stable* if for every $e \in \mathcal{Q}$ and $m \in \mathcal{M}$, with the same codomain, there exists a pullback of e along m in $\mathcal{Q}$.

Definition 2.3.8. A triple $(\mathcal{C}, \mathcal{Q}, \mathcal{M})$ is a *factorised category* if $(\mathcal{Q}, \mathcal{M})$ is a stable proper factorisation system of $\mathcal{C}$.

Throughout this thesis, we will use some properties of factorised categories which we enumerate in Proposition 2.3.9 below. This proposition is a restatement of [9, Lemma 2.5], but these are well-known results that also appear in reference material, e.g. [77] and [37].

Proposition 2.3.9. *Let $(\mathcal{C}, \mathcal{Q}, \mathcal{M})$ be a factorised category. The following statements hold:*

(a) $\mathcal{Q}$ *and* $\mathcal{M}$ *are closed under composition.*

(b) $\mathcal{Q} \cap \mathcal{M} = \{$*isomorphisms*$\}$

(c) The pullback of a morphism in $\mathcal{M}$ *along any morphism in* $\mathbf{Mor}(\mathcal{C})$*, 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* $f \in \mathcal{M}$*.*

We refer to members of $\mathcal{M}$ as embeddings (denoted by $\rightarrowtail$) and to members of $\mathcal{Q}$ as *quotients* (denoted by $\twoheadrightarrow$). Given two *embeddings* $m \colon S \rightarrowtail X$ and $n \colon T \rightarrowtail X$, we write $m \trianglelefteq n$ to denote that there is a morphism $i \colon S \to T$ such that $m = n \circ i$. Note that $\trianglelefteq$ induces a preorder on embeddings with the same codomain. The symmetrisation of $\trianglelefteq$ induces an equivalence relation $\sim$. The relation $\sim$ can be characterised as $m \sim n$ if there exists an isomorphism $i \colon S \to T$ such that $m = n \circ i$. Let $\mathbf{S}X$ denote the class of $\sim$-equivalence classes of embeddings with codomain X with the partial order $\leq$ induced by $\trianglelefteq$. The $\sim$-equivalence class with a representative $m \colon S \rightarrowtail X$ is denoted $[m]$.

2.3.4 Categories for model theory

The categories that we are primarily interested in are categories of relational structures for fixed signature σ.

Definition 2.3.10. Given a signature σ, a *σ-morphism* $h \colon \mathcal{A} \to \mathcal{B}$ is a set function $h \colon A \to B$ such that for every relational symbol $R \in \sigma$ of arity n,

$$R^{\mathcal{A}}(a_1, \ldots, a_n) \Rightarrow R^{\mathcal{B}}(h(a_1), \ldots, h(a_n)),$$

for all $a_1, \ldots, a_n \in A$.

Definition 2.3.11. A *σ-embedding* $h \colon \mathcal{A} \to \mathcal{B}$ is an injective set function $h \colon A \to B$ such that for every relational symbol $R \in \sigma$ of arity n,

$$R^{\mathcal{A}}(a_1, \ldots, a_n) \Leftrightarrow R^{\mathcal{B}}(h(a_1), \ldots, h(a_n)),$$

for all $a_1, \ldots, a_n \in A$.

We use $\mathsf{Struct}(\sigma)$ to denote the category of σ-structures and σ-morphisms. We let $\mathsf{Struct}^f(\sigma)$ denote the full subcategory of $\mathsf{Struct}(\sigma)$ containing only *finite* σ-structures. We also consider the category of pointed σ-structures and point-preserving σ-morphisms

$\mathsf{Struct}_\star(\sigma)$. In $\mathsf{Struct}_\star(\sigma)$, objects are pointed σ-structures $(\mathcal{A}, a_0)$. A morphism $h\colon (\mathcal{A}, a_0) \to (\mathcal{B}, b_0)$ in $\mathsf{Struct}_\star(\sigma)$ is a σ-morphism $h\colon \mathcal{A} \to \mathcal{B}$ such that $h(a_0) = b_0$. We let $\mathsf{Struct}^f_\star(\sigma)$ denote full subcategory of $\mathsf{Struct}^f_\star(\sigma)$ containing only finite pointed σ-structures. The signature of simple directed graphs contains one binary edge relation E and we will denote $\mathsf{Struct}(\sigma)$ in this case as **DiGraph**. The category of simple undirected graphs **Graph** can be seen as full subcategory of **DiGraph** whose objects are directed graphs with a symmetric edge relation. For every σ-structure $\mathcal{A}$, the *Gaifman Graph of* $\mathcal{A}$, denoted by $\mathcal{G}(\mathcal{A})$, is an undirected graph with vertex set A, where elements a, a' are adjacent if $a = a'$ or a, a' appear in the same tuple of $R^{\mathcal{A}}$ for some $R \in \sigma$. Since every σ-morphism $f\colon \mathcal{A} \to \mathcal{B}$ induces a graph homomorphism $\mathcal{G}(f)\colon \mathcal{G}(\mathcal{A}) \to \mathcal{G}(\mathcal{B})$, we see that $\mathcal{G}\colon \mathsf{Struct}(\sigma) \to \mathbf{Graph}$ is a functor.

When dealing with model-comparison games, we will need generalise the notion of σ-morphism to a binary relation between two σ-structures which preserves relational interpretations.

Definition 2.3.12. Given σ-structures $\mathcal{A}, \mathcal{B}$, a binary relation $\gamma \subseteq A \times B$ is a *σ-correspondence from $\mathcal{A}$ to $\mathcal{B}$* if for all $R \in \sigma$ with arity n,

$$R(a_1, \dots, a_n) \Rightarrow R(b_1, \dots, b_n)$$

if $(a_1, b_1), \dots, (a_n, b_n) \in \gamma$.

If γ is σ-correspondence from $\mathcal{A}$ to $\mathcal{B}$ and its converse

$$\gamma^c := \{(b, a) \mid (a, b) \in \gamma\}$$

is a σ-correspondence from $\mathcal{B}$ to $\mathcal{A}$, then we say that γ is *tight σ-correspondence from $\mathcal{A}$ to $\mathcal{B}$.* A σ-correspondence γ from $\mathcal{A}$ to $\mathcal{B}$ that is also a partial function is a *partial homomorphism.* If γ is partial homomorphism from $\mathcal{A}$ to $\mathcal{B}$ and γ^c is a partial homomorphism $\mathcal{B}$ to $\mathcal{A}$, then we say that γ *is a partial isomorphism.* Equivalently, γ is a partial isomorphism if γ is a tight σ-correspondence and a partial injective function.

Chapter 3

Arboreal Covers

The central topic of this thesis are the notions of arboreal category and arboreal cover first introduced by Abramsky and Reggio in [9] and its follow-up journal version [10]. In this chapter, we re-present the content of these papers and highlight the portions which will be needed for the contributions of this thesis. We begin by introducing the the following simple example of an arboreal category.

Example 3.0.1. Let $\mathcal{F}$ denote the category whose class of objects $\mathbf{Ob}(\mathcal{F})$ are the forests defined in Definition 2.1.1 and whose class of morphisms $\mathbf{Mor}(\mathcal{F})$ are the forest morphisms defined in Definition 2.1.2. Let $\mathcal{T}$ be the full subcategory of $\mathcal{F}$ whose class of objects $\mathbf{Ob}(\mathcal{T})$ are trees defined in Definition 2.1.3. The categories $\mathcal{F}$ and $\mathcal{T}$ are equipped with factorisation systems where

$$\begin{aligned} \mathcal{Q} &= \{q\colon (X, \leq_X) \to (Y, \leq_Y) \mid q\colon X \to Y \text{ surjective}\}; \text{ and} \\ \mathcal{M} &= \{m\colon (X, \leq_X) \to (Y, \leq_Y) \mid m\colon X \to Y \text{ injective}\}. \end{aligned}$$

$(\mathcal{F}, \mathcal{Q}, \mathcal{M})$ and $(\mathcal{T}, \mathcal{Q} \cap \mathbf{Mor}(\mathcal{T}), \mathcal{M} \cap \mathbf{Mor}(\mathcal{T}))$ are factorised categories as in Definition 2.3.8.

3.1 Arboreal categories

The definition of arboreal category is intended to formalise the intuitive notion of a category consisting of "tree shaped" objects. In order to accomplish this, the definition of arboreal category takes inspiration from the merelogical philosophy (that a whole is determined by its parts) of many other constructions in category theory. To illustrate what we mean by this philosophy, observe that in the category of **Set**, every set X is a directed colimit of its finite subsets and inclusions between them. More generally, in a locally finitely presentable category $\mathcal{D}$, objects are generated via filtered colimits from compact subobjects in a full subcategory $\mathcal{D}_{fp}$. In a similar vein, in order to capture the notion of a category of tree-shaped objects $\mathcal{C}$, objects of an arboreal category are generated as colimits of their branches, or more precisely, path-shaped subobjects. For this reason, we first need the intermediate notion of a path category which has enough structure to contain a full subcategory $\mathcal{C}_p$ of path objects.

3.1.1 Path categories

In this subsection, we will assume we have a factorised category $(\mathcal{C}, \mathcal{Q}, \mathcal{M})$ as in Definition 2.3.8. Recall that a factorisation system $(\mathcal{Q}, \mathcal{M})$ induces a partially ordered set of associated subobjects $\mathbf{S}X$ for every object $X \in \mathbf{Ob}(\mathcal{C})$.

Definition 3.1.1 (path)**.** An object P of $\mathcal{C}$ is a *path* if $\mathbf{S}P$ is a finite chain.

For every factorised category $(\mathcal{C}, \mathcal{Q}, \mathcal{M})$, we denote the full subcategory of path objects $\mathcal{C}_p$. An embedding $e\colon P \rightarrowtail X \in \mathcal{M}$ is a *path embedding* if P is a path. For an object $X \in \mathcal{C}$, let $\mathbf{P}X$ denote the subset of $\mathbf{S}X$ consisting of equivalence classes $[e]$ such that $e\colon P \rightarrowtail X$ is a path embedding.

Given a morphism $f\colon X \to Y$, this induces a set function $\mathbf{P}f\colon \mathbf{P}X \to \mathbf{P}Y$ where $[p] \mapsto [\exists_f p]$ from the factorisation

$$\begin{tikzcd} P \arrow[r, two heads, "e_p"] & \exists_f P \arrow[r, tail, "\exists_f p"] & X \end{tikzcd}$$

of $f \circ p$. By [10, Lemma 3.5(a)], it follows that $\exists_f P$ is a path, so $\exists_f p$ is a path embedding and $\mathbf{P}f$ is indeed a function.

Definition 3.1.2 (pathwise embedding)**.** A morphism $f\colon X \to Y \in \mathbf{Mor}(\mathcal{C})$ is a *pathwise embedding in* $\mathcal{C}$ if for all path embeddings $m\colon P \rightarrowtail X$, $f \circ m\colon P \to Y$ is a path embedding.

For factorised categories $\mathcal{C}$, we will use $X \rightharpoonup^{\mathcal{C}} Y$ to denote there exists a pathwise embedding $f\colon X \to Y \in \mathbf{Mor}(\mathcal{C})$. It is easy to show using the lifting property of Definition 2.3.6 that for pathwise embeddings $f\colon X \to Y$, $[f \circ m] = [\exists_f m]$.

Definition 3.1.3 (open map)**.** A morphism $f\colon X \to Y \in \mathbf{Mor}(\mathcal{C})$ is *open in* $\mathcal{C}$ if for every commutative square of the form:

$$\begin{tikzcd} P \arrow[r, tail] \arrow[d, tail] & Q \arrow[d, tail] \\ X \arrow[r, "f"'] & Y \end{tikzcd}$$

where $P, Q \in \mathbf{Ob}(\mathcal{C}_p)$, there exists a diagonal filler morphism $Q \to X$ such that the following two triangles commute:

$$\begin{tikzcd}
P \arrow[r, tail] \arrow[d, tail] & Q \arrow[d, tail] \arrow[dl, dashed] \\
X \arrow[r, "f"'] & Y
\end{tikzcd}$$

Since the triangle on the right commutes, by Proposition 2.3.9 (e), the diagonal filler $Q \to X$ is always an embedding.

Intuitively, an open map $f \colon X \to Y$ states that if a path P in the domain X, witnessed by $P \rightarrowtail X$, can be extended to a path Q in the codomain Y, witnessed by $Q \rightarrowtail Y$, then P can be extended to a path Q in the domain X, witnessed by the diagonal $Q \rightarrowtail X$. Moreover, this extension is the preimage of the path Q in Y under f.

Definition 3.1.4 (bisimulation)**.** A *bisimulation in* $\mathcal{C}$ is a span $X \xleftarrow{f} Z \xrightarrow{g} Y$ of open pathwise embeddings $f, g \in \mathbf{Mor}(\mathcal{C})$.

For factorised categories $\mathcal{C}$, we will use $X \leftrightarrow^{\mathcal{C}} Y$ to denote there exists a bisimulation $X \xleftarrow{f} Z \xrightarrow{g} Y$ in $\mathcal{C}$.

Definition 3.1.5. An object $X \in \mathcal{C}$ is *connected* if every morphism $X \to \bigsqcup_{i \in I} P_i$, where $\{P_i\}_{i \in I}$ is a small family of paths, factors through a coproduct injection $P_j \to \bigsqcup_{i \in I} P_i$.

This definition of connected differs slightly from the usual categorical formulation of connectedness in two ways. Firstly, we only consider coproducts of paths, and secondly, the coproduct injection $P_j \to \bigsqcup_{i \in I} P_i$ is not required to be unique. This second relaxation allows us to consider the initial object to be connected.

Definition 3.1.6 (path category)**.** A factorised category $(\mathcal{C}, \mathcal{Q}, \mathcal{M})$ is a *path category* if it satisfies the following axioms:

1. $\mathcal{C}$ has all coproducts $\coprod_{j \in J} X_j$ where $\{X_j\} \subseteq \mathbf{Ob}(\mathcal{C}_p)$ is small family of paths.
2. For all paths P, Q and R in $\mathcal{C}_p$, if $P \to Q \to R \in \mathcal{Q}$, then $P \to Q \in \mathcal{Q}$.
3. Every path $P \in \mathcal{C}_p$ of $\mathcal{C}$ is connected.

It is apparent from the axioms of a path category that the full subcategory $\mathcal{C}_p$ of path objects, induced by the factorisation system $(\mathcal{Q}, \mathcal{M})$, plays an important role for path categories $\mathcal{C}$.

Proposition 3.1.7. $(\mathcal{F}, \mathcal{Q}, \mathcal{M})$ *as in Example 3.2.1 is path category.*

Proof. For an object $(X, \leq) \in \mathbf{Ob}(\mathcal{F})$, each subobject $[m] \in \mathbf{S}(X)$ corresponds to the down set $\downarrow x$ of an element $x \in X$. Therefore, by the definition of path (Definition 3.1.1), the path objects $(P, \leq) \in \mathbf{Ob}(\mathcal{F}_p)$ are precisely the objects such that $\leq$ is a finite chain. With this definition of $\mathcal{F}_p$ in place, we verify each of the axioms of path category and show that $\mathcal{F}$ satisfies Definition 3.1.6.

To verify axiom (1), for any set S of finite chains $\{(Y_s, \leq_s)\}_{s \in S} \subseteq \mathbf{Ob}(\mathcal{F}_p)$, we define linear forest $(X, \leq)$ such that X is the disjoint union $\coprod_{s \in S} Y_s$ and $\leq$ is the disjoint union $\coprod_{s \in S} \leq_s$. The object $(X, \leq) \in \mathbf{Ob}(\mathcal{F})$ is the coproduct $\biguplus (Y_s, \leq_s)$ of the chains in S.

Intuitively, axiom (2) is capturing the fact that forest morphisms preserve covering chains, i.e. paths can be merged, but not shortened. To verify axiom (2), suppose that $P, Q, R \in \mathcal{F}_p$ with morphisms $u \colon P \to Q$ and $v \colon Q \to R$ such that $v \circ u \in \mathcal{Q}$. By the definition of $\mathcal{Q}$, we assume that $v \circ u$ is surjective, and we must show that u is surjective. Since Q is a finite chain, we can present Q as $q_1 \prec \cdots \prec q_m$ where q_1 is the root of Q and q_m is the maximal element of Q. Since $v \circ u$ preserves covering chains, $v(q_1) \prec \cdots \prec v(q_m)$ is a subchain of R. By $v \circ u$ being surjective, for every $i \in [m]$, there exists a $p_i \in P$ such that $v(u(p_i)) = v(q_i) \in R$. By u preserving roots and P being a chain with a unique root, $u(p_1) = q_1$ where p_1 is the root of P. By u preserving the covering relation $\prec$, for all $i \in [m]$, $u(p_i) = q_i$. Thus, every $q_i \in Q$ has preimage $p_i \in P$, and u is surjective.

To verify axiom (3), consider paths $Q, P_1, \ldots, P_n \in \mathcal{F}_p$ and morphism $f \colon Q \to \biguplus_{i \in [n]} P_i$. Since Q is a finite chain, we can present Q as $q_1 \prec \cdots \prec q_m$ where q_1 is the root of Q and q_m is the maximal element of Q. By f preserving roots, $f(q_1)$ is the root of P_j for some $j \in [n]$. By f preserving the covering relation, for all $i \in [m]$, $f(q_i) \in P_j$. Thus, f factors through the coproduct injection $i_j \colon P_j \to \biguplus_{i \in [n]} P_i$. $\square$

For every object $X \in \mathbf{Ob}(\mathcal{C})$ in a path category $\mathcal{C}$, the subobject relation restricted to $\mathbf{P}(X)$ is a tree-order. In particular, the root of $[\perp_X \colon \tilde{0} \rightarrowtail X] \in \mathbf{P}(X)$ is computed from the $(\mathcal{Q}, \mathcal{M})$ factorisation

$$\begin{tikzcd} 0 \arrow[r, two heads] & \tilde{0} \arrow[r, tail, "\perp_X"] & X \end{tikzcd}$$

of the unique map $!_X$ from the initial object $0 \in \mathbf{Ob}(\mathcal{C})$. The induced set function $\mathbf{P}(f) \colon \mathbf{P}(X) \to \mathbf{P}(Y)$ for an $f \in \mathcal{C}(X, Y)$ is a forest morphism. Thus, we restate [10, Theorem 3.11].

Theorem 3.1.8. *If $\mathcal{C}$ is a path category, $X \mapsto \mathbf{P}(X)$ induces a functor $\mathbf{P} \colon \mathcal{C} \to \mathcal{T}$.*

Definition 3.1.9. An object X in a path category $\mathcal{C}$ is *path-generated* if X is the colimit of the cocone consisting of all commuting triangles of the form

$$\begin{tikzcd}
P \arrow[rr, tail] \arrow[dr, tail] & & Q \arrow[dl, tail] \\
& X &
\end{tikzcd} .$$

where P and Q are paths.

Definition 3.1.10. An *arboreal category* is a path category in which all objects are path-generated.

Definition 3.1.11. An *arboreal category resource-indexed by parameter k* is an arboreal category $\mathcal{C}$ equipped with a full subcategory $\mathcal{C}_p^k$ of $\mathcal{C}_p$ for all $k > 0$, closed under embeddings with inclusions

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

This induces a corresponding chain of full subcategories $\mathcal{C}_k$ of $\mathcal{C}$

$$\mathcal{C}_1 \hookrightarrow \mathcal{C}_2 \hookrightarrow \mathcal{C}_3 \hookrightarrow \dots$$

with the objects of $\mathcal{C}_k$ being those whose cocone of path embeddings with domain in $\mathcal{C}_p^k$ is colimit cocone in $\mathcal{C}$. Since the chain $\mathcal{C}_1 \hookrightarrow \mathcal{C}_2 \hookrightarrow \mathcal{C}_3 \hookrightarrow \dots$ is determined uniquely by the choice of full subcategories $\mathcal{C}_p^k$ of $\mathcal{C}_p$, we can denote a resource-indexed arboreal category as $\{\mathcal{C}_k\}$.

A corollary of Theorem 3.1.8, is that every arboreal category $\mathcal{C}$ has a natural resource indexing $\{\mathcal{C}_k\}$ by height of the underlying branches. Explicitly, $\mathcal{C}_k$ is the full subcategory of $\mathcal{C}$ consisting of objects X such that $\mathsf{ht}(\mathbf{P}(X)) \leq k$. We call this resource-indexing the *depth-indexing of* $\mathcal{C}$. This default indexing reflects the fact that model comparison games are naturally parameterised by the number of rounds. For some arboreal categories $\mathcal{C}$, this is not the only way to resource index $\mathcal{C}$. For instance, Example 3.2.3 of k-pebble forest covers detailed below is a resource-indexing which is not a depth-indexing.

3.1.2 Games in arboreal categories

In the context of an arboreal category $(\mathcal{C}, \mathcal{Q}, \mathcal{M})$, a bisimulation $X \leftarrow Z \rightarrow Y$ in $\mathcal{C}$ can be formulated as an abstract 'path-matching' back-and-forth game from X to Y over $\mathcal{C}$. This fulfils the original motivation for the notion of arboreal category: to extract the common properties shared by the Eilenberg-Moore categories of game comonads which enable them to encode model comparison games in logic. The following is the definition of this abstract back-and-forth game:

Definition 3.1.12. Let X, Y be objects in an arboreal category $\mathcal{C}$. We define the *back-and-forth game* $\mathcal{G}(X, Y)$ *over* $\mathcal{C}$. Positions in $\mathcal{G}(X, Y)$ are given by pairs $([m], [n]) \in \mathbf{P}(X) \times \mathbf{P}(Y)$. The initial position is given by the pair of roots $([\perp_X], [\perp_Y])$. During each round of the game, Spoiler chooses one of the following moves:

- Spoiler chooses $[m] \prec [m']$ in $\mathbf{P}(X)$. Duplicator then responds with $[n] \prec [n']$ in $\mathbf{P}(Y)$.
- Spoiler chooses $[n] \prec [n']$ in $\mathbf{P}(Y)$. Duplicator then responds with $[m] \prec [: m']$ in $\mathbf{P}(X)$.

Duplicator wins the round if she has a response and the resulting position $([m'], [n'])$ is in the set:

$$\mathcal{W}(X, Y) := \{([m], [n]) \mid \mathsf{dom}(m) \cong \mathsf{dom}(n) \text{ in } \mathcal{C}\}.$$

Duplicator wins $\mathcal{G}(X, Y)$ if for every round $t \geq 0$ and every possible Spoiler choice, Duplicator has a winning response.

Each game comonad does not only encode a back-and-forth model comparison game, but also two forth variants. Thus in the abstract arboreal setting, the game $\mathcal{G}(X, Y)$ over $\mathcal{C}$ also has two variants.

- ${}^{\exists}\mathcal{G}(X, Y)$ where Spoiler is forced to play a path in $\mathbf{P}(X)$, and thus Duplicator is forced to respond with a path in $\mathbf{P}(Y)$.
- ${}^{\exists^+}\mathcal{G}(X, Y)$ has the same player restriction as ${}^{\exists}\mathcal{G}(X, Y)$, but the winning positions for Duplicator has been widened to the superset:

$$^{+}\mathcal{W}(X, Y) := \{([m], [n]) \mid \exists \ \mathsf{dom}(m) \to \mathsf{dom}(n) \in \mathbf{Mor}(\mathcal{C})\}.$$

The notation for these games hints at the intuition that if $\mathcal{G}(X, Y)$ over $\mathcal{C}$ "captures" equivalence in a logic $\mathcal{J}$, then ${}^{\exists^+}\mathcal{G}(X, Y)$ over $\mathcal{C}$ captures equivalence in the positive existential fragment $\exists^+\mathcal{J}$ of $\mathcal{J}$, ${}^{\exists}\mathcal{G}(X, Y)$ captures equivalence in the existential fragment $\exists\mathcal{J}$ of $\mathcal{J}$. The abstract existential game ${}^{\exists}\mathcal{G}(X, Y)$ and its connection with concrete model-comparison games characterising existential fragments stems from the work [6].

In an arboreal category $\mathcal{C}$, these abstract games provide an operational definition of morphisms, pathwise embeddings, bisimulation, and isomorphisms in $\mathcal{C}$.

Proposition 3.1.13. *Let $\mathcal{C}$ be an arboreal category and $X, Y \in \mathbf{Ob}(\mathcal{C})$. For the following two statements, assume that the categorical product $X \times Y$ exists in $\mathbf{Ob}(\mathcal{C})$:*

(1) Duplicator has winning strategy in ${}^{\exists}\mathcal{G}(X, Y)$ iff there exists a pathwise embedding of $X \to Y$.

(2) Duplicator has winning strategy in $\mathcal{G}(X, Y)$ iff there exists a bisimulation $X \leftarrow Z \rightarrow Y$.

For the following statement, assume that the paths functor $\mathbf{P}\colon \mathcal{C} \to \mathcal{T}$ is faithful:

(3) Duplicator has winning strategy in ${}^{\exists^{+}}\mathcal{G}(X, Y)$ iff there exists a morphism $X \to Y$.

Proof. Statement (1) is [10, Proposition 6.18]. Statement (2) is [10, Theorem 6.12]. Statement (3) is [10, Proposition 6.15]. □

These abstract games are the bridge for relating morphisms in $\mathcal{C}$ to the concrete model comparison games used in logic. This will be illustrated through the examples in Section 3.2. Further, by providing an operational definition of various morphisms in $\mathcal{C}$, these games allow for clean proofs for the implications between the different types of morphisms. This is illustrated in the proofs of the following propositions from [9, 10].

Proposition 3.1.14. *Let $\mathcal{C}$ be an arboreal category and $X, Y, Z \in \mathbf{Ob}(\mathcal{C})$ such that product of each pair exists in $\mathbf{Ob}(\mathcal{C})$. If Duplicator has a winning strategy in $\mathcal{G}(X, Y)$ and winning strategy in $\mathcal{G}(Y, Z)$ over $\mathcal{C}$, then Duplicator has a winning strategy in $\mathcal{G}(X, Z)$.*

Proof. This is [10, Corollary 6.13]. □

Proposition 3.1.15. *If $X \leftarrow Z \rightarrow Y$ is a bisimulation in $\mathcal{C}$, then there exists pathwise embeddings (and thus, morphisms) $X \to Y$ and $Y \to Z$.*

Proof. This is [10, Lemma 6.20]. □

Finally, we make the following observation that isomorphic objects in $\mathcal{C}$ are bisimilar.

Proposition 3.1.16. *If $X \cong Y$ in $\mathcal{C}$, then there exists a bisimulation $X \leftarrow Z \rightarrow Y$ in $\mathcal{C}$.*

Proof. If $f\colon X \to Y$ is an isomorphism, then we set $Z = X$. The bisimulation is witnessed by the span $X \xleftarrow{\mathsf{id}_X} X \xrightarrow{f}$ Y. It is straightforward to verify that any isomorphism, including id_X and f, is an open pathwise embedding. □

3.1.3 Arboreal covers and behavioural equivalences

As the previous section demonstrated, because objects in arboreal category $\mathcal{C}$ are tree-shaped and generated by their branches, the category $\mathcal{C}$ is a native setting for formulating dynamics notions, such as bisimulation, resource-indexing, and abstract path-matching games. In order to apply these notions to relational structures and other 'static' objects for any category $\mathcal{E}$, we use a comonadic adjunction $L \dashv R$ from $\mathcal{E}$ to $\mathcal{C}$. The comonadicity of the adjunctions allows us to view $R(X)$ as tree-shaped cover of $X \in \mathbf{Ob}(\mathcal{E})$ where the branches of $R(X)$ represent a process which searches through or builds-up portions of X.

Definition 3.1.17. An *arboreal cover of* $\mathcal{E}$ *by* $\mathcal{C}$ is a comonadic adjunction

$$\begin{tikzcd} \mathcal{C} \arrow[r, bend left, "L"] \arrow[r, phantom, "\perp"] & \mathcal{E} \arrow[l, bend left, "R"] \end{tikzcd} .$$

such that $\mathcal{C}$ is an arboreal category.

The following definition extends the notion of arboreal cover to resource-indexed arboreal categories.

Definition 3.1.18. An *resource-indexed arboreal cover of* $\mathcal{E}$ *by* $\{\mathcal{C}_k\}$ is family of comonadic adjunctions $L_k \dashv R_k$ for every $k > 0$ where $\{\mathcal{C}_k\}$ is a resource-indexed arboreal category, $L_k \colon \mathcal{C}_k \to \mathcal{E}$ and $R_k \colon \mathcal{E} \to \mathcal{C}_k$.

Definition 3.1.19. Consider a resource-indexed arboreal cover of $\mathcal{E}$ by $\{\mathcal{C}_k\}$, and two objects A, B of $\mathcal{E}$. For all $k > 0$, we can extend the preorders $\to_k^{\mathcal{C}}, \rightharpoonup_k^{\mathcal{C}}, \leftrightarrow_k^{\mathcal{C}}, \cong_k^{\mathcal{C}}$ on objects of $\mathcal{C}_k$ to apply to objects of $\mathcal{E}$ by using the functor R_k:

- $A \to_k^{\mathcal{C}} B$ if there exists a morphism $R_k(A) \to R_k(B)$ in $\mathcal{C}_k$.
- $A \rightharpoonup_k^{\mathcal{C}} B$ if there is a pathwise embedding $R_k(A) \to R_k(B)$ in $\mathcal{C}_k$.
- $A \leftrightarrow_k^{\mathcal{C}} B$ if there is a bisimulation $R_k(A) \leftarrow Z \to R_k(B)$ in $\mathcal{C}_k$.
- $A \cong_k^{\mathcal{C}} B$ if there is an isomorphism $R_k(A) \cong R_k(B)$ in $\mathcal{C}_k$.

A resource-indexed arboreal cover of $\mathcal{E}$ by $\{\mathcal{C}_k\}$ yields a family of comonads $\mathbb{C}_k \colon \mathcal{E} \to \mathcal{E}$ on $\mathcal{E}$ related by inclusions $\iota_{k \leq l} \colon \mathbb{C}_k \to \mathbb{C}_l$. Since an arboreal cover as a comonadic adjunction, for every $k > 0$, $\mathcal{C}_k$ is isomorphic to the category of coalgebras $\mathbf{EM}(\mathbb{C}_k)$ of the comonad $\mathbb{C}_k$.

Definition 3.1.20. Given an $A \in \mathbf{Ob}(\mathcal{E})$, and a resource-indexed arboreal cover $L_k \dashv R_k$ of $\mathcal{E}$ by $\{\mathcal{C}_k\}$ yielding the family of comonads $\{\mathbb{C}_k \colon \mathcal{E} \to \mathcal{E}\}$, the $\mathbb{C}$-*coalgebra number of* A, denoted $\chi^{\mathbb{C}}(A)$ is the least k such that there exists a $\mathbb{C}_k$-coalgebra $A \to \mathbb{C}_k(A)$.

Since $L_k \dashv R_k$ is a comonadic adjunction, $\mathcal{C}_k$ is isomorphic to the category $\mathbf{EM}(\mathbb{C}_k)$, so the $\mathbb{C}$-coalgebra number could also defined as the least k such there there exists an $X \in \mathbf{Ob}(\mathcal{C}_k)$ where $L_k(X) = A$.

3.2 Example arboreal covers

Every example arboreal category in this section related to game comonads is a subcategory or minor variant of the factorised category $\mathcal{R}(\sigma)$ of forest-ordered σ-structures.

Example 3.2.1. Given a signature σ, we define a category of *forest-ordered σ-structures* $\mathcal{R}(\sigma)$. In this category, objects are pairs $(\mathcal{A}, \leq_A)$ where $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma))$ and $(A, \leq_A) \in \mathbf{Ob}(\mathcal{F})$. A morphism $h\colon (\mathcal{A}, \leq_A) \to (\mathcal{B}, \leq_B)$ in this category is both a σ-morphism $h \in \mathsf{Struct}(\sigma)(\mathcal{A}, \mathcal{B})$ of the underlying σ-structures and a morphism of forests $h \in \mathcal{F}((A, \leq_A), (B, \leq_B))$. The category $\mathcal{R}(\sigma)$ is equipped with a factorisation system where

$$\mathcal{Q} = \{q\colon (\mathcal{A}, \leq_A) \to (\mathcal{B}, \leq_B) \mid q\colon A \to B \text{ surjective}\},$$
$$\mathcal{M} = \{m\colon (\mathcal{A}, \leq_A) \to (\mathcal{B}, \leq_B) \mid m\colon \mathcal{A} \to \mathcal{B} \text{ } \sigma\text{-embedding}\}.$$

The tuple $(\mathcal{R}(\sigma), \mathcal{Q}, \mathcal{M})$ is a factorised category as in Definition 2.3.8.

We also define a category $\mathcal{R}_\star(\sigma)$ which is the tree-analogue of $\mathcal{R}(\sigma)$. In this category, objects are triples $(\mathcal{A}, a_0, \leq)$ where $(\mathcal{A}, a_0) \in \mathbf{Ob}(\mathsf{Struct}_\star(\sigma))$, and $(A, \leq) \in \mathbf{Ob}(\mathcal{T})$ with root a_0. A morphism $h\colon (\mathcal{A}, a_0, \leq^A) \to (\mathcal{B}, b_0, \leq^B)$ is the same as in $\mathcal{R}(\sigma)$, but where root preservation implies that $h\colon (\mathcal{A}, a_0) \to (\mathcal{B}, b_0)$ is a morphism in $\mathsf{Struct}_\star(\sigma)$. Similarly to $\mathcal{R}(\sigma)$, $(\mathcal{R}_\star(\sigma), \mathcal{Q}_*, \mathcal{M}_*)$ is a factorised category where $\mathcal{Q}_*, \mathcal{M}_*$ are $\mathcal{Q}, \mathcal{M}$ restricted to functions which preserve the distinguished point.

Proposition 3.2.2. $(\mathcal{R}(\sigma), \mathcal{Q}, \mathcal{M})$ *and* $(\mathcal{R}_\star(\sigma), \mathcal{Q}_*, \mathcal{M}_*)$ *are path categories.*

Proof. The proof follows essentially the same lines as the proof of Proposition 3.1.7. The key difference is that $\mathcal{Q}$ is the class of surjective *σ-morphisms* and $\mathcal{M}$ is the class of σ-embeddings □

3.2.1 Prefix list

As a warm-up to the game comonad examples, we first exhibit the non-empty prefix list comonad, see e.g. [72], as arising from the arboreal cover of **Set** by the category of forests $\mathcal{F}$ from Example 3.0.1.

Proposition 3.2.3. *$\mathcal{F}$ is an arboreal category.*

Proof. By Proposition 3.1.7, $\mathcal{F}$ is a path category. Suppose $(X, \leq) \in \mathbf{Ob}(\mathcal{F})$, then each down set $\downarrow x$ is a path subobject of $(X, \leq)$. Observe that every forest $(X, \leq)$ is colimit cocone of consisting of inclusions from down sets $\downarrow x \hookrightarrow X$ for every $x \in X$ and inclusions between them, i.e. for $x \leq x'$, $\downarrow x \hookrightarrow \downarrow x'$. □

We derive the non-empty prefix list comonad $N^+\colon \mathbf{Set} \to \mathbf{Set}$ from an arboreal cover of $\mathbf{Set}$ by $\mathcal{F}$. As part of this derivation, we first define the object mapping $N^+\colon \mathbf{Ob}(\mathbf{Set}) \to \mathbf{Ob}(\mathbf{Set})$ which sends a set X to all non-empty finite sequences of elements in X, i.e.

$$N^+(X) := \{[x_1, \ldots, x_n] \mid \forall n \in \mathbb{N}, x_1, \ldots, x_n \in X\}.$$

For every $X \in \mathbf{Ob}(\mathbf{Set})$, there is a function $\varepsilon_X\colon N^+(X) \to X$ which maps a sequence $[x_1, \ldots, x_n]$ to its last element x_n. This mapping decomposes as $N^+ = L \circ R$ where L is the object mapping $(X, \leq) \mapsto X$ of the forgetful functor $L\colon \mathcal{F} \to \mathbf{Set}$ and R is the mapping $X \mapsto (N^+(X), \sqsubseteq)$. In the following proposition, we show that for every $Y \in \mathbf{Ob}(\mathbf{Set})$, $R(Y)$ and ε_Y is a universal morphism from L to Y.

Proposition 3.2.4. *$L \dashv R$ is an adjunction from $\mathcal{F}$ to* $\mathbf{Set}$*.*

Proof. Suppose $f\colon L((X, \leq)) \to Y \in \mathbf{Mor}(\mathbf{Set})$, we need to construct a unique forest morphism $g\colon (X, \leq) \to R(Y)$ such that $\varepsilon_X \circ L(g) = f$. Since $(X, \leq)$ is a tree order, for every $x \in X$, we can present $\downarrow x$ as a finite chain $x_1 \prec \cdots \prec x_n = x$ and define $g(x) = [f(x_1), \ldots, f(x_n)]$. The function g is a forest morphism as

$$\begin{aligned} x_1 \prec \cdots \prec x_n &\Rightarrow [f(x_1)] \sqsubset \cdots \sqsubset [f(x_1), \ldots, f(x_n)] \\ &\Rightarrow g(x_1) \sqsubset \cdots \sqsubset g(x_n). \end{aligned}$$

where we use $\sqsubset$ to denote the covering relation of $\sqsubseteq$. As the counit ε_X returns the last element of the sequence, for every $x \in X$, $(\varepsilon_X \circ L(g))(x) = f(x_n) = f(x)$.

To demonstrate uniqueness, suppose $g'\colon (X, \leq) \to R(Y)$ such that $\varepsilon_Y \circ L(g') = f$. For $x \in X$ such that $\downarrow x$ is the chain $x_1 \prec \cdots \prec x_n = x$, we know that $g'(x_i)$ is the i-th prefix of $g'(x)$ by g' being a forest morphism. Moreover, by the equation $\varepsilon_X \circ L(g') = f$, the last element of $g'(x_i)$ is $f(x_i)$. Since the i-th element of $g'(x)$ is the last element of its i-th prefix $g'(x_i)$, $g'(x) = [f(x_1), \ldots, f(x_n)] = g(x)$. $\square$

Thus, from the standard properties of adjunctions, we can compute that:

- The functor object mapping R extends to functor $R\colon \mathbf{Set} \to \mathcal{F}$ which maps a function $f\colon X \to Y$ to the unique morphism $R(f) = R_f$ satisfying $\varepsilon_Y \circ L(R_f) = f \circ \varepsilon_X$ defined as $R_f[x_1, \ldots, x_n] = [f(x_1), \ldots, f(x_n)]$.
- The counit maps ε_X of the universal morphism form a natural transformation and is the counit of the adjunction $L \dashv R$.
- The unit is a natural transformation $\eta\colon \mathsf{Id}_{\mathcal{F}} \to R \circ L$ where the component for $(X, \leq)$ maps an element $x \in X$ where $\downarrow x$ is the finite chain $x_1 \prec \cdots \prec x_n = x$ to the sequence $[x_1, \ldots, x_n]$.

We can conclude $N^+ \colon \mathbf{Set} \to \mathbf{Set}$ is a functor and obtain our derivation of the non-empty of prefix list comonad $(N^+, \varepsilon, \delta)$ where the comultiplication has components $\delta_X = L(\eta_{R(X)})$. Explicitly, for every sequence $s = [x_1, \dots, x_n] \in N^+(X)$,

$$\delta_X(s) = [[x_1], [x_1, x_2], \dots, [x_1, \dots, x_n]].$$

However, in order to show that $L \dashv R$ is an arboreal cover of $\mathbf{Set}$ by $\mathcal{F}$, we must strengthen Proposition 3.2.4 by demonstrating that $L \dashv R$ is a comonadic adjunction. This amounts to showing that that the comparison functor $K_! \colon \mathcal{F} \to \mathbf{EM}(N^+)$ is an isomorphism. In particular, $K_!$ induces a bijective correspondence on objects yielding the following concrete characterisation of N^+-coalgebras:

Proposition 3.2.5. *There is a bijective correspondence between:*

(1) N^+*-coalgebras* $\chi \colon X \to N^+(X)$

(2) Forests $(X, \leq)$, *i.e.* $(X, \leq) \in \mathbf{Ob}(\mathcal{F})$

Proof. Abstractly the definition of $K_!$ is that for all $F \in \mathbf{Ob}(\mathcal{F})$, $K_!(F) = (LF, L\eta_F \colon LF \to LRL(F))$. Concretely, $K_!$ maps a forest $(X, \leq)$ to (X, χ) such that for each $x \in X$ where $\downarrow x$ is the chain $x_1 \prec \cdots \prec x_j = x$, $\chi(x) = L\eta(x) = [x_1, \dots, x_j]$.

Conversely, suppose that $(X, \chi \colon X \to N^+(X))$ is a coalgebra. For $x \in X$, let $\chi(x) = [x_1, \dots, x_j]$, then the comultiplication δ_X applied to $\chi(x)$ is $[[x_1], [x_1, x_2], \dots, [x_1, \dots, x_j]]$. Hence, equation (2.1) states that $\chi(x_i) = [x_1, \dots, x_i]$, for all $i \in [j]$. The equation (2.2) states that $x_j = x$. Thus $\chi \colon X \to N^+(X)$ is an injective function whose image is a prefix-closed subset of $N^+(X)$. This allows us to define a forest order $\leq$ where $x \leq x'$ iff $\chi(x) \sqsubseteq \chi(x')$ Let $K_!^{-1} \colon \mathbf{Ob}(\mathbf{EM}(N^+)) \to \mathbf{Ob}(\mathcal{F})$ be defined as $K_!^{-1}(X, \chi \colon X \to N^+(X)) = (X, \leq)$. By construction, $K_! \circ K_!^{-1}$ and $K_!^{-1} \circ K_!$ are the identity object mappings. ☐

This bijective correspondence extends to a isomorphism of categories $\mathcal{F} \cong \mathbf{EM}(N^+)$. Thus, demonstrating that $L \dashv R$ is a comonadic adjunction and an arboreal cover of $\mathbf{Set}$ by $\mathcal{F}$.

Proposition 3.2.6. *$L \dashv R$ is an arboreal cover of* $\mathbf{Set}$ *by* $\mathcal{F}$.

Proof. Proposition 3.2.5 demonstrates that $K_! \colon \mathbf{Ob}(\mathcal{F}) \to \mathbf{EM}(N^+)$ is a bijective object mapping, we must show that this bijection extends to morphisms and demonstrate that the functor $K_! \colon \mathcal{F} \to \mathbf{EM}(N^+)$ is an isomorphism.

Abstractly, the definition of $K_!$ is such that for all $f \in \mathcal{F}(X,Y)$, $K_!(f)$ is the commutative diagram

$$\begin{tikzcd}
L(X) \arrow[r, "L(\eta_X)"] \arrow[d, "L(f)"'] & LRLX \arrow[d, "L(\phi(f))"] \\
L(Y) \arrow[r, "L(\eta_Y)"] & LRLY
\end{tikzcd} \tag{3.1}$$

where $\phi(f) \in \mathbf{Set}(X, RLY)$ resulting from the bijection induced by the adjunction $L \dashv R$. Observe that for the covering chain $x_1 \prec \cdots \prec x_n$ ending in $x_n = x$ of forest X, $\phi(f)(x) = [f(x_1), \ldots, f(x_n)]$. Diagram (3.1), for our case of the adjunction from $\mathcal{F}$ to **Set**, expresses that as f preserves roots and the covering relation $\prec$, it follows that the covering chain $x_1 \prec \cdots \prec x_n$ maps to the covering chain $f(x_1) \prec \cdots \prec f(x_n)$.

Conversely, given a commutative diagram D expressing that $\upsilon \circ g = N^+(g) \circ \chi$ for some N^+-coalgebras $(X, \chi\colon X \to N^+(X))$ and $(Y, \upsilon\colon Y \to N^+(Y))$. We observe that D commuting means that g preserves covering chains from $K^{-1}(X,\chi)$ to $K^{-1}(Y,\upsilon)$. Thus, g preserves roots, the covering relation, and lifts to a morphism $K_!^{-1}(g)$ of $\mathcal{F}$. $\square$

3.2.2 Ehrenfeucht-Fraïssé

The category of forest covers $\mathcal{R}^E(\sigma)$ is the full subcategory of $\mathcal{R}(\sigma)$ whose objects $(\mathcal{A}, \leq) \in \mathbf{Ob}(\mathcal{R}^E(\sigma))$ satisfy the following compatibility condition:

(E) If $a, a' \in A$ are adjacent in $\mathcal{G}(\mathcal{A})$, then $a \leq a'$ or $a' \leq a$.

Proposition 3.2.7. *$\mathcal{R}^E(\sigma)$ is an arboreal category.*

Proof. Since the paths of $\mathcal{R}(\sigma)$ are the paths of the full subcategory $\mathcal{R}^E(\sigma)$, $\mathcal{R}^E(\sigma)$ is a path category. The path subobjects of an $(\mathcal{A}, \leq) \in \mathcal{R}^E(\sigma)$ are the induced σ-structures on the down sets $\downarrow a$ for every $a \in A$. Condition (E) ensures that $(\mathcal{A}, \leq)$ can be realised as the colimit cocone of these down sets. $\square$

Let $\{\mathcal{R}^E_k(\sigma)\}$ be the depth-indexing of $\mathcal{R}^E(\sigma)$. Concretely, $\mathcal{R}^E_k(\sigma)$ is the category of forest covers of height $\leq k$.

For every $k > 0$, we derive the k-round EF comonad $\mathbb{E}_k$ from [12, 13] via the resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}^E_k(\sigma)$. Intuitively, we think of the elements of $\mathbb{E}_k(\mathcal{A})$ as Spoiler plays in the k-round positive-existential EF game. Thus, the universe $\mathbb{E}_k(A)$ of σ-structure $\mathbb{E}_k(\mathcal{A})$ is the subset of $N^+(A)$ consisting of sequences of length $\leq k$. Just as with N^+, the counit $\varepsilon\colon \mathbb{E}_k \to \mathsf{Id}_{\mathsf{Struct}(\sigma)}$ returns the last element. For every m-ary relation $R \in \sigma$, we define the interpretation

$$\begin{aligned} R^{\mathbb{E}_k(\mathcal{A})}(s_1, \dots, s_m) \Leftrightarrow \forall i, j \in [m], s_i \sqsubseteq s_j \text{ or } s_j \sqsubseteq s_i \qquad & \text{(pairwise comparability)} \\ \text{and } R^{\mathcal{A}}(\varepsilon(s_1), \dots, \varepsilon(s_m)) \qquad & \text{(compatibility).} \end{aligned}$$

Note that the compatibility condition ensures that $\varepsilon_{\mathcal{A}} \colon \mathbb{E}_k(\mathcal{A}) \to \mathcal{A}$ is indeed a σ-morphism. Following a similar story to N^+, for every $k > 0$ we factor the object mapping $\mathbb{E}_k$ as $\mathbb{E}_k = L_k \circ R_k$ where L_k is the object mapping $(\mathcal{A}, \leq) \mapsto \mathcal{A}$ of the forgetful functor $L_k \colon \mathcal{R}^E_k(\sigma) \to \mathsf{Struct}(\sigma)$ and R_k is the mapping $\mathcal{A} \mapsto (\mathbb{E}_k(A), \sqsubseteq)$. In the following proposition, we show that for every $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma))$, $R_k(\mathcal{A})$ and $\varepsilon_{\mathcal{A}}$ is a universal morphism from L_k to $\mathcal{A}$.

Proposition 3.2.8. *For every $k > 0$, $L_k \dashv R_k$ is an adjunction from $\mathcal{R}^E_k(\sigma)$ to $\mathsf{Struct}(\sigma)$.*

Proof. This proof is similar to the proof of Proposition 3.2.4, but with two additional verifications: we must take into account the resource parameter k denoting height, and we must check that the function constructed preserves relations. Suppose that $f \colon L_k((\mathcal{A}, \leq)) \to \mathcal{B} \in \mathbf{Mor}(\mathsf{Struct}(\sigma))$, we need to construct a forest-ordered σ-morphism $g \colon (\mathcal{A}, \leq) \to R_k(\mathcal{B})$. Since $(\mathcal{A}, \leq)$ is a forest order, for every $a \in A$, we can present $\downarrow a$ as a finite chain $a_1 \prec \cdots \prec a_n = a$ and define $g(a) = [f(a_1), \dots, f(a_n)]$. As $(\mathcal{A}, \leq)$ has height $\leq k$, $n \leq k$, and so $g(a) \in \mathbb{E}_k(\mathcal{A})$. The proof of Proposition 3.2.4 demonstrates that g is a forest morphism. It remains to check that this definition of g preserves relational interpretations, suppose $R \in \sigma$ is a relation of arity r and that $R^{\mathcal{A}}(a_{i_1}, \dots, a_{i_r})$. By Condition (E), $a_{i_1}, \dots, a_{i_r}$ are pairwise comparable in the tree order $\leq$, so there is maximal element $a_* \in \{a_{i_1}, \dots, a_{i_r}\}$ such that $\{a_{i_1}, \dots, a_{i_r}\} \subseteq \downarrow a_*$. By the definition of g, every $g(a_{i_j})$ is a prefix of $g(a_*)$, thus the pairwise comparability condition of $R^{R_k(\mathcal{B})}$ is satisfied. Since f is σ-morphism and the last element of $g(a_{i_j})$ is $f(a_{i_j})$, the compatibility condition of $R^{R_k(\mathcal{B})}$ is satisfied. Thus, g is a forest-ordered σ-morphism. The proof that g is the unique morphism such that $\varepsilon_{\mathcal{A}} \circ L_k(g) = f$ is no different than analogous step in the proof of Proposition 3.2.4. ☐

We compute the data of the adjunction $L_k \dashv R_k$ and observe its similarity with the corresponding data for the adjunction $L \dashv R$ from Section 3.2.1 yielding the prefix list comonad N^+.

- The functor $R_k \colon \mathsf{Struct}(\sigma) \to \mathcal{R}^E_k(\sigma)$ maps a σ-morphism $f \colon \mathcal{A} \to \mathcal{B}$ to the unique forest-ordered σ-morphism $R_k(f)$ satisfying $\varepsilon_{\mathcal{B}} \circ L_k(R_k(f)) = f \circ \varepsilon_{\mathcal{A}}$ defined as $R_k(f)[a_1, \dots, a_n] = [f(a_1), \dots, f(a_n)]$. As the length of $R_k(f)(s)$ is equal to the length of s, we can conclude that $R_k(f)$ is morphism of forest order $(A, \leq)$ to $(B, \leq)$. By f being σ-morphism and $R_k(f)$ preserving the prefix relation, we have that the pairwise comparability and compatibility conditions for the relational interpretations $R^{R_k\mathcal{B}}$ are satisfied, so $R_k(f)$ is a σ-morphism.

- The counit maps $\varepsilon_{\mathcal{A}}$ of the universal morphism form a natural transformation and is the counit of the adjunction $L_k \dashv R_k$.
- The unit is a natural transformation $\eta\colon \mathsf{Id}_{\mathcal{R}^E_k(\sigma)} \to R_k \circ L_k$ where the component $(\mathcal{A}, \leq)$ maps an element $a \in A$ where $\downarrow a$ is the finite chain $a_1 \prec \cdots \prec a_n = a$ to the sequence $[a_1, \ldots, a_n]$. By construction, for $a_j \in \downarrow a$, $\eta(a_j)$ is the j-th prefix $[a_1, \ldots, a_j]$ of $[a_1, \ldots, a_n]$, so $\eta_{\mathcal{A}}$ is indeed a forest morphism. Condition (E) and a_j being the last element of $\eta(a_j)$ correspond to the pairwise comparability and compatibility conditions which ensure that η is a σ-morphism.

We can conclude $\mathbb{E}_k\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma)$ is a functor and obtain our derivation of k-round EF comonad $(\mathbb{E}_k, \varepsilon, \delta)$ of [12, 13] where the comultiplication has components $\delta_X = L(\eta_{R(X)})$. However, in order to show that $L_k \dashv R_k$ is a resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}^E_k(\sigma)$, we must strengthen Proposition 3.2.8 by demonstrating that $L \dashv R$ is a comonadic adjunction. This amounts to showing that that the comparison functor $K_!\colon \mathcal{R}^E_k(\sigma) \to \mathbf{EM}(\mathbb{E}_k)$ is an isomorphism. In particular, $K_!$ induces a bijective correspondence on objects yielding the following concrete characterisation of $\mathbb{E}_k$-coalgebras. As a corollary, we demonstrate that tree-depth is equivalence to the $\mathbb{E}$-coalgebra number.

Theorem 3.2.9. *There is a bijective correspondence between:*

(1) $\mathbb{E}_k$*-coalgebras* $\alpha\colon \mathcal{A} \to \mathbb{E}_k\mathcal{A}$

(2) Forest covers of height $\leq k$, $(\mathcal{A}, \leq) \in \mathbf{Ob}(\mathcal{R}^E_k(\sigma))$.

Proof. Concretely, the object mapping $K_!\colon \mathcal{R}^E_k(\sigma) \to \mathbf{EM}(\mathbb{E}_k)$ can be described similarly to the object mapping of the comparison functor in the proof of Proposition 3.2.5. Given a forest cover $(\mathcal{A}, \leq)$, for each $a \in A$, its predecessors form a chain $a_1 \prec \cdots \prec a_j$, with $a_j = a$, and $j \leq k$. We see that $K_!(\mathcal{A}, \leq) = (\mathcal{A}, \alpha)$ where $\alpha(a) = L_k(\eta_{(\mathcal{A},\leq)})(a) = [a_1, \ldots, a_j]$.

Conversely, suppose that $\alpha\colon \mathcal{A} \to \mathbb{E}_k\mathcal{A}$ is a coalgebra. For $a \in A$, let $\alpha(a) = [a_1, \ldots, a_j]$, then the comultiplication $\delta_{\mathcal{A}}$ applied to $\alpha(a)$ is $[[a_1], [a_1, a_2], \ldots, [a_1, \ldots, a_j]]$. Hence, equation (2.1) states that $\alpha(a_i) = [a_1, \ldots, a_i]$, for all $i \in [j]$. The equation (2.2) states that $a_j = a$. Thus $\alpha\colon A \to \mathbb{E}_k\mathcal{A}$ is an injective map whose image is a prefix-closed subset of $\mathbb{E}_k$. Defining $a \leq a'$ iff $\alpha(a) \sqsubseteq \alpha(a')$ yields a forest order on A, of height $\leq k$. If a, a' are adjacent in $\mathcal{G}(\mathcal{A})$, then for some $a_1, \ldots, a_n$ with $a = a_i$, $a' = a_j$, we have $R^{\mathcal{A}}(a_1, \ldots, a_n)$.

Since α is a homomorphism, we must have $R^{\mathbb{E}_k\mathcal{A}}(\alpha(a_1),\dots,\alpha(a_n))$. By the pairwise comparability condition in the definition of $R^{\mathbb{E}_k\mathcal{A}}$, $\alpha(a_i),\alpha(a_j)$ are prefix comparable, so a_i,a_j are $\leq$-comparable and condition (E) is satisfied. Thus $(A,\leq)$ is a forest cover of $\mathcal{G}(\mathcal{A})$, of height $\leq k$. We can define the inverse object mapping $K_!^{-1}\colon \mathbf{EM}(\mathbb{E}_k) \to \mathcal{R}_k^E(\sigma)$ as $K_!^{-1}(\mathcal{A},\alpha) = (\mathcal{A},\leq)$. □

The combinatorial parameter of tree-depth for a σ-structure $\mathcal{A}$, denoted $\mathsf{td}(\mathcal{A})$ is defined as the minimum height of a forest cover of $\mathcal{A}$. Thus, as an immediate corollary of Theorem 3.2.9, we can use the $\mathbb{E}$-coalgebra number to obtain a new definition of the tree-depth of a σ-structure.

Corollary 3.2.10. *For every finite σ-structure $\mathcal{A}$, $\mathsf{td}(\mathcal{A}) = \chi^{\mathbb{E}}(\mathcal{A})$.*

The bijective correspondence of Theorem 3.2.9 extends to a isomorphism of categories $\mathcal{R}_k^E(\sigma) \cong \mathbf{EM}(\mathbb{E}_k)$. Thus, demonstrating that $L_k \dashv R_k$ is a comonadic adjunction and an arboreal cover of $\mathsf{Struct}(\sigma)$ by $\mathbb{E}_k$.

Proposition 3.2.11. *$L_k \dashv R_k$ is a resource-indexed cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}_k^E(\sigma)\}$.*

Proof. Theorem 3.2.9 demonstrates that $K_!\colon \mathbf{Ob}(\mathcal{R}_k^E(\sigma)) \to \mathbf{EM}(\mathbb{E}_k)$ is a bijective object mapping, we must show that this bijection extends to morphisms and demonstrate that the functor $K_!\colon \mathcal{R}_k^E(\sigma) \to \mathbf{EM}(\mathbb{E}_k)$ is an isomorphism.

Concretely, the action of $K_!\colon \mathcal{R}_k^E(\sigma) \to \mathbf{EM}(\mathbb{E}_k)$ on morphisms can be described similarly to the the comparison functor in the proof of Proposition 3.2.6. Given a morphism $f\colon (\mathcal{A},\leq) \to (\mathcal{B},\leq') \in \mathbf{Mor}(\mathcal{R}_k^E(\sigma))$ of forest covers, $K_!(f)$ is the diagram D expressing the equation $\beta \circ L_k(f) = \mathbb{E}_k(L_k(f)) \circ \alpha$ where $K_!(\mathcal{A},\leq) = (\mathcal{A},\alpha)$ and $K_!(\mathcal{B},\leq') = (\mathcal{B},\beta)$. As with the proof of Proposition 3.2.6, it follows from f being a forest morphism that f preserves covering chains so diagram D commutes. We also additionally observe, for the arboreal cover of $L_k \dashv R_k$ of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}_k^E(\sigma)$, that $L(f)$ must σ-morphism as f is a forest-ordered σ-morphism. Moreover, from $f \in \mathbf{Mor}(\mathcal{R}_k^E(\sigma))$ acting on chains of height $\leq k$ it must be that case that D involves $\mathbb{E}_k(L(f))$ as $\mathbb{E}_k(L(f))$ acts on sequences of length $\leq k$.

Conversely, suppose we have commuting diagram D expressing $\beta \circ g = \mathbb{E}_k(g) \circ \alpha$ for coalgebras $(\mathcal{A},\alpha),(\mathcal{B},\beta) \in \mathbf{Ob}(\mathbf{EM}(\mathbb{E}_k))$ and σ-morphism $g\colon \mathcal{A} \to \mathcal{B}$. We observe that D commuting means that g preserves covering chains, i.e. $a_1 \prec \cdots \prec a_n \mapsto g(a_1) \prec \cdots \prec g(a_n)$, of height $n \leq k$ from $K^{-1}(\mathcal{A},\alpha)$ to $K^{-1}(\mathcal{B},\beta)$. Thus, since g preserves roots, the covering relation and is a σ-morphism, it lifts to a morphism $K_!^{-1}(g)$ of $\mathcal{R}_k^E(\sigma)$. □

We can now turn to the original motivation of defining $\mathcal{R}_k^E(\sigma)$ and show how the behavioural relations in Definition 3.1.19, instantiated in this case, captures equivalence in variants of $\mathbf{FO}_k$. We begin defining the concrete k-round EF game in order to show $\mathcal{R}_k^E(\sigma)$ can encode this game.

Definition 3.2.12. Given σ-structures $\mathcal{A}, \mathcal{B}$ and $k \geq 0$, we define the *back-and-forth k-round Ehrenfeucht-Fraïssé game* $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ *from $\mathcal{A}$ to $\mathcal{B}$*. During each round $n \leq k$ of the game,

- Spoiler chooses an element $a_n \in A$ or $b_n \in B$.
- Duplicator responds with an element $b_n \in B$ or $a_n \in A$ in the opposite structure.

Duplicator wins round n if the relation $\gamma_n := \{(a_j, b_j) \mid j \leq n\}$ is a tight σ-correspondence. Duplicator wins in $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ if for every round $n \leq k$ and possible Spoiler choice, Duplicator has a winning move.

The abstract back-and-forth game from Definition 3.1.12 has three variants which are mirrored by variants of the concrete game $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$:

- ${}^{\exists}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ is the game where in every round, Spoiler is forced to play an $a \in A$ and Duplicator is forced to respond $b \in B$.
- ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ has the same player restrictions as ${}^{\exists}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$, but additionally Duplicator's winning condition has been weakened. Duplicator can win in ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ if the relation γ is a σ-correspondence.
- ${}^{\#}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ is the bijective variant of the game. In this game, at every round $n \leq k$, Duplicator first chooses a bijection $\psi_n \colon A \to B$. If there is no such bijection, Spoiler wins automatically. Spoiler then chooses an $a \in A$ and Duplicator must respond with $\psi_n(a)$. The winning condition is the same as $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$.

Remark 3.2.13. Recall from the discussion in Section 2.2.1, by default, the logics we are investigating in this thesis are equality-free. Thus, the game $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ differs from the usual Ehrenfeucht-Fraïssé game in that the winning condition is weakened from partial isomorphism to tight correspondence. Since the relation in the winning condition does not preserve equality, we obtain that Duplicator winning strategy in our $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ actually characterises equivalence in equality-free $\mathbf{FO}_k$ [23]. Similarly, for the $\exists^+\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$, the winning condition has been weakened from partial homomorphism to correspondence, and thus captures equality-free $\exists^+\mathbf{FO}_k$. The ${}^{\exists}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ and ${}^{\#}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ correspond to equality-free $\exists\mathbf{FO}_k$ and $\#\mathbf{FO}_k$. Extending to logics with equality will be part of our investigation in Section 3.3.

With the definitions of the EF games in place, we can now proceed with the proofs of the morphism, pathwise, bisimulation, and isomorphism power theorems for the resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}^E_k(\sigma)\}$.

Theorem 3.2.14. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \Rrightarrow^{\exists^+\mathbf{FO}_k} \mathcal{B}$

(2) Duplicator has a winning strategy in ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$.

(3) Duplicator has a winning strategy in ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ *over* $\mathcal{R}^E(\sigma)$.

(4) $\mathcal{A} \to_k^{\mathcal{R}^E(\sigma)} \mathcal{B}$

Proof. For the (1) ⇔ (2) equivalence, [58] demonstrated a similar equivalence for the case where $\exists^+\mathbf{FO}_k$ has equality and the winning condition for Duplicator in ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ is strengthened so that the chosen elements form a partial σ-morphism. Eliminating equality in $\exists^+\mathbf{FO}_k$ corresponds to weakening Duplicator's winning condition from partial σ-morphism to σ-correspondence.

For (2) ⇔ (3) equivalence, the proof is inductive. The base case is trivial, since γ is the empty set for $k = 0$, γ is (vacuously) a σ-correspondence. In the game ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^E(\sigma)$, this corresponds to the initial position $(\perp_{R_k(\mathcal{A})}, \perp_{R_k(\mathcal{B})}) \in {}^{+}\mathcal{W}(R_k(\mathcal{A}), R_k(\mathcal{B}))$. For the inductive step, in round $j + 1$, observe that the history of Spoiler's moves in ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ for the previous j rounds, is a sequence $[a_1, \ldots, a_j] \in R_k(\mathcal{A})$ which is in bijective correspondence with $[u]$ where $\mathsf{dom}(u)$ is a chain of length j and the i-th element of $\mathsf{dom}(u)$ is mapped to a_j. The relations on $\mathsf{dom}(u)$ are such that u is an embedding. Spoiler playing a_{j+1} in the next round of ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ corresponds to Spoiler extending $[u] \prec [u'] \in \mathbf{P}(R_k(\mathcal{A}))$ for the corresponding round in ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ where $\mathsf{dom}(u')$ has one additional element r such that $u'(r) = s' = [a_1, \ldots, a_{j+1}] \in R_k(\mathcal{A})$. Since u' is an embedding, $L_k(\mathsf{dom}(u'))$ is isomorphic to the induced substructure $[\downarrow s']$ of $\mathbb{E}_k(\mathcal{A})$ consisting of prefixes of s'. Similarly, Duplicator's history of responses $[b_1, \ldots, b_j]$ in ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ corresponds to a $[v] \in R_k(\mathcal{B})$. Duplicator's response b_{j+1} to Spoiler playing a_{j+1} in ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ results in Duplicator responding to Spoiler playing $[u'] \succ [u]$ in ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ with $[v'] \succ [v] \in \mathbf{P}(R_k(\mathcal{B}))$ where the additional element r' of $\mathsf{dom}(v')$ is mapped to $v'(r') = t' = [b_1, \ldots, b_{j+1}] \in R_k(\mathcal{B})$. As with $[u']$, $L_k(\mathsf{dom}(v'))$ is isomorphic to the induced substructure $[\downarrow t']$ of $\mathbb{E}_k(\mathcal{B})$ consisting of prefixes of t'. The winning condition for Duplicator in ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ is that $\gamma = \{(a_i, b_i) \mid i \in [j + 1]\}$ is a σ-correspondence which is equivalent to the existence of a σ-morphism $[\downarrow s'] \to [\downarrow t']$, and therefore a $\mathcal{R}^E(\sigma)$-morphism $\mathsf{dom}(u') \to \mathsf{dom}(v') \in {}^{+}\mathcal{W}(R_k(\mathcal{A}), R_k(\mathcal{B}))$. This is precisely the winning condition for Duplicator in ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^E(\sigma)$.

For (3) ⇔ (4) equivalence, we observe that, from morphisms in $\mathcal{R}^E(\sigma)$ preserving covering chains, $\mathbf{P}\colon \mathcal{R}^E_k(\sigma) \to \mathcal{T}$ is faithful. Applying Proposition 3.1.13(3) to the case where $\mathcal{C} = \mathcal{R}^E(\sigma)$, $X = R_k(\mathcal{A})$, and $Y = R_k(\mathcal{B})$ yields the result. □

Theorem 3.2.15. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \Rrightarrow^{\exists\mathbf{FO}_k} \mathcal{B}$

(2) Duplicator has a winning strategy in $^{\exists}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$.

(3) Duplicator has a winning strategy in $^{\exists}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ *over* $\mathcal{R}^E(\sigma)$.

(4) $\mathcal{A} \rightharpoonup^{\mathcal{R}^E(\sigma)}_k \mathcal{B}$

Proof. For the (1) ⇔ (2) equivalence, the proof is similar to the (1) ⇔ (2) equivalence in Theorem 3.2.14 above. Including all literals, in particular negative atomics, in the logic $\exists\mathbf{FO}_k$ corresponds to the fact that the relational interpretations restricted to the all chosen elements at round n in the game $^{\exists}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ must be reflected and preserved, and so Duplicator's winning condition at round n is a *tight* σ-correspondence.

For the (2) ⇔ (3) equivalence, the proof is similar to the (2) ⇔ (3) equivalence in Theorem 3.2.14 above. The key difference is that Duplicator's winning condition in $^{\exists}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ is that $\gamma = \{(a_i, b_i) \mid i \in [j+1]\}$ is a *tight* σ-correspondence. This is equivalent to the existence of a σ-isomorphism $[\downarrow s'] \to [\downarrow t']$, and therefore a $\mathcal{R}^E(\sigma)$-isomorphism $\mathsf{dom}(u') \cong \mathsf{dom}(v') \in \mathcal{W}(R_k(\mathcal{A}), R_k(\mathcal{B}))$. This is precisely the winning condition for Duplicator in $^{\exists}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^E(\sigma)$.

For the (3) ⇔ (4) equivalence, since the right adjoint R_k preserves the product $\mathcal{A} \times \mathcal{B} \in \mathsf{Struct}(\sigma)$, $R_k(\mathcal{A}) \times R_k(\mathcal{B}) \in \mathcal{R}^E(\sigma)$. Applying Proposition 3.1.13(1) to the case where $\mathcal{C} = \mathcal{R}^E(\sigma)$, $X = R_k(\mathcal{A})$, and $Y = R_k(\mathcal{B})$ yields the result. □

Theorem 3.2.16. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \equiv^{\mathbf{FO}_k} \mathcal{B}$

(2) Duplicator has a winning strategy in $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$.

(3) Duplicator has a winning strategy in $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ *over* $\mathcal{R}^E(\sigma)$.

(4) $\mathcal{A} \leftrightarrow^{\mathcal{R}^E(\sigma)}_k \mathcal{B}$

Proof. For the (1) ⇔ (2) equivalence, the paper [23] demonstrated that equivalence in equality-free $\mathbf{FO}_k$ is characterised by the game $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ where the winning condition is tight σ-correspondence. This is the equality-free variant of the classic Ehrenfeucht-Fraïssé theorem [34, 36]. For a textbook account of the classic Ehrenfeucht-Fraïssé theorem see e.g. [59].

For the (2) $\Leftrightarrow$ (3), the proof is inductive. The base case is trivial, since γ is the empty set for $k = 0$, γ is (vacuously) a tight σ-correspondence. In the game $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^E(\sigma)$, this corresponds to the initial position $(\perp_{R_k(\mathcal{A})}, \perp_{R_k(\mathcal{B})}) \in \mathcal{W}(R_k(\mathcal{A}), R_k(\mathcal{B}))$. The inductive step splits into two cases. For the case where Spoiler chooses to play in $a_{j+1} \in \mathcal{A}$, in the $j + 1$ round, the argument is the same as in the equivalence (2) $\Leftrightarrow$ (3) of Theorem 3.2.15. If instead, Spoiler chooses to play $b_{j+1} \in \mathcal{B}$ in $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ for the round $j + 1$, then he extends the play from $[v'] \succ [v] \in \mathbf{P}(R_k(\mathcal{B}))$ for the corresponding round in $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ and Duplicator's response a_{j+1} means she extends the play $[u'] \succ [u] \in \mathbf{P}(R_k(\mathcal{A}))$ for the corresponding round in $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$. Duplicator's winning condition for $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ is equivalent to Duplicator's winning condition for $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^E(\sigma)$ as was demonstrated in Theorem 3.2.15.

For the (3) $\Leftrightarrow$ (4) equivalence, since the right adjoint R_k preserves the product $\mathcal{A} \times \mathcal{B} \in \mathsf{Struct}(\sigma)$, $R_k(\mathcal{A}) \times R_k(\mathcal{B}) \in \mathcal{R}^E_k(\sigma) \subseteq \mathcal{R}^E(\sigma)$. Applying Proposition 3.1.13(2) to the case where $\mathcal{C} = \mathcal{R}^E(\sigma)$, $X = R_k(\mathcal{A})$, and $Y = R_k(\mathcal{B})$ yields the result. $\square$

For the morphism, pathwise, and bisimulation cases we used the abstract 'path matching' games over $\mathcal{R}^E(\sigma)$. By contrast, in order to prove the isomorphism power theorem, we introduce the notion of branch bijective morphisms in $\mathsf{Struct}(\sigma)$. Recall that by the definition of adjunction in Definition 2.3.5, that morphisms $f^*\colon R_k(\mathcal{A}) \to R_k(\mathcal{B}) \in \mathbf{Mor}(\mathcal{R}^E(\sigma))$ are in bijection with morphisms $f\colon L_k(R_k(\mathcal{A})) \to \mathcal{B} \in \mathbf{Mor}(\mathsf{Struct}(\sigma))$. By $\mathbb{E}_k = L_k \circ R_k$, f can be regarded as a coKleisli morphism $f\colon \mathbb{E}_k(\mathcal{A}) \to \mathcal{B}$. For every morphism $f\colon \mathbb{E}_k(\mathcal{A}) \to \mathcal{B}$ and (possibly empty) sequence $s \in \mathbb{E}_k(\mathcal{A}) \cup \{\epsilon\}$, we can define the *$f$-induced* **EF**-*branch function* $\psi_s\colon A \to B$ *at* s as $\psi_s(a) = f(s[a])$. Given a pair of sequences of $(s, t) \in R_k(\mathcal{A}) \times R_k(\mathcal{B})$ such that $|s| = |t|$, we define the *induced* **EF**-*position* $\Gamma(s, t)$ *of* (s, t) as

$$\Gamma(s, t) = \{(a_i, b_i) \mid s = [a_1, \ldots, a_n], t = [b_1, \ldots, b_n]\}.$$

Definition 3.2.17. A morphism $f\colon \mathbb{E}_k(\mathcal{A}) \to \mathcal{B} \in \mathbf{Mor}(\mathsf{Struct}(\sigma))$ is **EF**-*branch bijective* if for every $s \in \mathbb{E}_k(\mathcal{A}) \cup \{\epsilon\}$, the f-induced **EF**-branch function $\psi_s\colon A \to B$ at s is bijective and if s is non-empty, the induced position $\Gamma(s, f^*(s))$ is a tight σ-correspondence.

Theorem 3.2.18. *The following are equivalent for all finite σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \equiv^{\#\mathbf{FO}_k} \mathcal{B}$

(2) Duplicator has a winning strategy in $^{\#}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$*.*

(3) There exists a **EF***-branch bijective* $f\colon \mathbb{E}_k(\mathcal{A}) \to \mathcal{B}$*.*

(4) $\mathcal{A} \cong_k^{\mathcal{R}^E(\sigma)} \mathcal{B}$

Proof. For the (1) ⇔ (2) equivalence, [44] demonstrated a similar equivalence for the case where $\#\mathbf{FO}_k$ has equality and the winning condition for Duplicator in ${}^{\#}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ is strengthened so that the chosen elements form a partial σ-isomorphism. Eliminating equality in $\#\mathbf{FO}_k$ corresponds to weakening Duplicator's winning condition from partial σ-isomorphism to tight σ-correspondence.

For the (2) ⇔ (3) equivalence, this is immediate from definitions. At round $n \geq 0$ of ${}^{\#}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ with previous Spoiler's previous moves $s = [a_1, \ldots, a_n] \in \mathbb{E}_k(\mathcal{A}) \cup \{\epsilon\}$, Duplicator plays bijection $\psi_s\colon A \to B$. Conversely, we can define ψ_s as Duplicator response ψ_n in the n-th round of ${}^{\#}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ and then define $f\colon \mathbb{E}_k(\mathcal{A}) \to \mathcal{B}$ as $f(s[a]) = \psi_s(a)$. The winning condition of ${}^{\#}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ is exactly that $\gamma_n = \Gamma(s, f^*(s))$ is a tight σ-correspondence.

For the (3) ⇔ (4) equivalence, suppose $f\colon L_k(R_k(\mathcal{A})) \to \mathcal{B}$ is branch bijective, then via the adjunction we obtain a morphism $f^*\colon R_k(\mathcal{A}) \to R_k(\mathcal{B})$. By induction on the length of sequences $s \in \mathbb{E}_k(\mathcal{A})$, we can conclude that f^* is a bijection, by the fact that ψ_s is bijective. We can conclude that f^* is a σ-isomorphism, by the fact that $\Gamma(s, f^*(s))$ is a tight σ-correspondence. Conversely, given an isomorphism $f^*\colon R_k(\mathcal{A}) \to R_k(\mathcal{B})$ with inverse $g^*\colon R_k(\mathcal{B}) \to R_k(\mathcal{A})$, via the adjunction, we can produce $f\colon L_k(R_k(\mathcal{A})) \to \mathcal{B}$ and observe that this satisfies the conditions the f-induced **EF**-branch functions $\psi_s\colon A \to B$ are bijective. Further, $\Gamma(s, t)$ with $t = f^*(s)$ and $g^*(t) = s$ is a tight σ-correspondence by the fact that f^*, g^* are $\mathcal{R}_k^E(\sigma)$-morphisms. □

In summary, the behavioural relations of Definition 3.1.19 for the resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}_k^E(\sigma)\}$ characterise equivalence for variants of first order logic $\mathbf{FO}_k$ graded by quantifier rank.

Theorem 3.2.19. *The following equivalences hold:*

(1) $\mathcal{A} \to_k^{\mathcal{R}^E(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \Rrightarrow^{\exists^+\mathbf{FO}_k} \mathcal{B}$

(2) $\mathcal{A} \rightharpoonup_k^{\mathcal{R}^E(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \Rrightarrow^{\exists\mathbf{FO}_k} \mathcal{B}$

(3) $\mathcal{A} \leftrightarrow_k^{\mathcal{R}^E(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \equiv^{\mathbf{FO}_k} \mathcal{B}$

(4) $\mathcal{A} \cong_k^{\mathcal{R}^E(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \equiv^{\#\mathbf{FO}_k} \mathcal{B}$ *(assuming* $\mathcal{A}, \mathcal{B}$ *are finite)*

3.2.3 Pebbling

The category of k-pebble forest covers $\mathcal{R}^P_k(\sigma)$ has objects which are triples $(\mathcal{A}, \leq, \mathbf{p}\colon A \to [k])$ where $(\mathcal{A}, \leq) \in \mathbf{Ob}(\mathcal{R}^E(\sigma))$ and $p\colon A \to [k]$ is a 'pebbling' function satisfying an additional 'active pebble' condition:

(P) If $a, a' \in A$ are adjacent in $\mathcal{G}(\mathcal{A})$ and $a \leq a'$, then for all $b \in (a, a']$, $\mathbf{p}(b) \neq \mathbf{p}(a)$.

Note that since $(\mathcal{A}, \leq) \in \mathbf{Ob}(\mathcal{R}^E(\sigma))$, objects in this category $\mathcal{R}^P_k(\sigma)$ satisfy condition (E). Morphisms in $\mathcal{R}^P_k(\sigma)$ are forest-ordered σ-morphisms which preserve the pebbling function.

Proposition 3.2.20. *$\{\mathcal{R}^P_k(\sigma)\}$ is a resource-indexed arboreal category.*

Proof. Observe that for every $k > 0$, the objects and morphisms, without the pebbling function, are the objects and morphisms of $\mathcal{R}^E(\sigma)$. Thus, as in Proposition 3.2.7, $\mathcal{R}^P_k(\sigma)$ is arboreal by Condition (E) ensuring the each object $(\mathcal{A}, \leq, \mathbf{p})$ is colimit cocone of the induced σ-substructure on the down sets $\downarrow a$ for every $a \in A$.

Since a pebbling function $\mathbf{p}\colon [j] \to A$ can be extended to a pebbling function $\mathbf{p}'\colon [k] \to A$ where $j \leq k$, there is a inclusion from the path subcategory of $\mathcal{R}^P_j(\sigma)$ to the path subcategory of $\mathcal{R}^P_k(\sigma)$. Equipping the collection $\{\mathcal{R}^P_k(\sigma)\}$ with these inclusions yields a resource-indexed arboreal category. $\square$

We derive the family of pebbling comonads $\mathbb{P}_k$ from the seminal paper [5] via the resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}^P_k(\sigma)\}$. For every $k > 0$, the universe $\mathbb{P}_k(A)$ of σ-structure $\mathbb{P}_k(\mathcal{A})$ is the set of non-empty 'pebbled' sequences of A. Explicitly,

$$\mathbb{P}_k(A) := \{[(p_1, a_1), \ldots, (p_n, a_n)] \mid \forall n \in \mathbb{N} \text{ and } \forall i \in [n], (p_i, a_i) \in [k] \times A\}$$

Similar to N^+, the counit $\varepsilon\colon \mathbb{P}_k \to \mathsf{Id}_{\mathsf{Struct}(\sigma)}$ maps a pebble sequence to the last pebbled element, i.e. $[(p_1, a_1), \ldots, (p_n, a_n)] \mapsto a_n$. We also need a last pebble function $\rho_{\mathcal{A}}$ where $[(p_1, a_1), \ldots, (p_n, a_n)] \mapsto p_n$. For every m-relation $R \in \sigma$, $(s_1, \ldots, s_m) \in R^{\mathbb{P}_k(\mathcal{A})}$ if and only if

(1) $\forall i, j \in [m], s_i \sqsubseteq s_j$ or $s_j \sqsubseteq s_i$ *(pairwise comparability)*

(2) if $s_i \sqsubseteq s_j$, then $\rho_{\mathcal{A}}(s_i)$ does not appear in $[\mathsf{suffix}(s_i, s_*)]$ *(active pebble)*

(3) $(\varepsilon_{\mathcal{A}}(s_1), \ldots, \varepsilon_{\mathcal{A}}(s_m)) \in R^{\mathbb{P}_k(\mathcal{A})}$ *(compatibility)*

where $s_* = \max\{s_1, \ldots, s_m\}$ with respect to the prefix order $\sqsubseteq$ and for all $i \in [m]$, $\mathsf{suffix}(s_i, s_*)$ is such that $s_* = s_i \mathsf{suffix}(s_i, s_*)$. As in the example of Section 3.2.2, for every $k > 0$ we factor the object mapping $\mathbb{P}_k$ as $\mathbb{P}_k = L_k \circ R_k$ where L_k is the object mapping

$(\mathcal{A}, \leq, \mathbf{p}) \mapsto \mathcal{A}$ of the forgetful functor $L_k \colon \mathcal{R}_k^P(\sigma) \to \mathsf{Struct}(\sigma)$ and R_k is the mapping $\mathcal{A} \mapsto (\mathbb{P}_k(\mathcal{A}), \sqsubseteq, \rho)$. In the following proposition, we show that for every $\mathcal{A} \in \mathsf{Struct}(\sigma)$, $R_k(\mathcal{A})$ and $\varepsilon_{\mathcal{A}}$ is a universal morphism from L_k to $\mathcal{A}$.

Proposition 3.2.21. *For every $k > 0$, $L_k \dashv R_k$ is an adjunction from $\mathcal{R}_k^P(\sigma)$ to $\mathsf{Struct}(\sigma)$.*

Proof. This proof is similar to the proof of Proposition 3.2.4, but we must construct a morphism g which preserves relations and the pebbling function. Suppose that

$$f \colon L_k((\mathcal{A}, \leq, \mathbf{p})) \to \mathcal{B} \in \mathbf{Mor}(\mathsf{Struct}(\sigma)),$$

we need to construct a forest-ordered σ-morphism $g \colon (\mathcal{A}, \leq, \mathbf{p}) \to R_k(\mathcal{B})$. Since $(\mathcal{A}, \leq)$ is a tree order, for every $a \in A$, we can present $\downarrow a$ as a finite chain $a_1 \prec \cdots \prec a_n = a$ and define $g(a) = [(\mathbf{p}(a_1), f(a_1)), \ldots, (\mathbf{p}(a_n), f(a_n))]$. As the domain of pebbling function $\mathbf{p}$ is $[k]$, we observe that $g(a) \in R_k(\mathcal{A})$. The proof of Proposition 3.2.4 demonstrates that g is a forest morphism. By definition, $\rho(g(a))$, where ρ is the pebbling function in tuple $(\mathbb{P}_k(\mathcal{B}), \sqsubseteq, \rho)$, is the pebble $\mathbf{p}(a_n) = \mathbf{p}(a)$ on the last element $f(a)$ of the sequence $g(a)$ so g preserves the pebbling function. It remains to check that this definition of g preserves relational interpretations, suppose $R \in \sigma$ is a relation of arity m and that $R^{\mathcal{A}}(a_{i_1}, \ldots, a_{i_m})$. By Condition (E), $a_{i_1}, \ldots, a_{i_m}$ are pairwise comparable in the tree order $\leq$, so there is maximal element $a_* \in \{a_{i_1}, \ldots, a_{i_m}\}$ such that $\{a_{i_1}, \ldots, a_{i_m}\} \subseteq \downarrow a_*$. By the definition of g, $g(a_{i_j})$ is a prefix of $g(a_*)$, thus the pairwise comparability condition of $R^{R_k(\mathcal{B})}$ is satisfied. By Condition (P), for all $j \in [m]$ such that $a_{i_j} \neq a_*$, $\mathbf{p}(a_{i_j}) \neq \mathbf{p}(a_*)$. Thus, since $\rho(g(a_{i_j})) = \mathbf{p}(a_{i_j})$ is the last pebble of some prefix of the sequence $g(a)$, the active pebble condition of $R^{R_k(\mathcal{B})}$ is satisfied. Since f is σ-morphism and the last pebbled element of $g(a_{i_j})$ is $f(a_{i_j})$, the compatibility condition of $R^{R_k(\mathcal{B})}$ is satisfied. Thus, g is indeed a morphism of $\mathcal{R}_k^P(\sigma)$. The proof that g is the unique morphism such that $\varepsilon_{\mathcal{A}} \circ L_k(g) = f$ is similar to the analogous step in the proof of Proposition 3.2.4. $\square$

We compute the data of the adjunction $L_k \dashv R_k$ and observe its similarity with the corresponding data for the adjunction $L_k \dashv R_k$ from Section 3.2.2 associated to $\mathbb{E}_k$.

- The functor $R_k \colon \mathsf{Struct}(\sigma) \to \mathcal{R}_k^P(\sigma)$ maps a σ-morphism $f \colon \mathcal{A} \to \mathcal{B}$ to the unique forest ordered σ-morphism $R_k(f)$ satisfying $\varepsilon_{\mathcal{B}} \circ L_k(R_k(f)) = f \circ \varepsilon_{\mathcal{A}}$ defined as $R_k(f)([(p_1, a_1), \ldots, (p_n, a_n)]) = [(p_1, f(a_1)), \ldots, (p_n, f(a_n))]$. As the length of $R_k(f)(s)$ is equal to the length of s, we can conclude that $R_k(f)$ is morphism of forest order $(A, \leq)$ to $(B, \leq)$. By $R_k(f)$ preserving the prefix relation, pebble at each index, and f being σ-morphism, we have that the pairwise comparability, active pebble, and compatibility conditions for the relational interpretations $R^{R_k \mathcal{B}}$ are satisfied, so $R_k(f)$ is a σ-morphism.

- The counit maps $\varepsilon_{\mathcal{A}}$ of the universal morphism form a natural transformation and is the counit of the adjunction $L_k \dashv R_k$.
- The unit is a natural transformation $\eta\colon \mathsf{Id}_{\mathcal{R}^P_k(\sigma)} \to R_k \circ L_k$ where the component for $(\mathcal{A}, \leq, \mathbf{p})$ maps an element $a \in A$ where $\downarrow a$ is the finite chain $a_1 \prec \cdots \prec a_n = a$ to the pebble sequence $[(\mathbf{p}(a_1), a_1), \ldots, (\mathbf{p}(a_n), a_n)]$. By construction, for $a_j \in \downarrow a$, $\eta(a_j)$ is the j-th prefix $[(\mathbf{p}(a_1), a_1), \ldots, (\mathbf{p}(a_j), a_j)]$ of $[(\mathbf{p}(a_1), a_1), \ldots, (\mathbf{p}(a_n), a_n)]$, so η is indeed a forest morphism. Condition (E), condition (P), and a_j being the last pebbled element of $\eta(a_j)$ correspond to the pairwise comparability, active pebble, and compatibility conditions which ensure that η is indeed a σ-morphism

We can conclude $\mathbb{P}_k\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma)$ is a functor and obtain our derivation of k-pebble comonad $(\mathbb{P}_k, \varepsilon, \delta)$ of [5] where the comultiplication has components $\delta_{\mathcal{A}} = L(\eta_{R(\mathcal{A})})$. Explicitly,

$$\delta_{\mathcal{A}}([(p_1, a_1), \ldots, (p_n, a_n)]) = [(p_1, s_1), \ldots, (p_n, s_n)].$$

where for all $i \in [n]$, $s_i = s[1, i] = [(p_1, a_1), \ldots, (p_n, a_i)]$. However, as in the previous example sections, we must strengthen Proposition 3.2.21 to show that $L_k \dashv R_k$ is a resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}^P_k(\sigma)$. This amounts to showing that that the comparison functor $K_!\colon \mathcal{R}^P_k(\sigma) \to \mathbf{EM}(\mathbb{P}_k)$ is an isomorphism. In particular, $K_!$ induces a bijective correspondence on objects yielding the following concrete characterisation of $\mathbb{P}_k$-coalgebras and as a corollary define treewidth in terms of the $\mathbb{P}$-coalgebra number.

Theorem 3.2.22. *There is a bijective correspondence:*

(1) $\mathbb{P}_k$-coalgebras $\alpha\colon \mathcal{A} \to \mathbb{P}_k\mathcal{A}$

(2) k-pebble forest covers $(\mathcal{A}, \leq, \mathbf{p}) \in \mathbf{Ob}(\mathcal{R}^P_k(\sigma))$.

Proof. Concretely, the object mapping $K_!\colon \mathbf{Ob}(\mathcal{R}^P_k(\sigma)) \to \mathbf{Ob}(\mathbf{EM}(\mathbb{P}_k))$ can be described similarly to the object mapping of the comparison functor in the proof of Proposition 3.2.5. Given a k-pebble forest cover $(\mathcal{A}, \leq, \mathbf{p})$, for each $a \in A$, its predecessors form a chain $a_1 \prec \cdots \prec a_j$, with $a_j = a$. We see that $K_!(\mathcal{A}, \leq, \mathbf{p}) = (\mathcal{A}, \alpha)$ where $\alpha(a) = L_k(\eta_{(\mathcal{A},\leq)})(a) = [(a_1, \mathbf{p}(a_1)), \ldots, (a_j, \mathbf{p}(a_j))]$.

Conversely, suppose that $\alpha\colon \mathcal{A} \to \mathbb{P}_k\mathcal{A} \in \mathbf{EM}(\mathbb{P}_k)$ is a coalgebra. For $a \in A$, let $\alpha(a) = [(p_1, a_1), \ldots, (p_n, a_j)]$, then the comultiplication $\delta_{\mathcal{A}}$ applied to $\alpha(a)$ is $[(p_1, s_1), (p_2, s_2), \ldots, (p_j, s_j)]$ where $s_i = [(p_1, a_1), \ldots, (p_i, a_i)]$ for all $i \in [j]$. Hence, equation (2.1) states that $\alpha(a_i) = [a_1, \ldots, a_i]$, for all $i \in [j]$. The equation (2.2) states that $a_j = a$. Thus $\alpha\colon A \to \mathbb{P}_k\mathcal{A}$ is an injective map whose image is a prefix-closed subset of $\mathbb{P}_k$. Defining $a \leq a'$ iff

$\alpha(a) \sqsubseteq \alpha(a')$ yields a forest order on A. Defining $\mathbf{p}(a) = \rho(\alpha(a)) = p_n$ yields a pebbling function $\mathbf{p}\colon A \to [k]$. If a, a' are adjacent in $\mathcal{G}(\mathcal{A})$, then for some $a_1, \ldots, a_n$ with $a = a_i$, $a' = a_j$, we have $R^{\mathcal{A}}(a_1, \ldots, a_n)$. Since α is a homomorphism, we must have $R^{\mathbb{P}_k\mathcal{A}}(\alpha(a_1), \ldots, \alpha(a_n))$. By the pairwise comparability condition in the definition of $R^{\mathbb{P}_k\mathcal{A}}$, $\alpha(a_i), \alpha(a_j)$ are prefix comparable, so a_i, a_j are $\leq$-comparable and condition (E) is satisfied. By the active pebble condition in the definition of $R^{\mathbb{P}_k\mathcal{A}}$, $\mathbf{p}(\alpha(a_i)) \neq \mathbf{p}(\alpha(a_j))$. Thus $(A, \leq, \mathbf{p})$ is a k-pebble forest cover of $\mathcal{G}(\mathcal{A})$. We can define the inverse object mapping $K_!^{-1}\colon \mathbf{Ob}(\mathbf{EM}(\mathbb{P}_k)) \to \mathbf{Ob}(\mathcal{R}_k^P(\sigma))$ as $K_!^{-1}(\mathcal{A}, \alpha) = (\mathcal{A}, \leq, \mathbf{p})$. □

A k-pebble forest cover $(\mathcal{A}, \leq, \mathbf{p})$, equivalently a $\mathbb{P}_k$-coalgebra of type $\mathcal{A} \to \mathbb{P}_k\mathcal{A}$, encodes the data of of tree decomposition of width $< k$. Tree decompositions are a common way of defining the graph parameter treewidth. Treewidth is very ubiquitous parameter appearing throughout many fields within theoretical computer science. Intuitively, treewidth is a sparsity measure which captures how similar an arbitrary graph is to a tree. We now review the definition of tree decomposition of width $< k$ and treewidth in order to obtain its characterisation in terms of $\mathbb{P}_k$.

Definition 3.2.23. A tree decomposition of simple undirected graph $G = (V, \frown)$ is a tree $(T, \leq, l)$ with labelling function $l\colon T \to \mathcal{P}(V)$ that assigns a subset of vertices in G to each node in T satisfying the following conditions:

(TD1) For all $v \in V$, there exists some $x \in T$, $v \in l(x)$.

(TD2) If $v \frown v'$ in G, then for some $x \in T$, $\{v, v'\} \subseteq l(x)$

(TD3) If $v \in l(x) \cap l(x')$, then for all $y \in T$ along the unique path from x to x', $v \in l(y)$.

The width $\mathsf{width}(T)$ of a tree decomposition is given by $\max_{x \in T} |l(x)| - 1$. The treewidth $\mathsf{tw}(\mathcal{A})$ of a σ-structure $\mathcal{A}$ is $\min_T \mathsf{width}(T)$ where T ranges over tree decompositions of $\mathcal{G}(\mathcal{A})$. With this definition in place, we can show how a k-pebble forest cover $(\mathcal{A}, \leq, \mathbf{p})$ encodes the data of tree decomposition of $\mathcal{A}$ with width $< k$.

Proposition 3.2.24. *For all finite σ-structures $\mathcal{A}$, the following are equivalent*

(1) There exists a k-pebble forest cover $(\mathcal{A}, \leq, \mathbf{p}) \in \mathbf{Ob}(\mathcal{R}_k^P(\sigma))$

(2) $\mathcal{A}$ has a tree decomposition of width $< k$.

Proof. This is a consequence of [5, Theorem 6, Theorem 21], where the term k-traversal was used instead of k-pebble forest cover. A proof is also given directly in [13, Theorem 6.4]. □

Thus, as a corollary of Theorem 3.2.22 and Proposition 3.2.24, we can demonstrate that the $\mathbb{P}$-coalgebra number corresponds to treewidth.

Corollary 3.2.25. *For every finite σ-structure $\mathcal{A}$, $\chi^{\mathbb{P}}(\mathcal{A}) = \mathsf{tw}(\mathcal{A}) + 1$.*

The bijective correspondence of Theorem 3.2.22 extends to a isomorphism of categories $\mathcal{R}^P_k(\sigma) \cong \mathbf{EM}(\mathbb{P}_k)$. Thus, demonstrating that $L_k \dashv R_k$ is a comonadic adjunction and an arboreal cover of $\mathsf{Struct}(\sigma)$ by $\mathbb{P}_k$.

Proposition 3.2.26. *$L_k \dashv R_k$ is a resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}^P_k(\sigma)\}$.*

Proof. Theorem 3.2.22 demonstrates that $K_!\colon \mathbf{Ob}(\mathcal{R}^P_k(\sigma)) \to \mathbf{EM}(\mathbb{P}_k)$ is a bijective object mapping, we must show that this bijection extends to morphisms and demonstrate that the functor $K_!\colon \mathcal{R}^P_k(\sigma) \to \mathbf{EM}(\mathbb{P}_k)$ is an isomorphism.

Concretely, the action of $K_!\colon \mathcal{R}^P_k(\sigma) \to \mathbf{EM}(\mathbb{P}_k)$ on morphisms can be described similarly to the the comparison functor in the proof of Proposition 3.2.6. Given a morphism $f\colon (\mathcal{A}, \leq, \mathbf{p}) \to (\mathcal{B}, \leq', \mathbf{p}') \in \mathbf{Mor}(\mathcal{R}^P_k(\sigma))$ of k-pebble forest covers, $K_!(f)$ is the diagram D expressing the equation $\beta \circ L_k(f) = \mathbb{E}_k(L_k(f)) \circ \alpha$ where $K_!(\mathcal{A}, \leq, \mathbf{p}) = (\mathcal{A}, \alpha)$ and $K_!(\mathcal{B}, \leq', \mathbf{p}') = (\mathcal{B}, \beta)$. As with the proof of Proposition 3.2.6, it follows from f being a forest morphism that f preserves covering chains and diagram D commutes. We also additionally observe, for the arboreal cover of $L_k \dashv R_k$ of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}^P_k(\sigma)$, that $L_k(f)$ must be a σ-morphism as f is a forest-ordered σ-morphism. Moreover, from $f \in \mathbf{Mor}(\mathcal{R}^P_k(\sigma))$ preserving the pebbling function it must be the case that D involves $\mathbb{P}_k(L(f))$ as $\mathbb{P}_k(L(f))$ preserves the pebbles in a sequence $s \in \mathbb{P}_k(\mathcal{A})$, i.e. $p \in [k]$ being the i-th pebble of $s \in \mathbb{P}_k(\mathcal{A})$ implies p is the i-th pebble of $\mathbb{P}_k(L(f))(s)$.

Conversely, suppose we have commuting diagram D expressing $\beta \circ g = \mathbb{P}_k(g) \circ \alpha$ for coalgebras $(\mathcal{A}, \alpha), (\mathcal{B}, \beta) \in \mathbf{Ob}(\mathbf{EM}(\mathbb{P}_k))$ and σ-morphism $g\colon \mathcal{A} \to \mathcal{B}$. We observe that D commuting means that g preserves covering chains, i.e. $a_1 \prec \cdots \prec a_n \mapsto g(a_1) \prec \cdots \prec g(a_n)$ and pebbles, i.e $\rho(\alpha(a_i)) = \rho(\beta(g(a_i))$ for all $i \in [n]$, from $K^{-1}(\mathcal{A}, \alpha)$ to $K^{-1}(\mathcal{B}, \beta)$. Thus, since g preserves roots, the covering relation, pebbling function, and is a σ-morphism, it lifts to a morphism $K_!^{-1}(g)$ of $\mathcal{R}^P_k(\sigma)$. $\square$

With the resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}^P_k(\sigma)$, we can demonstrate how the behavioural relations in Definition 3.1.19, instantiated in this case, capture equivalence in variants of $\mathcal{L}^k$. This is an axiomatic recovery of the some of the results in the seminal pebbling comonad paper [5]. We begin defining the concrete k-pebble game in order to show $\mathcal{R}^P_k(\sigma)$ can encode this game.

Definition 3.2.27. Given σ-structures $\mathcal{A}, \mathcal{B}$ and $k > 0$, we define the *back-and-forth k-pebble game* $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ *from $\mathcal{A}$ to $\mathcal{B}$*. In this game, Spoiler and Duplicator each have a set of $[k] = \{1, \ldots, k\}$ pebbles placed on the structures $\mathcal{A}$ and $\mathcal{B}$. During each round n of the game,

- Spoiler places a pebble $p_n \in [k]$ on either $a_n \in A$ or $b_n \in B$. If the pebble p_n was previously placed on an element, Spoiler moves the pebble to their newly chosen element.
- Duplicator responds by placing the corresponding pebble $p_n \in [k]$ on an element of the opposite structure $b_n \in B$ or $a_n \in A$.

This game played for n rounds yields sequences of pebble placements $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ and $t = [(p_1, b_1), \ldots, (p_n, a_n)]$. Duplicator wins round n if the relation

$$\gamma_n := \{(a^p, b^p) \mid a^p = \mathsf{last}_p(s), b^p = \mathsf{last}_p(t)\}$$

is a tight correspondence. Duplicator wins the game $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ if she has a winning move for every possible Spoiler choice forever, i.e. for every round $n \in \mathbb{N}$.

As with the $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$, the game $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ has three variants obtained using analogous modifications:

- ${}^{\exists}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is the game where in every round, Spoiler is forced to place pebbles on elements $a \in A$ and Duplicator is forced to place her corresponding pebbles on elements $b \in B$.
- ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ has the same player restrictions as ${}^{\exists}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, but additionally Duplicator's winning condition has been weakened. Duplicator can win in ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ if the relation γ is a correspondence.
- ${}^{\#}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is the bijective variant of the game. At round 0, the starting position is the pair of empty sequences (s_0, t_0). At each round $n + 1$ of the game, assume that (s_n, t_n) is the previous position consisting of sequences of pebble placements on A and B.
  - Spoiler first chooses a pebble $p_{n+1} \in [k]$.
  - Duplicator then chooses a bijection $\psi_{n+1} \colon A \to B$ such that for all $p \neq p_{n+1}$, $\psi_{n+1}(a^p) = \psi_n(a^p) = \mathsf{last}_p(t_n)$ where $a^p = \mathsf{last}_p(s_n)$. If there is no such bijection, Spoiler wins automatically.

- Spoiler then places pebble p_{n+1} on $a_{n+1} \in A$ and Duplicator must respond by placing pebble p_{n+1} on $\psi_{n+1}(a_{n+1})$. The sequences are updated as $s_{n+1} = s_n[(p_{n+1}, a_{n+1})]$ and $t_{n+1} = t_n[(p_{n+1}, \psi_{n+1}(a_{n+1}))]$.

The winning condition is the same as $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$.

With the definitions of the pebble games in place, we can now proceed with the proofs of the morphism, pathwise, bisimulation, and isomorphism power theorems for the resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}^P_k(\sigma)\}$.

Theorem 3.2.28. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \Rrightarrow^{\exists^+\mathcal{L}^k} \mathcal{B}$

(2) Duplicator has a winning strategy in ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$.

(3) Duplicator has a winning strategy in ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ *over* $\mathcal{R}^P_k(\sigma)$.

(4) $\mathcal{A} \rightarrow^{\mathcal{R}^P(\sigma)}_k \mathcal{B}$

Proof. For the (1) $\Leftrightarrow$ (2) equivalence, [58] demonstrated a similar equivalence for the case where $\exists^+\mathcal{L}^k$ has equality and the winning condition for Duplicator in ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is strengthened so that the pebbled elements form a partial σ-morphism. Eliminating equality in $\exists^+\mathcal{L}^k$ corresponds to weakening Duplicator's winning condition from partial σ-morphism to σ-correspondence.

For the (2) $\Leftrightarrow$ (3) equivalence, the proof is inductive. The base case is trivial, since γ is the empty set for $j = 0$, so γ is (vacuously) a σ-correspondence. In the game ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^P(\sigma)$, this corresponds to the initial position $(\bot_{R_k(\mathcal{A})}, \bot_{R_k(\mathcal{B})}) \in {}^{+}\mathcal{W}(R_k(\mathcal{A}), R_k(\mathcal{B}))$. For the inductive step, in round $j + 1$, observe that the history of Spoiler's moves in ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ for the previous j rounds, is a sequence $[(p_1, a_1), \ldots, (p_j, a_j)] \in R_k(\mathcal{A})$ which is in bijective correspondence with $[u]$ where $\mathsf{dom}(u)$ is a chain of length j, the pebbling function of $\mathsf{dom}(u)$ maps the i-th element to p_i, and the i-th element of $\mathsf{dom}(u)$ is mapped to a_j by u. The relations on $\mathsf{dom}(u)$ are such that u is an embedding. Spoiler placing pebble p_{j+1} on a_{j+1} in the next round of ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ corresponds to Spoiler extending $[u] \prec [u'] \in \mathbf{P}(R_k(\mathcal{A}))$ for the corresponding round in ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ where $\mathsf{dom}(u')$ has one additional element r such that $u'(r) = s' = [(p_1, a_1), \ldots, (p_{j+1}, a_{j+1})] \in R_k(\mathcal{A})$ and the pebbling function of $\mathsf{dom}(u')$ maps r to $p_{j+1} \in [k]$. Similarly, Duplicator's history of responses $[(p_1, b_1), \ldots, (p_j, b_j)]$ in ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ corresponds to a $[v] \in R_k(\mathcal{B})$. Duplicator's response of placing p_{j+1} on b_{j+1} to Spoiler placing p_{j+1} on a_{j+1} in ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ results in Duplicator responding to

Spoiler playing $[u'] \succ [u]$ in ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ with $[v'] \succ [v] \in \mathbf{P}(R_k(\mathcal{B}))$. The additional element r' of $\mathsf{dom}(v')$ is mapped to $v'(r') = t' = [(p_1, b_1), \ldots, (p_{j+1}, b_{j+1})] \in R_k(\mathcal{B})$. The winning condition for Duplicator in the $j+1$-round of ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is that

$$\gamma_{j+1} := \{(a^p, b^p) \mid a^p = \mathsf{last}_p(s'), b^p = \mathsf{last}_p(t')\}$$

is a σ-correspondence. This is equivalent the existence of a $\mathcal{R}_k^P(\sigma)$-morphism $\mathsf{dom}(u') \to \mathsf{dom}(v') \in {}^{+}\mathcal{W}(R_k(\mathcal{A}), R_k(\mathcal{B}))$.

For the (3) $\Leftrightarrow$ (4) equivalence, we observe that, from morphisms in $\mathcal{R}_k^P(\sigma)$ preserving covering chains, $\mathbf{P}\colon \mathcal{R}_k^P(\sigma) \to \mathcal{T}$ is faithful. Applying Proposition 3.1.13(3) to the case where $\mathcal{C} = \mathcal{R}_k^P(\sigma)$, $X = R_k(\mathcal{A})$, and $Y = R_k(\mathcal{B})$ yields the result. □

Theorem 3.2.29. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \Rrightarrow^{\exists\mathcal{L}^k} \mathcal{B}$

(2) Duplicator has a winning strategy in ${}^{\exists}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$.

(3) Duplicator has a winning strategy in ${}^{\exists}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ *over* $\mathcal{R}_k^P(\sigma)$.

(4) $\mathcal{A} \rightharpoonup_k^{\mathcal{R}^P(\sigma)} \mathcal{B}$

Proof. For the (1) $\Leftrightarrow$ (2) equivalence, the proof is similar to the (1) $\Leftrightarrow$ (2) equivalence in Theorem 3.2.28 above. Including all literals, in particular negative atomics, in the logic $\exists\mathcal{L}^k$ corresponds to the fact that the relational interpretations restricted to the pebbled elements at every round in the game ${}^{\exists}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ must be reflected and preserved, and so Duplicator's winning condition at every round is a *tight* σ-correspondence.

For the (2) $\Leftrightarrow$ (3) equivalence, the proof is similar to the (2) $\Leftrightarrow$ (3) equivalence in Theorem 3.2.28 above. The key difference is that Duplicator's winning condition in the $j+1$ round of ${}^{\exists}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is that

$$\gamma_{j+1} := \{(a^p, b^p) \mid a^p = \mathsf{last}_p(s'), b^p = \mathsf{last}_p(t')\}$$

is a *tight* σ-correspondence. This is equivalent to the existence a a $\mathcal{R}_k^P(\sigma)$-isomorphism $\mathsf{dom}(u') \cong \mathsf{dom}(v') \in \mathcal{W}(R_k(\mathcal{A}), R_k(\mathcal{B}))$. This is precisely the winning condition for Duplicator in ${}^{\exists}\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^P(\sigma)$.

For the (3) $\Leftrightarrow$ (4) equivalence, since the right adjoint R_k preserves the product $\mathcal{A} \times \mathcal{B} \in \mathsf{Struct}(\sigma)$, $R_k(\mathcal{A}) \times R_k(\mathcal{B}) \in \mathcal{R}_k^P(\sigma)$. Applying Proposition 3.1.13(1) to the case where $\mathcal{C} = \mathcal{R}_k^P(\sigma)$, $X = R_k(\mathcal{A})$, and $Y = R_k(\mathcal{B})$ yields the result. □

Theorem 3.2.30. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \equiv^{\mathcal{L}^k} \mathcal{B}$

(2) Duplicator has a winning strategy in $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$*.*

(3) Duplicator has a winning strategy in $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ *over* $\mathcal{R}^P_k(\sigma)$*.*

(4) $\mathcal{A} \leftrightarrow_k^{\mathcal{R}^P(\sigma)} \mathcal{B}$

Proof. For the (1) ⇔ (2) equivalence, [47] demonstrated a similar equivalence for the case where $\mathcal{L}^k$ has equality and the winning condition for Duplicator in $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is strengthened so that the pebbled elements form a partial σ-isomorphism. Eliminating equality in $\mathcal{L}^k$ corresponds to weakening Duplicator's winning condition from partial σ-isomorphism to tight σ-correspondence.

For the (2) ⇔ (3), the proof is inductive. The base case is trivial, since γ is the empty set for $j = 0$, so γ is (vacuously) a tight σ-correspondence. In the game $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^P(\sigma)$, this corresponds to the initial position $(\bot_{R_k(\mathcal{A})}, \bot_{R_k(\mathcal{B})}) \in \mathcal{W}(R_k(\mathcal{A}), R_k(\mathcal{B}))$. The inductive step splits into two cases. For the case where Spoiler chooses to pebble p_{j+1} on $a_{j+1} \in \mathcal{A}$, in the $j + 1$ round of $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, the argument is the same as in the equivalence (2) ⇔ (3) of Theorem 3.2.29. If instead, Spoiler chooses to place pebble p_{j+1} on $b_{j+1} \in \mathcal{B}$ in $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ for the round $j + 1$, then he extends the play from $[v'] \succ [v] \in \mathbf{P}(R_k(\mathcal{B}))$ for the corresponding round in $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ and Duplicator's response a_{j+1} means she extends the play $[u'] \succ [u] \in \mathbf{P}(R_k(\mathcal{A}))$ for the corresponding round in $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^P(\sigma)$. Duplicator's winning condition for $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is equivalent to Duplicator's winning condition for $\mathcal{G}(R_k(\mathcal{A}), R_k(\mathcal{B}))$ over $\mathcal{R}^P(\sigma)$ as was demonstrated in 3.2.29.

For (3) ⇔ (4) equivalence, since the right adjoint R_k preserves the product $\mathcal{A} \times \mathcal{B} \in \mathsf{Struct}(\sigma)$, $R_k(\mathcal{A}) \times R_k(\mathcal{B}) \in \mathcal{R}^P_k(\sigma)$. Applying Proposition 3.1.13(2) to the case where $\mathcal{C} = \mathcal{R}^P_k(\sigma)$, $X = R_k(\mathcal{A})$, and $Y = R_k(\mathcal{B})$ yields the result. □

Mirroring the EF case in Section 3.2.2, we will use the notion of **Peb**-branch bijective in order to obtain an isomorphism power theorem. By the definition of adjunction, morphisms $f^* \colon R_k(\mathcal{A}) \to R_k(\mathcal{B}) \in \mathbf{Mor}(\mathcal{R}^P(\sigma))$ are in bijection with morphisms $f \colon \mathbb{P}_k(\mathcal{A})) \to \mathcal{B} \in \mathbf{Mor}(\mathsf{Struct}(\sigma))$ where we note that $\mathbb{P}_k = L_k \circ R_k$. For every morphism $f \colon \mathbb{P}_k(\mathcal{A}) \to \mathcal{B}$, (possibly empty) sequence $s \in \mathbb{P}_k(\mathcal{A}) \cup \{\epsilon\}$, and pebble $p \in [k]$, we can define the *f-induced* **Peb**-*branch function* $\psi_{s,p} \colon A \to B$ *at* (s, p) as $\psi_{s,p}(a) = f(s[(p, a]))$. Given a pair of sequences of $(s, t) \in R_k(\mathcal{A}) \times R_k(\mathcal{B})$, which are of the same length and have the same pebble at each index, we define the *induced* **Peb**-*position* $\Gamma(s, t)$ *of* (s, t) as

$$\Gamma(s, t) = \{(a^p, b^p) \mid p \in [k], a^p = \mathsf{last}_p(s), b^p = \mathsf{last}_p(t)\}.$$

Definition 3.2.31. A morphism $f\colon \mathbb{P}_k(\mathcal{A}) \to \mathcal{B} \in \mathbf{Mor}(\mathsf{Struct}(\sigma))$ is **Peb**-*branch bijective* if for every $s \in \mathbb{P}_k(\mathcal{A}) \cup \{\epsilon\}$ and pebble $p \in [k]$, the f-induced **Peb**-branch function $\psi_{s,p}\colon A \to B$ at (s, p) is bijective and if s is non-empty, the induced position $\Gamma(s, f^*(s))$ is a tight σ-correspondence.

Theorem 3.2.32. *The following are equivalent for all finite σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \equiv^{\#\mathcal{L}^k} \mathcal{B}$

(2) Duplicator has a winning strategy in ${}^{\#}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$*.*

(3) There exists a **Peb**-*branch bijective* $f\colon \mathbb{P}_k(\mathcal{A}) \to \mathcal{B}$*.*

(4) $\mathcal{A} \cong_k^{\mathcal{R}^P(\sigma)} \mathcal{B}$

Proof. For the (1) ⇔ (2) equivalence, [44] demonstrated a similar equivalence for the case where $\#\mathcal{L}^k$ has equality and the winning condition for Duplicator in ${}^{\#}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is strengthened so that the chosen elements form a partial σ-isomorphism. Eliminating equality in $\#\mathcal{L}^k$ corresponds to weakening Duplicator's winning condition from partial σ-isomorphism to tight σ-correspondence.

For the (2) ⇔ (3) equivalence, this is immediate from definitions. At round $n \geq 0$ of ${}^{\#}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, suppose Spoiler's previous moves is the sequence $s = [(p_1, a_1), \ldots, (p_n, a_n)] \in \mathbb{P}_k(\mathcal{A}) \cup \{\epsilon\}$ and Spoiler plays pebble p, Duplicator plays bijection $\psi_{s,p}\colon A \to B$. Conversely, we can define $\psi_{s,p}$ as Duplicator response bijection ψ_n to Spoiler playing p in the n-th round of ${}^{\#}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ and then define $f\colon \mathbb{P}_k(\mathcal{A}) \to \mathcal{B}$ as $f(s[(p, a)]) = \psi_{s,p}(a)$. The winning condition of ${}^{\#}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is exactly that $\gamma_n = \Gamma(s, f^*(s))$ is a tight σ-correspondence.

For the (3) ⇔ (4) equivalence, suppose $f\colon L_k(R_k(\mathcal{A})) \to \mathcal{B}$ is branch bijective, then via the adjunction we obtain a morphism $f^*\colon R_k(\mathcal{A}) \to R_k(\mathcal{B})$. By induction on the length of sequences $s \in \mathbb{P}_k(\mathcal{A})$, we can conclude that f^* is a bijection, by the fact that $\psi_{s,p}$ is bijective We conclude that f^* is a σ-isomorphism, by the fact that $\Gamma(s, f^*(s))$ is a tight σ-correspondence. Conversely, given an isomorphism $f^*\colon R_k(\mathcal{A}) \to R_k(\mathcal{B})$ with inverse $g^*\colon R_k(\mathcal{B}) \to R_k(\mathcal{A})$, via the adjunction, we can produce $f\colon L_k(R_k(\mathcal{A})) \to \mathcal{B}$ and observe that this satisfies the condition that the f-induced **Peb**-branch functions $\psi_{s,p}\colon A \to B$ are bijective. Further, $\Gamma(s, t)$ with $t = f^*(s)$ and $g^*(t) = s$ is a tight σ-correspondence follows from the fact that f^*, g^* are $\mathcal{R}_k^P(\sigma)$-morphisms. □

In summary, the behavioural relations of Definition 3.1.19 for the resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}_k^P(\sigma)\}$ characterise equivalence for variants of finite variable logic $\mathcal{L}^k$.

Theorem 3.2.33. *The following equivalences hold:*

(1) $\mathcal{A} \rightarrow_k^{\mathcal{R}^P(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \Rrightarrow^{\exists^+\mathcal{L}^k} \mathcal{B}$

(2) $\mathcal{A} \rightharpoonup_k^{\mathcal{R}^P(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \Rrightarrow^{\exists\mathcal{L}^k} \mathcal{B}$

(3) $\mathcal{A} \leftrightarrow_k^{\mathcal{R}^P(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \equiv^{\mathcal{L}^k} \mathcal{B}$

(4) $\mathcal{A} \cong_k^{\mathcal{R}^P(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \equiv^{\#\mathcal{L}^k} \mathcal{B}$ *(assuming* $\mathcal{A}, \mathcal{B}$ *are finite)*

3.2.4 Modal

Throughout this example, we will work in the category of pointed σ-structures $\mathsf{Struct}_\star(\sigma)$ where σ is a modal signature. Recall that since σ is a modal signature, the set of binary relations $\{R_\alpha \mid \alpha \in \mathsf{Act}\}$ is indexed by the transition alphabet Act, and the set of unary relations P encode the truth of propositional variables in PV.

The category $\mathcal{R}^M(\sigma)$ defined in this example is the subcategory of a category of pointed tree-ordered structures $\mathcal{R}_\star(\sigma)$. Intuitively, this subcategory $\mathcal{R}^M(\sigma)$ can be seen a category of tree models for $\Diamond^+\mathcal{M}$ theories.

The subcategory $\mathcal{R}^M(\sigma)$ of modal tree-ordered σ-structures has objects $(\mathcal{A}, \leq, a_0)$ whose tree-order $\leq$ and transition relations $\{R_\alpha^\mathcal{A} \mid \alpha \in \mathsf{Act}\}$ satisfy the following compatibility condition:

(M) $a \prec a'$ if and only if there is a unique binary relation symbol $R_\alpha \in \sigma$ such that $R_\alpha^\mathcal{A}(a, a')$.

Proposition 3.2.34. $\mathcal{R}^M(\sigma)$ *is an arboreal category.*

Proof. Since the paths of $\mathcal{R}_\star(\sigma)$ are the paths of the full subcategory $\mathcal{R}^M(\sigma)$, $\mathcal{R}^M(\sigma)$ is a path category. The path subobjects of an $(\mathcal{A}, a_0, \leq) \in \mathcal{R}^M(\sigma)$ are the induced σ-structures on the down sets $\downarrow a = a_0 \prec \cdots \prec a$ for every $a \in A$. Condition (M) ensures that $(\mathcal{A}, a_0, \leq)$ can be realised as the colimit cocone of these down sets. ☐

Let $\{\mathcal{R}_k^M(\sigma)\}$ be the depth-indexing of $\mathcal{R}^M(\sigma)$. Concretely, $\mathcal{R}_k^M(\sigma)$ is the category of k-height modal tree-ordered σ-structures.

For every $k > 0$, we derive the k-round modal comonad $\mathbb{M}_k$ from [12, 13]. The universe $\mathbb{M}_k(A, a_0)$ of $\mathbb{M}_k(\mathcal{A}, a_0)$ consists sequences of transitions with length $\leq k$ in $(\mathcal{A}, a_0)$. Explicitly,

$$\mathbb{M}_k(\mathcal{A}, a_0) := \{a_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} a_n \mid n \in \mathbb{N}, R_{\alpha_i}^\mathcal{A}(a_{i-1}, a_i)\}.$$

with distinguished point given by the sequence a_0. Similar to N^+, the counit $\varepsilon\colon \mathbb{M}_k \to \mathsf{Id}_{\mathsf{Struct}_\star(\sigma)}$ maps a transition sequence to the last state, i.e. $a_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} a_n \mapsto a_n$. For every binary relation $R_\alpha \in \sigma$, we define the interpretation

$$R_\alpha^{\mathbb{M}_k(\mathcal{A},a_0)}(s,t) \Leftrightarrow t \text{ is the extension of } s = a_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_{n-1}} a_{n-1} \text{ by } a_{n-1} \xrightarrow{\alpha} a$$

For every unary relation $P \in \sigma$, $P^{\mathbb{M}_k(\mathcal{A},a_0)}(s) \Leftrightarrow P^{\mathcal{A}}(\varepsilon(a))$. As in the previous examples of Section 3.2.2 and 3.2.3, for every $k > 0$, we factor the object mapping $\mathbb{M}_k$ as $\mathbb{M}_k = L_k \circ R_k$ where L_k is the object mapping $(\mathcal{A}, a_0, \leq) \mapsto (\mathcal{A}, a_0)$ of the forgetful functor $L_k\colon \mathcal{R}_k^M(\sigma) \to \mathsf{Struct}_\star(\sigma)$ and R_k is the mapping $(\mathcal{A}, a_0) \mapsto (\mathbb{M}_k(\mathcal{A}, a_0), a_0, \sqsubseteq)$ where $\sqsubseteq$ is the prefix relation on transition sequences, i.e. $s \sqsubseteq a_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} a_n$ if for some $i \in (n)$, $s = a_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_i} a_i$. In the following proposition, we show that for every $(\mathcal{A}, a_0) \in \mathsf{Struct}_\star(\sigma)$, $R_k(\mathcal{A}, a_0)$ and $\varepsilon_{(\mathcal{A},a_0)}$ is a universal morphism from L_k to $(\mathcal{A}, a_0)$.

Proposition 3.2.35. *For every $k > 0$, $L_k \dashv R_k$ is an adjunction from $\mathcal{R}_k^M(\sigma)$ to $\mathsf{Struct}_\star(\sigma)$.*

Proof. Suppose that $f\colon L_k((\mathcal{A}, \leq, a_0)) \to (\mathcal{B}, b_0) \in \mathbf{Mor}(\mathsf{Struct}_\star(\sigma))$, we need to construct a forest-ordered σ-morphism $g\colon (\mathcal{A}, \leq, a_0) \to R_k((\mathcal{B}, b_0))$. Since $(\mathcal{A}, \leq, a_0)$ is a tree order, for every $a \in A$, we can present $\downarrow a$ as a finite chain $a_0 \prec a_1 \prec \cdots \prec a_n = a$. By Condition (M), for every $i \in [n]$, there exists a unique transition relation $R_{\alpha_i} \in \sigma$ such that $R_{\alpha_i}^{\mathcal{A}}(a_{i-1}, a_i)$. We define $g(a) = f(a_0) \xrightarrow{\alpha_1} f(a_1) \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} f(a_n)$. As $(\mathcal{A}, \leq, a_0)$ has height $\leq k$, $n \leq k$ and so $g(a) \in \mathbb{M}_k(\mathcal{A}, a_0)$. The proof of Proposition 3.2.4 and $g(a_0) = f(a_0) = b_0$ demonstrates that g is a tree morphism. It remains to check that this definition of g preserves relational interpretations. Suppose $R_\alpha \in \sigma$ is a binary transition relation and $R_\alpha^{\mathcal{A}}(a, a')$. By Condition (M), $a \prec a'$, $g(a')$ is the extension of $g(a)$ by $a \xrightarrow{\alpha} a'$, so $R_\alpha^{R_k(\mathcal{B},b_0)}(a, a')$. Suppose $P \in \sigma$ is a unary relation, then since $f(a)$ is the last element of the transition sequence $g(a)$ and f is a σ-morphism, if $P^{\mathcal{A}}(a)$, then $P^{R_k(\mathcal{B},b_0)}(g(a))$.

To demonstrate uniqueness, suppose $g'\colon (\mathcal{A}, \leq, a_0) \to R_k(\mathcal{B}, b_0)$ such that $\varepsilon_{(\mathcal{B},b_0)} \circ L_k(g') = f$. For $a \in A$ such that $\downarrow a$ is the chain $a_0 \prec a_1 \prec \cdots \prec a_n = a$, we know by Condition (M) that this chain uniquely determines the transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n = a$. Moreover, by g' being a forest morphism $g'(a_i)$ is the i-th prefix of $g'(a)$, and since g' is σ-morphism, for all $i \in [n]$, $R_{\alpha_i}(g(a_{i-1}), g(a_i))$. Thus, the i-th transition in the transition sequence $g(a)$ is α_i. Further, by the equation $\varepsilon_X \circ L(g') = f$, the last element of $g'(a_i)$, i.e. the i-th element of $g(a)$ is $f(a_i)$. Combining these facts, we get that $g'(a) = f(a_0) \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} f(a_n) = g(a)$. $\square$

We compute the data of the adjunction $L_k \dashv R_k$ and observe its similarity with the corresponding data for the adjunction $L_k \dashv R_k$ from Section 3.2.2 yielding the k-round EF comonad $\mathbb{E}_k$.

- The functor $R_k\colon \mathsf{Struct}_\star(\sigma) \to \mathcal{R}^M_k(\sigma)$ maps a pointed σ-morphism $f\colon (\mathcal{A}, a_0) \to (\mathcal{B}, b_0)$ to the unique tree-ordered σ-morphism $R_k(f)$ satisfying $\varepsilon_{(\mathcal{B},b_0)} \circ L_k(R_k(f)) = f \circ \varepsilon_{(\mathcal{A},a_0)}$ defined as $R_k(f)(a_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} a_n) = f(a_0) \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} f(a_n)$. As the length of $R_k(f)(s)$ is equal to the length of s, we can conclude that $R_k(f)$ is morphism of tree order $(A, a_0, \leq)$ to $(B, b_0, \leq)$. By $R_k(f)$ preserving α-successors and f being σ-morphism, we have that the interpretations for binary relations $R_\alpha^{R_k(\mathcal{B},b_0)}$ and unary relations $P^{R_k(\mathcal{B},b_0)}$ are satisfied, so $R_k(f)$ is a σ-morphism.
- The counit maps $\varepsilon_{\mathcal{A}}$ of the universal morphism form a natural transformation and is the counit of the adjunction $L_k \dashv R_k$.
- The unit is a natural transformation $\eta\colon \mathsf{Id}_{\mathcal{R}^M_k(\sigma)} \to R_k \circ L_k$. To compute the component $(\mathcal{A}, a_0, \leq)$ of η observe that for an element $a \in A$ where $\downarrow a$ is the finite chain $a_1 \prec \cdots \prec a_n = a$, by condition (M), there is a unique transition sequence $a_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} a_n$ which we define as $\eta(a)$. By construction, for $a_j \in \downarrow a$, $\eta(a_j)$ is the j-th prefix $a_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_j} a_j$ of $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n$, so η is indeed a forest morphism. Condition (M) ensures that η preserves the binary relations. The fact that a_j is the last element of the transition sequence $\eta(a_j)$ ensures that unary relations are preserved. Hence, η is a σ-morphism.

We can conclude $\mathbb{M}_k\colon \mathsf{Struct}_\star(\sigma) \to \mathsf{Struct}_\star(\sigma)$ is a functor and obtain our derivation of k-round modal comonad $(\mathbb{M}_k, \varepsilon, \delta)$ of [12, 13] where the comultiplication has components $\delta_{(\mathcal{A},a_0)} = L(\eta_{R(\mathcal{A},a_0)})$. Explicitly, for a transition sequence $s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n \in \mathbb{M}_k(\mathcal{A}, a_0)$,

$$\delta_X(s) = s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} s_n.$$

where $s_i = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_i} a_i$ for all $i \in (n)$. However, in order to show that $L_k \dashv R_k$ is a resource-indexed arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\mathcal{R}^M_k(\sigma)$, we must strengthen Proposition 3.2.35 by demonstrating that $L \dashv R$ is a comonadic adjunction. This amounts to showing that that the comparison functor $K_!\colon \mathcal{R}^M_k(\sigma) \to \mathbf{EM}(\mathbb{M}_k)$ is an isomorphism. As with the previous cases of the k-round EF comonad and the k-pebble comonad, we first demonstrate that the object mapping of $K_!$ is a bijective correspondence.

Theorem 3.2.36. *There is a bijective correspondence:*

(1) $\mathbb{M}_k$*-coalgebras* $\gamma\colon (\mathcal{A}, a_0) \to \mathbb{M}_k(\mathcal{A}, a_0)$

(2) k*-height modal tree* σ*-structure* $(\mathcal{A}, a_0, \leq) \in \mathbf{Ob}(\mathcal{R}^M_k(\sigma))$

Proof. Concretely, the object mapping $K_!\colon \mathcal{R}_k^M(\sigma) \to \mathbf{EM}(\mathbb{M}_k)$ can be described similarly to the object mapping of the comparison functor in the proof of Proposition 3.2.5. Given a k-height modal tree tree σ-structure $(\mathcal{A}, a_0, \leq)$, for each $a \in A$, its predecessors form a chain $a_1 \prec \cdots \prec a_j$, with $a_j = a$, and $j \leq k$. By condition (M), this chain corresponds uniquely to a transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_j} a_j$. We see that $K_!(\mathcal{A}, \leq) = (\mathcal{A}, \gamma)$ where $\alpha(a) = L_k(\eta_{(\mathcal{A},\leq)})(a) = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_j} a_j$.

Conversely, suppose that $\gamma\colon (\mathcal{A}, a_0) \to \mathbb{M}_k(\mathcal{A}, a_0)$ is a coalgebra. For $a \in A$, let $\gamma(a) = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_j} a_j$, then the comultiplication $\delta_{\mathcal{A}}$ applied to $\gamma(a)$ is $s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_j} s_j$ where $s_i = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_j} a_i$ for all $i \in (j)$. Hence, equation (2.1) states that $\gamma(a_i) = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_j} a_i$, for all $i \in (j)$. The equation (2.2) states that $a_j = a$. Thus $\alpha\colon (A, a_0) \to \mathbb{M}_k(A, a_0)$ is an injective map whose image is a prefix-closed subset of $\mathbb{M}_k(A, a_0)$. Defining $a \leq a'$ if $\alpha(a) \sqsubseteq \alpha(a')$ yields a tree order A with root a_0 and height $\leq k$. Since α is a homomorphism, we have for all binary relations $R_\alpha \in \sigma$, $R_\alpha^{\mathbb{M}_k(\mathcal{A},a_0)}(s,t)$ if and only if t is the extension of $s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_j} a_i$ by $a_i \xrightarrow{\alpha_{i+1}} a_{i+1}$, so condition (M) is satisfied. Thus $(A, a_0, \leq)$ is a modal tree σ-structure of height $\leq k$. We can define the inverse object mapping $K_!^{-1}\colon \mathbf{EM}(\mathbb{M}_k) \to \mathcal{R}_k^M(\sigma)$ as $K_!^{-1}((\mathcal{A}, a_0), \gamma) = (\mathcal{A}, a_0, \leq)$. ☐

Theorem 3.2.36 was stated and proved to exhibit the uniformity of the argument with analogous theorems for the EF and pebbling comonads. However, we observe that, due to the uniqueness in condition (M), the unit of η is in fact an isomorphism and the modal adjunction $L_k \dashv R_k$ is a coreflection. Therefore, we can view $\mathcal{R}_k^M(\sigma)$ as isomorphic to a subcategory of $\mathsf{Struct}_\star(\sigma)$. The isomorphic subcategory is the full subcategory consisting of synchronisation trees, and we get a conditional statement of when a coalgebra exists for a modal structure $(\mathcal{A}, a_0)$. We say a pointed σ-structure $(\mathcal{A}, a_0)$ in a modal signature σ is a *synchronisation tree* if every element $a \in A$ is reachable from a_0 via a unique transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \ldots \xrightarrow{\alpha_n} a_n = a$. The height of a synchronisation tree $(\mathcal{A}, a_0)$ is the maximum length of any path from the root a_0.

Theorem 3.2.37. *The following are equivalent:*

(1) There exists a $\mathbb{M}_k$-coalgebra $\gamma\colon (\mathcal{A}, a_0) \to \mathbb{M}_k(\mathcal{A}, a_0)$

(2) $(\mathcal{A}, a_0)$ is a synchronisation tree of height $\leq k$

Proof. Suppose there is a coalgebra $\gamma\colon (\mathcal{A}, a) \to \mathbb{M}_k(\mathcal{A}, a)$. For $a' \in \mathcal{A}$, let $\gamma(a') = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_j} a_j$ where $a_0 = a$. The first coalgebra equation states that $\gamma(a_i) = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_i} a_i$ for $0 \leq i \leq j$. The second coalgebra equation states that $a_j = a'$. Therefore, γ is injective. By injectivity and since γ is a homomorphism and thus preserves transitions, the path of transitions determined by $\gamma(a')$, i.e. $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_j} a'$ is unique. Hence, every $a' \in \mathcal{A}$ has a unique path from a, so the submodel generated by a is a synchronisation tree. Since $j \leq k$, the height of the tree is at most k.

Conversely, suppose $(\mathcal{A}, a_0)$ is a synchronisation tree of height $\leq k$, then for every a' in this submodel, there exists a unique path of transitions $a \xrightarrow{\alpha_1} \cdots \xrightarrow{\alpha_j} a'$. We can define the morphism $\gamma\colon (\mathcal{A}, a) \to \mathbb{M}_k(\mathcal{A}, a)$ as $\gamma(a') = a_0 \xrightarrow{\alpha_1} a_1 \ldots \xrightarrow{\alpha_j} a'$. Since the path of transitions is unique, the first coalgebra equation is satisfied. Moreover, since the last element of $\gamma(a')$ is a', the second coalgebra equation is satisfied. $\square$

Given this concrete characterisation of $\mathbb{M}_k$-coalgebras, we can define a combinatorial parameter which coincides with the $\mathbb{M}$-coalgebra number. The modal depth $\mathsf{md}(\mathcal{A}, a_0)$ of a modal structure $(\mathcal{A}, a_0)$ is k if $(\mathcal{A}, a_0)$ is synchronisation tree of height $= k$. In contrast with the analogous results for the EF comonad and pebbling comonad, this result is conditional. This is because the definition of objects in $\mathcal{R}^M_k(\sigma)$ reflect information about the structure of possible transitions. Thus, having $\mathbb{M}_k$-coalgebra for any k at all puts a strong constraint, i.e. being a tree, on the structure of the transition system.

Corollary 3.2.38. *If $(\mathcal{A}, a_0)$ is a finite synchronisation tree in modal signature σ, then $\chi^{\mathbb{M}}(\mathcal{A}, a_0) = \mathsf{md}(\mathcal{A})$.*

Proposition 3.2.39. *$L_k \dashv R_k$ is a resource-indexed arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\{\mathcal{R}^M_k(\sigma)\}$.*

Proof. Theorem 3.2.36 demonstrates that $K_!\colon \mathbf{Ob}(\mathcal{R}^M_k(\sigma)) \to \mathbf{EM}(\mathbb{M}_k)$ is a bijective object mapping, we must show that this bijection extends to morphisms and demonstrate that the functor $K_!\colon \mathcal{R}^M_k(\sigma) \to \mathbf{EM}(\mathbb{M}_k)$ is an isomorphism.

Concretely, the action of $K_!\colon \mathcal{R}^M_k(\sigma) \to \mathbf{EM}(\mathbb{M}_k)$ on morphisms can be described similarly to the the comparison functor in the proof of Proposition 3.2.6. Given a morphism $f\colon (\mathcal{B}, b_0, \leq) \to (\mathcal{C}, c_0, \leq') \in \mathbf{Mor}(\mathcal{R}^M_k(\sigma))$ of tree-ordered modal structures, $K_!(f)$ is the diagram D expressing the equation $\gamma \circ L_k(f) = \mathbb{M}_k(L_k(f)) \circ \beta$ where $K_!(\mathcal{B}, b_0, \leq) = ((\mathcal{B}, b_0), \beta)$ and $K_!(\mathcal{C}, c_0, \leq') = ((\mathcal{C}, c_0), \gamma)$. As with the proof of Proposition 3.2.6, it follows from f being a tree morphism that f preserves covering chain and diagram D commutes. We also additionally observe, for the arboreal cover of $L_k \dashv R_k$ of $\mathsf{Struct}_\star(\sigma)$ by $\mathcal{R}^M_k(\sigma)$, that $L(f)$ must pointed σ-morphism as f is a tree-ordered σ-morphism which preserves transitions in a sequence, i.e. $b_0 \xrightarrow{\alpha_1} b_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} b_n$ maps to $c_0 \xrightarrow{\alpha_1} f(b_1) \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} f(b_n)$ by condition (M). Moreover, from $f \in \mathbf{Mor}(\mathcal{R}^M_k(\sigma))$ acting on chains of height $\leq k$ it must be that case that D involves $\mathbb{M}_k(L(f))$ as $\mathbb{M}_k(L(f))$ acts on transition sequences of length $\leq k$.

Conversely, suppose we have commuting diagram D expressing $\gamma \circ g = \mathbb{M}_k(g) \circ \beta$ for coalgebras $((\mathcal{B}, b_0), \beta), ((\mathcal{C}, c_0), \gamma) \in \mathbf{Ob}(\mathbf{EM}(\mathbb{M}_k))$ and σ-morphism $g \colon \mathcal{A} \to \mathcal{B}$. We observe that D commuting means that g preserves transition sequences, i.e. $b_0 \xrightarrow{\alpha_1} b_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} b_n \mapsto c_0 \xrightarrow{\alpha_1} g(b_0) \xrightarrow{\alpha_1} g(b_1) \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} (b_n)$, of height $n \leq k$ from $K^{-1}((\mathcal{B}, b_0), \beta)$ to $K^{-1}((\mathcal{C}, c_0), \gamma)$. Thus, since g preserves transition sequences and is a pointed σ-morphism, it lifts to a morphism $K_!^{-1}(g)$ of $\mathcal{R}^M_k(\sigma)$. $\square$

With the resource-indexed arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\mathcal{R}^M_k(\sigma)$, we can demonstrate how the behavioural relations in Definition 3.1.19, instantiated in this case, recover the notions of simulation, property-preserving simulation, and bisimulation they generalise. Thus, proving how this arboreal cover captures equivalence in modal logic. For the following definition, given a pointed σ-structure $(\mathcal{A}, a_0)$ in a modal signature σ, an element $a \in A$, and a transition relation $\alpha \in \mathsf{Act}$, we will denote the set of α-successors to a as

$$R^{\mathcal{A}}_\alpha(a) = \{a' \mid R^{\mathcal{A}}_\alpha(a, a')\}.$$

Definition 3.2.40. Given pointed σ-structures $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$ and $k > 0$, we define the *k-modal bisimulation game* $\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$. We start with position (a_0, b_0). At each round $n + 1$ of the game, with current position (a_n, b_n):

- Spoiler chooses a transition associated to a binary relation $R_\alpha \in \sigma$ in one of the structures $\mathcal{A}, \mathcal{B}$, e.g.
  - If Spoiler chooses $\mathcal{A}$, then Spoiler picks element $a_{n+1} \in R^{\mathcal{A}}_\alpha(a_n)$.
  - If Spoiler chooses $\mathcal{B}$, then Spoiler picks element $b_{n+1} \in R^{\mathcal{B}}_\alpha(b_n)$.
- Duplicator responds by choosing a corresponding transition in the opposite structure, e.g.
  - If Spoiler chose $\mathcal{A}$, the Duplicator responds with element $b_{n+1} \in R^{\mathcal{B}}_\alpha(b_n)$.
  - If Spoiler chose $\mathcal{B}$, the Duplicator responds with element $a_{n+1} \in R^{\mathcal{A}}_\alpha(a_n)$.
- The position after round $n + 1$ is (a_{n+1}, b_{n+1}).

Duplicator wins round $n \geq 0$ if for all unary predicates $P \in \sigma$, $P^{\mathcal{A}}(a_n) \Leftrightarrow P^{\mathcal{B}}(b_n)$. Duplicator wins the game $\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ if she has a winning move for every possible Spoiler choice and every round $0 \leq n \leq k$.

As with previous cases, e.g. $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ and $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, $\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ has three variants obtained using analogous modifications:

- ${}^{\exists}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ is the game where in every round, Spoiler is forced to place pebbles on elements $a \in A$ and Duplicator is forced to place her corresponding pebbles on elements $b \in B$.
- ${}^{\exists^+}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ has the same player restrictions as ${}^{\exists}\mathbf{M}_k(\mathcal{A}, \mathcal{B})$, but additionally Duplicator's winning condition has been weakened. Duplicator can win in ${}^{\exists^+}\mathbf{M}_k(\mathcal{A}, \mathcal{B})$ if for all unary predicates $P \in \sigma$, $P^{\mathcal{A}}(a_n) \Rightarrow P^{\mathcal{B}}(b_n)$.
- ${}^{\#}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ is the bijective variant of the game. At round 0 of the game, the starting position of the game is (a_0, b_0). At each round $n+1$ of the game, with previous position (a_n, b_n):
  - Spoiler first chooses an action $\alpha \in \mathsf{Act}$.
  - Duplicator then chooses a bijection $\psi_{n+1} \colon R^{\mathcal{A}}_{\alpha}(a_n) \to R^{\mathcal{B}}_{\alpha}(b_n)$. If there is no such bijection, Spoiler automatically wins.
  - Spoiler then chooses $a_{n+1} \in R^{\mathcal{A}}_{\alpha}(a)$ and Duplicator responds with $\psi_{n+1}(a_{n+1}) \in R^{\mathcal{B}}_{\alpha}(b)$.

  The winning condition is the same as $\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$.

With the definitions of the modal (bi)simulation games in place, we can now proceed with the proofs of the morphism, pathwise, bisimulation, and isomorphism power theorems for the resource-indexed arboreal cover of $\mathsf{Struct}_{\star}(\sigma)$ by $\{\mathcal{R}^M_k(\sigma)\}$.

Theorem 3.2.41. *The following are equivalent for all pointed σ-structures $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$:*

(1) $(\mathcal{A}, a_0) \Rrightarrow^{\Diamond^+\mathcal{M}_k} (\mathcal{B}, b_0)$

(2) Duplicator has a winning strategy in ${}^{\exists^+}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$.

(3) Duplicator has a winning strategy in ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}, a_0), R_k(\mathcal{B}, b_0))$ *over* $\mathcal{R}^M_k(\sigma)$.

(4) $(\mathcal{A}, a_0) \to^{\mathcal{R}^M(\sigma)}_k (\mathcal{B}, b_0)$

Proof. For the (1) ⇔ (2) equivalence, this is the old result that simulation between labelled transition systems, which can be phrased as the game ${}^{\exists^+}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ preserves $\Box_\alpha$-free positive formulas. For a textbook account, see e.g. [21].

For the (2) ⇔ (3) equivalence, observe that elements $[m] \in \mathbf{P}(R_k(\mathcal{A}, a_0))$ are in bijective correspondence with transition sequences $s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_j} a_j$ such that $m(r) = s$ where r is the maximal element of the chain $\mathsf{dom}(m)$; and conversely, up to equivalence, we can assume m is embedding mapping the transition sequence s into $R_k(\mathcal{A}, a_0)$, i.e.

for all $i \in (j)$, the i-th element of $\mathsf{dom}(m)$ is a_i. Therefore, Spoiler extending a transition sequence s by a transition $a_j \xrightarrow{\alpha} a_{j+1}$ in the $j+1$ round of ${}^{\exists^+}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ corresponds to Spoiler extending $[m]$ by $[m']$ in ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}, a_0), R_k(\mathcal{B}, b_0))$ over $\mathcal{R}_k^M(\sigma)$. This similarly holds for Duplicator responses. The winning condition of position (a_j, b_j) in ${}^{\exists^+}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ that $P^{\mathcal{A}}(a_j) \Rightarrow P^{\mathcal{B}}(b_j)$ is equivalent to the resulting position $([m], [n]) \in {}^{+}\mathcal{W}(R_k(\mathcal{A}, a_0), R_k(\mathcal{B}, b_0))$ for the game ${}^{\exists^+}\mathcal{G}(R_k(\mathcal{A}, a_0), R_k(\mathcal{B}, b_0))$ over $\mathcal{R}^M(\sigma)$.

For the (3) ⇔ (4) equivalence, we observe that, from morphisms in $\mathcal{R}_k^M(\sigma)$ preserving covering chains, $\mathbf{P} \colon \mathcal{R}_k^M(\sigma) \to \mathcal{T}$ is faithful. Applying Proposition 3.1.13(3) to the case where $\mathcal{C} = \mathcal{R}_k^M(\sigma)$, $X = R_k(\mathcal{A}, a_0)$, and $Y = R_k(\mathcal{B}, b_0)$ yields the result. □

Theorem 3.2.42. *The following are equivalent for all pointed σ-structures $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$:*

(1) $(\mathcal{A}, a_0) \Rrightarrow^{\Diamond\mathcal{M}_k} (\mathcal{B}, b_0)$

(2) Duplicator has a winning strategy in ${}^{\exists}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$.

(3) Duplicator has a winning strategy in ${}^{\exists}\mathcal{G}(R_k((\mathcal{A}, a_0)), R_k((\mathcal{B}, b_0)))$ *over* $\mathcal{R}_k^M(\sigma)$.

(4) $(\mathcal{A}, a_0) \rightharpoonup_k^{\mathcal{R}^M(\sigma)} (\mathcal{B}, b_0)$

Proof. For the (1) ⇔ (2) equivalence, the proof is similar to the (1) ⇔ (2) equivalence in Theorem 3.2.41 above. Including propositional variables and their negations in the logic $\mathcal{M}_k$ corresponds to the fact that the interpretations of unary relations on a_n, b_n at round n in the game ${}^{\exists}\mathbf{M}_k(\mathcal{A}, \mathcal{B})$ must be reflected and preserved, i.e. $P^{\mathcal{A}}(a_n) \Leftrightarrow P^{\mathcal{B}}(b_n)$.

For the (2) ⇔ (3) equivalence, the proof is similar to the (2) ⇔ (3) equivalence in Theorem 3.2.41 above. The key difference is that Duplicator's winning condition for the position (a_n, b_n) in the n-th round of ${}^{\exists}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ is for all unary $P \in \sigma$, $P^{\mathcal{A}}(a_n) \Leftrightarrow P^{\mathcal{B}}(b_n)$. This is equivalent to the resulting position $([m], [n]) \in \mathcal{W}(R_k(\mathcal{A}, a_0), R_k(\mathcal{B}, b_0))$ for the game ${}^{\exists}\mathcal{G}(R_k(\mathcal{A}, a_0), R_k(\mathcal{B}, b_0))$ over $\mathcal{R}^M(\sigma)$.

For the (3) ⇔ (4) equivalence, since the right adjoint R_k preserves the product $(\mathcal{A}, a_0) \times (\mathcal{B}, b_0) \in \mathsf{Struct}_\star(\sigma)$, $R_k(\mathcal{A}, a_0) \times R_k(\mathcal{B}, b_0) \in \mathcal{R}_k^M(\sigma)$. Applying Proposition 3.1.13(1) to the case where $\mathcal{C} = \mathcal{R}_k^M(\sigma)$, $X = R_k(\mathcal{A}, a_0)$, and $Y = R_k(\mathcal{B}, b_0)$ yields the result. □

Theorem 3.2.43. *The following are equivalent for all pointed σ-structures $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$:*

(1) $(\mathcal{A}, a_0) \equiv^{\mathcal{M}_k} (\mathcal{B}, b_0)$

(2) Duplicator has a winning strategy in $\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$.

(3) Duplicator has a winning strategy in $\mathcal{G}(R_k((\mathcal{A}, a_0)), R_k((\mathcal{B}, b_0)))$ *over* $\mathcal{R}_k^M(\sigma)$.

(4) $(\mathcal{A}, a_0) \leftrightarrow_k^{\mathcal{R}^M(\sigma)} (\mathcal{B}, b_0)$

Proof. For the (1) $\Leftrightarrow$ (2), this is the classical result that bisimulation, which can be phrased as the game $\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ characterises equivalence in $\mathcal{M}_k$. This results originates from [45]. For a textbook account, see e.g. [21].

For the (2) $\Leftrightarrow$ (3) equivalence, the proof is similar to the (2) $\Leftrightarrow$ (3) equivalence in Theorem 3.2.42 above. The key difference is that Spoiler can now extend a position (a_n, b_n) by choosing a transition $b_n \xrightarrow{\alpha} b_{n+1}$. Similarly, Spoiler can also extends position $([m], [n])$ by $[n'] \succ [n]$ in $\mathbf{P}(R_k(\mathcal{B}, b_0))$. Thus, we have a correspondence between plays in $\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ and $\mathcal{G}(R_k((\mathcal{A}, a_0)), R_k((\mathcal{B}, b_0)))$ over $\mathcal{R}_k^M(\sigma)$. Duplicator's winning condition in $\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ is equivalent to Duplicator's winning condition in $\mathcal{G}(R_k((\mathcal{A}, a_0)), R_k((\mathcal{B}, b_0)))$ over $\mathcal{R}_k^M(\sigma)$ as was demonstrated in Theorem 3.2.42.

For the (3) $\Leftrightarrow$ (4) equivalence, since the right adjoint R_k preserves the product $(\mathcal{A}, a_0) \times (\mathcal{B}, b_0) \in \mathsf{Struct}_\star(\sigma)$, $R_k(\mathcal{A}, a_0) \times R_k(\mathcal{B}, b_0) \in \mathcal{R}_k^M(\sigma)$. Applying Proposition 3.1.13(2) to the case where $\mathcal{C} = \mathcal{R}_k^M(\sigma)$, $X = R_k(\mathcal{A}, a_0)$, and $Y = R_k(\mathcal{B}, b_0)$ yields the result. $\square$

The isomorphism power theorems in Section 3.2.2 and Section 3.2.3 were conditional on $\mathcal{A}, \mathcal{B}$ being finite structures. This condition was related to the fact that injective branch maps $\psi \colon A \to B$ that Duplicator plays are bijective for finite sets. In the modal case, since Duplicator's winning strategy in ${}^{\#}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ relies on branch maps of the form $\psi \colon R_\alpha^{\mathcal{A}}(a) \to R_\alpha^{\mathcal{B}}(b)$ for $a \in A$ and $b \in \mathcal{B}$, we instead require that $(\mathcal{A}, a_0)$ and $(\mathcal{B}, b_0)$ are *image-finite*, i.e. for all $a \in \mathcal{A}$ and $R \in \sigma$, $R^{\mathcal{A}}(a)$ is finite.

By the definition of adjunction, morphisms $f^* \colon R_k(\mathcal{A}, a_0) \to R_k(\mathcal{B}, b_0) \in \mathbf{Mor}(\mathcal{R}^M(\sigma))$ are in bijection with morphisms $f \colon \mathbb{M}_k(\mathcal{A}, a_0) \to (\mathcal{B}, b_0) \in \mathbf{Mor}(\mathsf{Struct}_\star(\sigma))$ where we note that that $\mathbb{M}_k = L_k \circ R_k$. For every morphism $f \colon \mathbb{M}_k(\mathcal{A}, a_0) \to (\mathcal{B}, b_0)$, transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \ldots \xrightarrow{\alpha_n} a_n \in \mathbb{M}_k(\mathcal{A}, a_0)$, and action $\alpha \in \mathsf{Act}$, we can define the *f-induced* $\mathbf{M}$*-branch function* $\psi_{s,\alpha} \colon R_\alpha^{\mathcal{A}}(a_n) \to R_\alpha^{\mathcal{B}}(f(s))$ *at* (s, α) as

$$\psi_{s,\alpha}(a) = f(a_0 \xrightarrow{\alpha_1} a_1 \ldots \xrightarrow{\alpha_n} a_n \xrightarrow{\alpha} a).$$

Definition 3.2.44. A morphism $f \colon \mathbb{M}_k(\mathcal{A}, a_0) \to (\mathcal{B}, b_0) \in \mathbf{Mor}(\mathsf{Struct}_\star(\sigma))$ is $\mathbf{M}$*-branch bijective* if for every $s \in \mathbb{M}_k(\mathcal{A}, a_0)$ ending in state a_n, the f-induced $\mathbf{M}$-branch function $\psi_{s,\alpha} \colon R_\alpha^{\mathcal{A}}(a_n) \to R_\alpha^{\mathcal{B}}(f(s))$ at α is bijective, for all unary $P \in \sigma$, $P^{\mathcal{A}}(a_0) \Leftrightarrow P^{\mathcal{B}}(b_0)$, and for all $a_{n+1} \in R_\alpha^{\mathcal{A}}(a_n)$, $P^{\mathcal{A}}(a_{n+1}) \Leftrightarrow P^{\mathcal{B}}(\psi_{s,\alpha}(a_{n+1}))$.

Theorem 3.2.45. *The following are equivalent for all image-finite pointed σ-structures $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$:*

(1) $(\mathcal{A}, a_0) \equiv^{\#\mathcal{M}_k} (\mathcal{B}, b_0)$

(2) Duplicator has a winning strategy in $^{\#}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$*.*

(3) There exists a $\mathbf{M}$*-branch bijective* $f\colon \mathbb{M}_k(\mathcal{A}, a_0) \to (\mathcal{B}, b_0)$*.*

(4) $(\mathcal{A}, a_0) \cong_k^{\mathcal{R}^M(\sigma)} (\mathcal{B}, b_0)$

Proof. For the (1) ⇔ (2) equivalence, it was shown in [32] that a notion of graded bisimulation, which can be phrased as the game $^{\#}\mathbf{M}_k(\mathcal{A}, \mathcal{B})$ characterises equivalence in $\#\mathcal{M}_k$. Graded bisimulation is related to the notion of resource bisimulation from [15].

For the (2) ⇔ (3) equivalence, we proceed by induction on the number of rounds. At round 0, the position (a_0, b_0) is winning for Duplicator if for all unary $P \in \sigma$, $P^{\mathcal{A}}(a_0) \Leftrightarrow P^{\mathcal{B}}(b_0)$ which is the second item in the definition of $\mathbf{M}$-branch bijective. At round $n+1$ of $^{\#}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$, suppose Spoiler's previous moves is the sequence $s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n$ and Spoiler chooses action $\alpha \in \mathsf{Act}$, Duplicator plays bijection $\psi_{s,\alpha}\colon R_\alpha^{\mathcal{A}}(a_n) \to R_\alpha^{\mathcal{B}}(f(s))$. Conversely, we can define $\psi_{s,\alpha}$ as Duplicator response bijection $\psi_{n+1}\colon R_\alpha^{\mathcal{A}}(a_n) \to R_\alpha^{\mathcal{B}}(b_n)$ to Spoiler choosing $\alpha \in \mathsf{Act}$ in the $n+1$-th round of $^{\#}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ and then define $f\colon \mathbb{M}_k(\mathcal{A}, a_0) \to \mathcal{B}$ as $f(t) = \psi_{s,\alpha}(a_{n+1})$ where t is the extension of s by the transition $a_n \xrightarrow{\alpha_n} a_{n+1}$. The winning condition of $^{\#}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ is that for all unary $P \in \sigma$ and $a_{n+1} \in R_\alpha^{\mathcal{A}}(a_n)$, $P^{\mathcal{A}}(a_{n+1}) \Leftrightarrow P^{\mathcal{B}}(\psi_{s,\alpha}(a_{n+1}))$, i.e. the third item in the definition of $\mathbf{M}$-branch bijective.

For the (3) ⇔ (4) equivalence, suppose $f\colon L_k(R_k(\mathcal{A}, a_0)) \to (\mathcal{B}, b_0)$ is branch bijective, then via the adjunction we obtain a morphism $f^*\colon R_k(\mathcal{A}, a_0) \to R_k(\mathcal{B}, b_0)$. By induction on the length of sequences $s \in \mathbb{P}_k(\mathcal{A})$, we can conclude that f^* is a bijection, by the fact that $\psi_{s,\alpha}$ is bijective, and is a σ-isomorphism, since unary relations are preserved and reflection and the interpretation of binary relations R_α is via extension by transitions $a_n \xrightarrow{\alpha} a_{n+1}$. Conversely, given an isomorphism $f^*\colon R_k(\mathcal{A}, a_0) \to R_k(\mathcal{B}, b_0)$ with inverse $g^*\colon R_k(\mathcal{B}, a_0) \to R_k(\mathcal{A}, b_0)$, via the adjunction, we can produce $f\colon L_k(R_k(\mathcal{A}, a_0)) \to (\mathcal{B}, b_0)$ and observe that this satisfies the conditions the f-induced $\mathbf{M}$-branch functions $\psi_{s,\alpha}\colon R_\alpha^{\mathcal{A}}(\varepsilon(s)) \to R_\alpha^{\mathcal{B}}(f(s))$ are bijective. Further, by f^*, g^* being $\mathcal{R}_k^M(\sigma)$-morphisms, for all unary P, $P^{\mathcal{A}}(a_0) \Leftrightarrow P^{\mathcal{B}}(b_0)$ and for all $a_{n+1} \in R_\alpha^{\mathcal{A}}$, $P^{\mathcal{A}}(a_{n+1}) \Leftrightarrow P^{\mathcal{B}}(\psi_{s,\alpha}(a_{n+1}))$. ☐

In summary, the behavioural relations of Definition 3.1.19 for the resource-indexed arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}_k^M(\sigma)\}$ characterise equivalence for variants of modal logic $\mathcal{M}_k$ graded by nesting depth of modalities.

Theorem 3.2.46. *The following equivalences hold:*

(1) $(\mathcal{A}, a_0) \to_k^{\mathbb{C}} (\mathcal{B}, b_0)$ *if and only if* $(\mathcal{A}, a_0) \Rrightarrow^{\Diamond+\mathcal{M}_k} (\mathcal{B}, b_0)$

(2) $(\mathcal{A}, a_0) \rightharpoonup_k^{\mathbb{C}} (\mathcal{B}, b_0)$ *if and only if* $(\mathcal{A}, a_0) \Rrightarrow^{\Diamond\mathcal{M}_k} (\mathcal{B}, b_0)$

(3) $(\mathcal{A}, a_0) \leftrightarrow_k^{\mathbb{C}} (\mathcal{B}, b_0)$ *if and only if* $(\mathcal{A}, a_0) \equiv^{\mathcal{M}_k} (\mathcal{B}, b_0)$

(4) $(\mathcal{A}, a_0) \cong_k^{\mathbb{C}} (\mathcal{B}, b_0)$ *if and only if* $(\mathcal{A}, a_0) \equiv^{\#\mathcal{M}_k} (\mathcal{B}, b_0)$ *(assuming* $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$ *are image-finite).*

3.3 Interpreted symbols

In this section, we extend the arboreal approach to capture logics that have interpreted symbols. This addresses a gap in the use of arboreal covers. In particular, this extension allows us to use arboreal categories to capture equivalence in logics with equality. We use the machinery of relative comonads to address the gap, and show that relative-liftings of game comonads enable us to capture a wider range of logics. Namely, after adapting the definition of the Ehrenfeucht-Fraïssé comonad for multi-sorted structures, we are able to use this approach to capture equivalence in second-order logics.

3.3.1 Enriched logics

Throughout this chapter, we have only been able to use arboreal covers to give characterisations of equivalence for logics without equality. However, model theorists typically deal with logics where there is an atomic symbol '=' interpreted on a structure $\mathcal{A}$ as equality on the underlying universe A. To illustrate how to captures logics with equality using game comonads, we will first consider adding equality to $\mathbf{FO}_k$ making use of the Ehrenfeucht-Fraïssé comonad $\mathbb{E}_k$ from Section 3.2.2 capturing this logic.

In order to add equality to $\mathbf{FO}_k$, we make explicit the use extra atomic symbol by moving from the signature σ to the signature $\sigma^+ = \sigma \cup \{I\}$ with an extra binary relation symbol I. From there, in order to encapsulate the intended interpretation of equality, we consider a functor $J_= \colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma^+)$ where

- For objects, $\mathcal{A} = (A, \{R^{\mathcal{A}}\}_{R \in \sigma})$, we define $J_=(\mathcal{A}) = (A, \{R^{\mathcal{A}}\}_{R \in \sigma}, I^{\mathcal{A}})$ where $I^{\mathcal{A}}$ is the equality relation on A, i.e. $I^{\mathcal{A}} = \{(a, a) \mid \forall a \in A\}$.
- For morphisms $f \colon \mathcal{A} \to \mathcal{B}$, the morphism $J_=(f) \colon J_=(\mathcal{A}) \to J_=(\mathcal{B})$ has the same underlying set function as f. As f is a σ-morphism and set functions preserve equality, $J_=(f)$ is a σ^+-morphism.

Let $\mathbb{E}_k^{\sigma^+} : \mathsf{Struct}(\sigma^+) \to \mathsf{Struct}(\sigma^+)$ denote the Ehrenfeucht-Fraïssé comonad on $\mathsf{Struct}(\sigma^+)$.

Definition 3.3.1. Given σ-structures $\mathcal{A}, \mathcal{B}$ and $k > 0$, we define the *back-and-forth k-round Ehrenfeucht-Fraïssé game* $\mathbf{EF}^{=}_{k}(\mathcal{A}, \mathcal{B})$ *from* $\mathcal{A}$ *to* $\mathcal{B}$ *with equality.* During each round $n \leq k$ of the game,

- Spoiler chooses an element $a_n \in A$ or $b_n \in B$.
- Duplicator responds with an element $b_n \in B$ or $a_n \in A$ in the opposite structure.

Duplicator wins round n if the relation

$$\gamma_n := \{(a_j, b_j) \mid j \leq n\}$$

is a partial σ-isomorphism. Duplicator wins in $\mathbf{EF}^{=}_{k}(\mathcal{A}, \mathcal{B})$ if for every round $n \leq k$ and possible Spoiler choice, Duplicator has a winning move.

We can define the one-sided and bijective variants ${}^{\exists^{+}}\mathbf{EF}^{=}_{k}(\mathcal{A}, \mathcal{B})$, ${}^{\exists}\mathbf{EF}^{=}(\mathcal{A}, \mathcal{B})$, ${}^{\#}\mathbf{EF}^{=}(\mathcal{A}, \mathcal{B})$ of $\mathbf{EF}^{=}_{k}(\mathcal{A}, \mathcal{B})$ as we did for $\mathbf{EF}_{k}(\mathcal{A}, \mathcal{B})$. In particular, for the game ${}^{\exists^{+}}\mathbf{EF}_{k}(\mathcal{A}, \mathcal{B})$, Spoiler is forced to play in $\mathcal{A}$ and Duplicator's winning condition is that γ_n is a partial σ-morphism.

The key lemma relating the game $\mathbf{EF}^{=}_{k}(\mathcal{A}, \mathcal{B})$ with equality and the game without equality $\mathbf{EF}_{k}(\mathcal{A}, \mathcal{B})$ involves translating the winning condition of Duplicator 'along the functor $J_{=}$'. This is formalised in the following lemma:

Lemma 3.3.2. *For every finite relational signature σ and $\mathcal{A}, \mathcal{B} \in \mathbf{Ob}(\mathsf{Struct}(\sigma))$:*

- *γ is a partial σ-morphism from $\mathcal{A}$ to $\mathcal{B}$ if and only if γ is a σ^{+}-correspondence from $J_{=}(\mathcal{A})$ to $J_{=}(\mathcal{B})$.*
- *γ is a partial σ-isomorphism from $\mathcal{A}$ to $\mathcal{B}$ if and only if γ is a tight σ^{+}-correspondence from $J_{=}(\mathcal{A})$ to $J_{=}(\mathcal{B})$.*

Here, we are viewing γ as being specified by a subset of $A \times B$.

Proof. Suppose γ is a partial σ-morphism from $\mathcal{A}$ to $\mathcal{B}$, and consider pairs $(a, b), (a', b')$. By γ being a σ-morphism, $a = a' \Rightarrow b = b'$ By the definition of $I^{J(\mathcal{A})}, I^{J(\mathcal{B})}$ as equality on the underlying universes, this implication is equivalent to γ preserving the I-relation. As the interpretations of the relations are unchanged by J, we see that γ is σ-morphism if and only if γ is a σ^{+}-correspondence. The proof from the second item is the same, but using the fact that the bi-implication $a = a' \Leftrightarrow b = b'$ holds for partial σ-isomorphisms. $\square$

Theorem 3.3.3. *The following are equivalent:*

(1) $\mathcal{A} \equiv^{\mathbf{FO}^{=}_{k}} \mathcal{B}$

(2) Duplicator has a winning strategy in $\mathbf{EF}^{=}_{k}(\mathcal{A}, \mathcal{B})$.

(3) Duplicator has a winning strategy in $\mathbf{EF}_k(J_=(\mathcal{A}), J_=(\mathcal{B}))$*.*

(4) Duplicator has a winning strategy in $\mathcal{G}(R_k(J_=(\mathcal{A})), R_k(J_=(\mathcal{B})))$ *over* $\mathcal{R}_k^E(\sigma)$*.*

(5) $J_=(\mathcal{A}) \leftrightarrow_k^{\mathcal{C}} J_=(\mathcal{B})$

Proof. For the (1) ⇔ (2) equivalence, this is the classic Ehrenfeucht-Fraïssé theorem [34, 36]. For a textbook account of the classic Ehrenfeucht-Fraïssé theorem see e.g. [59].

For the (2) ⇔ (3) equivalence, for any Spoiler move $\mathbf{EF}^=(\mathcal{A}, \mathcal{B})$, Duplicator responds with the same element as she would in $\mathbf{EF}(J_=(\mathcal{A}), J_=(\mathcal{B}))$, and vice-versa for the converse. By Lemma 3.3.2, the relation γ generated by this strategy is winning for Duplicator in $\mathbf{EF}^=(\mathcal{A}, \mathcal{B})$, and similarly for the converse.

For the (3) ⇔ (4) ⇔ (5) equivalences, these follow from the corresponding equivalences in Theorem 3.2.16. □

Similar arguments can be repeated for the variants $\exists^+\mathbf{FO}_k^=$, $\exists\mathbf{FO}_k^=$, and $\#\mathbf{FO}_k^=$ of $\mathbf{FO}_k^=$ yielding the following theorem.

Theorem 3.3.4. *The following equivalences hold for all* $k > 0$*:*

(1) $J_=(\mathcal{A}) \to_k^{\mathcal{C}} J_=(\mathcal{B})$ *if and only if* $\mathcal{A} \Rrightarrow^{\exists^+\mathbf{FO}_k^=} \mathcal{B}$

(2) $J_=(\mathcal{A}) \rightharpoonup_k^{\mathcal{C}} J_=(\mathcal{B})$ *if and only if* $\mathcal{A} \Rrightarrow^{\exists\mathbf{FO}_k^=} \mathcal{B}$

(3) $J_=(\mathcal{A}) \leftrightarrow_k^{\mathcal{C}} J_=(\mathcal{B})$ *if and only if* $\mathcal{A} \equiv^{\mathbf{FO}_k^=} \mathcal{B}$

(4) $J_=(\mathcal{A}) \cong_k^{\mathcal{C}} J_=(\mathcal{B})$ *if and only if* $\mathcal{A} \equiv^{\#\mathbf{FO}_k^=} \mathcal{B}$

An additional example of an enriched logic is first-order logic with a connectivity predicate conn, $\mathbf{FO}^{\mathsf{conn}}$ which was explored in [80]. Connectivity cannot be expressed in bounded quantifier rank fragments of first-order logic, so $\mathbf{FO}_k^{\mathsf{conn}}$ is strictly more expressive than $\mathbf{FO}_k$. To encapsulate this example, we consider a signature $\sigma^{\mathsf{conn}} = \sigma \cup \{\mathsf{conn}, I\}$ which has two additional binary relation I and conn [1] and a functor $J_{\mathsf{conn}} \colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma^{\mathsf{conn}})$ where

- For objects $\mathcal{A}$ with universe A, we define $J_{\mathsf{conn}}(\mathcal{A})$ as the σ^{conn}-structure with universe A and relational interpretations:
  - $(a, a') \in I^{\mathcal{A}}$ if $a = a'$.

---

[1] Schirrmacher et. al. [80] introduce a logic with additional connectivity predicates conn_n of arity $n+2$ which express connectivity avoiding n many elements. For ease of presentation, we present the case where $\mathsf{conn} = \mathsf{conn}_0$.

  - $(a, a') \in \mathsf{conn}^{\mathcal{A}}$ if there exists a path $a_0 = a \frown a_1 \frown \cdots \frown a_n = a'$ in the Gaifman graph $\mathcal{G}(\mathcal{A})$.
  - $R^{\mathcal{A}} = R^{J_{\mathsf{conn}}(\mathcal{A})}$ for all $R \in \sigma$.

- For morphisms $f \colon \mathcal{A} \to \mathcal{B}$, the morphism $J_{\mathsf{conn}}(f) \colon J_{\mathsf{conn}}(\mathcal{A}) \to J_{\mathsf{conn}}(\mathcal{B})$ has the same underlying set function as f. As f is a σ-morphism, f preserves the edge relation of the Gaifman graph and thus the conn-relation, so $J_{\mathsf{conn}}(f)$ is a σ^{conn}-morphism.

If we use $\mathbf{FO}_k^{\mathsf{conn}}$ to denote first-order logic with connectivity predicate with quantifier rank $\leq k$, then it is easy to see that we can use an analogous argument to the equality case (Theorem 3.3.3) to obtain a characterisations of equivalence in this logic via the relative comonad $\mathbb{E}_k \circ J_{\mathsf{conn}}$.

Theorem 3.3.5. *$J_{\mathsf{conn}}(\mathcal{A}) \leftrightarrow_k^{\mathfrak{C}} J_{\mathsf{conn}}(\mathcal{B})$ if and only if $\mathcal{A} \equiv^{\mathbf{FO}_k^{\mathsf{conn}}} \mathcal{B}$*

Proof. By Theorem 3.2.16, $J_{\mathsf{conn}}(\mathcal{A}) \leftrightarrow_k^{\mathfrak{C}} J_{\mathsf{conn}}(\mathcal{B})$ is equivalent to Duplicator having a winning strategy in $\mathbf{EF}_k(J_{\mathsf{conn}}(\mathcal{A}), J_{\mathsf{conn}}(\mathcal{B}))$. By [80, Theorem 3.9], equivalence $\mathcal{A} \equiv^{\mathbf{FO}_k^{\mathsf{conn}}} \mathcal{B}$ is characterised by a modified version of the $\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$ game where a position for Duplicator is winning if, in addition to the chosen elements forming a partial isomorphism, the connectivity predicate must be reflected and preserved. This is precisely the same as a winning position for Duplicator in $\mathbf{EF}_k(J_{\mathsf{conn}}(\mathcal{A}), J_{\mathsf{conn}}(\mathcal{B}))$. $\square$

3.3.2 Second-order logic

Monadic second-order logic (MSO) is the extension of first-order logic with quantifiers $\exists X$ and $\forall X$ which quantify over subsets of the domain. A standard way of viewing monadic second-order logic is to restrict the semantics of two-sorted first-order logic with an additional symbol relating the two sorts. The semantics restriction to obtain monadic second-order logic involves interpreting the first sort as the domain A, the second sort as the powerset $\mathcal{P}(A)$ of A, and the additional binary relation symbol as the set membership relation $\in$.

Building from monadic second-order logic, general k-adic second-order logic extends first-order logic by also allowing quantifiers that range over subsets of $\mathcal{P}(A^n)$ for any $n \leq k$. Similar to the monadic case, the standard semantics involves restricting the semantics of (k)-sorted first-order logic where the n-th sort is $\mathcal{P}(A^n)$ for $n > 0$ and A for $n = 0$. Thus, we need to extend our comonadic analysis to the setting of multi-sorted model theory.

In the setting of multi-sorted model theory, we fix a set S of sorts. Signatures σ come up with relational symbols that have arities specified by vectors of sorts $\langle s_1, \ldots, s_m \rangle$ where $s_1, \ldots, s_m \in S$.

Universes of σ-structures are S-sorted sets rather than ordinary sets. An S-sorted set A is an S-indexed family of sets. For such a family A, the set A_s denotes the component of A at sort $s \in S$. An element $a \in A_s$ will be denoted $a{:}s \in A$. An *S-sorted σ-structure* $\mathcal{A}$ has universe given by an S-sorted set A, and interpretations

$$R^{\mathcal{A}} \subseteq A_{s_1} \times \cdots \times A_{s_m}$$

for every relational symbol $R \in \sigma$ with arity $\langle s_1, \ldots, s_m \rangle$. The morphism between σ-structures is an S-sorted function $h \colon A \to B$ that is specified by an S-indexed family of set functions $h_s \colon A_s \to B_s$. An *S-sorted σ-morphism* $h \colon \mathcal{A} \to \mathcal{B}$ is a S-sorted function $h \colon A \to B$ such that for every relational symbol $R \in \sigma$ of arity $\langle s_1, \ldots, s_m \rangle$,

$$R^{\mathcal{A}}(a_1{:}s_1, \ldots, a_m{:}s_m) \Rightarrow R^{\mathcal{B}}(h(a_1{:}s_1), \ldots, h(a_m{:}s_m))$$

for all $a_1{:}s_1, \ldots, a_m{:}s_m \in A$. We will write $\mathsf{Struct}^S(\sigma)$ for the category S-sorted σ-structures and σ-morphisms. Fixing a subset of $I \subseteq S$ of sorts with equality, we can define S-sorted $\mathbf{FO}^{I,S}$ with equality on I sorts. In $\mathbf{FO}^{I,S}$, variables x are equipped with a sort s, denoted $x{:}s$, atomic formulas are either of the form $R(x_1{:}s_1, \ldots, x_m{:}s_m)$ whenever $R \in \sigma$ has arity $\langle s_1, \ldots, s_m \rangle$ or of the form $x{:}s = y{:}s$ whenever $s \in I$. Let $\mathbf{FO}_k^{I,S}$ denote logic with formulas of $\mathbf{FO}^{I,S}$ with quantifier rank $\leq k$. For more details on multi-sorted model theory, see Appendix A.

We 'upgrade' the Ehrenfeucht-Fraïssé comonad to the S-sorted setting to capture equivalence in $\mathbf{FO}^{\varnothing,S}$. To begin, we first define the functor $\mathbb{E}_k^S \colon \mathsf{Struct}^S(\sigma) \to \mathsf{Struct}^S(\sigma)$, suppose $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}^S(\sigma))$, then we can define the universe $\mathbb{E}_k^S(A)$ of $\mathbb{E}_k^S(\mathcal{A})$ inductively as:

$$\frac{[a_1{:}s_1, \ldots, a_n{:}s_n] \in (\mathbb{E}_k^S A)_{s_n} \qquad n < k \qquad a_{n+1} \in A_{s_{n+1}}}{[a_1{:}s_1, \ldots, a_n{:}s_n, a_{n+1}{:}s_{n+1}] \in (\mathbb{E}_k^S A)_{s_{n+1}}}$$

Mirroring the single-sorted case, the counit $\varepsilon \colon \mathsf{Id}_{\mathsf{Struct}(\sigma^{\mathsf{e}})} \to \mathbb{E}_k^S$ has components $\varepsilon_{\mathcal{A}}$ specified by an S-sorted function such that for every $s \in S$ $\varepsilon_{\mathcal{A},s}$ returns the last element. The coextension $h^* \colon \mathbb{E}_k^S(\mathcal{A}) \to \mathbb{E}_k^S(\mathcal{B})$ of an σ-morphism $h \colon \mathbb{E}_k^S(\mathcal{A}) \to \mathcal{B}$ is defined component wise as

$$h_s([a_1{:}s_1, \ldots, a_n{:}s_n]) = [h_{s_1}(a_1), \ldots, h_{s_n}(a_n)]$$

where we note that $s = s_n$.

Proposition 3.3.6. *$(\mathbb{E}_k^S, \varepsilon, (\cdot)^*)$ is a comonad in Kleisli form.*

The coalgebras of $\mathbb{E}_k^S$ are forest covers $(\mathcal{A}, \leq)$ of σ-structures $\mathcal{A}$ where the forest order $\leq$ is over the set $\coprod_{s \in S} A_s$. The coalgebra morphisms in $\mathbf{EM}(\mathbb{E}_k^S)$ are equivalent to forest-ordered σ-morphisms which preserve sort. Thus, applying Proposition 3.2.11, $\{\mathbf{EM}(\mathbb{E}_k^S\}$ is an resource-indexed arboreal cover of $\mathsf{Struct}^S(\sigma)$. We now prove the generalisation of Theorem 3.2.16 for this multi-sorted EF comonad. Let $G^{\mathbb{E}_k^S} \colon \mathsf{Struct}^S(\sigma) \to \mathbf{EM}(\mathbb{E}_k)$ denote the cofree coalgebra sending a σ-structure $\mathcal{A}$ to $(\mathbb{E}_k^S(\mathcal{A}), \delta_{\mathcal{A}})$. Generalising from the 'equality interpretation' $J_=$ in the single-sorted case, we consider a signature $\sigma^+ = \sigma \cup \{I_s\}_{s \in I}$ where for every sort $s \in I$, there is relation I of arity $\langle s, s \rangle$. The functor $J_=^I \colon \mathsf{Struct}^S(\sigma) \to \mathsf{Struct}^S(\sigma^+)$ encodes the equality interpretation by mapping a σ-structure $\mathcal{A}$ to the σ^+-structure $J_=^I(\mathcal{A})$ where $I_s^{\mathcal{A}} = \{(a, a) \mid a \in A_s\}$.

Theorem 3.3.7. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B} \in \mathbf{Ob}(\mathsf{Struct}^S(\sigma))$:*

(1) $\mathcal{A} \equiv^{\mathbf{FO}_k^{I,S}} \mathcal{B}$.

(2) Duplicator has a winning strategy in $\mathbf{EF}_k^{I,S}(\mathcal{A}, \mathcal{B})$.

(3) There exists a bisimulation in $G^{\mathbb{E}_k^S}(J_=^I(\mathcal{A})) \leftarrow X \to G^{\mathbb{E}_k^S}(J_=^I(\mathcal{B}))$ in $\mathbf{EM}(\mathbb{E}_k^S)$.

Proof. For the (1) $\Leftrightarrow$ (2) equivalence, this is a straightforward adaption of the single-sorted proof. For details, see in Appendix A. For the (2) $\Leftrightarrow$ (3) equivalence, following the same lines as Theorem 3.2.16, a Duplicator winning strategy in $\mathbf{EF}_k^{I,S}(\mathcal{A}, \mathcal{B})$ is equivalent to a Duplicator winning strategy in $\mathcal{G}(G^{\mathbb{E}_k^S}(\mathcal{A}), G^{\mathbb{E}_k^S}(\mathcal{B}))$ over $\mathbf{EM}(\mathbb{E}_k^S)$. Thus, by Proposition 3.1.13(2), this equivalent to the existence of a bisimulation $G^{\mathbb{E}_k^S}(J_=(\mathcal{A})) \leftarrow X \to G^{\mathbb{E}_k^S}(J_=(\mathcal{B}))$ in $\mathbf{EM}(\mathbb{E}_k^S)$. $\square$

Using a relative-lifting of this multi-sorted EF comonad, we can capture equivalence in monadic second order logic $\mathbf{MSO}_k$ up to quantifier rank $< k$. Recall the syntax of $\mathbf{MSO}$ extends $\mathbf{FO}$ with set variables X, Y,etc. in addition to the ordinary variables x, y,etc., and atomic formulas of the form $x \in X$ relating these two sorts of variables. The standard semantics of $\mathbf{MSO}(\sigma)$ is to interpret the set variables in $\mathcal{P}(A)$ for a σ-structure $\mathcal{A}$ with universe A and $\in$ as the set membership relation.

For any single-sorted signature σ, we can encode this standard semantics by using a functor $J_\in \colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma^{\mathbf{e}})$ from a single-sorted signature σ to a multi-sorted signature $\sigma^{\mathbf{e}}$. The signature $\sigma^{\mathbf{e}}$ is a S-sorted signature with sorts $S = \{\mathbf{f}, \mathbf{s}\}$ and $\sigma^{\mathbf{e}} = \sigma \cup \{\mathbf{e}\}$. For every r-ary relation $R \in \sigma$, the identical symbol in $\sigma^{\mathbf{e}}$ has arity given by the r-length tuple $\langle \mathbf{f}, \dots, \mathbf{f} \rangle$, the additional symbol $\mathbf{e}$ has arity $\langle \mathbf{f}, \mathbf{s} \rangle$.

The functor $J_\in \colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma^{\mathbf{e}})$ encodes the standard semantics for $\mathbf{MSO}$.

- On objects, for $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma))$, $J_{\in}(\mathcal{A})$ is the $\sigma^{\mathbf{e}}$ structure with universe given by the $\{\mathbf{f}, \mathbf{s}\}$-set with $A_{\mathbf{f}} = A$ and $A_{\mathbf{s}} = \mathcal{P}(A)$ (powerset of A) and relational interpretations
  - $\mathbf{e}^{J_{\in}(\mathcal{A})} = \{(a, S) \mid a \in S\}$
  - $R^{J_{\in}(\mathcal{A})} = R^{\mathcal{A}}$ for all $R \in \sigma$.
- On morphisms, a σ-morphism $f\colon \mathcal{A} \to \mathcal{B}$ induces a $\sigma^{\mathbf{e}}$-morphism with underlying $\{\mathbf{f}, \mathbf{s}\}$-function where $f_{\mathbf{f}} = f$ and $f_{\mathbf{s}}$ is equal to the direct image function $\mathcal{P}(f)\colon \mathcal{P}(A) \to \mathcal{P}(B)$. By definition of direct image, for all $a \in S \subseteq A$, $f_{\mathbf{f}}(a) \in f_{\mathbf{s}}(S)$, so $J_{\in}(f)$ is a $\sigma^{\mathbf{e}}$-morphism.

Typically **MSO** has equality on ordinary variables, so the standard semantics is encoded by the composed functor $J\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma^{\mathbf{e},+})$ where $J = J^{\mathbf{f}}_{=} \circ J_{\in}$.

There is a bijective correspondence between formulas of $\mathbf{FO}^{I,S}_k(\sigma^{\mathbf{e}})$ and formulas of $\mathbf{MSO}_k(\sigma)$. This bijective correspondence allows us to obtain the following obvious proposition.

Using Theorem 3.3.7 and the bijective syntactic translation between formulas of $\mathbf{MSO}_k(\sigma)$ and $\mathbf{FO}^{I,S}_k(\sigma^{\mathbf{e}})$, we obtain the following theorem demonstrating that the relative comonad $\mathbb{E}^S_k \circ J$ on $J\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma^{\mathbf{e},+})$ captures equivalence in $\mathbf{MSO}_k$.

Theorem 3.3.8. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B} \in \mathbf{Ob}(\mathsf{Struct}(\sigma))$:*

(1) $\mathcal{A} \equiv^{\mathbf{MSO}_k(\sigma)} \mathcal{B}$.

(2) Duplicator has a winning strategy in $\mathbf{EF}^{I,S}_k(J_{\in}(\mathcal{A}), J_{\in}(\mathcal{B}))$.

(3) There exists a bisimulation in $G^{\mathbb{E}^S_k}(J(\mathcal{A})) \leftarrow X \to G^{\mathbb{E}^S_k}(J(\mathcal{B}))$ *in* $\mathbf{EM}(\mathbb{E}^S_k)$.

Proof. For the (1) ⇔ (2) equivalence, there are syntactic translation maps $F\colon \mathbf{MSO}(\sigma) \to \mathbf{FO}^{I,S}(\sigma^{\mathbf{e}})$ and $M\colon \mathbf{FO}^{I,S}(\sigma^{\mathbf{e}}) \to \mathbf{MSO}(\sigma)$ such that

- F and M are inverses,
- for all $\varphi \in \mathbf{MSO}(\sigma)$, $\mathcal{A} \vDash \varphi \Leftrightarrow J(\mathcal{A}) \vDash F(\varphi)$, and conversely,
- for all $\varphi \in \mathbf{FO}^{I,S}_k(\sigma^{\mathbf{e}})$, $\mathcal{A} \vDash M(\varphi) \Leftrightarrow J(\mathcal{A}) \vDash \varphi$.

For more details on this translation, see Proposition A.0.14. Thus, $\mathcal{A} \equiv^{\mathbf{MSO}_k(\sigma)} \mathcal{B} \Leftrightarrow J_{\in}(\mathcal{A}) \equiv^{\mathbf{FO}^{I,S}_k(\sigma^{\mathbf{e}})} J_{\in}(\mathcal{B})$ where $I = \{\mathbf{f}\}$ and $S = \{\mathbf{f}, \mathbf{s}\}$. The desired equivalence then follows from the (1) ⇔ (2) equivalence of Theorem 3.3.7.

For the (2) $\Leftrightarrow$ (3) equivalence, this follows from the (2) $\Leftrightarrow$ (3) equivalence of Theorem 3.3.7 and $J = J^{\mathbf{f}}_{=} \circ J_{\in}$. $\square$

Chapter 4

Linear arboreal covers

In this chapter, we strengthen the notion of arboreal category to linear arboreal category. The notion of linear arboreal category adds two axioms to the definition of arboreal category that ensure that the branches of an object in our category, i.e. the path subobjects, are ‘full paths’ or ‘traces’.

With this notion in place, we demonstrate that any arboreal category $\mathcal{C}$ whose objects are strongly connected has a linear arboreal subcategory $\mathcal{C}^L$ related to $\mathcal{C}$ via an adjunction. Composing an arboreal cover of $\mathcal{E}$ by $\mathcal{C}$ with this adjunction, yields a linear arboreal cover of $\mathcal{E}$ by $\mathcal{C}^L$. The behavioural equivalences of Definition 3.1.19 when instantiated for one of these linear arboreal covers yields linear-time variants of the corresponding branching-time equivalence.

As the previous chapter demonstrated, many model-comparison games are instances of a notion of bisimulation. In this chapter, we generalise the relationship between bisimulation and its linear variant–trace equivalence. In particular, we show that the linear variant of each of these model comparison games can be phrased as an ‘all-in-one’ game.

4.1 Linear arboreal categories

In an arboreal category, since every object X is path-generated, X is a colimit cocone of its branches. For the notion of linear arboreal category, we would like to impose additional conditions to exclude ‘non-trivial branching’. Branching appears in two forms for an arboreal category $\mathcal{C}$. The first form is evident in the objects of $\mathcal{C}$. Namely, objects in $\mathcal{C}$ are colimits of their branch subobjects. Thus to exclude branching in the objects of $\mathcal{C}$, we additionally require that every object $X \in \mathbf{Ob}(\mathcal{C})$ is a coproduct of paths. The second form is evident in the morphisms of $\mathcal{C}$. In particular, path embeddings $p\colon P \rightarrowtail X$ in $\mathcal{C}$ isolate a partial branch of X that can be extended in possibly multiple ways, via extensions $j\colon P \rightarrowtail Q$, to longer partial branches $q\colon Q \rightarrowtail X$ where $p = q \circ j$. Thus to exclude branching in the morphisms of $\mathcal{C}$, we need a path embedding $P \rightarrowtail X$ to isolate a 'full path' or 'trace' (rather than partial path). To accomplish this, we add an axiom to ensure that paths in $\mathcal{C}$ have only trivial extension. We formalise these exclusions in the following axiomatic definition.

Definition 4.1.1. A *linear arboreal category* is an arboreal category $\mathcal{C}$ such that the following two axioms are satisfied:

(L1) Every object $X \in \mathbf{Ob}(\mathcal{C})$ is a coproduct

$$\biguplus_{P_i \in \mathbf{Ob}(\mathcal{C}_p)} P_i$$

of path objects P_i.

(L2) For any two non-initial paths P, Q if $j\colon P \rightarrowtail Q$ is an embedding, then $j\colon P \cong Q$ is an isomorphism.

We say an arboreal category $\mathcal{C}$ is a *quasi-linear arboreal category* if (L1) is satisfied.

One counterintuitive consequence of condition (L2) is that the path objects in a bijective or linear arboreal category, appear 'externally', i.e. from the perspective of the paths $\mathbf{P}\colon \mathcal{C} \to \mathcal{T}$ functor, rather trivial. This captured in the the following proposition:

Proposition 4.1.2. *If $\mathcal{C}$ is an arboreal category satisfying (L2) and $P \in \mathbf{Ob}(\mathcal{C})$ is a path, then $\mathsf{ht}(P) \leq 1$.*

Proof. If P is the initial path, then by the definition of height, $\mathsf{ht}(P) = 0$. If P is not the initial path, then we show that $\mathbf{P}(P)$ has a unique non-root element. Suppose $[j\colon P' \rightarrowtail P] \in \mathbf{P}(P)$ is a non-root element, then P' is not the initial path. Thus, by Axiom (L2), $j\colon P' \rightarrowtail P$ is an isomorphism, and $[j] = [\mathsf{id}_P]$. The equivalence class $[\mathsf{id}_P] \in \mathbf{P}(P)$ is the unique non-root element of $\mathbf{P}(P)$. Since $\mathbf{P}(P)$ has only one non-root element, $\mathsf{ht}(P) = 1$. $\square$

However, as will be demonstrated in the examples of Section 4.2, Proposition 4.1.2 hints at that idea that 'paths' in the linear arboreal setting can be seen has types for full behaviours, i.e. traces and, in terms of model comparison games, produce 'all-in-one' variants.

Another counterintuitive consequence of these axioms is that, bisimulations in $\mathcal{C}$ trivialise to bidirectional pathwise embeddings.

Corollary 4.1.3. *Let $\mathcal{C}$ be a linear arboreal category, and suppose $X, Y \in \mathcal{C}$, the categorical product $X \times Y$ exists in $\mathcal{C}$. The following are equivalent:*

(1) There exists pathwise embeddings $h\colon X \to Y$ and $h'\colon Y \to X$ in $\mathbf{Mor}(\mathcal{C})$.

(2) There exists a bisimulation $X \xleftarrow{f} Z \xrightarrow{g} Y$ in $\mathcal{C}$.

Proof. For the (1) ⇒ (2) direction, by Proposition 3.1.13(2), the existence of a bisimulation $X \xleftarrow{f} Z \xrightarrow{g} Y$ in $\mathcal{C}$ is equivalent to Duplicator having a winning strategy in the abstract back-and-forth game $\mathcal{G}(X, Y)$ over $\mathcal{C}$ described in Definition 3.1.12. Given pathwise embeddings $h\colon X \to Y$ and $h'\colon Y \to X$ in $\mathbf{Mor}(\mathcal{C})$, we can construct such a strategy. The game starts with roots $([\bot_X], [\bot_Y])$. If Spoiler chooses to extend $[\bot_X]$ with $[m\colon P \rightarrowtail X] \succ [\bot_X]$, then Duplicator responds with $[h \circ m\colon P \rightarrowtail Y]$. Clearly, $\mathsf{dom}(h \circ m) \cong \mathsf{dom}(m) \cong P$, so $([m], [h \circ m]) \in \mathcal{W}(X, Y)$ and Duplicator wins the round. Moreover, from Proposition 4.1.2 we see that every path in $\mathbf{P}(X)$ is of height at most 1, so Spoiler has no more extensions of $[m]$ and Duplicator wins the whole game $\mathcal{G}(X, Y)$ over $\mathcal{C}$. Similarly, if Spoiler chooses to extend $[\bot_Y]$ with $[n\colon P \rightarrowtail Y] \succ [\bot_Y]$, then Duplicator responds with $[h' \circ n\colon P \to X] \succ [\bot_X]$. As with the previous case, we can conclude that Duplicator wins the whole game $\mathcal{G}(X, Y)$.

Conversely, the (2) ⇒ (1) is just the statement of Proposition 3.1.15. ☐

4.1.1 The linear adjunction

In this section, we construct a linear arboreal subcategory $\mathcal{C}^L$ from any arboreal category $\mathcal{C}$ which satisfies a linearisable condition. The linear arboreal subcategory $\mathcal{C}^L$ is related to $\mathcal{C}$ via a right adjoint to the inclusion functor $I\colon \mathcal{C}^L \to \mathcal{C}$. Using this adjunction, we can derive linear variants of many of the corresponding constructions in $\mathcal{C}$.

To define the linearisable condition, we simply strengthen to notion of connectedness in Axiom (3) of Definition 3.1.6 to the usual categorical notion.

Definition 4.1.4. An object $X \in \mathbf{Ob}(\mathcal{C})$ is *strongly connected* if every morphism $X \to \biguplus_{i \in I} P_i$, where $\{P_i\}$ is a family of paths, factors through a unique coproduct injection $P_j \to \biguplus_{i \in I} P_i$.

The following condition, which we can define for any path category, is sufficient to guarantee that an arboreal category $\mathcal{C}$ contains an linear arboreal subcategory $\mathcal{C}^L$.

Definition 4.1.5. A path category is *linearisable* if every non-initial path $P \in \mathcal{C}$ is strongly connected.

The linearisable condition is independent from the axioms of a path category. The following is an example, suggested by Luca Reggio to the author, of a path category which is not linearisable.

Counterexample 4.1.6. Suppose $k \geq 2$, we define a full subcategory $\mathcal{HT}_k$ of $\mathcal{T}$ where the objects $(T, \leq)$ of $\mathcal{HT}_k$ have height $\leq k$ and satisfy the following condition:

(HT) For all $x, y \in T$, if $\mathsf{ht}(x) = \mathsf{ht}(y)$, then $x = y$.

In other words, for every $h \leq k$, there is at most one $x \in T$ such that $\mathsf{ht}(x) = h$. The coproduct in $\mathcal{HT}_k$ maintains the (HT) condition by merging elements of equal height from each component. As $\mathcal{HT}_k$ is a full subcategory of $\mathcal{T}$ closed under quotients, $\mathcal{HT}_k$ inherits the property of being a path category from $\mathcal{T}$. However, $\mathcal{HT}_3$ is not a linearisable path category. Namely, observe that if P is a path of height 2 and Q is a path of height 3, then there is a tree morphism $P \to P \uplus Q$ which factors both the through the injection $P \to P$ and the injection $P \to Q$. This is because in $P \uplus Q$ elements of Q of height ≤ 2 are merged with the elements of P.

For every linearisable arboreal category $\mathcal{C}$, we construct a subcategory $\mathcal{C}^L$ which is a linear arboreal category. To construct $\mathcal{C}^L$, we first a construct the quasi-linear arboreal $\mathcal{C}^{qL}$ from $\mathcal{C}$ by restricting the objects of $\mathcal{C}$ to those which are generated by their maximal paths.

Definition 4.1.7. A path embedding $m \colon P \rightarrowtail X$ is maximal if P is not initial and for all $n \colon P' \rightarrowtail X$ such that $m \trianglelefteq n$, $m \sim n$.

Equivalently, $m \colon P \rightarrowtail X$ is maximal if the $\sim$-equivalence class $[m]$ is a maximal element in the poset $\mathbf{P}(X)$. For every object X in a path category, let $\mathbf{P}^{\top}(X)$ be the subset of maximal elements in $\mathbf{P}(X)$. With the definition of maximal in place, we obtain $\mathcal{C}^{qL}$ an objects which are generated by their maximal paths:

Definition 4.1.8. An object $X \in \mathcal{C}$ of a path category $\mathcal{C}$ is *linearly path-generated* if

$$X \cong \biguplus_{[m] \in \mathbf{P}^{\top}(X)} \mathsf{dom}(m).$$

The following lemmas will be useful when working with linearly path-generated objects in a linearisable arboreal category $\mathcal{C}$.

Lemma 4.1.9. *If $\mathcal{C}$ is a linearisable arboreal category and $X \in \mathbf{Ob}(\mathcal{C})$ is linearly path-generated, then for every maximal $[m] \in \mathbf{P}^{\top}(X)$, $m = i_{[m]}$ where $i_{[m]} \colon \mathsf{dom}(m) \to X$ is the coproduct injection.*

Proof. By the universal property of the coproduct X, there exists a unique morphism $u\colon X \to X$ such that for all components $\mathsf{dom}(m)$ of X indexed by $[m] \in \mathbf{P}^{\top}(X)$, $u \circ i_{[m]} = m$. Since $u \circ i_{[m]} = m$ and $m\colon \mathsf{dom}(m) \to X$ is an embedding, by Proposition 2.3.9(e), $i_{[m]}\colon \mathsf{dom}(m) \to X$ is an embedding for every $[m] \in \mathbf{P}^{\top}(X)$. As m is a maximal path embedding, $\mathsf{dom}(m)$ is a non-initial path. By $\mathcal{C}$ being linearisable, $\mathsf{dom}(m)$ is strongly connected, so $m\colon \mathsf{dom}(m) \to X$ factors through a unique coproduct injection $i_{[n]}\colon \mathsf{dom}(n) \rightarrowtail X$ obtaining the equality $m = i_{[n]} \circ j_m$. The morphism $j_m\colon \mathsf{dom}(m) \to \mathsf{dom}(n)$ is an embedding by Proposition 2.3.9(e). Since m is a maximal path and $i_{[n]}\colon \mathsf{dom}(n) \rightarrowtail X$ is an embedding, j_m is an isomorphism, and $[m] = [i_{[n]}]$. Up to isomorphism, we can replace the component $\mathsf{dom}(n)$ of X with $\mathsf{dom}(m)$ and conclude that $m = i_{[m]}$. ☐

Lemma 4.1.10. *If $\mathcal{C}$ is a linearisable arboreal category, $X \in \mathbf{Ob}(\mathcal{C})$ is linearly path-generated, and $[p\colon P \rightarrowtail X] \in \mathbf{P}(X)$ is not the root, then there exists a unique maximal element $[m] \in \mathbf{P}^{\top}(X)$, such that $[p] \leq [m]$.*

Proof. Since $[p\colon P \rightarrowtail X] \in \mathbf{P}(X)$ is not the root, P is a non-initial path. By $\mathcal{C}$ being linearisable, P is strongly connected. Since X is linearly path-generated, the path embedding $p\colon P \rightarrowtail X$ factors through a unique coproduct injection $p = i_{[m]} \circ z$ for some $z\colon P \to Q$. By Lemma 4.1.9, $m = i_{[m]}$, so $p = m \circ z$. By p being a path embedding and Proposition 2.3.9(e), z is an embedding. Therefore, $[p] \leq [m]$. ☐

The next step in our construction of $\mathcal{C}^L$ is to obtain $\mathcal{C}^L$ from $\mathcal{C}^{qL}$ by restricting to those morphisms of $\mathcal{C}^{qL}$ that preserve maximal elements. Formally, such morphisms are defined as follows:

Definition 4.1.11. A morphism $f\colon X \to Y \in \mathbf{Mor}(\mathcal{C})$ is a *leaf morphism in* $\mathcal{C}$ if for every $[m] \in \mathbf{P}^{\top}(X)$, $\mathbf{P}(f)([m]) = [\exists_f m] \in \mathbf{P}^{\top}(Y)$.

Thus, $\mathcal{C}^L$ is the subcategory of $\mathcal{C}$ where the objects are the linearly-path generated objects of $\mathcal{C}$ and the morphisms are the leaf morphisms of $\mathcal{C}$.

Proposition 4.1.12. *If $\mathcal{C}$ is a linearisable arboreal category, then $\mathcal{C}^L$ is a linear arboreal category. In particular, $\mathcal{C}^{qL}$ is a quasi-linear arboreal category.*

Proof. By definition, the objects $\mathcal{C}^{qL}$ and $\mathcal{C}^L$ are the objects of $\mathcal{C}$ which are coproducts of their maximal paths. Thus, $\mathcal{C}^{qL}$ and $\mathcal{C}^L$ satisfy Axiom (L1) of Definition 4.1.1.

To show that $\mathcal{C}^L$ satisfies Axiom (L2) of Definition 4.1.1, suppose P, Q are non-initial paths and $m\colon P \rightarrowtail Q \in \mathbf{Mor}(\mathcal{C}^L) \subseteq \mathbf{Mor}(\mathcal{C})$, we need to show that $P \cong Q$. Consider the $(\mathcal{Q}, \mathcal{M})$ factorisation of $m \circ \mathsf{id}_P$ into quotient $q\colon P \twoheadrightarrow \exists_m P$ and embedding $\exists_m \mathsf{id}_P\colon \exists_m P \to$

Q. Since $m\colon P \rightarrowtail Q \in \mathbf{Mor}(\mathcal{C}^L)$ is a leaf morphism, $[\exists_m \mathsf{id}_P \colon \exists_m P \to Q] \in \mathbf{P}^\top(Q)$. As Q is a path, $[\mathsf{id}_Q]$ is the unique maximal element in $\mathbf{P}(Q)$, so there exists an isomorphism i such that i such that $\exists_m \mathsf{id}_P = i \circ \mathsf{id}_Q$. Hence, $\exists_m \mathsf{id}_P = i$ is an isomorphism $\exists_m P \cong Q$.

On the other hand, since $m = m \circ \mathsf{id}_P = \exists_m \mathsf{id}_P \circ q$ is an embedding, by Proposition 2.3.9(e), the quotient q is also an embedding. By Proposition 2.3.9(b), q is an isomorphism $P \cong \exists_m P$.

Composing these two isomorphisms allows us to conclude that $P \cong Q$ as desired. ☐

Since leaf morphisms preserve maximal elements, the object mapping $X \mapsto \mathbf{P}^\top(X)$ extends to a functor $\mathbf{P}^\top \colon \mathcal{C}^L \to \mathbf{Set}$.

Proposition 4.1.13. *If $\mathcal{C}$ is a linearisable arboreal category, $X \mapsto \mathbf{P}^\top(X)$ induces a functor $\mathbf{P}^\top \colon \mathcal{C}^L \to \mathbf{Set}$ from the linear arboreal subcategory $\mathcal{C}^L$ of $\mathcal{C}$.*

Proof. Follows immediately from the definition of leaf morphism. ☐

For the linear arboreal subcategory $\mathcal{C}^L$ of an arboreal category $\mathcal{C}$, there is an inclusion $I\colon \mathcal{C}^L \hookrightarrow \mathcal{C}$. Paths of an object in $X \in \mathbf{Ob}(\mathcal{C}^L)$ are essentially the maximal paths of $I(X) \in \mathbf{Ob}(\mathcal{C})$.

Proposition 4.1.14. *For every object $X \in \mathbf{Ob}(\mathcal{C}^L)$ and non-initial path P, if $[m\colon P \rightarrowtail X] \in \mathbf{P}(X)$, then $[I(m)] \in \mathbf{P}^\top(I(X))$.*

Proof. The function $\mathbf{P}(I(m))$ maps the unique maximal path $[\mathsf{id}_{I(P)}]$ in $\mathbf{P}^\top(I(P))$ to

$$[\exists_{I(m)} \mathsf{id}_{I(P)}] \overset{*}{=} [I(m) \circ \mathsf{id}_{I(P)}] = [I(m)]$$

where the starred equality follows from $I(m)$ being a pathwise embedding. Since $m\colon P \rightarrowtail X \in \mathbf{Mor}(\mathcal{C}^L)$, $I(m)$ is leaf morphism of $\mathcal{C}$. By definition of leaf morphism, $\mathbf{P}(I(m)) = [I(m)]$ is a maximal element in $\mathbf{P}^\top(I(X))$. ☐

The inclusion $I\colon \mathcal{C}^L \to \mathcal{C}$ has a right adjoint $T\colon \mathcal{C} \to \mathcal{C}^L$. The right adjoint T will be central for computing from a game comonad its linear variant. To construct the object mapping of T, suppose X is an object of arboreal category $\mathcal{C}$, we define $T(X)$ as the coproduct in $\mathcal{C}$ of all paths of X. In notation,

$$T(X) = \biguplus_{[p] \in \mathbf{P}(X)} \mathsf{dom}(p).$$

For every $X \in \mathcal{C}$, by the universal property of the coproduct $I(T(X))$, there exists a unique morphism $\varepsilon^L_X \colon I(T(X)) \to X$ such that for all $[p] \in \mathbf{P}(X)$, $p = \varepsilon^L_X \circ i_{[p]}$ where

$i_{[p]}\colon \mathsf{dom}(p) \to I(T(X))$ is the coproduct injection. The following theorem demonstrates that for every object $X \in \mathcal{C}$, $T(X)$ and ε^L_X is a universal morphism from I to X. Thus, T extends to a functor $T\colon \mathcal{C} \to \mathcal{C}^L$ and is right adjoint to I.

Theorem 4.1.15. *For every linearisable arboreal category $\mathcal{C}$, $I \dashv T$ is an adjunction from $\mathcal{C}^L$ to $\mathcal{C}$.*

Proof. Suppose $f\colon I(X) \to Y \in \mathbf{Mor}(\mathcal{C})$, we need to construct a unique morphism $g\colon X \to T(Y) \in \mathbf{Mor}(\mathcal{C}^L)$ satisfying $\varepsilon^L_Y \circ I(g) = f$. By the definition of $\mathbf{Mor}(\mathcal{C}^L)$, it suffices to construct a unique leaf morphism $I(g)\colon I(X) \to I(T(Y))$ in $\mathbf{Mor}(\mathcal{C})$ such that $\varepsilon^L_Y \circ I(g) = f$. Since $X \in \mathbf{Ob}(\mathcal{C}^L)$, we know that $I(X)$ is linearly path-generated in $\mathcal{C}$ and thus, a coproduct of $Q = \mathsf{dom}(m)$ ranging over maximal path embeddings $[m\colon Q \rightarrowtail I(X)]$. Recall that for every such maximal element $[m]$, we have the factorisation of $f \circ m$ into:

$$Q \xrightarrow{q_m} \mathrel{\!\!\twoheadrightarrow} \exists_f Q \overset{n_m}{\rightarrowtail} Y \tag{4.1}$$

such that $[n_m] = \mathbf{P}f([m]) \in \mathbf{P}(Y)$. Since $[n_m] \in \mathbf{P}(Y)$, we have corresponding component of $T(Y)$ and coproduct injection $j_m\colon \exists_f Q \to T(Y)$ such that $\varepsilon^L_Y \circ j_m = n_m$. Thus, for every component Q of the coproduct $I(X)$, we have morphism $j_m \circ q_m\colon Q \to T(Y)$. By the universal property of the coproduct $I(X)$, we have a unique morphism $I(g)\colon I(X) \to I(T(Y))$ such that $I(g) \circ i_{[m]} = j_m \circ q_m$.

To verify that $\varepsilon^L_Y \circ I(g) = f$, we compute the following equality for every component $\mathsf{dom}(m) = Q$ of the coproduct $I(X)$:

$$\begin{aligned}
(\varepsilon^L_Y \circ I(g)) \circ i_{[m]} &= \varepsilon^L_Y \circ (I(g) \circ i_{[m]}) && \\
&= \varepsilon^L_Y \circ (j_m \circ q_m) && I(g) \circ i_{[m]} = j_m \circ q_m \\
&= (\varepsilon^L_Y \circ j_m) \circ q_m && \\
&= n_m \circ q_m && \varepsilon^L_Y \circ j_m = n_m \\
&= f \circ m && \text{Equation (4.1)} \\
&= f \circ i_{[m]} && \text{Lemma 4.1.9}
\end{aligned}$$

Thus, by the universal property of coproduct $I(X)$, $\varepsilon^L_Y \circ I(g) = f$.

Combining this equation, the universal property of $I(g)$, and $m = i_{[m]}$ from Lemma 4.1.9, we obtain the following commutative diagram

$$\begin{tikzcd}
Q \arrow[r, two heads, "q_m"] \arrow[d, tail, "m"] & \exists_f Q \arrow[d, "j_m"] \arrow[dr, tail, "n_m"] & \\
I(X) \arrow[r, "I(g)"] \arrow[rr, bend right, "f"'] & I(T(Y)) \arrow[r, "\varepsilon^L_Y"] & Y
\end{tikzcd} \tag{4.2}$$

To verify that $I(g)$ is a leaf morphism, consider maximal path embedding $m\colon Q \to I(X)$ and suppose there exists an $n\colon Q' \to I(T(Y))$ such that $\mathbf{P}(I(g))([m]) \leq [n]$. By equation (4.2), n_m being an embedding, and Proposition 2.3.9 (e), j_m is an embedding and $\mathbf{P}(I(g))([m]) = [j_m\colon \exists_f Q \rightarrowtail I(T(Y))]$. To prove maximality of $\mathbf{P}(I(g))([m])$, we need to show that $[j_m] = [n]$. By the supposition $[j_m] \leq [n]$, there exists an embedding $z\colon \exists_{I(g)}Q \rightarrowtail Q'$ such that the following triangle commutes:

$$\begin{tikzcd}
\exists_{I(g)}Q \arrow[rr, tail, "z"] \arrow[dr, tail, "j_m"'] & & Q' \arrow[dl, tail, "n"] \\
& I(T(Y)) &
\end{tikzcd}$$

By $\mathcal{C}$ being linearisable and Q is a non-initial path, $\exists_f Q$ and Q' are strongly connected. By Q' being strongly connected and $I(T(X))$ being a coproduct of paths, there exists a unique coproduct injection $j_n\colon R \to I(T(X))$ such that $n = j_n \circ w$ for some $w\colon Q' \to R$. However, from the equations $n = j_n \circ w$, $j_m = n \circ z$, we must conclude that $j_m = j_n \circ w \circ z$ and so the coproduct injection j_m factors through the coproduct injection j_n. By $\exists_f Q$ being strongly connected, we conclude that $j_n = j_m$. Thus, we have that $n = j_m \circ w$, $R = \exists_f Q$, and $w\colon Q' \to \exists_f Q$. By n being embedding and Proposition 2.3.9 (e), w is embedding and $[j_m] \geq [n]$. Combining with our supposition $[j_m] \leq [n]$, we obtain that $[j_m] = [n]$.

To verify the uniqueness of g, suppose there exists a $g'\colon X \to T(Y) \in \mathbf{Mor}(\mathcal{C}^L)$ such that $\varepsilon^L_Y \circ I(g') = f$. For every path $[p] \in \mathbf{P}(Y)$, $p = \varepsilon^L_Y \circ i_{[p]}$. In particular, for every $[m] \in \mathbf{P}^\top(I(X))$, for the path

$$[n] = \mathbf{P}(\varepsilon^L_Y \circ I(g'))([m]) = \mathbf{P}(\varepsilon^L_Y \circ I(g))([m]) = \mathbf{P}(f)([m]) \in \mathbf{P}(Y),$$

we have that $n = \varepsilon^L_Y \circ i_{[n]}$ where $i_{[n]}\colon \mathsf{dom}(n) \to I(T(Y))$ is the coproduct injection. Thus, $I(g)$ and $I(g')$ map m onto the same component $\mathsf{dom}(n)$ of $I(T(Y))$. Since $I(g)$ and $I(g')$ are leaf morphisms, and $[m] \in \mathbf{P}^\top(X)$ is a maximal path, g and g' must map $[m]$ to the maximal path $[i_{[n]}] \in \mathbf{P}(I(T(Y)))$ corresponding to the full component $\mathsf{dom}(n)$ of $I(T(Y))$. By the universal property of the coproduct $I(X)$, $I(g) = I(g')$. As I is a faithful functor, $g = g'$. $\square$

Throughout the rest of this chapter, $\eta^L\colon \mathsf{Id} \to T \circ I$ and $\varepsilon^L\colon I \circ T \to \mathsf{Id}$ denote the unit and counit, respectively, of the adjunction $I \dashv T$.

As demonstrated in Chapter 3 the objective of arboreal categories $\mathcal{C}$ is to view its tree shaped objects as covers of objects in another extensional category $\mathcal{E}$. Consequently, we would like to use the linear adjunction $I \dashv T$ to produce an linear arboreal cover of $\mathcal{E}$ by $\mathcal{C}^L$ from an arboreal cover $\mathcal{E}$ by $\mathcal{C}$.

Given the arboreal cover $L \dashv R$ of $\mathcal{E}$ by a linearisable arboreal category $\mathcal{C}$, we can compose this adjunction with the adjunction $I \dashv T$ to obtain an adjunction $L \circ I \dashv T \circ R$ between $\mathcal{E}$ and $\mathcal{C}^L$. However, a priori, this adjunction may not be comonadic, and therefore may not necessarily yield a linear arboreal cover of $\mathcal{E}$ by $\mathcal{C}^L$. In Theorem 4.1.19, we show that this composed adjunction is indeed comonadic. The proof uses Beck's comonadicity theorem which gives necessary and sufficient conditions for when a functor F yields a comonadic adjunction. One of the conditions in Beck's theorem involves demonstrating that equalisers of split pairs are created by F. For this reason, we first prove some lemmas about equalisers in arboreal categories.

Lemma 4.1.16. *Let $(\mathcal{C}, \mathcal{Q}, \mathcal{M})$ be an factorised category. If $o\colon O \to X$ equalises $f, g\colon X \to Y$ and $(e_o\colon O \twoheadrightarrow \hat{O}, m_o\colon \hat{O} \rightarrowtail X)$ is the $(\mathcal{Q}, \mathcal{M})$-factorisation of o, then $m_o\colon \hat{O} \rightarrowtail X$ equalises f, g.*

Proof.

$$\begin{aligned}
f \circ o &= g \circ o && o \text{ equalises } f, g\\
f \circ (m_o \circ e_o) &= g \circ (m_o \circ e_o) && o = m_o \circ e_o \text{ factorisation}\\
(f \circ m_o) \circ e_o &= (g \circ m_o) \circ e_o &&\\
f \circ m_o &= g \circ m_o && \text{quotient } e_o \text{ is an epimorphism}
\end{aligned}$$

□

For the next two lemmas, we consider for all objects X in a linear arboreal category, the subset $Z \subseteq \mathbf{P}(X)$ of elements whose representatives equalise f, g, i.e. $z = [p\colon P \rightarrowtail X] \in Z$ if $f \circ p = g \circ p$. Let $Z^\top$ be the elements which are maximal in the induced sub-poset of Z in $\mathbf{P}(X)$, i.e. $y \in Z^\top$ if $y \in Z$ and for all $z \in Z$ such that $y \leq z$, $y = z$.

Lemma 4.1.17. *For every $x \in Z$ such that $x \neq [\bot\colon \tilde{0} \rightarrowtail X]$ there exists a unique $t \in Z^\top$ such that $x \leq t$.*

Proof. By Lemma 4.1.10, every such $x \in Z$ is contained within a unique maximal element $m \in \mathbf{P}^\top(X)$. This allows us to partition Z into non-empty sets $Z_m = \downarrow m \cap Z$ where $m \in \mathbf{P}^\top(X)$ and there exists an $x \in Z$ with $x \leq m$. As $\downarrow m$ is a finite chain, Z_m is a finite subset of a chain, and there exists a maximal element $t_m \in Z_m$. To verify that $t_m \in Z^\top$, suppose that $z \in Z$ and $t_m \leq z$. It must be the case that $z \leq m$, if not then $t_m \leq z \leq m'$ for $m' \neq m \in \mathbf{P}^\top(X)$ which would contradict Lemma 4.1.10. Therefore, $z \in Z_m$, and by the maximality of t_m, $z = t_m$. As $t_m \in Z^\top$, we set $t = t_m$ and the statement of the theorem holds. $\square$

Lemma 4.1.18. *If (E, e) is the equaliser of morphisms $I(f), I(g)\colon I(X) \to I(Y)$ in $\mathcal{C}$ for $X, Y \in \mathcal{C}^L$, then*

(1) there is an equaliser (E', e') of $f, g\colon X \to Y$ in $\mathcal{C}^L$, and

(2) $(E, e) \cong (I(E'), I(e'))$ in $\mathcal{C}$.

Proof. Take E' to be the coproduct:

$$E' = \biguplus_{[q\colon Q \rightarrowtail X] \in Z^\top} Q$$

Since E' is a coproduct of paths, $E' \in \mathcal{C}^L$. As every component of E' has a morphism $Q \rightarrowtail X$, by the universal property of the coproduct E', there exists a unique morphism $e'\colon E' \to X$.

Both statement (1) and (2) will rely on showing that $(I(E'), I(e'))$ is an equaliser of $I(f), I(g)\colon I(X) \to I(Y)$ in $\mathcal{C}$. Since $q\colon Q \rightarrowtail X$ equalises f, g for all $[q\colon Q \rightarrowtail X] \in Z^\top$, it follows from the universal morphism property of coproducts that $e'\colon E' \rightarrowtail X$ equalises f, g. By applying I, we see that $I(e')$ equalises $I(f), I(g)$.

Now suppose that $o\colon O \to I(X)$ also equalises $I(f), I(g)$, we need to construct a unique morphism $u\colon O \to I(E')$. We will accomplish this by constructing a cocone over $I(E')$ from the diagram of path embeddings into O. To begin, we will first consider the case where the embedding $p\colon P \rightarrowtail O$ is not the root, i.e. P is non-initial. Recall that $o\colon O \to I(X)$ induces a mapping which sends a path embedding $p\colon P \rightarrowtail O$ to a path embedding $m_p\colon \exists_o P \rightarrowtail I(X)$ originating from the factorisation $(e_p\colon P \twoheadrightarrow \exists_o P, m_p\colon \exists_o P \rightarrowtail I(X))$ of $o \circ p\colon P \to I(X)$. Since o equalises $I(f), I(g)$, $o \circ p$ equalises $I(f), I(g)$, and by Lemma 4.1.16, m_p equalises $I(f), I(g)$. Since $[p] \neq [\bot]$ and $[m_p] \neq [\bot]$, by Lemma 4.1.17, there exists a unique $[q\colon Q \rightarrowtail X] \in Z^\top$ such that $[m_p] \leq [q]$ in $\mathbf{P}X$, i.e. there exists an embedding $j_p\colon \exists_o P \rightarrowtail Q$ such that $m_p = q \circ j_p$. Let $z_p = i_{[q]} \circ j_p \circ e_p$ where $i_{[q]}$ is the unique coproduct injection from the component $[q] \in Z^\top$ to $I(E')$. Further, if $p'\colon P' \rightarrowtail O$ is a path embedding such that $p = p' \circ v$ for some embedding $v\colon P \rightarrowtail P'$, then by the function $\mathbf{P}o\colon \mathbf{P}O \to \mathbf{P}X$ being monotone, there exists an embedding $d_v\colon \exists_o P \rightarrowtail \exists_o P'$

and $[m_p] \leq [m'_p] \leq [q]$. It follows that $z_p = z_{p'} \circ v$:

$$\begin{aligned}
p &= p' \circ v \\
o \circ p &= o \circ p' \circ v \\
m_p \circ e_p &= m_{p'} \circ e_{p'} \circ v && \text{factorisation of } o \circ p, o \circ p' \\
(q \circ j_p) \circ e_p &= (q \circ j_{p'}) \circ e_{p'} \circ v && m_p = q \circ j_p, m_{p'} = q \circ j_{p'} \\
q \circ (j_p \circ e_p) &= q \circ (j_{p'} \circ e_{p'} \circ v) \\
j_p \circ e_p &= j_{p'} \circ e_{p'} \circ v && \text{embedding } q \text{ is a monomorphism} \\
i_{[q]} \circ j_p \circ e_p &= i_{[q]} \circ j_{p'} \circ e_{p'} \circ v \\
z_p &= z_{p'} \circ v && \text{definition of } z_p, z_{p'}
\end{aligned}$$

For the root path embedding $\bot\colon 0 \rightarrowtail O$, we take $z_\bot\colon 0 \to I(E')$ to be the unique morphism from the initial object 0. Hence, from the cocone of path embeddings over O, we have constructed a cocone over $I(E')$. Since O is path-generated, O is a colimit cocone and there exists a unique morphism $u\colon O \to I(E')$.

Thus, we have shown that $I(E')$ is an equaliser of $I(f), I(g)$ in $\mathcal{C}$. Statement (1) that E' is an equaliser of f, g is a special case of the above proof where O is assumed to be linearly path-generated. Statement (2) that $(I(E'), e') \cong (E, e)$ follows from the hypothesis that (E, e) is an equaliser and that equalisers are unique up to isomorphism. □

Finally, with Lemmas 4.1.16, 4.1.17 and 4.1.18 in place, we can prove that the composed adjunction $L \circ I \dashv T \circ R$ is comonadic.

Theorem 4.1.19. *If $L \dashv R$ is an arboreal cover of $\mathcal{E}$ by $\mathcal{C}$, then $L \circ I \dashv T \circ R$ is a linear arboreal cover of $\mathcal{E}$ by $\mathcal{C}^L$.*

Proof. For this theorem, it suffices to show that if $L\colon \mathcal{C} \to \mathcal{E}$ is a comonadic functor, then $L \circ I\colon \mathcal{C}^L \to \mathcal{E}$ is a comonadic functor. By the dual of Beck's monadicity theorem, a functor $F\colon \mathcal{A} \to \mathcal{B}$ is comonadic if and only if:

(1) $F\colon \mathcal{A} \to \mathcal{B}$ has a right adjoint.

(2) F reflects isomorphisms.

(3) $\mathcal{A}$ has equalisers of F-split pairs, and F preserves those equalisers. More precisely, for every diagram

$$\begin{tikzcd} X \arrow[r, "e"] & F(B) \arrow[r, shift left, "F(f)"] \arrow[r, shift right, "F(g)"'] & F(C) \end{tikzcd}$$

in $\mathcal{B}$ that can be embedded in the commutative diagram

$$\begin{tikzcd} X \arrow[r, "e"] & F(B) \arrow[r, shift left, "F(f)"] \arrow[r, shift right, "F(g)"'] \arrow[l, bend left, "s"] & F(C) \arrow[l, bend left, "t"] \end{tikzcd} ,$$

the following diagram is an equaliser in $\mathcal{A}$:

$$\begin{tikzcd} Y \arrow[r, "e'"] & B \arrow[r, shift left, "f"] \arrow[r, shift right, "g"'] & C \end{tikzcd} .$$

Moreover, $(F(Y), F(e')) \cong (X, e)$.

To obtain our proposition, we must show that if L satisfies conditions (1)-(3), then $L \circ I$ satisfies conditions (1)-(3). We verify that $L \circ I$ satisfies these conditions:

(1) The composition of adjoint $L \dashv R$ with adjoint $I \dashv T$ yields an adjoint $L \circ I \dashv T \circ R$. Therefore, $L \circ I$ has a right adjoint $T \circ R$.

(2) Since L satisfies condition (2), L reflects isomorphisms. Further, as I is a fully-faithful functor, I reflects isomorphisms. Therefore, $L \circ I$ reflects isomorphisms.

(3) Consider the following split equaliser in $\mathcal{E}$:

$$\begin{tikzcd} X \arrow[r, "e"] & L(I(B)) \arrow[r, shift left, "L(I(f))"] \arrow[r, shift right, "L(I(g))"'] \arrow[l, bend left, "s"] & L(I(C)) \arrow[l, bend left, "t"] \end{tikzcd} .$$

By L satisfying condition (3), the following is an equaliser in $\mathcal{C}$:

$$\begin{tikzcd} Y \arrow[r, "e'"] & I(B) \arrow[r, shift left, "I(f)"] \arrow[r, shift right, "I(g)"'] & I(C) \end{tikzcd} .$$

Moreover, $(L(Y), L(e')) \cong (X, e)$. By Lemma 4.1.18, there is $Z \in \mathcal{C}^L$ such that the following is an equaliser in $\mathcal{C}^L$:

$$\begin{tikzcd} Z \arrow[r, "e''"] & B \arrow[r, shift left, "f"] \arrow[r, shift right, "g"'] & C \end{tikzcd} .$$

and $(I(Z), I(e''))) \cong (Y, e')$. Therefore $(L(I(Z)), L(I(e''))) \cong (L(Y), L(e')) \cong (X, e)$. Hence, $L \circ I$ satisfies condition (3).

By Beck's comonadicity theorem, $L \circ I$ is comonadic. □

The following proposition allows us to transfer paths from X to their corresponding paths in $T(X)$.

Proposition 4.1.20. *For every object $X \in \mathcal{C}$ in a linearisable path category $\mathcal{C}$, $\varepsilon^L_X \colon I(T(X)) \to X$ is a pathwise embedding.*

Proof. Suppose $p \colon P \rightarrowtail I(T(X))$ is a path embedding, we must show that $\varepsilon^L_X \circ p \colon P \rightarrowtail X$ is a path embedding. By construction $I(T(X))$ is a coproduct of small family of paths, and so by connectedness axiom of path categories (Axiom (3) of Definition 3.1.6), $p \colon P \rightarrowtail I(T(X))$ factors through a coproduct injection $i_{[m]} \colon Q \to I(T(X))$. That is, there exists a morphism $j \colon P \to Q$ such that $i_{[m]} \circ j = p$. Since p is an embedding, by Proposition 2.3.9(e), j is an embedding.

$$\begin{aligned} \varepsilon^L_X \circ p &= \varepsilon^L_X \circ (i_{[m]} \circ j) && p = i_{[m]} \circ j \\ &= (\varepsilon^L_X \circ i_{[m]}) \circ j && \\ &= m \circ j && \varepsilon^L_X \circ i_{[m]} = m \end{aligned}$$

As m and j are embeddings, by Proposition 2.3.9 (a), $m \circ j \colon P \rightarrowtail X$ is an embedding. □

Proposition 4.1.21. *If $g \colon X \to Y \in \mathbf{Mor}(\mathcal{C})$ is a pathwise embedding, then $T(g) \colon T(X) \to T(Y) \in \mathbf{Mor}(\mathcal{C}^L)$ is a pathwise embedding.*

Proof. To show that $T(g)$ is a pathwise embedding in $\mathcal{C}^L$, consider a path embedding $p \colon P \rightarrowtail T(X) \in \mathbf{Mor}(\mathcal{C}^L)$. Using the ε^L-naturality square of g and precomposing with path embedding $I(p)$, we obtain the equation

$$\varepsilon^L_Y \circ I(T(g)) \circ I(p) = g \circ \varepsilon^L_X \circ I(p)$$

Since g is a pathwise embedding, by hypothesis, and ε^L_X is a pathwise embedding, by Proposition 4.1.20, the composition on the right hand side is a path embedding. Thus, $\varepsilon^L_Y \circ I(T(g)) \circ I(p)$ is a path embedding. By Proposition 2.3.9(e), $I(T(g)) \circ I(p) = I(T(g) \circ p)$ is a path embedding. As I is faithful, $T(g) \circ p$ is a path embedding. By definition, $T(g)$ is a pathwise embedding. □

Proposition 4.1.22. *If $g \colon X \to Y \in \mathbf{Mor}(\mathcal{C})$ is a open map, then $T(g) \colon T(X) \to T(Y) \in \mathbf{Mor}(\mathcal{C}^L)$ is an open map.*

Proof. To show that $T(g)$ is an open map, consider the following commutative square in $\mathcal{C}^L$:

$$\begin{tikzcd} P \arrow[r, tail, "j"] \arrow[d, tail, "p"'] & Q \arrow[d, tail, "q"] \\ T(X) \arrow[r, "T(g)"'] & T(Y) \end{tikzcd} \tag{4.3}$$

Applying the inclusion I and composing with the ε^L-naturality square of g, we obtain the commutative square (4.4)-left in $\mathcal{C}$.

$$\begin{tikzcd}
I(P) \arrow[r, tail, "I(j)"] \arrow[d, tail, "I(p)"'] & I(Q) \arrow[d, tail, "I(q)"] \\
I(T(X)) \arrow[r, "I(T(g))"'] \arrow[d, "\varepsilon^L_X"'] & I(T(Y)) \arrow[d, "\varepsilon^L_Y"] \\
X \arrow[r, "g"'] & Y
\end{tikzcd}
\qquad
\begin{tikzcd}
I(P) \arrow[r, tail, "I(j)"] \arrow[d, tail, "I(p)"'] & I(Q) \arrow[d, tail, "I(q)"] \arrow[ddl, dashed, tail, "d"] \\
I(T(X)) \arrow[d, "\varepsilon^L_X"'] & I(T(Y)) \arrow[d, "\varepsilon^L_Y"] \\
X \arrow[r, "g"'] & Y
\end{tikzcd}
\tag{4.4}$$

By Proposition 4.1.20, ε^L_X and ε^L_Y are pathwise embeddings, so the two vertical sides of the square (4.4)-right are path embeddings. By the hypothesis that g is an open map, there exists a diagonal filler arrow $d\colon I(Q) \rightarrowtail X$ obtaining the commuting triangles of (4.4)-right. By the universal property of T being right adjoint to I, there exists a unique morphism $d'\colon Q \to T(X)$ such that $\varepsilon^L_X \circ I(d') = d$.

We claim that $d'\colon Q \to T(X)$ is the diagonal filler of diagram (4.3). To show the top triangle commutes:

$$\begin{aligned}
\varepsilon^L_X \circ I(p) &= d \circ I(j) && \text{top triangle of (4.4)-right} \\
&= \varepsilon^L_X \circ I(d') \circ I(j) && \varepsilon^L_X \circ I(d') = d \\
&= \varepsilon^L_X \circ I(d' \circ j) && \text{functorality of } I \\
p &= d' \circ j && \text{universal property of counit } \varepsilon^L_X \text{ for } I \dashv T
\end{aligned}$$

To show the bottom triangle commutes:

$$\begin{aligned}
\varepsilon^L_Y \circ I(q) &= g \circ d && \text{bottom triangle of (4.4)-right} \\
&= g \circ \varepsilon^L_X \circ I(d') && \varepsilon^L_X \circ I(d') = d \\
&= \varepsilon^L_Y \circ I(T(g)) \circ I(d') && \varepsilon^L\text{-naturality square of } g \\
&= \varepsilon^L_Y \circ I(T(g) \circ d') && \text{functorality of } I \\
q &= T(g) \circ d' && \text{universal property of counit } \varepsilon^L_X \text{ for } I \dashv T
\end{aligned}$$

By definition, $T(g)$ is an open map. □

Theorem 4.1.23. *Given a resource-indexed arboreal cover $L_k \dashv R_k$ of $\mathcal{E}$ by $\{\mathcal{C}_k\}$ and two objects X, Y of $\mathcal{E}$, for all $k > 0$,*

(1) $X \to^{\mathcal{C}}_k Y$ implies $X \to^{\mathcal{C}^L}_k Y$.

(2) $X \rightharpoonup^{\mathcal{C}}_k Y$ implies $X \rightharpoonup^{\mathcal{C}^L}_k Y$.

(3) $X \leftrightarrow_k^{\mathcal{C}} Y$ *implies* $X \leftrightarrow_k^{\mathcal{C}^L} Y$*.*

(4) $X \cong_k^{\mathcal{C}} Y$ *implies* $X \cong_k^{\mathcal{C}^L} Y$*.*

Proof. For (1), we note that $X \to_k^{\mathcal{C}} Y$ is witnessed by a morphism $f\colon R_k(X) \to R_k(Y)$ in $\mathcal{C}$. From this morphism we obtain a morphism $T_k(f)\colon T_k(R_k(X)) \to T_k(R_k(Y))$ witnessing that $X \to_k^{\mathcal{C}^L} Y$.

For (2), the result follows from the fact that if f is a pathwise embedding in $\mathcal{C}_k$, then $T_k(f)$ is a pathwise embedding in $\mathcal{C}_k^L$ which holds by Proposition 4.1.21.

For (3), by the definition of $X \leftrightarrow_k^{\mathcal{C}} Y$, there exists a span of open pathwise embeddings

$$R_k(X) \xleftarrow{g} Z \xrightarrow{h} R_k(Y)$$

in $\mathcal{C}_k$. Applying T_k, we obtain the span

$$T_k(R_k(X)) \xleftarrow{T_k(g)} T_k(Z) \xrightarrow{T_k(h)} T_k(R_k(Y))$$

in $\mathcal{C}_k^L$. By Proposition 4.1.21 and Proposition 4.1.22, both $T_k(g)$ and $T_k(h)$ are open pathwise embeddings; so $X \leftrightarrow_k^{\mathcal{C}^L} Y$. For (4), the result follows from the fact that if f is an isomorphism in $\mathcal{C}_k$, then $T_k(f)$ is an isomorphism in $\mathcal{C}_k^L$ which holds as every functor preserves isomorphisms. □

4.1.2 A quasi-linear weak adjunction

In this section, we extend our analysis from the linear arboreal subcategory $\mathcal{C}^L$ to the quasi-linear arboreal category $\mathcal{C}^{qL}$. Recall that $\mathcal{C}^{qL}$ is the full subcategory of $\mathcal{C}$ consisting of only linear path-generated objects. Since $\mathcal{C}^{qL}$ is a full subcategory of $\mathcal{C}$, unlike $\mathcal{C}^L$, we do not require that the morphisms between objects in $\mathcal{C}^{qL}$ to be leaf morphisms.

One key difference between $\mathcal{C}^{qL}$ and $\mathcal{C}^L$ is that the inclusion functor $I^q\colon \mathcal{C}^{qL} \to \mathcal{C}$ is not a left adjoint. Intuitively, this is because without the requirement that morphisms preserve maximal elements, a partial branch of X represented by a path embedding $P \rightarrowtail X$ in $\mathcal{C}^{qL}$ can embed into multiple full branches of X. However, I^q can be paired with a weak right adjoint $T^q\colon \mathcal{C} \to \mathcal{C}^{qL}$. Weak right adjoints, originating from [54], are defined by dropping the uniqueness requirement on factorisations:

Definition 4.1.24. If $F\colon \mathcal{D} \to \mathcal{E}$ is a functor, then a functor $G\colon \mathcal{E} \to \mathcal{D}$ is a *weak right adjoint of* F if for all objects $Y \in \mathbf{Ob}(\mathcal{E})$, there exists a quasi-counit $\zeta_Y\colon F(G(Y)) \to Y$ such that for every morphism $F(X) \to Y \in \mathbf{Mor}(\mathcal{E})$, there exists a (not necessarily unique) morphism $g\colon X \to G(Y) \in \mathbf{Mor}(\mathcal{D})$ such that $f = \zeta_Y \circ F(g)$.

For every linearisable arboreal category $\mathcal{C}$, the inclusion $I^q\colon \mathcal{C}^{qL} \to \mathcal{C}$ has weak right adjoint $T^q\colon \mathcal{C} \to \mathcal{C}^{qL}$ defined in a similar manner as $T\colon \mathcal{C} \to \mathcal{C}^L$.

- For every object $X \in \mathbf{Ob}(\mathcal{C})$,

$$T^q(X) = \biguplus_{[p]\in \mathbf{P}(X)} \mathsf{dom}(p).$$

- For every morphism $f\colon X \to Y$, by the universal property of the coproduct $T^q(X)$, there exists a unique morphism $T^q(f)\colon T^q(X) \to T^q(Y)$ such that $T^q(f) \circ i_{[p]} = i_{[\exists_f p]} \circ q_p$ where

$$P \xrightarrow{q_p} \mathrel{\mkern-14mu}\rightarrow \exists_f P \xrightarrow{\exists_f p} Y \tag{4.5}$$

  is the factorisation of $f \circ p\colon P \to Y \in \mathbf{Mor}(\mathcal{C})$.

By the universal property of the coproduct, for every $X \in \mathbf{Ob}(\mathcal{C})$, we have a unique morphism $\zeta_X\colon I^q(T^q(X)) \to X$ such that for all $[p] \in \mathbf{P}(X)$, $\zeta \circ i_{[p]} = p$.

Proposition 4.1.25. *T^q is a weak right adjoint of I^q.*

Proof. The proof follows a similar argument to Theorem 4.1.15. Suppose $f\colon I^q(X) \to Y \in \mathbf{Mor}(\mathcal{C})$, we need to construct a morphism $g\colon X \to T^q(Y) \in \mathbf{Mor}(\mathcal{C}^{qL})$ satisfying $\zeta_Y \circ I^q(g) = f$. Since $X \in \mathbf{Ob}(\mathcal{C}^{qL})$, we know that $I^q(X)$ is linearly path-generated in $\mathcal{C}$ and thus, a coproduct of $Q = \mathsf{dom}(m)$ ranging over maximal path embeddings $[m\colon Q \rightarrowtail I(X)]$. Recall by equation (4.5) that for every such maximal element $[m]$, we have the factorisation of $f \circ m$ into:

$$Q \xrightarrow{q_m} \mathrel{\mkern-14mu}\rightarrow \exists_f Q \xrightarrow{n_m} Y \tag{4.6}$$

such that $[n_m] = \mathbf{P}f([m]) \in \mathbf{P}(Y)$. Since $[n_m] \in \mathbf{P}(Y)$, we have corresponding component of $T^q(Y)$ and coproduct injection $j_m\colon \exists_f Q \to T^q(Y)$ such that $\zeta_Y \circ j_m = n_m$. Thus, for every component Q of the coproduct $I(X)$, we have morphism $j_m \circ q_m\colon Q \to T^q(Y)$. By the universal property of the coproduct $I(X)$, we have a unique morphism $I^q(g)\colon I^q(X) \to I^q(T^q(Y))$ such that $I^q(g) \circ i_{[m]} = j_m \circ q_m$.

To verify that $\zeta_Y \circ I^q(g) = f$, we compute the following equality for every component

$\mathsf{dom}(m) = Q$ of the coproduct $I^q(X)$:

$$\begin{aligned}
(\zeta_Y \circ I^q(g)) \circ i_{[m]} &= \zeta_Y \circ (I^q(g) \circ i_{[m]}) && \\
&= \zeta_Y \circ (j_m \circ q_m) && I^q(g) \circ i_{[m]} = j_m \circ q_m \\
&= (\zeta_Y \circ j_m) \circ q_m && \\
&= n_m \circ q_m && \zeta_Y \circ j_m = n_m \\
&= f \circ m && \text{Equation (4.5)} \\
&= f \circ i_{[m]} && \text{Lemma 4.1.9}
\end{aligned}$$

Thus, by the universal property of coproduct $I^q(X)$, $\varepsilon^L_Y \circ I^q(g) = f$. □

We now discuss the how the contrast between $\mathcal{C}^{qL}$ and $\mathcal{C}^L$ means that T^q is not right adjoint to I^q. Every path P can be seen as coproduct over the singleton family $\{P\}$ where we consider P as the domain of the maximal path $[\mathsf{id}_P]$ in $\mathbf{P}(P)$, so every path is linearly path generated and $\mathbf{Ob}(\mathcal{C}^{qL}_p) = \mathbf{Ob}(\mathcal{C}_p)$. Consider a path embedding $p\colon P \rightarrowtail X$ into a object X where there are at least two $[b] \in \mathbf{P}(X)$ such that $[b] \leq [p]$. The embedding p represents a partial branch of X which can be extended to at least two full branches $b\colon Q \rightarrowtail X$. Indeed, for each $[b] \geq [p]$ in $\mathbf{P}(X)$ we have an embedding m_b such that $b = p \circ m_b$, so we can construct a $g_b\colon P \rightarrowtail T^q(X)$ which maps P into the $\mathsf{dom}(b)$ component of $T^q(X)$ defined as $g_b = i_{[b]} \circ m_b$. Thus, there exists multiple morphisms $g_b\colon P \rightarrowtail T^q(X)$ such that $\zeta_X \circ I^q(p) = g_b$. Note that of these g_b for every $[b] \geq [p]$, only the unique g_p is a leaf morphism which exists in $\mathcal{C}^L$.

Despite the pair (I^q, T^q) being merely a weak adjunction, we are able to use this construction fruitfully. Example 4.2.3 will illustrate how this pairing can be used to give a categorical definition of complete labelled trace equivalence.

4.2 Example linear arboreal covers

In this section, we examine the linear arboreal subcategories of the arboreal categories from Section 3.2. Recall that many of the examples from Section 3.2 are minor variants of the path category $\mathcal{R}(\sigma)$ from Example 3.2.1, thus we first prove that this category is linearisable.

Proposition 4.2.1. *$\mathcal{R}(\sigma)$ is a linearisable path category.*

Proof. Suppose $P, P_1, \ldots, P_n$ are paths in $\mathbf{Ob}(\mathcal{R}(\sigma))$, P is non-initial and there exists a morphism $f\colon P \to \biguplus_{i\in[n]} P_i$. Since morphisms in $\mathcal{R}(\sigma)$ are forest morphisms, f preserves roots. The coproduct in $\mathcal{R}(\sigma)$ does not merge roots, so the minimal element of P is mapped to the minimal element of P_j for some $j \in [n]$. As f preserves the covering relation and coproducts in $\mathcal{F}$ do not merge paths, the non-initial z-th element of P must be mapped to the z-element of P_j. Thus, f factors uniquely through the coproduct injection $i_j \colon P_j \to \biguplus_{i \in [n]} P_i$ and $\mathcal{R}(\sigma)$ is linearisable. $\square$

4.2.1 Pointed list

Mirroring the presentation of the examples in Section 3.2 of the previous Chapter 3, we first warm-up by deriving the pointed list comonad N_* from a linear arboreal cover of **Set** by $\mathcal{F}^L$. Thus, observing that the pointed list comonad N_* is the 'linear' variant of the prefix list comonad N^+ from Section 3.2.1.

Our first step is to the observe that the category of forests $\mathcal{F}$ from Example 3.0.1 is linearisable.

Proposition 4.2.2. *$\mathcal{F}$ is linearisable.*

Proof. This is a special case of Proposition 4.2.1 since $\mathcal{R}(\sigma) = \mathcal{F}$ when σ is the empty signature. $\square$

Thus, by Proposition 4.1.12, $\mathcal{F}$ has a linear arboreal subcategory $\mathcal{F}^L$. Concretely, the objects of $\mathcal{F}^L$ are linear forests as in Definition 2.1.4 and the morphisms, in addition to preserving roots and successors of elements as in $\mathcal{F}$, also preserve 'leaves', i.e. maximal elements. Intuitively, the objects X of $\mathcal{F}^L$ are simply disjoint unions of over a family of chains P_i and morphisms $f \colon X \to Y$ can only merge two components P_i, P_j of X into a component Q of Y if P_i, P_j and Q all have equal height.

We exhibit the pointed list comonad N_* as arising from a linear arboreal cover of **Set** by $\mathcal{F}^L$. The underlying functor $N_* \colon \mathbf{Set} \to \mathbf{Set}$ maps a set X to the set of non-empty sequences of elements in X paired with an index into the sequence, i.e.

$$N_*(X) = \{([x_1, \ldots, x_n], i) \mid n \in \mathbb{N}, x_1, \ldots, x_n \in X \text{ and } i \in [n]\}.$$

and maps a function $f \colon X \to Y$ to the function $N_*(f) \colon N_*(X) \to N_*(Y)$ which applies f component-wise, i.e. $N(f)([x_1, \ldots, x_n], i) = ([f(x_1), \ldots, f(x_n)], i)$.

In order to derive the comonad on N_* from a linear arboreal cover, we recall from Section 3.2.1 the arboreal cover $L \dashv R$ of **Set** by $\mathcal{F}$ which yielded the prefix list comonad N^+. In this construction, $L \colon \mathcal{F} \to \mathbf{Set}$ is the forgetful functor and $R \colon \mathbf{Set} \to \mathcal{F}$ is the cofree functor such that for $X \in \mathbf{Ob}(\mathbf{Set})$, $R(X)$ is the forest of non-empty sequences of elements of X and for $f \colon X \to Y$, $R(f)$ applies f component-wise, i.e.

$R(f)[x_1, \ldots, x_n] = [f(x_1), \ldots, f(x_n)]$. On the other hand, by Proposition 4.2.2 and Theorem 4.1.15, there is an adjunction $I \dashv T$ from $\mathcal{F}^L$ to $\mathcal{F}$. Concretely, $I \colon \mathcal{F}^L \to \mathcal{F}$ is simply the inclusion from the subcategory $\mathcal{F}^L$ to $\mathcal{F}$ and $T \colon \mathcal{F} \to \mathcal{F}^L$ maps each branch $\downarrow x = x_1 \prec \cdots \prec x_n = x$ of a forest $(X, \leq)$ it its own component $\downarrow x$ consisting of the same elements, i.e.

$$T(X) = \bigsqcup_{x \in X} \downarrow x.$$

As a set, the elements of $T(X)$ are pairs of the form $(\downarrow x, x_i)$ where $x_i \in \downarrow x$ is the i-th element of the chain $\downarrow x = x_1 \prec \cdots \prec x_n = x$. To make this 'underlining' of the i-th element x_i explicit, we will consider $T(X)$ as consisting of pairs (x, i).

We compose the arboreal cover $L \dashv R$ with the adjunction $I \dashv T$ to obtain the adjunction $L \circ I \dashv T \circ R$. In the following, we assume that η, ε are the unit, counit of the arboreal cover $L \dashv R$ and η^L, ε^L are the unit,counit of the linear adjunction $I \dashv T$.

- The left adjoint $L \circ I \colon \mathcal{F}^L \to \mathbf{Set}$ is the forgetful functor mapping a linear forest $(X, \leq)$ to X and a forest leaf morphism $f \colon (X, \leq) \to (Y, \leq)$ to its underlying set function $L(f) \colon X \to Y$.
- The right adjoint $T \circ R \colon \mathbf{Set} \to \mathcal{F}^L$ is the functor mapping X to the linear forest $(N_*(X), \leq)$ where $\leq$ defined by:

  $$(s, i) \leq (t, j) \text{ if } s = t \text{ and } i \leq j$$

  and a function $f \colon X \to Y$ to a forest leaf morphism with underlying function $N_*(f)$. The functorial action of $N_*(f)$ preserves the distinguished index in the list, i.e. $N_*(f)(s, i) = (t, i)$ where $t = R(f)(s)$ and n is the length of s and t. Assuming that n is the length of s and t, we can conclude that $N_*(f) \colon T(R(X)) \to T(R(Y))$ defines a leaf forest morphism which maps element (s, n) in the component $\downarrow s$ of $T(R(X))$ corresponding to sequence $s \in R(X)$ to the element (t, n) in the component $\downarrow t$ of $T(R(Y))$ corresponding to the sequence $t \in R(Y)$.
- The counit $\varepsilon^* \colon L \circ I \circ T \circ R \to \mathsf{Id}_{\mathbf{Set}}$ is the composition $\varepsilon \circ L\varepsilon^L R$. Explicitly, the component of ε^* for a set $X \in \mathbf{Ob}(\mathbf{Set})$ first maps (s, i) to the i-th prefix $s[1, i]$ of s via the $R(X)$ component of $L\varepsilon^L$ and then maps $s[1, i]$ to its last element via the X component of ε. Thus, for $s = [x_1, \ldots, x_n] \in R(X)$, $\varepsilon(s, i) = x_i$.
- The unit $\eta^* \colon \mathsf{Id}_{\mathcal{F}^L} \to T \circ R \circ L \circ I$ is the composition $T\eta I \circ \eta^L$. Explicitly, the component of η^* for a linear forest $(X, \leq) \in \mathbf{Ob}(\mathcal{F}^L)$ first maps the i-th element x_i in the component $x_1 \prec \cdots \prec x_n = x$ corresponding to maximal path $\downarrow x$ of X to (x, i) via η^L and then maps (x, i) to the pair $([x_1, \ldots, x_n], i)$ via the $I(X, \leq)$

component of $T\eta$. Thus, for $x_i \in (X, \leq)$ in the maximal path $x_1 \prec \cdots \prec x_n = x$ of X, $\eta^*(x_i) = ([x_1, \dots, x_n], i)$.

The composition $T \circ R$: **Set** $\to \mathcal{F}^L$ maps X to $(N_*(X), \leq)$ and f to the forest leaf morphism with underlying function $N_*(f)$. Since $L \circ I \to \mathcal{F}^L \to$ **Set** is a forgetful functor, we see that $N_* = L \circ I \circ T \circ R$. Therefore, we have obtained our derivation, from the adjunction $L \circ I \dashv T \circ R$ of the pointed list comonad $(N_*, \varepsilon_*, \delta)$ where the comultiplication has components $\delta_X = (L \circ I)(\eta_{(T \circ R)(X)})$. Explicitly, for every pair $(s, i) \in N_*(X)$,

$$\delta_X(s, i) = ([(s, 1), \dots, (s, n)], i).$$

As a corollary of Theorem 4.1.19 and Proposition 3.2.6, we show that $L \circ I \dashv T \circ R$ is and comonadic adjunction, and therefore, a linear arboreal cover of **Set** by $\mathcal{F}^L$. Thus, $\mathcal{F}^L \cong \mathbf{EM}(N_*)$.

Proposition 4.2.3. *$L \circ I \dashv T \circ R$ is a linear arboreal cover of* **Set** *by $\mathcal{F}^L$.*

4.2.2 Linear pebbling

In this section, we derive a linear variant of the pebbling comonad from Section 3.2.3 from a resource-indexed linear arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}_k^{PL}(\sigma)\}$. This linear variant is the so-called pebble-relation comonad from [68]. Generalising the relationship between the pebbling comonad and the pebble-relation comonad was the motivation for axiomatising linear arboreal categories. Studying the pebble-relation comonad itself was motivated by Victor Dalmau's paper [29]. In [29], Dalmau links bounded pathwidth classes to a restricted conjunction fragment of positive-existential finite variable logic 'linearising' the well-known connection between bounded treewidth classes and positive-existential finite variable logic from [58].

Our first step is to the observe that the category of k-pebble forest covers $\mathcal{R}_k^P(\sigma)$ is linearisable.

Proposition 4.2.4. *For every $k > 0$, $\mathcal{R}_k^P(\sigma)$ is linearisable.*

Proof. Coproducts in $\mathcal{R}_k^P(\sigma)$ are computed as coproducts of the underlying forest-ordered σ-structures and the pebbling functions extended in the obvious way. The category of $\mathcal{R}(\sigma)$ is linearisable by Proposition 4.2.1, so we obtain that $\mathcal{R}_k^P(\sigma)$ is linearisable. □

Thus, by Proposition 4.1.12, $\mathcal{R}_k^P(\sigma)$ has a linear arboreal subcategory $\mathcal{R}_k^{PL}(\sigma)$. Concretely, the objects of $\mathcal{R}_k^{PL}(\sigma)$ are are the k-pebble linear forest covers $(\mathcal{A}, \leq, \mathbf{p})$ where $(A, \leq)$ is a linear forest. The morphisms $f\colon (\mathcal{A}, \leq, \mathbf{p}) \to (\mathcal{B}, \leq', \mathbf{p}')$ are leaf forest-ordered σ-morphisms which preserve the pebbling function.

We exhibit the family of linear pebble comonads $\mathbb{P}_k^L$ as arising from the linear arboreal cover of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}_k^{PL}(\sigma)$. For every $k > 0$, the universe $\mathbb{P}_k^L(A)$ of σ-structure $\mathbb{P}_k^L(\mathcal{A})$ is the set of underlined k-pebble sequences of A. Explicitly,

$$\mathbb{P}_k^L(A) := \{([(p_1, a_1), \ldots, (p_n, a_n)], i) \mid \forall n \in \mathbb{N}, i \in [n] \text{ and } \forall j \in [n], (p_j, a_j) \in [k] \times A\}$$

where the index i in the th pair $(s, i) \in \mathbb{P}_k^L(A)$ indicates that (p_i, a_i) is underlined in the sequence (s, i). Similar to N_*, the counit $\varepsilon\colon \mathbb{P}_k^L \to \mathsf{Id}_{\mathsf{Struct}(\sigma)}$ maps a underlined pebble sequence to the underlined element, i.e. for $s = [(p_1, a_1), \ldots, (p_n, a_n)]$, $(s, i) \mapsto a_i$. We also need a 'underlined' pebble function $\rho_{\mathcal{A}}$ where $(s, i) \mapsto p_i$. For every m-relation $R \in \sigma$, we define the interpretation $((s_1, i_1), \ldots, (s_m, i_m)) \in R^{\mathbb{P}_k^L(\mathcal{A})}$ if and only if:

(1) $\forall j \in [m], s_j = s$ *(same component)*

(2) $\rho_{\mathcal{A}}(s, i_j)$ does not appear in $s(i_j, i]$ *(active pebble)*

(3) $(\varepsilon(s, i_1), \ldots, \varepsilon(s, i_m)) \in R^{\mathcal{A}}$ *(compatibility)*

where $i = \max\{i_1, \ldots, i_m\}$. The functor $\mathbb{P}_k^L$ maps the σ-morphism $f\colon \mathcal{A} \to \mathcal{B}$ to the σ-morphism $\mathbb{P}_k^L(f)\colon \mathbb{P}_k^L(\mathcal{A}) \to \mathbb{P}_k^L(\mathcal{B})$ which applies f component-wise to the pebbled elements, i.e. $\mathbb{P}_k^L(f)([(p_1, a_1), \ldots, (p_n, a_n], i) = ([(p_1, f(a_1)), \ldots, (p_n, f(a_n))], i)$.

In order to derive the comonad on $\mathbb{P}_k^L\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma)$ from a linear arboreal cover, we recall from Section 3.2.3 the resource-indexed arboreal cover $L_k \dashv R_k$ of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}_k^P(\sigma)$ which yielded the family of comonads $\mathbb{P}_k$. In this construction, $L_k\colon \mathcal{R}_k^P(\sigma) \to \mathsf{Struct}(\sigma)$ is the forgetful functor, throwing away the order and pebbling function, and $R_k\colon \mathsf{Struct}(\sigma) \to \mathcal{R}_k^P(\sigma)$ is the cofree functor such that for $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma))$, $R(\mathcal{A})$ is the k-pebble forest cover $(\mathbb{P}_k(\mathcal{A}), \sqsubseteq, \rho_{\mathcal{A}})$ on pebble sequences and for $f\colon \mathcal{A} \to \mathcal{B} \in \mathbf{Mor}(\mathsf{Struct}(\sigma))$, $R_k(f)$ is the forest-ordered σ-morphism whose underlying set function applies f to each pebbled element. By Proposition 4.2.4 and Theorem 4.1.15, there is an adjunction for every $k > 0$, $I_k \dashv T_k$ from $\mathcal{R}_k^{PL}(\sigma)$ to $\mathcal{R}_k^P(\sigma)$.

We compose the arboreal cover $L_k \dashv R_k$ with the adjunction $I_k \dashv T_k$ to obtain the adjunction $L_k \circ I_k \dashv T_k \circ R_k$. In the following, we assume that η, ε are the unit, counit of the arboreal cover $L_k \dashv R_k$ and η^L, ε^L are the unit,counit of the adjunction $I_k \dashv T_k$.

- The left adjoint $L_k \circ I_k\colon \mathcal{R}_k^{PL}(\sigma) \to \mathsf{Struct}(\sigma)$ is the forgetful functor mapping a k-pebble linear forest cover $(\mathcal{A}, \leq, \mathbf{p})$ to $\mathcal{A}$ and a leaf, pebble-preserving, forest-ordered σ morphism $f\colon (\mathcal{A}, \leq, \mathbf{p}) \to (\mathcal{B}, \leq', \mathbf{p}')$ to its underlying σ-morphism function $L_k(f)\colon \mathcal{A} \to \mathcal{B}$.

- The right adjoint $T_k \circ R_k \colon \mathsf{Struct}(\sigma) \to \mathcal{R}_k^{PL}(\sigma)$ is the functor mapping $\mathcal{A}$ to the k-pebble linear forest cover $(\mathbb{P}_k^L(\mathcal{A}), \leq, \rho_{\mathcal{A}})$ where $\leq$ defined, as in the N_* case, by:

  $$(s,i) \leq (t,j) \text{ if } s = t \text{ and } i \leq j$$

  and $\rho_{\mathcal{A}}$ is the function which returns the pebble p_i on the underlined element i in the pair (s,i). For a σ-morphism $f \colon \mathcal{A} \to \mathcal{B}$, $(T_k \circ R_k)(f)$ has underlying σ-morphism $\mathbb{P}_k^L(f) \colon \mathbb{P}_k^L(\mathcal{A}) \to \mathbb{P}_k^L(\mathcal{B})$. The functorial action of $\mathbb{P}_k^L(f)$ preserves the distinguished index in the list, i.e. $\mathbb{P}_k^L(f)(s,i) = (t,i)$ where $t = R_k(f)(s)$, and the pebbles at each position in the list, i.e. $\rho_{\mathcal{A}}(s,i) = \rho_{\mathcal{B}}(t,i)$ for all $i \in [n]$. Thus, $\mathbb{P}_k^L(f)$ lifts to a morphism $T_k(R_k(f)) \in \mathbf{Mor}(\mathcal{R}_k^{PL}(\sigma))$.

- The counit $\varepsilon^* \colon L_k \circ I_k \circ T_k \circ R_k \to \mathsf{Id}_{\mathsf{Struct}(\sigma)}$ is the composition $\varepsilon \circ L\varepsilon^L R$. Explicitly, as in the N_* case, the component of ε^* for a σ-structure $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma))$ first maps (s,i) to the i-th prefix $s[1,i]$ of s via the $R_k(X)$ component of $L_k\varepsilon^L$ and then maps $s[1,i]$ to its last element via the X component of ε. Thus, for $s = [(p_1,a_1),\dots,(p_n,a_n)] \in R_k(\mathcal{A})$, $\varepsilon(s,i) = a_i$.

- The unit $\eta^* \colon \mathsf{Id}_{\mathcal{R}_k^{PL}(\sigma)} \to T_k \circ R_k \circ L_k \circ I_k$ is the composition $T\eta I \circ \eta^L$. Explicitly, the component of η^* for a k-pebble linear forest cover $(\mathcal{A}, \leq, \mathbf{p}) \in \mathbf{Ob}(\mathcal{R}_k^{PL}(\sigma))$ first maps the i-th element a_i in the component $a_1 \prec \cdots \prec a_n = a$ corresponding to maximal path $\downarrow a$ of $(\mathcal{A}, \leq, \mathbf{p})$ to (a,i) via η^L and then maps (a,i) to the pair $([(\mathbf{p}(a_1),a_1),\dots,(\mathbf{p}(a_n),a_n)],i)$ via the $I(\mathcal{A},\leq,\mathbf{p})$ component of $T\eta$. Thus, for $a_i \in (\mathcal{A},\leq,\mathbf{p})$ in the maximal path $a_1 \prec \cdots \prec a_n = a$ of A, $\eta^*(a_i) = ([(\mathbf{p}(a_1),a_1,\dots,(\mathbf{p}(a_n),a_n)],i)$.

The composition $T_k \circ R_k \colon \mathsf{Struct}(\sigma) \to \mathcal{R}_k^{PL}(\sigma)$ maps $\mathcal{A}$ to $(\mathbb{P}_k^L(\mathcal{A}), \leq, \mathbf{p})$ and f to the leaf, pebble preserving, forest-ordered σ-morphism with underlying σ-morphism $\mathbb{P}_k^L(f)$. Since $L_k \circ I_k \colon \mathcal{R}_k^{PL}(\sigma) \to \mathsf{Struct}(\sigma)$ is a forgetful functor, we see that $\mathbb{P}_k^L = L_k \circ I_k \circ T_k \circ R_k$. Therefore, we have obtained our derivation, from the adjunction $L_k \circ I_k \dashv T_k \circ R_k$ of the linear k-pebble comonad $(\mathbb{P}_k^L, \varepsilon^*, \delta)$ where the comultiplication has components $\delta_X = (L_k \circ I_k)(\eta^*_{(T_k \circ R_k)(X)})$. Explicitly,

$$\delta_X(s,i) = ([(p_1,(s,1)),\dots,(p_n,(s,n))],i).$$

for $s = [(p_1,a_1),\dots,(p_n,a_n)]$. As a corollary of Theorem 4.1.19 and Proposition 3.2.26, we conclude that $L_k \circ I_k \dashv T_k \circ R_k$ is a comonadic adjunction, and therefore, a linear arboreal cover of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}_k^{PL}(\sigma)$. Thus, $\mathcal{R}_k^{PL}(\sigma) \cong \mathbf{EM}(\mathbb{P}_k^L)$.

Proposition 4.2.5. *$L_k \circ I_k \dashv T_k \circ R_k$ is a linear arboreal cover of* $\mathsf{Struct}(\sigma)$ *by* $\{\mathcal{R}_k^{PL}(\sigma)\}$.

A k-pebble linear forest cover $(\mathcal{A}, \leq, \mathbf{p})$, equivalently a $\mathbb{P}_k^L$-coalgebra of type $\mathcal{A} \to \mathbb{P}_k^L$, encodes the data of a path decomposition of width $< k$. Path decompositions of graph G are simply tree decompositions $(T, \leq, l)$ where T is assumed to be a path. Pathwidth of G is defined similarly treewidth, as the minimum width of a path decomposition of G. Just as a treewidth is a sparsity measure which captures how similar a graph G is to a tree, pathwidth is a sparsity measure which captures how similar G is to a path.

We can define path decomposition of a graph G as simply tree decomposition $(T, \leq, \lambda)$ as in Definition 3.2.23 where T is a path. However, it will be useful for our proofs to spell this out explicitly.

Definition 4.2.6. Given a finite graph G, a *path decomposition of* $\mathcal{A}$ is a triple $(X, \leq_X, l)$, where $(X, \leq_X)$ is a linearly ordered set and $l \colon X \to \mathcal{P}(A)$ is a function satisfying the following conditions:

(PD1) For every $a \in A$, there exists $x \in X$ such that $a \in l(x)$.

(PD2) If $a \frown a' \in \mathcal{A}$, then $a, a' \in l(x)$ for some $x \in X$.

(PD3) For all $y \in [x, x']$, $l(x) \cap l(x') \subseteq l(y)$, where $[x, x']$ is an interval with respect to $\leq_X$.

The width $\mathsf{width}(T)$ of a path decomposition is given by $\max_{x \in T} |l(x)| - 1$. As with the definition of treewidth, the pathwidth $\mathsf{pw}(\mathcal{A})$ of a σ-structure $\mathcal{A}$ is $\min_T \mathsf{width}(T)$ where T ranges over path decompositions of $\mathcal{G}(\mathcal{A})$.

We obtain the following proposition demonstrating that linear k-pebble forest covers encode the data of a path decomposition of width $< k$–essentially establishing a linear variant of Proposition 3.2.24.

Proposition 4.2.7. *The following are equivalent for all finite σ-structures $\mathcal{A}$:*

(1) $\mathcal{A}$ has a path decomposition of width $< k$.

(2) There exists a a k-pebble linear forest cover $(\mathcal{A}, \leq, \mathbf{p}) \in \mathbf{Ob}(\mathcal{R}_k^{PL}(\sigma))$

In order to construct the pebbling function inductively from a path decomposition, we will use the following intermediate data structure.

Definition 4.2.8 (k-pebbling section family)**.** Given a path decomposition $(X, \leq_X, \lambda)$ of width $< k$ for $\mathcal{A}$, we define a *k-pebbling section family for* $(X, \leq_X, \lambda)$ as a family of functions $\{\tau_x : \lambda(x) \to [k]\}$ indexed by $x \in X$, such that the following conditions hold:

(1) (Locally-injective) For every $x \in X$, τ_x is an injective function.

(2) (Glueability) For every $x, x' \in X$,

$$\tau_x|_{\lambda(x)\cap\lambda(x')} = \tau_{x'}|_{\lambda(x)\cap\lambda(x')}.$$

From a path decomposition $< k$, we can construct a k-pebbling section family.

Lemma 4.2.9. *If $(X, \leq_X, \lambda)$ is a path decomposition of width $< k$, then $(X, \leq_X, \lambda)$ has a k-pebbling section family $\{\tau_x\}_{x\in X}$.*

Proof. By induction on the linear order $\leq_X$ for every $x \in X$, we construct a k-pebbling section family $\{\tau_z\}_{z\in\downarrow x}$ for $(Y, \leq_Y, \lambda|_Y)$, where $Y = \downarrow x$ and $\leq_Y = \leq_X \cap (Y \times Y)$.

Base Case: Suppose r is the $\leq_X$-least element. By the path decomposition X having width $< k$, the cardinality of $\lambda(r)$ is $\leq k$, thus we can enumerate the elements via an injective function $\tau_r : \lambda(r) \to [k]$. By construction, τ_r is injective, so $\{\tau_r\}$ is locally-injective. Glueability follows trivially as every $x \in \downarrow r$ is equal to r.

Inductive Step: Let x' be the immediate $\leq_X$-successor of x. By the induction hypothesis, there exists a k-pebbling section family $\{\tau_y\}_{y\in\downarrow x}$. Let $V_{x'}$ denote the subset of 'new' elements $a \in \lambda(x')$ such that $a \notin \lambda(y)$ for every $y <_X x'$.

Claim: For every $y <_X x'$, $\lambda(y) \cap \lambda(x') \subseteq \lambda(x) \cap \lambda(x')$. We show that this claim holds: By (PD3), for all $z \in [y, x']$, $\lambda(y) \cap \lambda(x') \subseteq \lambda(z)$. In particular, since $y \leq_X x <_X x'$, $\lambda(y) \cap \lambda(x') \subseteq \lambda(x)$. Therefore, $\lambda(y) \cap \lambda(x') \subseteq \lambda(x) \cap \lambda(x')$.

From the claim and the definition of $V_{x'}$, we have that $\lambda(x') = (\lambda(x) \cap \lambda(x')) \sqcup V_{x'}$, where $\sqcup$ denotes a disjoint union. This allows us to define $\tau_{x'} : \lambda(x') \to [k]$ by cases on each of these parts. Fix an injective function $\upsilon_{x'} : V_{x'} \to [k]$ enumerating $V_{x'}$. Let $\{i_1, \dots, i_m\}$ enumerate the elements of $[k]$ not in the image of $\tau_x|_{\lambda(x)\cap\lambda(x')}$. Define $\tau_{x'} : \lambda(x') \to [k]$ as

$$\tau_{x'}(a) = \begin{cases} \tau_x(a) & \text{if } a \in \lambda(x) \cap \lambda(x'); \\ i_j & \text{if } a \in V_{x'} \text{ and } \upsilon_{x'}(a) = j. \end{cases}$$

Injectivity of $\tau_{x'}$ follows from the injectivity of $\tau_x|_{\lambda(x)\cap\lambda(x')}$ and $\upsilon_{x'}$. To verify glueability, it suffices to check that $\tau_{x'}|_{\lambda(y)\cap\lambda(x')} = \tau_y|_{\lambda(y)\cap\lambda(x')}$ for $y \in \downarrow x'$. Since $\{\tau_y\}_{y\in\downarrow x}$ is a k-pebbling section family, for all $y \in \downarrow x$, $\tau_y|_{\lambda(y)\cap\lambda(x)} = \tau_x|_{\lambda(y)\cap\lambda(x)}$. By construction, $\tau_x|_{\lambda(x)\cap\lambda(x')} = \tau_{x'}|_{\lambda(x)\cap\lambda(x')}$. By the claim, we have that $\lambda(y) \cap \lambda(x') \subseteq \lambda(x) \cap \lambda(x')$. Therefore, $\tau_y|_{\lambda(y)\cap\lambda(x')} = \tau_{x'}|_{\lambda(y)\cap\lambda(x')}$. □

Proof of Proposition 4.2.7. (1) ⇒ (2) Suppose $(X, \leq_X, \lambda)$ is a path decomposition of $\mathcal{A}$ of width $< k$. We define a family of linearly ordered sets $\{(S_i, \leq_i)\}$, where each S_i is the vertex set of a connected component of $\mathcal{G}(\mathcal{A})$. To define the order $\leq_i$, we define an order on $\leq_A$ and realise $\leq_i$ as the restriction of $\leq_A$ to S_i. For every $a \in A$, let $x_a \in X$ denote the $\leq_X$-least element in X such that $a \in \lambda(x_a)$. Such an x_a always exists by (PD1). By lemma 4.2.9 there exists a k-pebbling section family $\{\tau_x\}_{x \in X}$. We then define $\leq_A$ as follows:

$$a \leq_A a' \Leftrightarrow x_a <_X x_{a'} \text{ or } \tau_x(a) \leq \tau_x(a') \text{ if } x_a = x_{a'} = x.$$

The glueability condition on k-pebble section family $\{\tau_x\}$ allows us to obtain a well-defined pebbling function $p : A \to [k]$ from τ_x. Explicitly, thinking of functions as their sets of ordered pairs, $p = \bigcup_{x \in X} \tau_x$. The tuple $(\{(S_i, \leq_i)\}, p)$ is a k-pebble linear forest cover.

To verify that $\{(S_i, \leq_i)\}$ is a partition of A into linearly ordered subsets, we observe that by construction each S_i is a connected component of A, and so $\{S_i\}$ partitions A. Suppose $a, a' \in S_i$, then by $\leq_X$ being a linear order, either $x_a <_X x_{a'}$, $x_a >_X x_{a'}$, or $x_a = x_{a'}$. If $x_a <_X x_{a'}$ or $x_a >_X x_{a'}$, then $a <_i a'$ or $a >_i a'$ by the definition of $\leq_i$. If $x_a = x_{a'} = x$, then either $\tau_x(a) \leq \tau_x(a')$ or $\tau_x(a) \geq \tau_x(a')$ by the linear ordering $\leq$ on $[k]$. Hence, in either case, $a \leq_i a'$ or $a \geq_i a'$, so $\leq_i$ is a linear ordering.

To verify Condition (E), suppose $a \frown a' \in \mathcal{A}$. This means a, a' are connected in $\mathcal{G}(\mathcal{A})$, and so are in the same connected component S_i of $\mathcal{G}(\mathcal{A})$.

To verify Condition (P), suppose $a \frown a' \in S_i$, $a \leq_i a'$, and $b \in (a, a']_i$. By definition of $\leq_i$, $x_a \leq_X x_b \leq_X x_{a'}$. Since $a \frown a'$, by (PD2) there exists $x \in X$ such that $a, a' \in \lambda(x)$. By the definition of $x_{a'}$ as the $\leq_X$-least element of X containing a', we have $x_{a'} \leq_X x$. By transitivity of $\leq_X$ and $x_a \leq_X x_{a'} \leq_X x$, we have that $x_a \leq_X x$. By (PD3), for every $y \in [x_a, x]_X$, $\lambda(x_a) \cap \lambda(x) \subseteq \lambda(y)$. In particular, for $x_b \in [x_a, x_{a'}]_X \subseteq [x_a, x]_X$, $a \in \lambda(x_b)$. Hence, $a, b \in \lambda(x_b)$ and by the injectivity of τ_{x_b}, $\tau_{x_b}(a) \neq \tau_{x_b}(b)$. It follows that $p(a) \neq p(b)$.

(2) $\Rightarrow$ (1) Suppose $\mathcal{A}$ has k-pebble linear forest cover given by the partition $\{(S_i, \leq_i)\}_{i \in [n]}$ and pebbling function $p : A \to [k]$. We define a linearly ordered set $(A, \leq_A)$, where $\leq_A$ is the ordered sum of the family $\{(S_i, \leq_i)\}_{i \in [n]}$. Explicitly, $a \leq_A a'$ iff $a \in S_i$, $a' \in S_j$ for $i < j$ or $a \leq_i a'$ for $i = j$. We say that an element a is an *active predecessor* of a' if $a \leq_A a'$ and for all $b \in (a, a']_A$, $p(b) \neq p(a)$. Let $\lambda(a)$ be the set of active predecessors of a. The triple $(A, \leq_A, \lambda)$ is a path decomposition of $\mathcal{A}$ of width $< k$.

To verify (PD1), observe that, for every $a \in A$, a is an active predecessor of itself since $a \leq_A a$ and $(a, a]_A = \varnothing$. Hence, $a \in \lambda(a)$.

To verify (PD2), suppose $a \frown a' \in \mathcal{A}$. By condition (E), there exists an S_i where $a, a' \in S_i$. Without loss of generality, assume $a \leq_i a'$. By condition (P), for all $b \in (a, a']_i$, $p(b) \neq p(a)$. Therefore, a is an active predecessor of a', so $a, a' \in \lambda(a')$.

To verify (PD3), suppose $b \in [a, a']_A$ and that $c \in \lambda(a) \cap \lambda(a')$. By $c \in \lambda(a)$ and $b \in [a, a']_A$, we have that $c \leq_A a$ and $a \leq_A b$, so $c \leq_A b$. By $c \in \lambda(a')$, for all $d \in (c, a']_A$, $p(c) \neq p(d)$. In particular, for all $d \in (c, b]_A$, $p(c) \neq p(d)$. By definition, c is an active predecessor of b, so $c \in \lambda(b)$.

To verify the width of the decomposition $< k$, we need to show that for every $a' \in A$, $|\lambda(a')| \leq k$. Assume for contradiction that $|\lambda(a')| > k$ for some $a' \in A$. Consider the pebbling function restricted to $\lambda(a')$, $p|_{\lambda(a')} : \lambda(a') \to [k]$. By the Pigeonhole Principle, there must exist $a, c \in \lambda(a')$ with $a \neq c$, such that $p(a) = p(c)$. Without loss of generality assume that $a <_A c$. Since $a \in \lambda(a')$, a is an active predecessor of a', i.e. for all $b \in (a, a']_A$, $p(b) \neq p(a)$. In particular, since $c \in (a, a']_A$, as $a <_A c$ and $c \in \lambda(a')$, then $p(c) \neq p(a)$. This yields a contradiction. □

In analogy to $\mathbb{P}_k$, where k-pebble forest covers correspond to $\mathbb{P}_k$-coalgebras, we obtain as a corollary of Proposition 4.2.5 and Proposition 4.2.7, that $\mathbb{P}^L_k$-coalgebras. Thus, we have reproved the coalgebra characterisation theorem of [68] and obtain that the $\mathbb{P}^L$-coalgebra number corresponds to pathwidth

Corollary 4.2.10. *For all finite σ-structures $\mathcal{A}$, $\chi^{\mathbb{P}^L}(\mathcal{A}) = \mathsf{pw}(\mathcal{A}) + 1$.*

With the resource-indexed linear arboreal cover of $\mathsf{Struct}(\sigma)$ by $\mathcal{R}^{PL}_k(\sigma)$, we can demonstrate how the behavioural relations in Definition 3.1.19, instantiated in this case, capture equivalence in restricted conjunction fragments $\exists^+ \curlywedge \mathcal{L}^k$ of $\exists^+\mathcal{L}^k$, and $\exists \curlywedge \mathcal{L}^k$ of $\exists\mathcal{L}^k$. This will recover axiomatically recover some of the results of the pebble-relation comonad We begin by defining these restricted conjunction fragments.

Definition 4.2.11. A *restricted conjunction* is a conjunction $\bigwedge \Psi$ where Ψ is a set of formulas satisfying the following conditions:

(R) At most one formula $\psi \in \Psi$ with quantifiers is not a sentence.

The logics $\exists^+ \curlywedge \mathcal{L}^k$ and $\exists \curlywedge \mathcal{L}^k$ are the fragments of $\exists^+\mathcal{L}^k$ and $\exists\mathcal{L}^k$ in where each conjunction is a restricted conjunction. Preservation and equivalence in $\exists^+ \curlywedge \mathcal{L}^k, \exists \curlywedge \mathcal{L}^k$ are characterised by all-in-one ${}^{\exists^+} \curlywedge \mathbf{Peb}_k(\mathcal{A}, \mathcal{B}), {}^{\exists} \curlywedge \mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ variants of the one-sided k-pebble games ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B}), {}^{\exists}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ from Section 3.2.3, respectively.

Definition 4.2.12. Given σ-structures $\mathcal{A}, \mathcal{B}$ and $k > 0$, we define the all-in-one k-pebble game ${}^{\exists}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ from $\mathcal{A}$ to $\mathcal{B}$. In this game, Spoiler and Duplicator each have a set $[k] = \{1, \ldots, k\}$ of pebbles. In the first and only round of the game,

- Spoiler announces a full play in the game ${}^{\exists}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$:

$$s = [(p_1, a_1), \ldots, (p_n, a_n)]$$

- Duplicator responds with a compatible (same length and pebble at each index) full play in the game ${}^{\exists}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$:

$$t = [(p_1, b_1), \ldots, (p_n, b_n)]$$

Duplicator wins the game if for all $i \in [n]$, the relation

$$\gamma_i = \{(a^p, b^p) \mid a^p = \mathsf{last}_p(s[1, i]), b^p = \mathsf{last}_p(t[1, i])\}$$

is a tight σ-correspondence.

The all-in-one positive k-pebble game ${}^{\exists^+}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is ${}^{\exists}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ where Duplicator's winning condition is weakened so that for every $i \in [n]$, γ_i is only required to be σ-correspondence.

In order to connect the games ${}^{\exists^+}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ and ${}^{\exists}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ with their corresponding logics, we will need to construct the 'maximal' sentences associated with Spoiler plays s in these games.

We say a collection $v = \{(p_1, a_1), \ldots, (p_n, a_n)\}$ of pebble placements on $\mathcal{A}$ is an *active window* if the set of pebbles $\{p_1, \ldots, p_n\}$ is pairwise distinct. For each active window v on $\mathcal{A}$, we define the associated atomic $\mathbf{diag}^+_{\mathcal{A}}(v)$ and literal diagrams $\mathbf{diag}_{\mathcal{A}}(v)$ as:

$$\begin{aligned}
\mathbf{diag}^+_{\mathcal{A}}(v) &= \{R(\mathbf{x}(p_{i_1}), \ldots, \mathbf{x}(p_{i_r})) \mid (p_{i_1}, a_{i_1}), \ldots, (p_{i_r}, a_{i_r}) \in v, R \in \sigma, \text{ and } (a_{i_1}, \ldots, a_{i_r}) \in R^{\mathcal{A}}\} \\
\mathbf{diag}_{\mathcal{A}}(v) &= \mathbf{diag}^+_{\mathcal{A}}(v) \\
&\quad \cup \{\neg R(\mathbf{x}(p_{i_1}), \ldots, \mathbf{x}(p_{i_r})) \mid (p_{i_1}, a_{i_1}), \ldots, (p_{i_r}, a_{i_r}) \in v, R \in \sigma, \text{ and } (a_{i_1}, \ldots, a_{i_r}) \notin R^{\mathcal{A}}\}
\end{aligned}$$

where for every pebble $p \in [k]$, $\mathbf{x}(p) = x_p$ is the associated variable in $\mathcal{L}^k$. Next, given a sequence of pebble placements s, we define

$$\mathsf{Active}(s) = \{(p_i, a_i) \mid a_i = \mathsf{last}_{p_i}(s)\}.$$

Using these definitions, we are able to associate to every (possibly empty) pebble sequence s on $\mathcal{A}$ a sentence $\mathbf{tp}_s \in \exists\curlywedge\mathcal{L}^k$ encoding all the literals that are satisfied in every window of pebbled elemetns on $\mathcal{A}$ along Spoiler's play s. We first construct the formulas $\mathbf{tp}'_s$ by induction on the length of s. For the base case, $s = \epsilon$ is the empty sequence, and we set $\mathbf{tp}'_\epsilon = \top$ where $\top$ is true in every structure. For the inductive step, suppose $t = s[(p, a)]$, then

$$\mathbf{tp}'_t(\vec{x'}) = \exists x_p(\bigwedge \mathbf{diag}(\mathsf{Active}(t)) \wedge \mathbf{tp}'_s(\vec{x})).$$

where we note that $[\vec{x'}]$ is the set of variables $\mathbf{x}(q)$ such that $q \neq p$ and q is active pebble in t. Finally, for every s, let $\mathbf{tp}_s = \exists\vec{x}\mathbf{tp}'_s(\vec{x})$. We can similarly define the positive variant $\mathbf{tp}^+_s \in \exists^+\Yup\mathcal{L}^k$ by replacing $\mathbf{diag}_\mathcal{A}(\cdot)$ in the above definition with $\mathbf{diag}^+_\mathcal{A}(\cdot)$.

We now prove the morphism power theorem relating equivalence in $\exists^+\Yup\mathcal{L}^k$, Duplicator winning strategies in ${}^{\exists^+}\Yup\mathbf{Peb}(\mathcal{A}, \mathcal{B})$, and the relation $\rightarrow_k^{\mathcal{R}^{PL}(\sigma)}$

Theorem 4.2.13. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \Rrightarrow^{\exists^+\Yup\mathcal{L}^k} \mathcal{B}$

(2) Duplicator has a winning strategy in ${}^{\exists^+}\Yup\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$*.*

(3) Duplicator has a winning strategy in ${}^{\exists^+}\mathcal{G}(T_k(R_k(\mathcal{A})), T_k(R_k(\mathcal{B})))$ *over* $\mathcal{R}^{PL}_k(\sigma)$*.*

(4) $\mathcal{A} \rightarrow_k^{\mathcal{R}^{PL}(\sigma)} \mathcal{B}$

Proof. For the (1) ⇒ (2) implication, suppose $\mathcal{A} \Rrightarrow^{\exists^+\Yup\mathcal{L}^k} \mathcal{B}$ and consider Spoiler playing pebble sequence s of length n in the one and only around of ${}^{\exists^+}\Yup\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$. By construction, the pebble sequence s on $\mathcal{A}$ witnesses all the existential quantifiers in $\mathbf{tp}^+_s$ and $\mathcal{A} \vDash \mathbf{tp}^+_s$. By $\mathcal{A} \Rrightarrow^{\exists^+\Yup\mathcal{L}^k} \mathcal{B}$, $\mathcal{B} \vDash \mathbf{tp}^+_s$. Thus, there exists a corresponding pebble sequence t on $\mathcal{B}$ which we take to be Duplicator's response in ${}^{\exists^+}\Yup\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$. Since $\mathbf{tp}^+_s$ contains for every $i \in [n]$, the positive diagram $\mathbf{diag}^+(\mathsf{Active}(s[1, i]))$, we can conclude that for every $i \in [n]$, $\mathbf{diag}^+(\mathsf{Active}(s[1, i])) \subseteq \mathbf{diag}^+(\mathsf{Active}(t[1, i]))$ and the induced γ_i is partial homomorphisms. Thus, t is a winning move for Duplicator in ${}^{\exists^+}\Yup\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$.

For the (1) ⇐ (2) implication, suppose Duplicator has a winning strategy in ${}^{\exists^+}\Yup\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ and consider a sentence $\varphi \in \exists^+\Yup\mathcal{L}^k$ such that $\mathcal{A} \vDash \varphi$. We first assume that φ is primitive, i.e. contains no disjunctions and every restricted conjunctions contains no quantified sentences. Under these assumptions, by the restricted conjunction condition, every conjunction only involves atomics formulas and at most one existentially quantified formula. Thus, by $\mathcal{A} \vDash \varphi$, there exists a pebble sequence s on $\mathcal{A}$ witnessing the existential quantifiers in φ. Suppose Spoiler plays s in the one and only round of ${}^{\exists^+}\Yup\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$. By hypothesis, Duplicator has winning response pebble sequence t on $\mathcal{B}$. By the winning condition of ${}^{\exists^+}\Yup\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, the pebbled elements in t satisfy the atomic formulas in φ. Therefore, the elements of t witness the existential quantifiers of φ and $\mathcal{B} \vDash \varphi$. For the case where φ is not primitive, we can inductively apply rewrite rules to pull all the disjunctions outside the scope of quantifiers and all quantified sentences out of restricted conjunctions. Thus, every sentence is a positive Boolean combination of primitive sentences.

For the (2) $\Leftrightarrow$ (3) equivalence, observe that every Spoiler play s of length j determines a path $[m\colon P \rightarrowtail T_k(R_k(\mathcal{A}))]$ whose image are elements $(s, i) \in T_k(R_k(\mathcal{A}))$ for all $i \in [n]$. Conversely, every $[m\colon P \rightarrowtail T_k(R_k(\mathcal{A}))]$ contains the embedding of the substructure induced by the elements (s, i) for all $i \in [n]$. This bijective correspondence similarly holds for Duplicator responses t and $[n\colon P \rightarrowtail T_k(R_k(\mathcal{B}))]$. Duplicator's winning condition of ${}^{\exists^+}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, i.e. partial homomorphism at every index in the chosen plays (s, t), translates to the winning condition $([m], [n]) \in {}^{+}\mathcal{W}(T_k(R_k(\mathcal{A})), T_k(R_k(\mathcal{B})))$ in ${}^{\exists^+}\mathcal{G}(T_k(R_k(\mathcal{A})), T_k(R_k(\mathcal{B})))$.

For the (3) $\Leftrightarrow$ (4) equivalence, we observe that, from morphisms in $\mathcal{R}_k^{PL}(\sigma)$ preserving covering chains, $\mathbf{P}\colon \mathcal{R}_k^{PL}(\sigma) \to \mathcal{T}$ is faithful. Applying Proposition 3.1.13(3) to the case where $\mathcal{C} = \mathcal{R}_k^{PL}(\sigma)$, $X = T_k(R_k(\mathcal{A}))$, and $Y = T_k(R_k(\mathcal{B}))$ yields the result. ☐

We now prove the pathwise power theorem relating equivalence in $\exists^+\curlywedge\mathcal{L}^k$, Duplicator winning strategies in ${}^{\exists^+}\curlywedge\mathbf{Peb}(\mathcal{A}, \mathcal{B})$, and the relation $\rightarrow_k^{\mathcal{R}^{PL}(\sigma)}$

Theorem 4.2.14. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B}$:*

(1) $\mathcal{A} \Rrightarrow^{\exists\curlywedge\mathcal{L}^k} \mathcal{B}$

(2) Duplicator has a winning strategy in ${}^{\exists}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$.

(3) Duplicator has a winning strategy in ${}^{\exists}\mathcal{G}(T_k(R_k(\mathcal{A})), T_k(R_k(\mathcal{B})))$ *over* $\mathcal{R}_k^{PL}(\sigma)$.

(4) $\mathcal{A} \rightharpoonup_k^{\mathcal{R}^{PL}(\sigma)} \mathcal{B}$

Proof. For the (1) $\Rightarrow$ (2) implication, the proof follows the same lines as in Theorem 4.2.14. The key difference is to use the type formulas $\mathbf{tp}_s$ instead of $\mathbf{tp}_s$. Thus, for every Spoiler play s of length n in ${}^{\exists}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, Duplicator's response t is such that for every $i \in [n]$, $\mathbf{diag}(\mathsf{Active}(s[1, i])) = \mathbf{diag}(\mathsf{Active}(t[1, i]))$. This allows us to conclude that the induced γ_i are partial isomorphisms and t is a winning move for Duplicator in ${}^{\exists}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$.

For the (1) $\Leftarrow$ (2) implication, the proof follows the same lines as in Theorem 4.2.14, noting that Duplicator's stronger winning condition in ${}^{\exists}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ of partial isomorphism means that the pebbled elements in the play t satisfy the atomic and negated atomic formulas in a primitive sentence $\varphi \in \exists\curlywedge\mathcal{L}^k$.

For the (2) $\Leftrightarrow$ (3) equivalence, similarly to the analogous equivalence in Theorem 4.2.13, we can associate Spoiler play and Duplicator response pairs (s,t) in ${}^{\exists^+}\curlywedge\mathbf{Peb}_k(\mathcal{A},\mathcal{B})$ with positions $([m],[n])$ in ${}^{\exists}\mathcal{G}(T_k(R_k(\mathcal{A})),T_k(R_k(\mathcal{B})))$ over $\mathcal{R}^{PL}_k(\sigma)$. The key difference is to observe that the Duplicator's winning condition ${}^{\exists}\curlywedge\mathbf{Peb}_k(\mathcal{A},\mathcal{B})$, i.e. partial isomorphism at every index in the chosen plays (s,t), translates to the stronger winning condition $([m],[n]) \in \mathcal{W}(T_k(R_k(\mathcal{A})),T_k(R_k(\mathcal{B})))$ in ${}^{\exists}\mathcal{G}(T_k(R_k(\mathcal{A})),T_k(R_k(\mathcal{B})))$.

For the (3) $\Leftrightarrow$ (4) equivalence, since the right adjoint $T_k \circ R_k$ preserves the product $\mathcal{A} \times \mathcal{B} \in \mathsf{Struct}(\sigma)$, $T_k(R_k(\mathcal{A})) \times T_k(R_k(\mathcal{B})) \in \mathcal{R}^{P}_k(\sigma)$. Applying Proposition 3.1.13(1) to the case where $\mathcal{C} = \mathcal{R}^{PL}_k(\sigma)$, $X = T_k(R_k(\mathcal{A}))$, and $Y = T_k(R_k(\mathcal{B}))$ yields the result. $\square$

As with the previous cases of arboreal covers/games comonads, we would like to obtain a theorem demonstrating that isomorphism in the linear arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}^{PL}_k(\sigma)\}$ captures logical equivalence in a logic with counting quantifiers. However, counting quantifiers actually exhibit a hidden form of branching or unrestricted conjunction. To illustrate, consider the naive resource-increasing translation of $\exists^{\leq n}x\varphi(x)$ in the counting logic to

$$\exists x_1 \dots x_n \bigwedge_{i \neq j \in [n]} x_i \neq x_j \wedge \bigwedge \varphi(x_j)$$

in the full logic. For quantified formulas $\varphi(x)$, this translation results in an unrestricted conjunct. Thus, the "non-branching" fragment $\#\curlywedge\mathcal{L}^k$ of $\#\mathcal{L}^k$ not only restriction conjunctions, but also restricts to only counting 'walks' in a structure

Remark 4.2.15. Unfortunately, in the conference paper [68], we made a mistake in our reasoning. The necessity of restricting the logic to use walk counting quantifiers rather than (element-wise) counting quantifiers was overlooked. This was corrected in the subsequent journal version [70].

We will formalise this idea by including in our logic $\#\curlywedge\mathcal{L}^k$ special 'walk pointer' variables and 'walk counting quantifiers'.

In order to define our logic $\#\curlywedge\mathcal{L}^k$, we first define a logic $\exists\curlywedge\mathcal{L}^k(Z)$. This logic is the same as $\exists\curlywedge\mathcal{L}^k$, but in addition to the k-many ordinary variables $X = \{x_1, \dots, x_k\}$, we have an another sort of infinitely-many variables $Z = \{z_1, z_2, \dots\}$ which we call 'walk pointer variables'. We also require that every existential quantifier $\exists x_j$ is guarded with an equality $x_j = z_i$ equating a ordinary variable x_j with a walk pointer variable z_i. Explicitly, the formulas of $\exists\curlywedge\mathcal{L}^k(Z)$ can be defined recursively as follows:

$$\psi(\vec{x},\vec{z}) ::= R(x_{i_1},\dots,x_{i_r}) \mid \neg p \mid \bigwedge \Psi \mid \bigvee \Psi \mid \exists x_j(x_j = z_i \wedge \psi(\vec{x'},\vec{z'})),$$

where $R \in \sigma$ is an r-ary relation such that $\{x_{i_1}, \dots, x_{i_r}\} \in [\vec{x}] \subseteq \{x_1, \dots, x_k\}$, p is an atomic formula, $\bigwedge \Psi$ is a conjunction of formulas $\psi(\vec{x}, \vec{z})$ satisfying (R), $\bigvee \Psi$ is a disjunction of formulas $\psi(\vec{x}, \vec{z})$, $x_j \in [\vec{x'}]$ and $z_i \notin \vec{z'}$.

We can now define sentences of $\#\curlywedge\mathcal{L}^k$ as disjunctions of formulas that involve counting tuples of z-variables. Explicitly, the sentence of $\#\curlywedge\mathcal{L}^k$ can be defined recursively as follows:

$$\phi ::= \bigvee \Phi \mid \exists^j(z_{i_1}, \dots, z_{i_m})\psi(z_{i_1}, \dots, z_{i_m}),$$

where Φ is a collection of sentences ϕ, $j \in \mathbb{N}$ and $\psi(z_{i_1}, \dots, z_{i_m}) \in \exists\curlywedge\mathcal{L}^k(Z)$.

Observe that in the syntax, the $\exists^j \vec{z}$ walk counting quantifier only binds to formula $\psi(\vec{z}) \in \exists\curlywedge\mathcal{L}^k(Z)$ where every free variable is from the sort Z of walk pointer variables. For every $j \in \mathbb{N}$, we define the semantics of $\exists^j$ quantifier as:

- $\mathcal{A} \vDash \exists^j(z_{i_1}, \dots, z_{i_m})\psi(z_{i_1}, \dots, z_{i_m})$ if there exist exactly j-many tuples $\vec{a} = (a_1, \dots, a_m)$ such that $\mathcal{A}, \vec{a} \vDash \psi(z_{i_1}, \dots, z_{i_m})$.

The quantifier $\exists^j(z_{i_1}, \dots, z_{i_m})$ in the logic $\#\curlywedge\mathcal{L}^k$ is inspired by a similar k-walk quantifier used in the logic defined in [60].

Recall that, a formula of $\varphi(\vec{x}) \in \exists^+\curlywedge\mathcal{L}^k$ is primitive if it contains no disjunctive subformula and every restricted conjunction subformula $\bigvee \Psi$ does not contain a sentence. From definition of primitive formula and the condition (R), it follows that every conjunctive subformula $\bigvee \Psi$ of a primitive formula $\varphi(x)$ contains only literals and at most one quantified formula with a free variable. We similarly can define primitive formulas of $\exists\curlywedge\mathcal{L}^k(Z)$. Thus, we extend this definition to $\#\curlywedge\mathcal{L}^k$, and say a sentence of $\phi \in \#\curlywedge\mathcal{L}^k$ is *primitive* if it is of the form $\phi = \exists^j(z_{i_1}, \dots, z_{i_m})\psi(z_{i_1}, \dots, z_{i_m})$ if $\psi(\vec{z}) \in \exists\curlywedge\mathcal{L}^k(Z)$ is primitive.

Proposition 4.2.16. *Every sentence $\phi \in \#\curlywedge\mathcal{L}^k$ is a disjunction of primitive subformulas.*

Proof. Given a sentence ϕ, we can inductively apply the standard rewrite rule

$$\exists y(\bigvee_{i \in I} \varphi(y, \vec{x})) \mapsto \bigvee_{i \in I} \exists y \varphi(y, \vec{x})$$

along with commutativity of $\bigwedge, \bigvee$ to obtain an equivalent formula. From the semantics it is clear that this standard rewrite rule also works for the walk counting quantifiers $\exists^j(z_{i_1}, \dots, z_{i_m})$. $\square$

The equivalence relation $\equiv_{\#\curlywedge\mathcal{L}^k}$ is characterised by Duplicator winning strategies in an all-in-one version ${}^{\#}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ of the k-pebble bijective game ${}^{\#}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ and thus, isomorphism in $\mathbf{Kl}(\mathbb{P}^L_k)$.

Definition 4.2.17. Given σ-structures $\mathcal{A}, \mathcal{B}$ and $k > 0$, we define the all-in-one bijective k-pebble game ${}^{\#}\!\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$. During the first and only round,

- Spoiler chooses a sequence of pebbles:

$$\vec{p} = [p_1, \dots, p_n].$$

- Duplicator chooses a bijection $h_{\vec{p}} \colon A^n \to B^n$.
- Spoiler chooses a sequence of pebble placements respecting $\vec{p}$:

$$s = [(p_1, a_1), \dots, (p_n, a_n)]$$

- Duplicator must respond with the list of pebble placements:

$$t = [(p_1, b_1), \dots, (p_n, b_n)].$$

  where $h_{\vec{p}}([a_1, \dots, a_n]) = [b_1, \dots, b_n]$.

Duplicator wins the round if for all $i \in [n]$, the relation

$$\gamma_i = \{(a^p, b^p) \mid p \in [k], a^p = \mathsf{last}_p(s[1, i]), b^p = \mathsf{last}_p(t[1, i])\}$$

is a partial isomorphism.

In order to connect the game ${}^{\#}\!\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ with the logic $\#\curlywedge\mathcal{L}^k$ we need to construct, as with the previous cases, the 'maximal' sentences associated to a Spoiler strategy in this game. We retain the definitions of $\mathbf{diag}(v)$ and $\mathsf{Active}(s)$, and define new formulas $\mathbf{tp}'_s(\vec{z}, \vec{x}) \in \exists\curlywedge\mathcal{L}^k(Z)$ which guard quantifiers with equalities of the form $x_p = z_j$. As before, we construct these sentences by induction on the length of s. For the base case, $s = \epsilon$ is the empty sequence, and we set $\mathbf{tp}'_\epsilon = \top$ where $\top$ is true in every structure. For the inductive step, suppose $t = s[(p, a)]$, then

$$\mathbf{tp}'_t(\vec{x'}, z'\vec{z}) = \exists x_p(x_p = z' \wedge \bigwedge \mathbf{diag}(\mathsf{Active}(t)) \wedge \mathbf{tp}'_s(\vec{x}, \vec{z})).$$

Finally, we need quantify over the last active window of s. For every s of length n, we define the $\mathbf{tp}_s(z_1, \dots, z_n, \vec{z'}) = \mathbf{tp}^n_s(z_1, \dots, z_n, \vec{z'})$ by induction up to n as follows:

$$\mathbf{tp}^1_s(z_1, x_{p_2}, \dots, x_{p_n}, \vec{z'}) = \exists x_{p_1}(x_{p_1} = z_1 \wedge \mathbf{tp}'_s(x_{p_1}, \dots, x_{p_n}, \vec{z'}))$$
$$\mathbf{tp}^{i+1}_s(z_1, \dots, z_i, x_{p_{i+1}}, \dots, x_{p_n}, \vec{z'}) = \exists x_{p_i}(x_{p_i} = z_i \wedge \mathbf{tp}^i_s(z_1, \dots, z_{i-1}, x_{p_i}, \dots, x_{p_n}, \vec{z'}))$$

We now prove the isomorphism power theorem relating equivalence in $\#\curlywedge\mathcal{L}^k$, the game ${}^{\#}\!\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, and the relation $\cong^{\mathfrak{R}^{PL}(\sigma)}_k$.

Theorem 4.2.18. *The following are equivalent for all finite σ-structures $\mathcal{A}$ and $\mathcal{B}$:*

(1) $\mathcal{A} \equiv^{\#\curlywedge\mathcal{L}^k} \mathcal{B}$.

(2) Duplicator has a winning strategy in ${}^{\#}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$.

(3) $\mathcal{A} \cong_k^{\mathfrak{R}^{PL}(\sigma)} \mathcal{B}$.

Proof. For (1) $\Rightarrow$ (2), suppose $\mathcal{A} \equiv_{\#\curlywedge\mathcal{L}^k} \mathcal{B}$. In particular, for every sentence of the form $\phi = \exists^j(z_1, \dots, z_n)\mathbf{tp}_s^Z(z_1, \dots, z_n)$, $\mathcal{A} \vDash \phi \Leftrightarrow \mathcal{B} \vDash \phi$. This determines a bijection $h \colon A^n \to B^n$ such that:

$$\mathcal{A}, (a_1, \dots, a_n) \vDash \mathbf{tp}_s^Z(z_1, \dots, z_n) \Leftrightarrow \mathcal{B}, (h(b_1), \dots, h(b_n)) \vDash \mathbf{tp}_s^Z(z_1, \dots, z_n)$$

By construction of $\mathbf{tp}_s^Z(z_1, \dots, z_n)$, the bijections h determine a Duplicator winning strategy in ${}^{\#}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$.

For (1) $\Leftarrow$ (2), suppose Duplicator has a winning strategy in ${}^{\#}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$. Consider a primitive sentence ϕ of $\#\curlywedge\mathcal{L}^k$ such that $\mathcal{A} \vDash \phi$. As ϕ is primitive, it is of the form $\exists^j(z_{i_1}, \dots, z_{i_m})\psi(z_{i_1}, \dots, z_{i_m})$ for some $j \in \mathbb{N}$ and primitive formula $\psi(z_{i_1}, \dots, z_{i_m}) \in \exists\curlywedge\mathcal{L}^k(Z)$. From ψ, we construct Spoiler's sequence of pebbles $\vec{p} = [p_1, \dots, p_n]$ in the game ${}^{\#}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ such that for all $l \in [m]$, $x_{p_l} = z_{i_l}$ is subformula of ψ guarding a existential quantifier. Let $h_{\vec{p}} \colon A^m \to B^m$ denote Duplicator response in her winning strategy for ${}^{\#}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$. By $\mathcal{A} \vDash \phi$, there exist exactly j-many tuples $\vec{a} = (a_1, \dots, a_m)$ such that $\mathcal{A}, \vec{a} \vDash \psi(z_{i_1}, \dots, z_{i_m})$. Thus, for each of these tuples, we examine the play where Spoiler plays the sequence $[(p_1, a_1), \dots, (p_m, a_m)]$. Duplicator responds with the sequence $[(p_1, b_1), \dots, (p_m, b_m)]$ such that $h_{\vec{p}}(\vec{a}) = \vec{b} = [b_1, \dots, b_m]$. From the winning condition it follows that $\mathcal{B}, \vec{b} \vDash \psi(z_{i_1}, \dots, z_{i_m})$. Since $h_{\vec{p}}$ is a bijection there are j-many tuples $\vec{b}$ such that $\mathcal{B}, \vec{b} \vDash \psi(z_{i_1}, \dots, z_{i_m})$. Thus, $\mathcal{B} \vDash \phi$. A similar argument shows that if $\mathcal{B} \vDash \phi$, then $\mathcal{A} \vDash \phi$. By Proposition 4.2.16, every sentence in $\#\curlywedge\mathcal{L}^k$ is a disjunction of primitive sentences ϕ. Therefore, $\mathcal{A} \equiv_{\#\curlywedge\mathcal{L}^k} \mathcal{B}$.

Let $P_{\vec{p}}(\mathcal{A})$ denote the induced substructure of $\mathbb{P}_k^L(\mathcal{A})$ on the subset of plays which follow $\vec{p}$:

$$P_{\vec{p}}(A) = \{[(p_1, a_1), \dots, (p_n, a_n)] \mid \forall i \in [n], a_i \in A\}.$$

We can similarly define the substructure $P_{\vec{p}}(\mathcal{B})$ of $\mathbb{P}_k^L(\mathcal{B})$.

For (2) $\Rightarrow$ (3), suppose Duplicator has a winning strategy in ${}^{\#}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$. Suppose Spoiler plays the pebble sequence $\vec{p} = [p_1, \dots, p_n]$ and Duplicator responds with $\psi_{\vec{p}} \colon A^n \to B^n$. We define the σ-isomorphism $f_{\vec{p}} \colon P_{\vec{p}}(\mathcal{A}) \to P_{\vec{p}}(\mathcal{B})$ such that if $\psi_{\vec{p}}[a_1, \dots, a_n] = [b_1, \dots, b_n]$, then $f_{\vec{p}}([(p_1, a_1), \dots, (p_n, a_n)], i) = ([(p_1, b_1), \dots, (p_n, b_n)], i)$ for all $i \in [n]$.

Taking the coproduct of $f_{\vec{p}} \colon P_{\vec{p}}(\mathcal{A}) \to P_{\vec{p}}(\mathcal{B})$ over all $\vec{p} \in [k]^{<\omega}$ we obtain the isomorphism $f^* \colon \mathbb{P}_k^L\mathcal{A} \to \mathbb{P}_k^L\mathcal{B}$. Composing with $\varepsilon_{\mathcal{B}}$ yields the existence of the Kleisli isomorphism $f \colon \mathbb{P}_k^L\mathcal{A} \to \mathcal{B}$.

For (2) $\Leftarrow$ (3), assume there exists a Kleisli isomorphism $f \colon \mathbb{P}_k^L\mathcal{A} \to \mathcal{B}$. We need construct a winning strategy for Duplicator in ${}^{\#}\curlywedge\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, so suppose Spoiler plays the sequence $\vec{p} = [p_1, \ldots, p_n]$. The coextension f^* of the Kleisli isomorphism f when restricted to $P_{\vec{p}}(\mathcal{A})$ is the isomorphism $f_{\vec{p}} \colon P_{\vec{p}}(\mathcal{A}) \to P_{\vec{p}}(\mathcal{B})$. For every $[a_1, \ldots, a_n]$, there exists an $[b_1, \ldots, b_n]$, such that for all $i \in [n]$, $f_{\vec{p}}(([(p_1, a_1), \ldots, (p_n, a_n)], i)) = ([(p_1, b_1), \ldots, (p_n, b_n)], i)$. We define Duplicator's response to be the function $\psi \colon A^n \to B^n$ which maps $[a_1, \ldots, a_n]$ to $[b_1, \ldots, b_n]$. It follows from $f_{\vec{p}}$ being an isomorphism that ψ is a bijection. Namely. consider $g_{\vec{p}} \colon P_{\vec{p}}(\mathcal{B}) \to P_{\vec{p}}(\mathcal{A})$ to generated from the inverse $g^* \colon \mathbb{P}_k^L\mathcal{B} \to \mathbb{P}_k^L\mathcal{A}$ of f^*. There exists a $\phi \colon B^n \to A^n$ such that $\phi[b_1, \ldots, b_n] = [a_1, \ldots, a_n]$ if and only for all $i \in [n]$, $g_{\vec{p}}([(p_1, b_1), \ldots, (p_n, b_n)], i) = ([(p_1, a_1), \ldots, (p_n, a_n)], i)$. Clearly, by construction $\psi^{-1} = \phi$. Moreover, if Spoiler plays the sequence $s = [(p_1, a_1), \ldots, (p_n, a_n)]$, then Duplicator must respond with $t = [(p_1, b_1), \ldots, (p_n, b_n)]$ as $\psi[a_1, \ldots, a_n] = [b_1, \ldots, b_n]$. These sequences determine relation γ_i for all $i \in [n]$. By $f_{\vec{p}}$ being an isomorphism, we can show that γ_i is a partial isomorphism. Suppose $R \in \sigma$ and consider the pairs $(a^{q_1}, b^{q_1}), \ldots, (a^{q_r}, b^{q_r}) \in \gamma_i$ such that $R^{\mathcal{A}}(a^{q_1}, \ldots, a^{q_r})$. By definition of γ_i, each $q_j = \pi_{\mathcal{A}}(s, z_j)$ and $a^{q_j} = \varepsilon_{\mathcal{A}}(s, j)$ for some $z_j \leq i \leq n$. Thus, by $R^{\mathcal{A}}(a^{q_1}, \ldots, a^{q_r})$, we have that $R^{\mathcal{A}}(\varepsilon(s, z_1), \ldots, \varepsilon(s, z_r))$ and $R^{\mathbb{P}_k^L(\mathcal{A})}((s, z_1), \ldots, (s, z_r))$. By $f_{\vec{p}}$ being a morphism, $R^{\mathbb{P}_k^L(\mathcal{B})}(f_{\vec{p}}(s, z_1), \ldots, f_{\vec{p}}(s, z_r))$ and $R^{\mathbb{P}_k^L(\mathcal{B})}((t, z_1), \ldots, (t, z_r))$. Applying $\varepsilon_{\mathcal{B}}$, we obtain that $R^{\mathbb{P}_k^L(\mathcal{B})}((t, z_1), \ldots, (t, z_r))$ and $R^{\mathcal{B}}(b^{q_1}, \ldots, b^{q_r})$. A similar proof, using the inverse of $f_{\vec{p}}$, show that if $R^{\mathcal{B}}(b^{q_1}, \ldots, b^{q_r})$, then $R^{\mathcal{A}}(a^{q_1}, \ldots, a^{q_r})$. Thus, γ_i is a tight σ-correspondence. □

Theorem 4.2.19. *The following equivalences hold:*

(1) $\mathcal{A} \to_k^{\mathcal{R}^{PL}(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \Rrightarrow^{\exists^+\curlywedge\mathcal{L}^k} \mathcal{B}$

(2) $\mathcal{A} \rightharpoonup_k^{\mathcal{R}^{PL}(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \Rrightarrow^{\exists\curlywedge\mathcal{L}^k} \mathcal{B}$

(3) $\mathcal{A} \leftrightarrow_k^{\mathcal{R}^{PL}(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \equiv^{\exists\curlywedge\mathcal{L}^k} \mathcal{B}$

(4) $\mathcal{A} \cong_k^{\mathcal{R}^{PL}(\sigma)} \mathcal{B}$ *if and only if* $\mathcal{A} \equiv^{\#\curlywedge\mathcal{L}^k} \mathcal{B}$ *(assuming* $\mathcal{A}, \mathcal{B}$ *are finite)*

Proof. Statements (1),(2), and (4) are Theorem 4.2.13, Theorem 4.2.14, and Theorem 4.2.18, respectively. By Corollary 4.1.3 bisimulations are the same as bidirectional pathwise embeddings in a linear arboreal category $\mathcal{R}^{PL}(\sigma)$, so Statement (3) follows from Theorem 4.2.14. □

4.2.3 Linear modal

In this section, we derive a linear variant of the modal comonad from Section 3.2.4 from a resource-indexed linear arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\mathcal{R}_k^{ML}(\sigma)$.

Our first step is to the observe that the category of k-height tree-ordered modal structures $\mathcal{R}_k^{M}(\sigma)$ is linearisable.

Proposition 4.2.20. *For every $k > 0$, $\mathcal{R}_k^{M}(\sigma)$ is linearisable.*

Proof. The proof is essentially the same as the proof of Proposition 4.2.1. $\square$

Thus, by Proposition 4.1.12, $\mathcal{R}_k^{M}(\sigma)$ has a linear arboreal subcategory $\mathcal{R}_k^{ML}(\sigma)$. Concretely, the objects of $\mathcal{R}_k^{ML}(\sigma)$ are are the k-height tree modal structures $(\mathcal{A}, a_0, \leq)$ where $(A, \leq)$ is a linear tree with root a_0. The morphisms $f\colon (\mathcal{A}, a_0, \leq) \to (\mathcal{B}, b_0, \leq')$ are leaf tree-ordered pointed σ-morphisms.

We exhibit the family of linear modal comonads $\mathbb{M}_k^L$ as arising from the linear arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\mathcal{R}_k^{ML}(\sigma)$. For every $k > 0$, the universe $\mathbb{M}_k^L(A, a_0)$ of pointed σ-structure $\mathbb{M}_k^L(\mathcal{A}, a_0)$ is the set of underlined k-length transition sequence in modal structure $(\mathcal{A}, a_0)$ with the additional distinguished point a_0. Explicitly,

$$\begin{aligned}\mathbb{M}_k^L(A, a_0) &:= \{a_0\} \\ &\qquad \cup \{(a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n, i) \mid n \in \mathbb{N}, i \in [n] \text{ and } \forall j \in [n], R^{\mathcal{A}}_{\alpha_i}(a_{j-1}, a_j)\}\end{aligned}$$

Similar to N_*, the counit $\varepsilon\colon \mathbb{M}_k^L \to \mathsf{Id}_{\mathsf{Struct}_\star(\sigma)}$ maps a underlined transition sequence to the underlined state, i.e. for $s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n$, $(s, i) \mapsto a_i$. For every binary transition relation $R_\alpha \in \sigma$, we define the interpretation

$$\begin{aligned}R_\alpha^{\mathbb{M}_k^L(\mathcal{A}, a_0)}(x, y) \Leftrightarrow &\text{ either } x = a_0, y = (s, 1) \text{ and the first transition of } s \text{ is } \alpha\text{; or} \\ &x = (s, i), y = (s, i+1) \text{ and the } i+1 \text{ transition appearing in } s \text{ is } \alpha\end{aligned}$$

The functor $\mathbb{M}_k^L$ maps the pointed σ-morphism $f\colon (\mathcal{A}, a_0) \to (\mathcal{B}, b_0)$ to the pointed σ-morphism $\mathbb{M}_k^L(f)\colon \mathbb{M}_k^L(\mathcal{A}) \to \mathbb{M}_k^L(\mathcal{B})$ which applies f component-wise to the transition sequence, i.e. $\mathbb{M}_k^L(f)(a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n, i) = (f(a_0) \xrightarrow{\alpha_1} f(a_1) \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} f(a_n), i)$.

In order to derive the comonad on $\mathbb{M}_k^L\colon \mathsf{Struct}_\star(\sigma) \to \mathsf{Struct}_\star(\sigma)$ from a linear arboreal cover, we recall from Section 3.2.4 the resource-indexed arboreal cover $L_k \dashv R_k$ of $\mathsf{Struct}_\star(\sigma)$ by $\mathcal{R}_k^{M}(\sigma)$ which yielded the family of comonads $\mathbb{M}_k$. In this construction, $L_k\colon \mathcal{R}_k^{M}(\sigma) \to \mathsf{Struct}_\star(\sigma)$ is the forgetful functor, throwing away the order, and $R_k\colon \mathsf{Struct}_\star(\sigma) \to \mathcal{R}_k^{M}(\sigma)$ is the cofree functor such that for $(\mathcal{A}, a_0) \in \mathbf{Ob}(\mathsf{Struct}_\star(\sigma))$, $R(\mathcal{A}, a_0)$ is the k-height modal tree σ-structure $(\mathbb{M}_k(\mathcal{A}, a_0), a_0, \sqsubseteq)$ on transition sequences and for $f\colon (\mathcal{A}, a_0) \to (\mathcal{B}, b_0) \in \mathbf{Mor}(\mathsf{Struct}_\star(\sigma))$, $R_k(f)$ is the tree-ordered pointed σ-morphism whose underlying set function applies f to each state in a transition sequence. By Proposition 4.2.20 and Theorem 4.1.15, there is an adjunction for every $k > 0$, $I_k \dashv T_k$ from $\mathcal{R}_k^{ML}(\sigma)$ to $\mathcal{R}_k^{M}(\sigma)$.

We compose the arboreal cover $L_k \dashv R_k$ with the adjunction $I_k \dashv T_k$ to obtain the adjunction $L_k \circ I_k \dashv T_k \circ R_k$. In the following, we assume that η, ε are the unit, counit of the arboreal cover $L_k \dashv R_k$ and η^L, ε^L are the unit,counit of the adjunction $I_k \dashv T_k$.

- The left adjoint $L_k \circ I_k \colon \mathcal{R}_k^{ML}(\sigma) \to \mathsf{Struct}_\star(\sigma)$ is the forgetful functor mapping a k-height tree modal structure $(\mathcal{A}, a_0, \leq)$ to $\mathcal{A}$ and a leaf, tree-ordered σ-morphism $f\colon (\mathcal{A}, a_0, \leq) \to (\mathcal{B}, b_0, \leq')$ to its underlying pointed σ-morphism $L(f)\colon (\mathcal{A}, a_0) \to (\mathcal{B}, b_0)$.
- The right adjoint $T \circ R\colon \mathsf{Struct}_\star(\sigma) \to \mathcal{R}_k^{ML}(\sigma)$ is the functor mapping $\mathcal{A}$ to the k-height linear tree modal structure $(\mathbb{M}_k^L(\mathcal{A}, a_0), a_0, \leq)$ where $\leq$ defined, as in the N^*, by:
  $$(s, i) \leq (t, j) \text{ if } s = t \text{ and } i \leq j$$
  For a σ-morphism $f\colon (\mathcal{A}, a_0) \to (\mathcal{B}, b_0)$, $(T_k \circ R_k)(f)$ has underlying σ-morphism $\mathbb{M}_k^L(f)$. The functorial action of $\mathbb{M}_k^L(f)$ preserves the distinguished index in the list, i.e. $\mathbb{M}_k^L(f)(s, i) = (t, i)$ where $t = R_k(f)(s)$ and the transitions in the sequence s. Thus, $\mathbb{M}_k^L(f)$ lifts to a morphism $T_k(R_k(f)) \in \mathbf{Mor}(\mathcal{R}_k^{ML}(\sigma))$.
- The counit $\varepsilon^*\colon L \circ I \circ T \circ R \to \mathsf{Id}_{\mathsf{Struct}_\star(\sigma)}$ is the composition $\varepsilon \circ L\varepsilon^L R$. Explicitly, as in the N^* case, the component of ε^* for a set $(\mathcal{A}, a_0) \in \mathbf{Ob}(\mathsf{Struct}_\star(\sigma))$ first maps (s, i) to the i-th prefix $s[1, i]$ of s via the $R_k(\mathcal{A}, a_0)$ component of $L\varepsilon^L$ and then maps $s[1, i]$ to its last element via the $(\mathcal{A}, a_0)$ component of ε. Thus, for $s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n \in R_k(\mathcal{A}, a_0)$, $\varepsilon(s, i) = a_i$.
- The unit $\eta^*\colon \mathsf{Id}_{\mathcal{R}_k^{ML}(\sigma)} \to T_k \circ R_k \circ L_k \circ I_k$ is the composition $T_k\eta I_k \circ \eta^L$. Explicitly, the component of η^* for a k-height modal linear tree structure $(\mathcal{A}, a_0, \leq) \in \mathbf{Ob}(\mathcal{R}_k^{ML}(\sigma))$ first maps the i-th element a_i in the component $a_0 \prec a_1 \prec \dots \prec a_n = a$ corresponding to maximal path $\downarrow a$ of X to (a, i) via η^L and then maps (a, i) to the pair $(a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n, i)$ via the $I(\mathcal{A}, a_0, \leq)$ component of $T\eta$ where by Condition (M), $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n$ is the unique sequence corresponding to chain $a_0 \prec a_1 \prec \dots \prec a_n$ in $I(\mathcal{A}, a_0, \leq)$. Thus, for $a_i \in (\mathcal{A}, a_0, \leq)$ in the maximal path $a_1 \prec \dots \prec a_n = a$ of X, $\eta^*(a_i) = (a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n, i)$.

The composition $T_k \circ R_k\colon \mathsf{Struct}_\star(\sigma) \to \mathcal{R}_k^{ML}(\sigma)$ maps $(\mathcal{A}, a_0)$ to $(\mathbb{M}_k^L(\mathcal{A}, a_0), a_0, \leq)$ and f to the leaf tree morphism with underlying σ-morphism $\mathbb{M}_k^L(f)$. Since $L \circ I \to \mathcal{R}_k^{ML}(\sigma) \to$

$\mathsf{Struct}_\star(\sigma)$ is a forgetful functor, we see that $\mathbb{M}^L_k = L_k \circ I_k \circ T_k \circ R_k$. Therefore, we have obtained our derivation, from the adjunction $L_k \circ I_k \dashv T_k \circ R_k$ of the linear k-height modal comonad $(\mathbb{M}^L_k, \varepsilon_*, \delta)$ where the comultiplication has components $\delta_X = (L \circ I)(\eta_{(T\circ R)(X)})$. Explicitly, for every pair $(s, i) \in \mathbb{M}^L_k(\mathcal{A}, a_0)$,

$$\delta_X(s, i) = (a_0 \xrightarrow{\alpha_1} (s, 1) \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} (s, n), i).$$

for $s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n$. As a corollary of Theorem 4.1.19 and Proposition 3.2.39, we show that $L_k \circ I_k \dashv T_k \circ R_k$ is and comonadic adjunction, and therefore, a linear arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\mathcal{R}^{ML}_k(\sigma)$. Thus, $\mathcal{R}^{ML}_k(\sigma) \cong \mathbf{EM}(\mathbb{M}^L_k)$.

Proposition 4.2.21. *$L_k \circ I_k \dashv T_k \circ R_k$ is a linear arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\{\mathcal{R}^{ML}_k(\sigma)\}$.*

With the resource-indexed linear arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\mathcal{R}^{ML}_k(\sigma)$, we can demonstrate how the behavioural relations in Definition 3.1.19, instantiated in this case, recover their linear variants–trace inclusion, labelled trace inclusion, and trace equivalence. Thus, proving how this arboreal cover captures equivalence in linear fragments of modal logic.

A formula $\varphi \in \Diamond\mathcal{M}_k$ is called *linear* if each conjunction in every proper subformula of φ contains at most one formula with modal operations. Explicitly, the language $\Diamond\curlywedge\mathcal{M}_k$ consists of the formulas φ in the recursive grammar:

$$\varphi ::= p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond_\alpha \psi\ , \tag{4.7}$$

$$\psi ::= \top \mid \bot \mid \psi \vee \psi \mid p \wedge \psi \mid \neg p \wedge \psi \mid \Diamond_\alpha \psi\ ,\ \ p \in \mathsf{PV}, \alpha \in \mathsf{Act}. \tag{4.8}$$

This entails for example that $\Diamond(\Diamond p \wedge \Diamond q)$ is not a linear modal formula. Accordingly, let $\Diamond\curlywedge\mathcal{M}_k$ denote the fragment of $\Diamond\mathcal{M}_k$ in which every formula is linear.

The fragments are related to linear-time behavioural equivalences which will exhibit as arising from the linear arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\{\mathcal{R}^{ML}_k(\sigma)\}$. A pointed structure $(\mathcal{A}, a_0)$ in a modal signature σ with binary relations $\{R_\alpha\}_{\alpha\in\mathsf{Act}}$ indexed by transition alphabet Act and unary relations P associated with propositional variables $p \in \mathsf{PV}$ can be re-encoded as a multi-modal Kripke model $(A, a_0, \{R^{\mathcal{A}}_\alpha\}_{\alpha\in\mathsf{Act}}, V)$ where elements of A are states, a_0 is a distinguished initial state, and a valuation map $V\colon A \to \mathcal{P}(\mathsf{PV})$ where $p \in V(a)$ if and only if $P^{\mathcal{A}}(a)$. With this re-encoding, we can define the set of *labelled traces of* $(\mathcal{A}, a_0)$:

$$\mathsf{ltraces}(a_0) = \{V(a_0)\alpha_1 V(a_1)\alpha_2 \dots \alpha_n V(a_n) \mid n \in \mathbb{N}, \forall i \in [n](R_{\alpha_i}(a_{i-1}, a_i))\}.$$

Let $\mathsf{ltraces}_k(a_0)$ denote the set of traces $V(a_0)\alpha_1 V(a_1)\alpha_2 \dots \alpha_n V(a_n)$ of length $n \leq k$.

Definition 4.2.22. Given two objects $(\mathcal{A}, a_0)$ and $(\mathcal{B}, b_0)$ in $\mathsf{Struct}_\star(\sigma)$, we define

- $a_0 \subseteq^{\mathsf{tr}} b_0$ if for all labelled traces $V(a_0)\alpha_1 V(a_1)\alpha_2 \ldots \alpha_n V(a_n) \in \mathsf{ltraces}(a_0)$, there exists a trace $V(b_0)\alpha_1 V(b_1)\alpha_2 \ldots \alpha_n V(b_n) \in \mathsf{ltraces}(b_0)$ such that $V(a_i) \subseteq V(b_i)$ for all $i \in [n]$.
- $a_0 \subseteq^{\mathsf{ltr}} b_0$ if $\mathsf{ltraces}(a_0) \subseteq \mathsf{ltraces}(b_0)$. This equivalent is to $a_0 \subseteq^{\mathsf{tr}} b_0$ but with the stronger condition that $V(a_i) = V(b_i)$.

If $a_0 \subseteq^{\mathsf{tr}} b_0$, we say that there exists a *trace inclusion* from a_0 to b_0 and that a_0 is *trace included* in b_0. If both $a_0 \subseteq^{\mathsf{ltr}} b_0$ and $b_0 \subseteq^{\mathsf{ltr}} a_0$, we write $a_0 \sim^{\mathsf{ltr}} b_0$ and say that a_0 is *label trace equivalent* to b_0.

Each of these relations can be graded by a resource parameter $k > 0$, e.g. $\subseteq^{\mathsf{tr}}_k$ for a grading of $\subseteq^{\mathsf{tr}}$, where the definitions are restricted to traces of length $\leq k$, i.e. $\mathsf{ltraces}_k(a_0)$ and $\mathsf{ltraces}_k(b_0)$.

Though the behavioural relations $\subseteq^{\mathsf{tr}}$, $\subseteq^{\mathsf{ltr}}$, and $\sim^{\mathsf{ltr}}$ are rather simple to define, to exhibit uniformity with the previous examples of arboreal covers, we observe that these relations can be phrased as the existence of Duplicator winning strategies in all-in-one modal games.

Definition 4.2.23. Given pointed σ-structures $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$ and $k \geq 0$, we define the all-in-one k-depth modal game ${}^{\exists}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ from $(\mathcal{A}, a_0)$ to $(\mathcal{B}, b_0)$. In the first and only round of the game,

- Spoiler choose a transition sequence of length $n \leq k$ in $(\mathcal{A}, a_0)$:

$$s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} a_n$$

- Duplicator responds with a corresponding transition sequence in $(\mathcal{B}, b_0)$. If there is no such sequence, Spoiler wins.

$$t = b_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} a_n$$

Duplicator wins the game if

$$P^{\mathcal{A}}(a_i) \Leftrightarrow P^{\mathcal{B}}(b_i)$$

for all $i \in (n)$ and unary $P \in \sigma$.

The all-in-one positive k-depth modal game ${}^{\exists^+}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ is similar to ${}^{\exists}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ where Duplicator's winning condition is weakened so that for every $i \in (n)$ and unary $P \in \sigma$, $P^{\mathcal{A}}(a_i) \Rightarrow P^{\mathcal{B}}(b_i)$.

In order to connect to the games ${}^{\exists^+}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ and ${}^{\exists}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$, we construct 'maximal' sentences associated to Spoiler plays in these games.

For each transition sequence $s \in (\mathcal{A}, a_0)$, we construct 'maximal' sentences $\mathbf{tp}_s \in \Diamond \curlywedge \mathcal{M}_k$ and $\mathbf{tp}_s^+ \in \Diamond^+ \curlywedge \mathcal{M}_k$ such that $(\mathcal{A}, a_0) \vDash \mathbf{tp}_s$ and $(\mathcal{A}, a_0) \vDash \mathbf{tp}_s^+$ is witnessed by the states in the transition sequence s. We define these sentence $\mathbf{tp}_s^+ \in \Diamond^+ \curlywedge \mathcal{M}_k$, by induction on the length n of transition sequences s. For the base case $s = a_0$, we define

$$\begin{aligned} \mathbf{tp}_s^+ &= \bigwedge_{p \in V(a_0)} p \\ \mathbf{tp}_s &= \bigwedge_{p \in V(a_0)} p \wedge \bigwedge_{q \notin V(a_0)} \neg q. \end{aligned}$$

For the inductive step, assume s' is the extension of s by the transition $a_n \xrightarrow{\alpha} a_{n+1}$, we define

$$\begin{aligned} \mathbf{tp}_{s'}^+ &= \Diamond_\alpha \mathbf{tp}_s^+ \wedge \bigwedge_{p \in V(a_{n+1})} p \\ \mathbf{tp}_{s'} &= \Diamond_\alpha \mathbf{tp}_s \wedge \bigwedge_{p \in V(a_{n+1})} p \wedge \bigwedge_{q \notin V(a_{n+1})} \neg q. \end{aligned}$$

Intuitively, these sentences represent complete types (in the model-theoretic sense) associated to transitions sequences s in $(\mathcal{A}, a_0)$. The following proposition solidifies this intuition.

Proposition 4.2.24. *For every $\varphi \in \Diamond \curlywedge \mathcal{M}_k$ or $\varphi \in \Diamond^+ \curlywedge \mathcal{M}_k$, there exists a finite collection of sentences W of the form $\phi = \bigwedge \mathbf{tp}_s$ or $\phi = \bigwedge \mathbf{tp}_s^+$, respectively, such that φ is equivalent $\bigvee W$.*

Proof. For every structure $(\mathcal{A}, a_0)$, we consider the conjunction of all $\mathbf{tp}_s$ for every at most k-length reachable transition sequence s in $(\mathcal{A}, a_0)$:

$$\Phi(\mathcal{A}, a_0) = \bigwedge \{ \mathbf{tp}_s \mid s \in \mathbb{M}_k(\mathcal{A}, a_0) \}.$$

Consider the collection $Z(\varphi) = \{ (\Phi(\mathcal{A}, a_0) \mid (\mathcal{A}, a_0) \vDash \varphi \}$. The sentence $\varphi \in \Diamond \curlywedge \mathcal{M}_k$ is equivalent to the (ostensibly) infinitary disjunction $\bigvee Z(\varphi)$. However, since the relational signature σ is finite and the modal depth is bounded by k, by induction, there are only finitely many inequivalent sentences in the collection $Z(\varphi)$. Let W be a finite collection of inequivalent sentences in $Z(\varphi)$ such that every sentence in $Z(\varphi)$ is equivalent one of the sentences in $Z(\varphi)$. By construction, φ is equivalent to $\bigvee W$. The proof of the case where $\varphi \in \Diamond^+ \curlywedge \mathcal{M}_k$ is similar. $\square$

We now prove the morphism and pathwise power theorems demonstrating how the relations $\rightarrow_k^{\mathcal{R}^{ML}(\sigma)}$ and $\rightharpoonup_k^{\mathcal{R}^{ML}(\sigma)}$ capture equivalence in the linear fragments $\Diamond^+ \curlywedge \mathcal{M}_k$ and $\Diamond \curlywedge \mathcal{M}_k$, respectively, of $\Diamond \mathcal{M}_k$. Observe that these theorems also show that $\rightarrow_k^{\mathcal{C}^L}$ and $\rightharpoonup_k^{\mathcal{C}^L}$ are generalisations of the linear time behavioural relations $\subseteq_k^{\mathsf{tr}}$ and $\subseteq_k^{\mathsf{ltr}}$ as they coincide in the case that $\mathcal{C} = \mathcal{R}^M(\sigma)$.

Theorem 4.2.25. *The following are equivalent for all σ-structures $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$:*

(1) $(\mathcal{A}, a_0) \Rrightarrow^{\Diamond^+ \curlywedge \mathcal{M}_k} (\mathcal{B}, b_0)$

(2) $a_0 \subseteq_k^{\mathsf{tr}} b_0$

(3) Duplicator has a winning strategy in ${}^{\exists^+}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$.

(4) Duplicator has a winning strategy in ${}^{\exists^+}\mathcal{G}(T_k(R_k(\mathcal{A}, a_0)), T_k(R_k(\mathcal{B}, b_0)))$ *over* $\mathcal{R}_k^{ML}(\sigma)$.

(5) $(\mathcal{A}, a_0) \to_k^{\mathcal{R}^{ML}(\sigma)} (\mathcal{B}, b_0)$

Proof. For the (1) $\Rightarrow$ (2) implication, suppose $(\mathcal{A}, a_0) \Rrightarrow^{\Diamond^+ \curlywedge \mathcal{M}_k} (\mathcal{B}, b_0)$. In particular for every $s \in (\mathcal{A}, a_0)$, we have that $(\mathcal{A}, a_0) \vDash \mathbf{tp}_s^+ \Rightarrow (\mathcal{B}, b_0) \vDash \mathbf{tp}_s^+$. By construction the transition sequence $s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n$ witnesses a label trace $V(a_0)\alpha_1 \dots \alpha_n V(a_n)$ and $(\mathcal{A}, a_0) \vDash \mathbf{tp}_s^+$. Therefore, $(\mathcal{B}, b_0) \vDash \mathbf{tp}_s^+$, and there exists a corresponding transition sequence $t = b_0 \xrightarrow{\alpha_1} b_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} b_n$ witnessing a label trace $V(b_0)\alpha_1 \dots \alpha_n V(b_n)$ such that $V(a_i) \subseteq V(b_i)$. Thus, $a_0 \subseteq_k^{\mathsf{tr}} b_0$.

For the (1) $\Leftarrow$ (2) implication, suppose $a_0 \subseteq_k^{\mathsf{tr}} b_0$ and that $(\mathcal{A}, a_0) \vDash \varphi$ for some $\varphi \in \Diamond^+ \curlywedge \mathcal{M}_k$. By Proposition 4.2.24, φ is equivalent to a sentence $\bigvee S$ where S consists of sentences of the form $\phi = \bigwedge \mathbf{tp}_w^+$. Since $(\mathcal{A}, a_0) \vDash \varphi$ is equivalent to the disjunction $\bigvee S$ for at least one $\phi \in S$, $(\mathcal{A}, a_0) \vDash \phi$. Thus, $(\mathcal{A}, a_0) \vDash \bigwedge \mathbf{tp}_w^+$ and there exists a collection of transition sequences $s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n$ of $(\mathcal{A}, a_0)$ witnessing label traces $L \subseteq \mathsf{ltraces}_k(a_0)$ corresponding to each conjunct $\mathbf{tp}_w^+$. By the supposition, $a_0 \subseteq_k^{\mathsf{tr}} b_0$, we have that $L \subseteq \mathsf{ltraces}_k(b_0)$, so there also exists a collection of transition sequences $t = b_0 \xrightarrow{\alpha_1} b_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} b_n$ in $(\mathcal{B}, b_0)$ witnessing the same label traces in L, and also corresponding to each conjunct $\mathbf{tp}_w^+$. It follows that $(\mathcal{B}, b_0) \vDash \bigwedge \mathbf{tp}_w^+$ and $(\mathcal{B}, b_0) \vDash \varphi$.

For the (2) $\Leftrightarrow$ (3) $\Leftrightarrow$ (4) equivalences, observe every pair (s, t) of transition sequences, with s in $(\mathcal{A}, a_0)$ and t in $(\mathcal{B}, b_0)$ witnessing the same label trace is precisely a Spoiler play and Duplicator response in ${}^{\exists^+}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ which is in bijective correspondence with a Spoiler play $[m]$ and Duplicator response $[n]$ in ${}^{\exists^+}\mathcal{G}(T_k(R_k(\mathcal{A}, a_0)), T_k(R_k(\mathcal{B}, b_0)))$. The condition in the definition of $\subseteq_k^{\mathsf{tr}}$ that for all $i \in [n]$, $V(a_i) \subseteq V(b_i)$ is equivalent to Duplicator's winning condition in ${}^{\exists^+}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ that for all unary $P \in \sigma$, $P^{\mathcal{A}}(a_i) \Rightarrow P^{\mathcal{B}}(b_i)$ and Duplicator's winning condition in ${}^{\exists^+}\mathcal{G}(T_k(R_k(\mathcal{A}, a_0)), T_k(R_k(\mathcal{B}, b_0)))$ that $([m], [n]) \in {}^{+}\mathcal{W}(T_k(R_k(\mathcal{A}, a_0)), T_k(R_k(\mathcal{B}, b_0)))$.

For the (4) $\Leftrightarrow$ (5) equivalence, we observe that, from morphisms in $\mathcal{R}_k^{ML}(\sigma)$ preserving covering chains, $\mathbf{P} \colon \mathcal{R}_k^{ML}(\sigma) \to \mathcal{T}$ is faithful. Applying Proposition 3.1.13(3) to the case where $\mathcal{C} = \mathcal{R}_k^{ML}(\sigma)$, $X = T_k(R_k(\mathcal{A}, a_0))$, and $Y = T_k(R_k(\mathcal{B}, b_0))$ yields the result. $\square$

Theorem 4.2.26. *The following are equivalent for all* σ*-structures* $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$*:*

(1) $(\mathcal{A}, a_0) \Rrightarrow^{\Diamond \Yup \mathcal{M}_k} (\mathcal{B}, b_0)$

(2) $a_0 \subseteq_k^{\mathsf{ltr}} b_0$

(3) Duplicator has a winning strategy in ${}^{\exists}\Yup\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$.

(4) Duplicator has a winning strategy in ${}^{\exists}\mathcal{G}(T_k(R_k(\mathcal{A}, a_0)), T_k(R_k(\mathcal{B}, b_0)))$ *over* $\mathcal{R}_k^{ML}(\sigma)$.

(5) $(\mathcal{A}, a_0) \rightharpoonup_k^{\mathcal{R}^{ML}(\sigma)} (\mathcal{B}, b_0)$

Proof. The proof of the (1) ⇔ (2) equivalence follows same lines as the analogous equivalence in Theorem 4.2.25 by taking into account negations on propositional variables $p \in \mathsf{PV}$. This is done by replacing every occurrence of $\mathbf{tp}^+$ with $\mathbf{tp}$ in the proof.

For the (2) ⇔ (3) ⇔ (4) equivalences, we have the same correspondences between pairs of sequences, Spoiler and Duplicator plays in and in as is the proof of Theorem 4.2.25. The condition in the definition of $\subseteq_k^{\mathsf{ltr}}$ that for all $i \in [n]$, $V(a_i) = V(b_i)$ is equivalent to Duplicator's winning condition in ${}^{\exists^+}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ that for all unary $P \in \sigma$, $P^{\mathcal{A}}(a_i) \Leftrightarrow P^{\mathcal{B}}(b_i)$ and Duplicator's winning condition in ${}^{\exists^+}\mathcal{G}(T_k(R_k(\mathcal{A}, a_0)), T_k(R_k(\mathcal{B}, b_0)))$ that $([m], [n]) \in \mathcal{W}(T_k(R_k(\mathcal{A}, a_0)), T_k(R_k(\mathcal{B}, b_0)))$.

For the (4) ⇔ (5) equivalence, since the right adjoint $T_k \circ R_k$ preserves the product $\mathcal{A} \times \mathcal{B} \in \mathsf{Struct}(\sigma)$, $T_k(R_k(\mathcal{A})) \times T_k(R_k(\mathcal{B})) \in \mathcal{R}_k^P(\sigma)$. Applying Proposition 3.1.13(1) to the case where $\mathcal{C} = \mathcal{R}_k^{ML}(\sigma)$, $X = T_k(R_k(\mathcal{A}))$, and $Y = T_k(R_k(\mathcal{B}))$ yields the result. □

Below, given a list $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_n] \in \mathsf{Act}^*$ and pointed σ-structure $(\mathcal{A}, a_0)$, we define the set of processes which follows this list of actions:

$$S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0) := \{a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} a_n \mid \forall i \in [n], a_i \in A\}.$$

Let $\#\Yup\mathcal{M}_k$ denote the extension of $\Diamond\Yup\mathcal{M}_k$ with walk counting modalities $\Diamond_{\boldsymbol{\alpha}}^m$, where $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_n] \in \mathsf{Act}^*$ and m is a positive integer. Given a pointed Kripke structure $(\mathcal{A}, a_0)$, the semantics for this $\Diamond_{\boldsymbol{\alpha}}^m$ walk counting modality is defined by induction on the length of $\boldsymbol{\alpha}$:

- $\mathcal{A}, a \vDash \Diamond_{\epsilon}^1 \varphi$ if and only if $\mathcal{A}, a \vDash \varphi$
- $\mathcal{A}, a \vDash \Diamond_{\boldsymbol{\alpha}[\beta]}^m \varphi$ if for every $m_{\boldsymbol{\alpha}}, m_\beta \in \mathbb{N}$ such that $m = m_{\boldsymbol{\alpha}} m_\beta$, there exists m_β-many $a' \in R_\beta^{\mathcal{A}}(a)$ where $\mathcal{A}, a' \vDash \Diamond_{\boldsymbol{\alpha}}^{m_{\boldsymbol{\alpha}}} \varphi$.

where $\boldsymbol{\alpha}[\beta]$ denotes the concatenation of $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_n] \in \mathsf{Act}$ and $[\beta]$ for $\beta \in \mathsf{Act}$. The syntax of this logic is the same as of $\Diamond\curlywedge\mathcal{M}_k$ but with the addition of formulas of the form $\Diamond^m_{\boldsymbol{\alpha}}\psi$ to the φ grammar of equation (4.7).

Definition 4.2.27. Given two pointed Kripke structures $(\mathcal{A}, a_0)$ and $(\mathcal{B}, b_0)$ in $\mathsf{Struct}_\star(\sigma)$, we write $a_0 \sim^{\mathsf{bltr}} b_0$ if for every $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_n]$, there exists a bijection

$$f_{\boldsymbol{\alpha}} \colon S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0) \to S_{\boldsymbol{\alpha}}(\mathcal{B}, b_0),$$

such that for each $s = a_0 \xrightarrow{\alpha_1} a_1 \ldots \xrightarrow{\alpha_n} a_n$ and $t = b_0 \xrightarrow{\alpha_1} b_1 \ldots \xrightarrow{\alpha_n} b_n$ with $f_{\boldsymbol{\alpha}}(s) = t$, we have $V(a_i) = V(b_i)$, for all $i \geq 0$. If $a_0 \sim^{\mathsf{bltr}} b_0$, we say a_0 is *bijective trace equivalent* to b_0

As with $\subseteq^{\mathsf{tr}}$ and $\sim^{\mathsf{ltr}}$, we can also obtain the relation $\sim^{\mathsf{bltr}}_k$ by restricting the definition of $\sim^{\mathsf{bltr}}$ to use action sequences $\boldsymbol{\alpha} \in \mathsf{Act}^{\leq k}$

The relations $\subseteq^{\mathsf{tr}}_k$ and $\subseteq^{\mathsf{ltr}}_k$ could be formulated as the all-in-one k-depth modal games ${}^{\exists^+}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ and ${}^{\exists}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$. Similarly we can phrase $\sim^{\mathsf{bltr}}_k$ as an all-in-one bijective k-depth modal game ${}^{\#}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$.

Definition 4.2.28. Given pointed σ-structures $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$ and $k \geq 0$, we define the all-in-one k-round modal game ${}^{\#}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$ from $(\mathcal{A}, a_0)$ to $(\mathcal{B}, b_0)$. In the first and only round of the game,

- Spoiler choose an action sequence of length $n \leq k$, $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_n]$
- Duplicator responds with a bijection $f_{\boldsymbol{\alpha}} \colon S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0) \to S_{\boldsymbol{\alpha}}(\mathcal{B}, b_0)$
- Spoiler chooses a transition sequence

$$s = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} a_n \in S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0)$$

- Duplicator is forced to respond with the transition sequence

$$f(s) = b_0 \xrightarrow{\alpha_1} b_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} b_n \in S_{\boldsymbol{\alpha}}(\mathcal{B}, b_0)$$

Duplicator wins the game if

$$P^{\mathcal{A}}(a_i) \Leftrightarrow P^{\mathcal{B}}(b_i)$$

for all $i \in (n)$ and unary $P \in \sigma$.

Theorem 4.2.29. *The following are equivalent for all image-finite σ-structures $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$:*

(1) $(\mathcal{A}, a_0) \equiv^{\#\curlywedge\mathcal{M}_k} (\mathcal{B}, b_0)$

(2) $a_0 \sim^{\mathsf{bltr}}_k b_0$

(3) Duplicator has a winning strategy in $^{\#}\curlywedge\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$*.*

(4) $(\mathcal{A}, a_0) \cong_k^{\mathfrak{R}^{ML}(\sigma)} (\mathcal{B}, b_0)$

Proof. For the (1) $\Rightarrow$ (2) implication, we associate to each action sequence $\alpha = [\alpha_1, \ldots, \alpha_n] \in \mathsf{Act}^{\leq k}$, a 'maximal' sentence $\mathbf{tp}^{\#}_{\mathcal{A},\boldsymbol{\alpha}} \in \Diamond\curlywedge\mathcal{M}_k$ such that $(\mathcal{A}, a_0) \vDash \mathbf{tp}^{\#}_{\mathcal{A},\boldsymbol{\alpha}}$ defined as:

$$\mathbf{tp}^{\#}_{\mathcal{A},\boldsymbol{\alpha}} = \Diamond^m_{\boldsymbol{\alpha}} \left(\bigvee_{s \in S_{\boldsymbol{\alpha}}(\mathcal{A},a_0)} \mathbf{tp}_s \right) \wedge \bigwedge_{s \in S_{\boldsymbol{\alpha}}(\mathcal{A},a_0)} \mathbf{tp}_s$$

By hypothesis, $(\mathcal{A}, a_0) \equiv^{\#\curlywedge\mathcal{M}_k} (\mathcal{B}, b_0)$. In particular, from $(\mathcal{A}, a_0) \vDash \mathbf{tp}^{\#}_{\mathcal{A},\boldsymbol{\alpha}}$, we can conclude that $(\mathcal{B}, b_0) \vDash \mathbf{tp}^{\#}_{\mathcal{A},\boldsymbol{\alpha}}$. Since $(\mathcal{B}, b_0)$ satisfies the first conjunct of $\mathbf{tp}^{\#}_{\mathcal{A},\boldsymbol{\alpha}}$, $m = |S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0)| = |S_{\boldsymbol{\alpha}}(\mathcal{B}, b_0)|$ and there is a bijection $f_{\boldsymbol{\alpha}} \colon S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0) \to S_{\boldsymbol{\alpha}}(\mathcal{B}, b_0)$. Since $(\mathcal{B}, b_0)$ satisfies the second conjunct of $\mathbf{tp}^{\#}_{\mathcal{A},\boldsymbol{\alpha}}$, by the construction of $\mathbf{tp}_s$, if $f_{\boldsymbol{\alpha}}(s) = t$ for $s = a_0 \xrightarrow{\alpha_1} a_1 \ldots \xrightarrow{\alpha_n} a_n$ and $t = b_0 \xrightarrow{\alpha_1} b_1 \ldots \xrightarrow{\alpha_n} b_n$, then $V(a_i) = V(b_i)$, for all $i \geq 0$. Therefore, $a_0 \sim_k^{\mathsf{bltr}} b_0$.

For the (1) $\Leftarrow$ (2) implication, suppose $a_0 \sim_k^{\mathsf{bltr}} b_0$, we need to show that for all $\varphi \in \#\curlywedge\mathcal{M}_k$, $(\mathcal{A}, a_0) \vDash \varphi \Leftrightarrow (\mathcal{B}, b_0) \vDash \varphi$. By the definition of $\#\curlywedge\mathcal{M}_k$, every φ is a positive Boolean combination of sentences of the form $\Diamond^m_{\boldsymbol{\alpha}}\psi$ where ψ is formula in the ψ grammar of equation (4.8). Thus, it suffices to show that $(\mathcal{A}, a_0) \vDash \Diamond^m_{\boldsymbol{\alpha}}\psi \Leftrightarrow (\mathcal{B}, b_0) \vDash \Diamond^m_{\boldsymbol{\alpha}}\psi$ for all action sequences $\boldsymbol{\alpha} \in \mathsf{Act}^{\leq k}$. By Proposition 4.2.24, there is a collection W of inequivalent sentences ϕ such that each $\phi = \bigwedge \mathbf{tp}_w$ and ψ is equivalent to $\bigvee \phi$. The sentence $\Diamond^m_{\boldsymbol{\alpha}}\psi$ is equivalent to $\bigvee \Diamond^{n(\phi)}_{\boldsymbol{\alpha}}\phi$ where $\sum_{\phi \in W} n(\phi) = m$. We can partition the set of sequences $S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0)$ into sets

$$Z(\mathcal{A}, a_0, \phi) = \{s \in S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0) \mid (\mathcal{A}, a_0) \vDash \phi^j \to \mathbf{tp}_s\}$$

Intuitively, $Z(\mathcal{A}, a_0, \phi)$ consists of the sequences s whose type $\mathbf{tp}_s$ is in the conjuncts, up to logical equivalence, of $\phi = \bigwedge \mathbf{tp}_w$. By the supposition $a_0 \sim_k^{\mathsf{bltr}} b_0$ there exists a bijection $f \colon S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0) \to S_{\boldsymbol{\alpha}}(\mathcal{B}, b_0)$. Moreover, if $f_{\boldsymbol{\alpha}}(s) = t$ for $s = a_0 \xrightarrow{\alpha_1} a_1 \ldots \xrightarrow{\alpha_n} a_n$ and $t = b_0 \xrightarrow{\alpha_1} b_1 \ldots \xrightarrow{\alpha_n} b_n$, then $V(a_i) = V(b_i)$ for all $i \geq 0$, we can conclude that $f_{\boldsymbol{\alpha}}$ maps $Z(\mathcal{A}, a_0, \phi)$ to $Z(\mathcal{B}, b_0, \phi)$. Thus, $(\mathcal{A}, a_0) \vDash \bigvee \Diamond^{n(\phi)}_{\boldsymbol{\alpha}}\phi \Leftrightarrow (\mathcal{B}, b_0) \vDash \bigvee \Diamond^{n(\phi)}_{\boldsymbol{\alpha}}\phi$ Therefore, $(\mathcal{A}, a_0) \vDash \Diamond^m_{\boldsymbol{\alpha}}\psi \Leftrightarrow (\mathcal{B}, b_0) \vDash \Diamond^m_{\boldsymbol{\alpha}}\psi$.

For the (2) $\Leftrightarrow$ (3) $\Leftrightarrow$ (4) equivalences, observe that given isomorphism $f \colon R_k(\mathcal{A}, a_0) \to R_k(\mathcal{B}, b_0)$, we obtain for every $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_n] \in \mathsf{Act}^{\leq k}$, a bijection $f_{\boldsymbol{\alpha}} \colon S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0) \to S_{\boldsymbol{\alpha}}(\mathcal{B}, b_0)$ defined as $f_{\boldsymbol{\alpha}}(s) = t$ where $(t, n) = f(s, n)$. Conversely, a collection of bijections

$$\{f_{\boldsymbol{\alpha}} \colon S_{\boldsymbol{\alpha}}(\mathcal{A}, a_0) \to S_{\boldsymbol{\alpha}}(\mathcal{B}, b_0)\}_{\boldsymbol{\alpha} \in \mathsf{Act}^{\leq k}}$$

induces an morphism $f\colon R_k(\mathcal{A}, a_0) \to R_k(\mathcal{B}, b_0)$ defined as $f(s, i) = (f_{\boldsymbol{\alpha}}(s), i)$. If f is an isomorphism, then we have that $V(a_i) = V(b_i)$ for all $i \geq 0$ whenever the induced bijection $f_{\boldsymbol{\alpha}}(a_0 \xrightarrow{\alpha_1} a_1 \dots \xrightarrow{\alpha_n} a_n) = b_0 \xrightarrow{\alpha_1} b_1 \dots \xrightarrow{\alpha_n} b_n$. Conversely, if every bijection $\{f_{\boldsymbol{\alpha}}\}_{\alpha \in \mathsf{Act}^{\leq k}}$ satisfies the condition that labels are reflected and preserved, i.e. $V(a_i) = V(b_i)$, then the induced morphism $f\colon R_k(\mathcal{A}, a_0) \to R_k(\mathcal{B}, b_0)$ is an isomorphism. These definition of $\sim^{\mathsf{bltr}}$ can be phrased directly or as the game ${}^{\#}{\curlywedge}\mathbf{M}_k((\mathcal{A}, a_0), (\mathcal{B}, b_0))$. $\square$

In summary, the behavioural relations of Definition 3.1.19 for the resource-indexed arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\{\mathcal{R}^{ML}_k(\sigma)\}$ characterise equivalence various restricted conjunction, non-branching variants of modal logic $\mathcal{M}_k$ graded by nesting depth of modalities.

Theorem 4.2.30. *The following equivalences hold:*

(1) $(\mathcal{A}, a_0) \to_k^{\mathcal{R}^{ML}(\sigma)} (\mathcal{B}, b_0)$ *if and only if* $(\mathcal{A}, a_0) \Rrightarrow^{\Diamond^+ {\curlywedge}\mathcal{M}_k} (\mathcal{B}, b_0)$.

(2) $(\mathcal{A}, a_0) \rightharpoonup_k^{\mathcal{R}^{ML}(\sigma)} (\mathcal{B}, b_0)$ *if and only if* $(\mathcal{A}, a_0) \Rrightarrow^{\Diamond {\curlywedge}\mathcal{M}_k} (\mathcal{B}, b_0)$.

(3) $(\mathcal{A}, a_0) \leftrightarrow_k^{\mathcal{R}^{ML}(\sigma)} (\mathcal{B}, b_0)$ *if and only if* $(\mathcal{A}, a_0) \equiv^{\Diamond {\curlywedge}\mathcal{M}_k} (\mathcal{B}, b_0)$.

(4) $(\mathcal{A}, a_0) \cong_k^{\mathcal{R}^{ML}(\sigma)} (\mathcal{B}, b_0)$ *if and only if* $(\mathcal{A}, a_0) \equiv^{\# {\curlywedge}\mathcal{M}_k} (\mathcal{B}, b_0)$ *(assuming* $(\mathcal{A}, a_0), (\mathcal{B}, b_0)$ *are image-finite).*

Proof. Statements (1),(2), and (4) are Theorem 4.2.25, Theorem 4.2.26, and Theorem 4.2.29, respectively. By Corollary 4.1.3 bisimulations are the same as bidirectional pathwise embeddings in a linear arboreal category $\mathcal{R}^{ML}(\sigma)$, so Statement (3) follows from Theorem 4.2.26. $\square$

4.2.4 Quasi-linear modal

We now show how to the abstract definition of the quasi-linear weak adjunction from Section 4.1.2 can be applied to the resource-indexed modal arboreal cover of $\mathsf{Struct}_\star(\sigma)$ by $\{\mathcal{R}^M_k(\sigma)\}$ from Section 3.2.4 in order to obtain a categorical characterisation of complete labelled trace equivalence.

Complete labelled trace equivalence is a linear-time behavioural relation that refines labelled trace equivalence $\sim^{\mathsf{ltr}}$ by being able to distinguish between terminating processes (complete traces) and other partial processes (traces). The set $\mathsf{cltraces}(a_0) \subseteq \mathsf{ltraces}(a_0)$ of *complete labelled traces* of a_0 consists of all labelled traces $V(a_0)\alpha_1 \dots \alpha_n V(a_n)$ such that a_n is terminal, i.e. there exists no a_{n+1} such that $(a_n, a_{n+1}) \in R^{\mathcal{A}}_\alpha$ for some $\alpha \in \mathsf{Act}$.

Definition 4.2.31. Given two objects $(\mathcal{A}, a_0)$ and $(\mathcal{B}, b_0)$ in $\mathsf{Struct}_\star(\sigma)$, we write $a_0 \sim^{\mathsf{cltr}} b_0$ if $\mathsf{cltraces}(a_0) = \mathsf{cltraces}(b_0)$ and $\mathsf{ltraces}(a_0) = \mathsf{ltraces}(b_0)$ and say a_0 is *complete trace equivalent* to b_0.

As with label trace equivalence, we can obtain a grading $\sim_k^{\mathsf{cltr}}$ of $\sim^{\mathsf{cltr}}$. Let $\mathsf{cltraces}_k(a_0) = \mathsf{ltraces}_k(a_0) \cap \mathsf{cltraces}(a_0)$, then we define $a_0 \sim_k^{\mathsf{cltr}} b_0$ as $\mathsf{cltraces}_k(a_0) = \mathsf{cltraces}_k(b_0)$ and $\mathsf{ltraces}_k(a_0) = \mathsf{ltraces}_k(b_0)$.

Theorem 4.2.32. *The following are equivalent for all* $(\mathcal{A}, a_0), (\mathcal{B}, b_0) \in \mathsf{Struct}_\star(\sigma)$*:*

(1) $a_0 \sim_k^{\mathsf{cltr}} b_0$

(2) There exists a bisimulation $R_k^q(\mathcal{A}, a_0) \leftarrow Z \rightarrow R_k^q(\mathcal{B}, b_0)$ *in* $\mathcal{R}_k^M(\sigma)$

where $R_k^q = I_k^q \circ T_k^q \circ R_k \colon \mathsf{Struct}_\star(\sigma) \rightarrow \mathcal{R}_k^M(\sigma)$.

Proof. Throughout both directions of this proof, we will use the fact that an element $[m] \in \mathbf{P}(R_k^q(\mathcal{A}, a_0))$ corresponds to a transition sequence $m(r) = a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} a_n$ in $(\mathcal{A}, a_0)$ where r is the maximal element of the chain $\mathsf{dom}(m)$; and moreover, from a transition sequence s, we can construct an element $[m] \in \mathbf{P}(R_k^q(\mathcal{A}, a_0))$ where $\mathsf{dom}(m)$ is the chain whose i-th element is a_i and $m \colon \mathsf{dom}(m) \rightarrow R_k^q(\mathcal{A}, a_0)$ is the obvious embedding.

For the (1) $\Rightarrow$ (2) implication, we proceed by contrapositive and suppose that there is no bisimulation $R_k^q(\mathcal{A}, a_0) \leftarrow Z \rightarrow R_k^q(\mathcal{B}, b_0)$ in $\mathcal{R}_k^M(\sigma)$. By Proposition 3.1.13(2), Spoiler has a winning strategy in the game $\mathcal{G}(X, Y)$ over $\mathcal{R}_k^M(\sigma)$ with $X = R_k^q(\mathcal{A}, a_0)$ and $Y = R_k^q(\mathcal{B}, b_0)$. There are three cases to consider:

- Spoiler wins at round 0, thus the domains of the roots are not isomorphic. This implies that $V(a_0) \neq V(b_0)$. Therefore, $\mathsf{ltraces}_k(a_0) \neq \mathsf{ltraces}_k(b_0)$, and $a_0 \not\sim_k^{\mathsf{cltr}} b_0$.
- Spoiler wins at a subsequent round $t > 0$. where the previous position is given by $([m], [n]) \in \mathcal{W}(X, Y)$, by choosing an extension $[m' \colon P \rightarrowtail X]$ of $[m] \in \mathbf{P}(X)$, i.e. $[m'] \succ [m]$, such that there is no winning Duplicator response. That is, for every $[n'] \in \mathbf{P}(Y)$ either it does not hold that $[n'] \succ [n]$ or $([m'], [n']) \notin \mathcal{W}(X, Y)$.
  - In the case where there is no $[n'] \succ [n]$, then $[n \colon Q \rightarrowtail Y]$ is a maximal element in $\mathbf{P}(Y)$. Suppose n corresponds to the transition sequence $b_0 \xrightarrow{\alpha_1} b_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{t-1}} b_{t-1}$ in $\mathcal{B}$. Since $[n]$ is a maximal element, then $V(b_0)\alpha_1 V(b_1) \dots \alpha_{t-1} V(b_{t-1})$ is a complete trace in $\mathsf{cltraces}_k(b_0)$. On the hand, suppose m corresponds to the transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{t-1}} a_{t-1}$, then by $([m], [n]) \in \mathcal{W}(X, Y)$, $V(a_0)\alpha_1 V(a_1) \dots \alpha_{t-1} V(a_{t-1}) \in \mathsf{ltraces}_k(a_0)$ where $V(b_i) = V(a_i)$ for all $i \in$

    $\{0, \dots, t-1\}$. However, since $[m']$ properly extends $[m]$,

    $$V(a_0)\alpha_1 V(a_1)\dots\alpha_{t-1}V(a_{t-1}) = V(b_0)\alpha_1 V(b_1)\dots\alpha_{t-1}V(b_{t-1}) \notin \mathsf{cltraces}_k(a_0).$$

    Hence, $\mathsf{cltraces}_k(a_0) \neq \mathsf{cltraces}_k(b_0)$ and $a_0 \not\sim^{\mathsf{cltr}}_k b_0$.

  - In the case where $([m'],[n']) \notin \mathcal{W}(X,Y)$, suppose that m corresponds to the transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{t-1}} a_{t-1}$ of $\mathcal{A}$ and n corresponds to the to the transition sequence $b_0 \xrightarrow{\alpha_1} b_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{t-1}} b_{t-1}$ of $\mathcal{B}$. As $([m],[n]) \in \mathcal{W}(X,Y)$, we have that $V(b_i) = V(a_i)$ for all $i \in \{0,\dots,t-1\}$. Suppose $[m']$, properly extending $[m]$, corresponds to the transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{t-1}} a_{t-1} \xrightarrow{\alpha_t} a_t$. However, since every proper $[n'] \succ [n]$ extension is such that $([m'],[n']) \notin \mathcal{W}(X,Y)$, there is no b_t such that $b_0 \xrightarrow{\alpha_1} b_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{t-1}} b_{t-1} \xrightarrow{\alpha_t} b_t$ and $V(b_t) = V(a_t)$. Hence, $\mathsf{ltraces}_k(a_0) \neq \mathsf{ltraces}_k(b_0)$ and $a_0 \not\sim^{\mathsf{cltr}}_k b_0$.

- Spoiler wins at a subsequent round $t > 0$ by choosing an extension $[n']$ of $[n] \in \mathbf{P}(Y)$ such that there is no winning Duplicator response This is symmetric to the previous case.

For the (1) $\Leftarrow$ (2) implication, we proceed by contrapositive and suppose that $a_0 \not\sim^{\mathsf{cltr}}_k b_0$. By Proposition 3.1.13(2), to show there is no bisimulation $R^q_k(\mathcal{A},a_0) \leftarrow Z \rightarrow R^q_k(\mathcal{B},b_0)$ it suffices to show that Spoiler has a winning strategy in $\mathcal{G}(X,Y)$ over $\mathcal{R}^M_k(\sigma)$ where $X = R^q_k(\mathcal{A},a_0)$ and $Y = R^q_k(\mathcal{B},b_0)$. By definition of $\sim^{\mathsf{cltr}}_k$, either $\mathsf{cltraces}_k(a_0) \neq \mathsf{cltraces}_k(b_0)$ or $\mathsf{ltraces}_k(a_0) \neq \mathsf{ltraces}_k(b_0)$.

- In the case where $\mathsf{cltraces}_k(a_0) \neq \mathsf{cltraces}_k(b_0)$, we can assume without loss of generality, there is some transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_t} a_t$ in $\mathcal{A}$ witnessing the complete labelled trace $V(a_0)\alpha_1 V(a_1)\dots\alpha_t V(a_t)$, but no corresponding transition sequence $b_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_t} b_t$ in $\mathcal{B}$ witnessing the same complete trace. This allows us to construct a Spoiler winning strategy in which Spoiler plays along the path $[m] \in \mathbf{P}(X)$ corresponding to the transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{t-1}} a_{t-1}$. This results in the position $([m],[n])$ in $\mathcal{G}(X,Y)$. Spoiler can extend $[m]$ to the $[m']$ corresponding to the transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_t} a_t$, but there is is no such extension $[n']$ of $[n]$. Thus, Spoiler wins $\mathcal{G}(X,Y)$ over $\mathcal{R}^M_k(\sigma)$.

- In the case where $\mathsf{ltraces}_k(a_0) \neq \mathsf{ltraces}_k(b_0)$, we can assume without loss of generality, there is some transition sequence $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_t} a_t$ in $\mathcal{A}$ witnessing the labelled trace $V(a_0)\alpha_1 V(a_1)\dots\alpha_t V(a_t)$, but no corresponding transition sequence $b_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_t} b_t$ in $\mathcal{B}$ witnessing the same labelled trace. If consider the corresponding path elements of transition sequences, we see that that for some

$[m] \in \mathbf{P}(X)$ corresponding to transition sequence s in $\mathcal{A}$, and every $[n] \in \mathbf{P}(Y)$, $([m],[n]) \notin \mathcal{W}(X,Y)$. Thus, Spoiler wins the $\mathcal{G}(X,Y)$ over $\mathcal{R}_k^M(\sigma)$.

$\square$

Chapter 5

CoKleisli laws

Though model comparison games are a central tool for proving when two σ-structures are equivalent in a logic, finding Duplicator winning strategies in concrete cases can often involve intricate combinatorial arguments. For this reason, (finite) model theorists use higher-level methods that employ model-comparison games as a low-level tool, e.g. locality, zero-one laws. Another such method involves inductively applying operations F, on some base structures, to build-up a composite structure. For this method to be useful in establishing inexpressiblity, the operations F must be shown to preserve equivalence $\equiv^{\mathcal{J}}$ in the logic $\mathcal{J}$ of interest; or more generally, an n-ary operation F may map equivalence(s) $\equiv^{\mathcal{J}_i}$ between input structures to the equivalence $\equiv^{\mathcal{J}}$ in the output structure:

$$\text{if } \forall i \in [n] \text{ and } \mathcal{A}_i, \mathcal{B}_i \in \mathbf{Ob}(\mathcal{C}_i), \mathcal{A}_i \equiv^{\mathcal{J}_i} \mathcal{B}_i, \text{ then } F(\mathcal{A}_1, \ldots, \mathcal{A}_n) \equiv^{\mathcal{J}} F(\mathcal{B}_1, \ldots, \mathcal{B}_n)$$

The first result of this type was theorem of Mostoswki [71] stating that binary Cartesian products of structures preserve equivalence in first-order logic. Later, Feferman and Vaught in [35] showed that infinite disjoint unions, infinite products, and a host of other product operations also preserved logical equivalence. Thus, results of this type became known as Feferman-Vaught theorems. We will refer these results as Feferman-Vaught-Mostoswki, or FVM theorems.

5.1 Operations

Since we will be dealing with operations of arbitrary arity, we will need some notation navigating product categories and conveying the types for these operations. Let $\mathcal{C}_i$ be a collection of categories with indices $i \in I$, we write $\prod_{i \in I} \mathcal{C}_i$ for the product category where each object $\vec{A} \in \mathbf{Ob}(\prod_{i \in I} \mathcal{C}_i)$ is I-indexed collection of objects $A_i \in \mathcal{C}_i$ and each morphism $\vec{f} \colon \vec{A} \to \vec{B}$ is an I-indexed collection of morphisms $f_i \colon A_i \to B_i \in \mathbf{Mor}(\mathcal{C}_i)$. If I is a finite set $[n]$, we use $\mathcal{C}_1 \times \cdots \times \mathcal{C}_n$ to denote $\prod_{i \in I} \mathcal{C}_i$. When the indexing set I is clear from context, we will use $\vec{\mathcal{C}} = \prod_{i \in I} \mathcal{C}_i$. If $F_i \colon \mathcal{C}_i \to \mathcal{C}_i$ is an I-indexed collection of endofunctors, then this induces a endofunctor $\vec{F} \colon \vec{\mathcal{C}} \to \vec{\mathcal{C}}$. If $(\mathbb{S}_i \colon \mathcal{C}_i \to \mathcal{C}_i, \varepsilon_i, \delta_i)$ is an I-indexed collection of comonads, then this induces a product comonad $(\vec{\mathbb{S}}, \vec{\varepsilon}, \vec{\delta})$ on $\vec{\mathcal{C}}$ which acts component-wise. From routine calculations, observe that $\mathbf{Kl}(\vec{\mathbb{S}}) = \prod \mathbf{Kl}(\mathbb{S}_i)$ and for every $i \in I$, if $U_{\mathbb{S}_i} \dashv G_{\mathbb{S}_i}$ is the Kleisli resolution of $\mathbb{S}_i \colon \mathcal{C}_i \to \mathcal{C}_i$, then we obtain the Kleisli resolution $U_{\vec{\mathbb{S}}} \dashv G_{\vec{\mathbb{S}}}$ of comonad $\vec{\mathbb{S}} \colon \vec{\mathcal{C}} \to \vec{\mathcal{C}}$.

To illustrate these FVM theorems, we will use the binary disjoint union of σ-structures, which is the categorical coproduct $\uplus \colon \mathsf{Struct}(\sigma) \times \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma)$ in $\mathsf{Struct}(\sigma)$, as a running example. Recall that given two σ-structures $\mathcal{A}_1, \mathcal{A}_2$, the coproduct $\mathcal{A}_1 \uplus \mathcal{A}_2$ has universe given by the disjoint union

$$A_1 \sqcup A_2 = \{(z_i, a_i) \mid z_i \in \{1,2\}, a_i \in A_i\}; \text{ and}$$

for every m-ary relation $R \in \sigma$, the interpretation

$$R^{\mathcal{A}_1 \uplus \mathcal{A}_2}((z_1, a_1), \dots, (z_r, a_r)) \Leftrightarrow \exists z \in \{1,2\} \forall j \in [m], z_j = z \\ \text{and } R^{\mathcal{A}_i}(a_1, \dots, a_n).$$

The coproduct operation on σ-structures extends to a functor $\uplus \colon \mathsf{Struct}(\sigma) \times \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma)$ mapping morphisms $f_1 \colon A_1 \to B_1$, $f_2 \colon A_2 \to B_2$ where the morphism $f_1 \uplus f_2$ where

$$(f_1 \uplus f_2)(z, a) = (i, f_z(a)).$$

5.2 Discrete transformations and coKleisli laws

In this section, we show how to utilise the comonadic characterisation of logical preservation and equivalence, detailed in the examples in Section 3.2, to recover FVM theorems for the positive-existential $\exists^+\mathcal{J}_k$ and counting variants $\#\mathcal{J}_k$ of the logic $\mathcal{J}_k$ associated to a game comonad $\mathbb{C}_k$.

The statement of the coproduct FVM theorem for $\exists^+\mathbf{FO}_k$ establishes that

Proposition 5.2.1. *if* $\mathcal{A}_1 \Rrightarrow^{\exists^+\mathbf{FO}_k} \mathcal{B}_1$ *and* $\mathcal{A}_2 \Rrightarrow^{\mathbf{FO}_k} \mathcal{B}_2$*, then* $\mathcal{A}_1 \uplus \mathcal{A}_2 \Rrightarrow^{\exists^+\mathbf{FO}_k} \mathcal{B}_1 \uplus \mathcal{B}_2$.

To prove this statement using the k-round EF arboreal cover of $\mathsf{Struct}(\sigma)$ by $\{\mathcal{R}^E_k(\sigma)\}$, recall from Theorem 3.2.19 that $\mathcal{A} \Rrightarrow^{\exists^+\mathbf{FO}_k} \mathcal{B}$ is characterised by the existence of $R_k(\mathcal{A}) \to R_k(\mathcal{B}) \in \mathbf{Mor}(\mathcal{R}^E_k(\sigma))$. By the adjunction $L_k \dashv R_k$ witnessing the arboreal cover $\mathsf{Struct}(\sigma)$ by $\mathcal{R}^E_k(\sigma)$, this is equivalent to the existence of a morphism $L_k(R_k(\mathcal{A})) \to \mathcal{B} \in$

$\mathbf{Mor}(\mathsf{Struct}(\sigma))$, which by $L_k R_k = \mathbb{E}_k$ we can write as a coKleisli morphism $\mathbb{E}_k(\mathcal{A}) \to \mathcal{B}$. If we assume, that for $i \in \{1,2\}$, $\mathcal{A}_i \Rrightarrow^{\exists^+\mathbf{FO}_k} \mathcal{B}_i$, then this witnessed by a pair of morphisms $f_i \colon \mathbb{E}_k(\mathcal{A}_i) \to \mathcal{B}_i$. To prove $\mathcal{A}_1 \uplus \mathcal{A}_2 \Rrightarrow^{\exists^+\mathbf{FO}_k} \mathcal{B}_1 \uplus \mathcal{B}_2$, we must construct a morphism of type $\mathbb{E}_k(\mathcal{A}_1 \uplus \mathcal{A}_2) \to \mathcal{B}_1 \uplus \mathcal{B}_2$. Observe that for every $\mathcal{A}_1, \mathcal{A}_2$, we can define a morphism $\kappa_{\mathcal{A}_1,\mathcal{A}_2} \colon \mathbb{E}_k(\mathcal{A}_1 \uplus \mathcal{A}_2) \to \mathbb{E}_k(\mathcal{A}_1) \uplus \mathbb{E}_k(\mathcal{A}_2)$ defined as:

$$\kappa_{\mathcal{A}_1,\mathcal{A}_2}([(z_1,a_1),\ldots,(z_n,a_n)]) = (z_n, [a_j \mid z_j = z_n])^1 \tag{5.1}$$

If $\kappa(s)$ is in the z-th component of $\mathbb{E}_k(\mathcal{A}_1) \uplus \mathbb{E}_k(\mathcal{A}_2)$, let $\tau(s)$ be defined such that $\kappa(s) = (z, \tau(s))$. Composing with the coproduct $f_1 \uplus f_2 \colon \mathbb{E}_k(\mathcal{A}_1) \uplus \mathbb{E}_k(\mathcal{A}_2) \to \mathcal{B}_1 \uplus \mathcal{B}_2$, we can construct a morphism $f_1 \uplus f_2 \circ \kappa_{\mathcal{A}_1,\mathcal{A}_2} \colon \mathbb{E}_k(\mathcal{A}_1 \uplus \mathcal{A}_2) \to \mathcal{B}_1 \uplus \mathcal{B}_2$ as desired. Abstracting from this example to any operation F, we say that family of maps $\kappa_{\vec{A}} \colon \mathbb{T}F(\vec{A}) \to F(\vec{\mathbb{S}}(\vec{A}))$ is *a discrete transformation* $\kappa \colon \mathbb{T} \circ F \to F \circ \vec{\mathbb{S}}$. Using this terminology, the argument we gave for coproducts clearly generalises to any collection of comonads and operation F of arbitrary arity, allowing us to interpret discrete transformations as witnessing an abstract morphism FVM theorem:

Theorem 5.2.2. *Suppose for every $i \in I$, we have a comonad $\mathbb{S}_i \colon \mathcal{C}_i \to \mathcal{C}_i$, a comonad $\mathbb{T} \colon \mathcal{D} \to \mathcal{D}$, and a functor $F \colon \vec{\mathcal{C}} \to \mathcal{D}$. If for every $\vec{A} \in \mathbf{Ob}(\vec{\mathcal{C}})$, there exists a morphism*

$$\kappa_{\vec{A}} \colon \mathbb{T}F(\vec{A}) \to F\vec{\mathbb{S}}(\vec{A})$$

then for all $i \in I$, $A_i, B_i \in \mathbf{Ob}(\mathcal{C}_i)$, $A_i \to_{\mathbf{Kl}(\mathbb{S}_i)} B_i$ implies $F(\vec{A}) \to_{\mathbf{Kl}(\mathbb{T})} F(\vec{B})$.

Proof. Suppose for all $i \in I$ and $A_i, B_i \in \mathbf{Ob}(\mathcal{C}_i)$, there exists coKleisli morphisms $f_i \colon \mathbb{S}_i(A_i) \to B_i$. These induce a coKleisli morphism $\vec{f} \colon \vec{\mathbb{S}}(\vec{A}) \to \vec{B}$. Applying F and composing with $\kappa_{\vec{A}}$, we obtain coKleisli morphism $\kappa_{\vec{A}} \circ F(\vec{f}) \colon \mathbb{T}(F(\vec{A})) \to F(\vec{B})$ for the comonad $\mathbb{T}$. □

Thus, by observing that $\{\kappa_{\mathcal{A}_1,\mathcal{A}_2}\}_{(\mathcal{A}_1,\mathcal{A}_1) \in \mathsf{Struct}(\sigma) \times \mathsf{Struct}(\sigma)}$ is discrete transformation, we may obtain the coproduct FVM theorem for $\exists^+\mathbf{FO}_k$ as a corollary of Theorem 5.2.2 and Theorem 3.2.14.

Proposition 5.2.3. *κ is a discrete transformation.*

Proof. We must verify that for every $\vec{\mathcal{A}} = (\mathcal{A}_1, \mathcal{A}_2) \in \mathbf{Ob}(\mathsf{Struct}(\sigma) \times \mathsf{Struct}(\sigma))$, $\kappa_{\vec{\mathcal{A}}}$ as defined in equation (5.1) is a σ-morphism. Suppose $R \in \sigma$ is a r-ary relation and $(s_1, \ldots, s_r) \in R^{\mathbb{E}_k(\mathcal{A}_1 \uplus \mathcal{A}_2)}$. By the pairwise comparability condition, there exists some

---

[1] Here, we are using 'Haskell-style' list comprehension notation, where $[x_j \mid \Phi(j)]$ is the sublist of $[x_1, \ldots, x_n]$ which satisfies the predicate $\Phi(j)$. For example, given the list $[(i_1,a_1),(i_2,a_2),(i_3,a_3)] = [(1,a_1),(2,a_2),(1,a_3)]$, the notation $[a_j \mid i_j = 1]$ denotes $[a_1, a_3]$ and $[a_j \mid i_j = 2]$ denotes $[a_2]$.

maximal element $s_* = \max\{s_1, \dots, s_r\}$ with respect to the prefix order. By the compatibility condition, $(\varepsilon(s_1), \dots, \varepsilon(s_r)) \in R^{\mathcal{A}_1 \uplus \mathcal{A}_2}$ and so, by the definition of $R^{\mathcal{A}_1 \uplus \mathcal{A}_2}$, there exists a $z \in \{1, 2\}$, such that for all $j \in [r]$, $\varepsilon(s_j) \in \mathcal{A}_z$. Since all of the $\varepsilon(s_j)$ are in the z-th component of $\mathcal{A}_1 \uplus \mathcal{A}_2$ and $s_j \sqsubseteq s_*$, all of the $\kappa(s_j)$ are in the same component of $\mathbb{E}_k(\mathcal{A}_1) \uplus \mathbb{E}_k(\mathcal{A}_2)$ and we can assume $\kappa(s_j) = (z, t_j)$. By $t_j \sqsubseteq t_*$ $\varepsilon(s_j) = \varepsilon(t_j)$, we have that $(t_1, \dots, t_r) \in R^{\mathbb{E}_k(\mathcal{A}_z)}$ and $(\kappa(s_1), \dots, \kappa(s_r)) \in R^{\mathbb{E}_k(\mathcal{A}_1) \uplus \mathbb{E}_k(\mathcal{A}_2)}$. □

Thus, we have recovered using game comonads the statement of the coproduct FVM theorem for $\exists^+\mathbf{FO}_k$ (Proposition 5.2.1).

We now turn to proving a similar abstract isomorphism FVM theorem for counting logics.

The statement of the coproduct FVM theorem for $\#\mathbf{FO}_k$ establishes that

Proposition 5.2.4. . *if* $\mathcal{A}_1 \equiv^{\#\mathbf{FO}_k} \mathcal{B}_1$ *and* $\mathcal{A}_2 \equiv^{\#\mathbf{FO}_k} \mathcal{B}_2$*, then* $\mathcal{A}_1 \uplus \mathcal{A}_2 \equiv^{\#\mathbf{FO}_k} \mathcal{B}_1 \uplus \mathcal{B}_2$*,*

Just as coKleisli morphism $\to_{\mathbf{Kl}(\mathbb{E}_k)}$ characterises the preservation relation $\Rrightarrow^{\exists^+\mathbf{FO}_k}$ for positive-existential $\exists^+\mathbf{FO}$, coKleisli isomorphism $\cong_{\mathbf{Kl}(\mathbb{E}_k)}$ characterises the equivalence relation $\equiv^{\#\mathbf{FO}}$, on finite structures, for $\mathbf{FO}_k$ extended with counting quantifiers $\#\mathbf{FO}_k$. Since $\#\mathbf{FO}_k$ is characterised by isomorphism in $\mathbf{Kl}(\mathbb{E}_k)$, we must that if $\mathcal{A}_i \cong_{\mathbf{Kl}(\mathbb{E}_k)} \mathcal{B}_i$ for $i \in \{1, 2\}$, then $\mathcal{A}_1 \uplus \mathcal{A}_2 \cong_{\mathbf{Kl}(\mathbb{E}_k)} \mathcal{B}_1 \uplus \mathcal{B}_2$. Since $\mathcal{A}_i \cong_{\mathbf{Kl}(\mathbb{E}_k)} \mathcal{B}_i$, for each $i \in I$, we have a pair of morphisms in $f_i \colon \mathcal{A}_i \to \mathcal{B}_i \in \mathbf{Mor}(\mathbf{Kl}(\mathbb{E}_k))$ and $g_i \colon \mathcal{B}_i \to \mathcal{A}_i \in \mathbf{Mor}(\mathbf{Kl}(\mathbb{E}_k))$, which are in fact morphisms $f_i \colon \mathbb{E}_k(\mathcal{A}_i) \to \mathcal{B}_i \in \mathbf{Mor}(\mathsf{Struct}(\sigma))$ and $g_i \colon \mathbb{E}_k(\mathcal{B}_i) \to \mathcal{A}_i \in \mathbf{Mor}(\mathsf{Struct}(\sigma))$ such that $f_i \circ g_i = \varepsilon_{\mathcal{A}_i}$ and $g_i \circ f_i = \varepsilon_{\mathcal{B}_i}$ since ε, as a morphism in $\mathsf{Struct}(\sigma)$, is the identity in $\mathbf{Kl}(\mathbb{E}_k)$. In the previous case, we saw how to produce a morphisms $\kappa_{\mathcal{A}_1, \mathcal{A}_2} \circ f_1 \uplus f_2$ and $\kappa_{\mathcal{B}_1, \mathcal{B}_2} \circ g_1 \uplus g_2$. We can notate these morphisms as $f_1 \bar{\uplus} f_2 = \kappa_{\mathcal{A}_1, \mathcal{A}_2} \circ f_1 \uplus f_2 \in \mathbf{Mor}(\mathbf{Kl}(\mathbb{E}_k))$ and $g_1 \bar{\uplus} g_2 = \kappa_{\mathcal{B}_1, \mathcal{B}_2} \circ g_1 \uplus g_2$ To prove the coproduct FVM theorem, we must show that '$\bar{\uplus}$' operation preserves isomorphisms in $\mathbf{Kl}(\mathbb{E}_k)$. This follows from requiring the operation $\bar{\uplus}$ on morphisms of $\mathbf{Kl}(\mathbb{E}_k)$ to distribute over composition and preserve the identity in $\mathbf{Kl}(\mathbb{E}_k)$. That is, we require $\bar{\uplus} \colon \mathbf{Kl}(\mathbb{E}_k) \times \mathbf{Kl}(\mathbb{E}_k) \to \mathbf{Kl}(\mathbb{E}_k)$ to be a functor which extends $\uplus \colon \mathsf{Struct}(\sigma) \times \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma)$.

Definition 5.2.5. Given for all $i \in I$, a comonad $\mathbb{S}_i \colon \mathcal{C}_i \to \mathcal{C}_i$, and a comonad $\mathbb{T} \colon \mathcal{D} \to \mathcal{D}$, a functor $\bar{F} \colon \mathbf{Kl}(\vec{\mathbb{S}}) \to \mathbf{Kl}(T)$ *is a coKleisli extension of functor* $F \colon \vec{\mathcal{C}} \to \mathcal{D}$ if the following diagram commutes in $\mathbf{Cat}$:

$$\begin{tikzcd}
\mathbf{Kl}(\vec{\mathbb{S}}) \arrow[r, "\bar{F}"] & \mathbf{Kl}(\mathbb{T}) \\
\vec{\mathcal{C}} \arrow[u, "G_{\vec{\mathbb{S}}}"] \arrow[r, "F"] & \mathcal{D} \arrow[u, "G_{\mathbb{T}}"']
\end{tikzcd} \tag{5.2}$$

where for a comonad $\mathbb{W}$ on $\mathcal{E}$, $G_{\mathbb{W}} \colon \mathcal{E} \to \mathbf{Kl}(\mathbb{W})$ denotes the functor that is the identity on objects and for $f \colon X \to Y \in \mathbf{Mor}(\mathcal{E})$, $G_{\mathbb{W}}(f) = \varepsilon^{\mathbb{W}}_Y \circ \mathbb{W}(f)$.

Extensions of functors F to coKleisli categories are in bijective correspondence with natural transformations which distribute F over the underlying comonads. These are called coKleisli laws.

Definition 5.2.6. Given a comonad for all $i \in I$, a comonad $(\mathbb{S}_i\colon \mathcal{C}_i \to \mathcal{C}_i, \varepsilon_i, \delta_i)$, a comonad $(\mathbb{T}, \varepsilon, \delta)$, and a functor $F\colon \prod_{i\in I} \mathcal{C}_i \to \mathcal{D}$, then a natural transformation

$$\kappa\colon \mathbb{T}F \to F\vec{\mathbb{S}}$$

is *a coKleisli law* if it satisfies the following two diagrams:

$$\begin{tikzcd}
 & \mathbb{T}F \arrow[dl, "\kappa"'] \arrow[dr, "\varepsilon F"] & \\
F\vec{\mathbb{S}} \arrow[rr, "F\vec{\varepsilon}"] & & F
\end{tikzcd} \tag{5.3}$$

$$\begin{tikzcd}
\mathbb{T}F \arrow[rr, "\kappa"] \arrow[d, "\delta F"] & & F\vec{\mathbb{S}} \arrow[d, "F\vec{\delta}"'] \\
\mathbb{T}^2F \arrow[r, "\mathbb{T}\kappa"] & \mathbb{T}F\vec{\mathbb{S}} \arrow[r, "\kappa\vec{\mathbb{S}}"] & F\vec{\mathbb{S}}\vec{\mathbb{S}}
\end{tikzcd} \tag{5.4}$$

For reference, we state explicitly the characterisation of coKleisli extensions via coKleisli laws.

Proposition 5.2.7. *There is a bijective correspondence between:*

(1) coKleisli extensions $\bar{F}\colon \mathbf{Kl}(\vec{\mathbb{S}}) \to \mathbf{Kl}(\mathbb{T})$ *of functor* $F\colon \vec{\mathcal{C}} \to \mathcal{D}$

(2) coKleisli laws $\kappa\colon \mathbb{T}F \to F\vec{\mathbb{S}}$

Proof. The characterisation of coKleisli extensions via coKleisli laws is well-known folkloric fact. This proposition has been reproved either implicitly or as part of general statements see e.g. [19, 74, 64, 49] for details. For the sake of completeness, we give a proof sketch here. Given a coKleisli extension $\bar{F}\colon \mathbf{Kl}(\vec{\mathbb{S}}) \to \mathbf{Kl}(\mathbb{T})$ consider the identity morphism $\mathsf{id}_{\vec{\mathbb{S}}(X)}\colon \vec{\mathbb{S}}(X) \to \vec{\mathbb{S}}(X)$.

$$\begin{aligned}
\mathsf{id}_{\vec{\mathbb{S}}(X)} &\in \vec{\mathcal{C}}(\vec{\mathbb{S}}(X), \vec{\mathbb{S}}(X)) \\
&\in \mathbf{Kl}(\vec{\mathbb{S}})(X, \vec{\mathbb{S}}(X)) \\
\bar{F}(\mathsf{id}_{\vec{\mathbb{S}}(X)}) &\in \mathbf{Kl}(\mathbb{T})(\bar{F}(X), \bar{F}\vec{\mathbb{S}}(X)) \\
&\in \mathbf{Kl}(\mathbb{T})(\bar{F}(G_{\vec{\mathbb{S}}}(X)), \bar{F}\vec{\mathbb{S}}(G_{\vec{\mathbb{S}}}(X))) && G_{\vec{\mathbb{S}}} \text{ identity on objects} \\
&\in \mathbf{Kl}(\mathbb{T})(G_{\mathbb{T}}(F(X)), G_{\mathbb{T}}(F\vec{\mathbb{S}}(X))) && \text{diagram (5.2)} \\
&\in \mathbf{Kl}(\mathbb{T})(F(X), F\vec{\mathbb{S}}(X)) && G_{\mathbb{T}} \text{ identity on objects} \\
&\in \mathcal{D}(\mathbb{T}F(X), F\vec{\mathbb{S}}(X))
\end{aligned}$$

Thus, we can define $\kappa_X = \bar{F}(\mathsf{id}_{\vec{\mathbb{S}}(X)})$. From $\bar{F}$ being a functor and satisfying diagram (5.2), it follows that κ satisfies the counit diagram (5.3) and the comultiplication diagram (5.4) in Definition 5.2.6 of coKleisli law.

Conversely, given a coKleisli law $\kappa\colon \mathbb{T}F \to F\vec{\mathbb{S}}$ of functor $F\colon \vec{\mathcal{C}} \to \mathcal{D}$, we can define functor $\bar{F}\colon \mathbf{Kl}(\vec{\mathbb{S}}) \to \mathbf{Kl}(\mathbb{T})$ where $\bar{F}(X) = \bar{F}(X)$ for objects $X \in \mathbf{Ob}(\mathbf{Kl}(\vec{\mathbb{S}})) = \mathbf{Ob}(\vec{\mathcal{C}})$ and for morphisms $f \in \mathbf{Kl}(\vec{\mathbb{S}}(X), Y) = \vec{\mathcal{C}}(\vec{\mathbb{S}}(X), Y)$, $\bar{F}(f) = F(f) \circ \kappa_X$. The verification the F preserves identity and composition follows from the counit diagram (5.3) and the comultiplication diagram (5.4). $\square$

Using this correspondence, we can interpret coKleisli laws as witnessing an abstract isomorphism FVM theorem.

Theorem 5.2.8. *For every $i \in I$, we have a comonad $\mathbb{S}_i\colon \mathcal{C}_i \to \mathcal{C}_i$, a comonad $\mathbb{T}\colon \mathcal{D} \to \mathcal{D}$, and a functor $F\colon \vec{\mathcal{C}} \to \mathcal{D}$. If there exists a coKleisli law*

$$\kappa\colon \mathbb{T}F \to F\vec{\mathbb{S}}$$

then for all $i \in I$, $A_i, B_i \in \mathbf{Ob}(\mathcal{C}_i)$ $A_i \cong_{\mathbf{Kl}(\mathbb{S}_i)} B_i$ implies $AF(\vec{A}) \cong_{\mathbf{Kl}(\mathbb{T})} F(\vec{B})$.

Proof. By hypothesis, there exists coKleisli law $\kappa\colon \mathbb{T} \circ F \to F \circ \vec{\mathbb{S}}$. By Proposition 5.2.7, there exists a functor $\bar{F}\colon \mathbf{Kl}(\vec{\mathbb{S}}) \to \mathbf{Kl}(\mathbb{T})$ satisfying diagram 5.2. Suppose $i \in I$, $A_i, B_i \in \mathbf{Ob}(\mathcal{C}_i)$, $A_i \cong_{\mathbf{Kl}(\mathbb{S}_i)} B_i$, then

$$\begin{aligned}
\forall i \in I, A_i &\cong_{\mathbf{Kl}(\mathbb{S}_i)} B_i && \\
\vec{A} &\cong_{\mathbf{Kl}(\vec{\mathbb{S}})} \vec{B} && \\
G_{\vec{\mathbb{S}}}(\vec{A}) &\cong_{\mathbf{Kl}(\vec{\mathbb{S}})} G_{\vec{\mathbb{S}}}(\vec{B}) && G_{\vec{\mathbb{S}}} \text{ identity on objects} \\
\bar{F}(G_{\vec{\mathbb{S}}}(\vec{A})) &\cong_{\mathbf{Kl}(\mathbb{T})} \bar{F}(G_{\vec{\mathbb{S}}}(\vec{B})) && \hat{F} \text{ preserves isomorphism} \\
G_{\mathbb{T}}(F(\vec{A})) &\cong_{\mathbf{Kl}(\mathbb{T})} G_{\mathbb{T}}(F(\vec{B})) && \text{equation (5.2)} \\
F(\vec{A}) &\cong_{\mathbf{Kl}(\mathbb{T})} F(\vec{B}) && G_{\mathbb{T}} \text{ identity on objects}
\end{aligned}$$

$\square$

Thus, returning to the coproduct example, if we show the family of maps κ defined in Equation (5.1) is a natural transformation and satisfies diagrams (5.3) and (5.4) in Definition 5.2.6, we can obtain the coproduct FVM theorem for $\#\mathbf{FO}_k$ as corollary of Theorem 5.2.8 and Theorem 3.2.18.

Proposition 5.2.9. *$\kappa\colon \mathbb{E}_k \circ \uplus \to \uplus \circ (\mathbb{E}_k \times \mathbb{E}_k)$ is a coKleisli law.*

Proof. By Proposition 5.2.3, for every $\vec{\mathcal{A}} \in \mathbf{Ob}(\mathsf{Struct}(\sigma) \times \mathsf{Struct}(\sigma))$, $\kappa_{\vec{\mathcal{A}}}$ is a σ-morphism. To verify naturality, suppose $f_1 \colon \mathcal{A}_1 \to \mathcal{B}_1$, $f_2 \colon \mathcal{A}_2 \to \mathcal{B}_2$ are σ-morphisms, and $\vec{\mathcal{A}} = (\mathcal{A}_1, \mathcal{A}_2), \vec{\mathcal{B}} = (\mathcal{B}_1, \mathcal{B}_2)$ are objects in $\mathsf{Struct}(\sigma) \times \mathsf{Struct}(\sigma)$. We check that $\mathbb{E}_k(f_1) \uplus \mathbb{E}_k(f_2) \circ \kappa_{\vec{\mathcal{A}}} = \kappa_{\vec{\mathcal{B}}} \circ \mathbb{E}_k(f_1 \uplus f_2)$ is satisfied. Consider $[(z_1, a_1), \ldots, (z_n, a_n)] \in \mathbb{E}_k(A_1 \uplus A_2)$ and let $z = z_n$:

$$\begin{aligned}
&\mathbb{E}_k(f_1) \uplus \mathbb{E}_k(f_2) \circ \kappa_{\vec{\mathcal{A}}}([(z_1, a_1), \ldots, (z_n, a_n)]) \\
&\qquad = \mathbb{E}_k(f_1) \uplus \mathbb{E}_k(f_2)(z, [a_j \mid z_j = z)] \\
&\qquad = (z, \mathbb{E}_k(f_z)([a_j \mid z_j = z)])) \\
&\qquad = (z, [f_z(a_j) \mid z_j = z)]) \\
&\qquad = \kappa_{\vec{\mathcal{B}}}([(z_1, f_{z_1}(a_1)), \ldots, (z_n, f_{z_n}(a_n))]) \\
&\qquad = \kappa_{\vec{\mathcal{B}}}([(f_1 \uplus f_2)(z_1, a_1), \ldots, (f_1 \uplus f_2)(z_n, a_n)]) \\
&\qquad = \kappa_{\vec{\mathcal{B}}} \circ \mathbb{E}_k(f_1 \uplus f_2)([(z_1, a_1), \ldots, (z_n, a_n)])
\end{aligned}$$

The counit axiom (5.3) states that if the last element s of $\mathbb{E}_k(\mathcal{A}_1 \uplus \mathcal{A}_n)$ is (z, a) for $a \in A_z$, then the last element of $\tau(s)$ in $\kappa(s) = (z, \tau(s))$ is a. This is clear by the definition of κ. We have that $\kappa \circ \mathbb{E}_k(\kappa) \circ \delta(s) = (z, [s_j \mid S_z(s)])$ where S_z is the set of prefixes which end in an element of A_z. The comultiplication axiom (5.4) states that taking prefixes of the restriction $\tau(s)$ of s to elements in the component of S_z is the same as listing the A_z component restrictions $\tau(s_j)$ of each of the prefixes s_j in $S_z(s)$. This is clear by inspection. □

Thus, we have recovered using game comonads the statement of the FVM coproduct theorem for $\#\mathbf{FO}_k$ (Proposition 5.2.4).

Remark 5.2.10. In the paper [52], abstract bisimulation and pathwise embedding FVM theorems were also obtained. This involved using the fact that for game comonads $\mathbb{T}$, $\mathbf{EM}(\mathbb{T})$ has equalisers. For any operation F, we could compute a specific equaliser involving a coKleisli law $\kappa \colon \mathbb{T}F \to F\vec{\mathbb{S}}$ in order to define an extension $\hat{F} \colon \mathbf{EM}(\vec{\mathbb{S}}) \to \mathbf{EM}(\mathbb{T})$. The extended operation $\hat{F}$ must satisfy two axioms in order to show it preserved bisimulations. To present these axioms in terms of F instead of $\hat{F}$, we tailored the standard machinery of bilinear maps over strong commutative monads and documented these adaptations in [51]. Though I had contributed ideas and candidate constructions for these results, this was primarily the work of my co-authors and so is not detailed in this thesis.

5.3 Examples of coKleisli laws

Theorem 5.2.2 and Theorem 5.2.8 re-interpret of discrete transformations and coKleisli laws in category theory as witnesses to abstract FVM theorems in logic.g In this section, we apply this re-interpretation to game comonads recovering some results about logic.

5.3.1 Comonad morphisms and translations

One of the important special cases of coKleisli laws $\kappa\colon \mathbb{T}F \to F\mathbb{S}$ is the case where $F = \mathsf{Id}_{\mathcal{C}}$ and $\mathbb{S}, \mathbb{T}$ are comonads on the same category $\mathcal{C}$. A coKleisli laws of this form is a *comonad morphism* $\kappa\colon \mathbb{T} \to \mathbb{S}$. In this section, we demonstrate how comonad morphisms between games comonads correspond to translations between the logics they capture.

Proposition 5.3.1. *The natural transformation $\kappa\colon \mathbb{E}_k \to \mathbb{P}_k$ with components defined as:*

$$\kappa_{\mathcal{A}}([a_1, \dots, a_n]) = [(1, a_1), \dots (n, a_n)]$$

is a comonad morphism.

Proof. To verify $\kappa_{\mathcal{A}}$ is a σ-morphism, suppose $(s_1, \dots, s_r) \in R^{\mathbb{E}_k(\mathcal{A})}$ for r-ary relation $R \in \sigma$. Since κ preserves prefixes, $(\kappa(s_1), \dots, \kappa(s_r))$ satisfies the pairwise comparable condition of $R^{\mathbb{P}_k(\mathcal{A})}$ by $(s_1, \dots, s_r)$ satisfying the pairwise comparable condition of $R^{\mathbb{E}_k(\mathcal{A})}$. Moreover, since the last element of s is the last pebbled element of $\kappa(s)$, $(\kappa(s_1), \dots, \kappa(s_r))$ satisfies compatibility condition of $R^{\mathbb{P}_k(\mathcal{A})}$ by $(s_1, \dots, s_r)$ satisfying the compatibility condition of $R^{\mathbb{P}_k(\mathcal{A})}$. By definition of κ and the pairwise comparable condition of $R^{\mathbb{E}_k(\mathcal{A})}$, the active pebble condition is satisfied by $(\kappa(s_1), \dots, \kappa(s_r))$. Thus, $(\kappa(s_1), \dots, \kappa(s_r)) \in R^{\mathbb{P}_k(\mathcal{A})}$ and $\kappa_{\mathcal{A}}$ is a σ-morphism.

To verify naturality, suppose $f\colon \mathcal{A} \to \mathcal{B}$ is a σ-morphism and consider $[a_1, \dots, a_n] \in \mathbb{E}_k(\mathcal{A})$:

$$\begin{aligned}
\kappa_{\mathcal{B}} \circ \mathbb{E}_k(f)([a_1, \dots, a_n]) &= \kappa_{\mathcal{B}}([f(a_1), \dots, f(a_n)]) \\
&= [(1, f(a_1)), \dots, (n, f(a_n))] \\
&= \mathbb{P}_k(f)([(1, a_1), \dots, (n, a_n)]) \\
&= \mathbb{P}_k(f) \circ \kappa_{\mathcal{A}}([a_1, \dots, a_n])
\end{aligned}$$

Since the last element of $s \in \mathbb{E}_k(\mathcal{A})$ is the last pebbled element of $\kappa(s) \in \mathbb{P}_k(\mathcal{A})$, the counit axiom (5.3) is satisfied. For every $i \in [n]$, the i-th prefix of $\kappa(s)$ is $\kappa(s_i)$ where s_i is the i-th prefix of s. Thus, the the comultiplication axiom (5.4) and κ is coKleisli law. $\square$

Thus as a corollary of this proposition, Theorem 3.2.19 and Theorem 3.2.33, we recover semantically the following statements about translating $\exists^+\mathbf{FO}_k$ and $\#\mathbf{FO}_k$ into $\exists^+\mathcal{L}^k$ and $\#\mathcal{L}^k$, respectively:

Corollary. *For all $k > 0$ and σ-structures $\mathcal{A}, \mathcal{B}$,*

- *if $\mathcal{A} \Rrightarrow^{\exists^+\mathcal{L}^k} \mathcal{B}$, then $\mathcal{A} \equiv^{\exists^+\mathbf{FO}_k} \mathcal{B}$.*
- *if $\mathcal{A} \equiv^{\#\mathcal{L}^k} \mathcal{B}$, then $\mathcal{A} \equiv^{\#\mathbf{FO}_k} \mathcal{B}$.*

For a pointed σ-structures, we can view $\mathbb{P}_2$ as a comonad over $\mathsf{Struct}_\star(\sigma)$ where $\mathbb{P}_2(\mathcal{A}, a_0)$ has distinguishing point $[(1, a_0)]$.

Proposition 5.3.2. *Let σ be a modal signature. For every $k > 0$, the natural transformation $\kappa \colon \mathbb{M}_k \to \mathbb{P}_2$ with components defined as:*

$$\kappa_{\mathcal{A}}(a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} a_n) = [(1, a_0), (2, a_1), \dots, (1 + (n \bmod 2), a_n)]$$

is a comonad morphism.

Proof. To verify κ is a σ-morphism, suppose $a \xrightarrow{\alpha_n} a'$, i.e. $(a, a') \in R_\alpha$, then by definition $\kappa(a') = \kappa(a)[(m, a')]$ where m is not equal to the pebble on the last element of $\kappa(a)$. Thus, the pairwise comparable, active pebble, and compatibility conditions of $R^{\mathbb{P}_k(\mathcal{A})}$ are satisfied by $(\kappa(a), \kappa(a'))$. For unary relations, $P^{\mathbb{M}_k(\mathcal{A}, a_0)}(s)$ if and only if $P^{(\mathcal{A}, a_0)}(\varepsilon^{\mathbb{M}_k}(s)$ if and only if $P^{\mathbb{P}_k(\mathcal{A}, a_0)}(\varepsilon^{\mathbb{P}_k}(\kappa(s))$. Thus, κ is a σ-morphism.

To verify naturality, suppose $f \colon (\mathcal{A}, a_0) \to (\mathcal{B}, b_0)$ is a pointed σ-morphism and consider $a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} a_n \in \mathbb{M}_k(\mathcal{A}, a_0)$.

$$\begin{aligned} \kappa_{\mathcal{B}} \circ \mathbb{M}_k(f)(a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} a_n) &= \kappa_{\mathcal{B}}((b_0 \xrightarrow{\alpha_1} f(a_1) \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} f(a_n)) \\ &= [(1, a_0), (2, a_1), \dots, (m, a_n)] \\ &= \mathbb{P}_k(f)([(1, a_0), (2, a_1), \dots, (m, a_n)]) \\ &= \mathbb{P}_k(f) \circ \kappa_{\mathcal{A}}(a_0 \xrightarrow{\alpha_1} a_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} a_n) \end{aligned}$$

where $m = 1 + (n \bmod 2)$. Since the last element of transition sequence $s \in \mathbb{M}_k(\mathcal{A})$ is the last pebbled element of $\kappa(s)$, the counit axiom (5.3) is satisfied. For every $i \in [n]$, the i-th prefix of $\kappa(s)$ is $\kappa(s_i)$ where s_i is the i-th prefix of s. Thus, the the comultiplication axiom (5.4) and κ is coKleisli law. ☐

Thus as a corollary of this proposition, Theorem 3.2.33 and Theorem 3.2.46, we recover semantically the following statements about translating $\Diamond^+\mathcal{M}_k$ and $\#\mathcal{M}_k$ into $\exists^+\mathcal{L}^2$ and $\#\mathcal{L}^2$, respectively:

Corollary. *For all σ-structures $\mathcal{A}, \mathcal{B}$,*

- *if $\mathcal{A} \Rrightarrow^{\exists^{+}\mathcal{L}^{2}} \mathcal{B}$, then for all $k > 0$, $\mathcal{A} \equiv^{\mathcal{M}_k} \mathcal{B}$.*
- *if $\mathcal{A} \equiv^{\#\mathcal{L}^{2}} \mathcal{B}$, then for all $k > 0$, $\mathcal{A} \equiv^{\#\mathcal{M}_k} \mathcal{B}$.*

5.3.2 Clique-width operations

One important application of FVM theorems is to algorithmic meta-theorems such as Courcelle's clique-width theorem [27]. Courcelle's theorem proved that deciding **MSO**-definable properties over graphs of bounded clique-width is in linear time. Clique-width, like treewidth, is another important combinatorial parameter in graph theory. In fact, clique-width can be seen as the 'dense' analogue of the sparse parameter treewidth. One way of proving this meta-theorem is via FVM theorems that demonstrate that each of the operation in the grammar for building bounded clique-width graphs preserves equivalence in **MSO**. In this section, we use coKleisli laws to prove FVM theorems that demonstrate that each of the operations in a clique-width grammar over $\mathsf{Struct}(\sigma)$ preserve equivalence in $\exists^{+}\mathbf{FO}_k$, and $\#\mathbf{FO}_k$. We define a naïve generalisation of the clique-width graph grammar to $\mathsf{Struct}(\sigma)$. These operations work in the setting of m-coloured σ-structures. Thus, we extend our signature to $\sigma^m = \sigma \cup \{C_i \mid i \in [m]\}$ where each C_i is a fresh unary symbol.

Definition 5.3.3. Given a set of $[m] = \{1, \ldots, m\}$ labels, we define the following *m-clique width expression grammar* to produce relational structures $\mathsf{Struct}(\sigma^m)$.

- For every $i \in [m]$, a creation of a σ^m-structure $\mathcal{C}_i$ with singleton universe whose only element has colour C_i
- For every $i, j \in [m]$ and $i \neq j$, a recolouring operation $\rho_{i \to j}$ which produces a structure that recolours all the elements of colour C_i to colour C_j
- For every $R \in \sigma$ with arity r and $\bar{i} = (i_1, \ldots, i_r) \in [m]^r$, a R-creation operation $\eta^R_{\bar{i}}$ which produces a structure adding all tuples $C_{i_1} \times \cdots \times C_{i_r}$ to the interpretation of R.
- Coproduct of structures $\mathcal{A}_1 \uplus \mathcal{A}_2$

We say a σ^m-structure $\mathcal{B}$ is a *m-clique-width witness* if it can be obtained by finite application of the operations $\rho_{i \to j}, \eta^R_{\bar{i}}$ and $\uplus$ applied to the base structures $\mathcal{C}_1, \ldots, \mathcal{C}_m$.

Given a σ-structure $\mathcal{A}$, the clique-width $\mathsf{cw}(\mathcal{A})$ is the least m such that $\mathcal{A}$ is the σ-reduct of an m-clique-width witness $\mathcal{A}'$.

Remark 5.3.4. This naive generalisation of the clique-width grammar on graphs to $\mathsf{Struct}(\sigma)$ is used, for instance in [75]. However, as pointed out by Courcelle in [28], this definition of clique-width breaks some of the desirable properties of clique-width over graphs. In particular, in **Graph**, bounded treewidth classes of graphs are also of bounded clique-width (but not conversely). However, under this naive definition, there exist classes of σ-structures which are of bounded treewidth, but unbounded clique-width. Courcelle in [28] and Blumensath in [22] provide alternative generalisations which retain the desirable properties of clique-width.

Since Proposition 5.2.9 established the coproduct $\uplus$ operation has a coKleisli law over $\mathbb{E}_k$, we have already obtained an FVM theorem for this operation of the grammar involving $\exists^+\mathbf{FO}_k$ and $\#\mathbf{FO}_k$. It remains to prove the FVM theorems for $\rho_{i\to j}$ and $\eta_{\vec{i}}^R$.

In order to use the coKleisli method to obtain an FVM theorem for $\rho_{i\to j}$ involving $\exists^+\mathbf{FO}_k$ and $\#\mathbf{FO}_k$, we formulate the operation $\rho_{i\to j}$ as functor. For every $i, j \in [m]$, define the functor $\rho_{i\to j} \colon \mathsf{Struct}(\sigma^m) \to \mathsf{Struct}(\sigma^m)$ where

- On objects, for $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma^m))$ with universe A, $\rho_{i\to j}(\mathcal{A})$ has the same universe $\mathcal{A}$, and relational interpretations defined as:

$$\begin{aligned} C_i^{\rho_{i\to j}(\mathcal{A})} &= \varnothing \\ C_j^{\rho_{i\to j}(\mathcal{A})} &= C_i^{\mathcal{A}} \cup C_j^{\mathcal{A}} \\ Q_j^{\rho_{i\to j}(\mathcal{A})} &= Q^{\mathcal{A}} \text{ for all } Q \in \sigma^m \backslash \{C_i, C_j\} \end{aligned}$$

- On morphisms, note that a morphism $f \colon \mathcal{A} \to \mathcal{B} \in \mathbf{Ob}(\mathsf{Struct}(\sigma^m))$ induces a morphism $\rho_{i\to j}(f) \in \mathbf{Mor}(\mathsf{Struct}(\sigma^m))$ with the same underlying set function.

For every $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma^m))$, we have a σ^m-morphism

$$\kappa_{\mathcal{A}} \colon \mathbb{E}_k(\rho_{i\to j}(\mathcal{A})) \to \rho_{i\to j}(\mathbb{E}_k(\mathcal{A}))$$

which as a set function is the identity on the universe $\mathbb{E}_k(A)$ of $\mathbb{E}_k(\mathcal{A})$. These family of maps specify a coKleisli law $\kappa \colon \mathbb{E}_k \circ \rho_{i\to j} \to \rho_{i\to j} \circ \mathbb{E}_k$

Proposition 5.3.5. $\kappa_{\mathcal{A}} \colon \mathbb{E}_k \circ \rho_{i\to j} \to \rho_{i\to j} \circ \mathbb{E}_k$ *is a coKleisli law.*

Proof. Since the underlying set function is the identity, verifying that κ is natural and satisfies diagrams (5.3) and (5.4) is trivial. All that remains to show is that for every $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma^m))$, $\kappa_{\mathcal{A}}$ is a σ^m-morphism. For clarity of notation, we use $\rho = \rho_{i\to j}$.

Suppose $Q \in \sigma^m$, then we must show that $Q^{\mathbb{E}_k\rho(\mathcal{A})} \subseteq Q^{\rho(\mathbb{E}_k\mathcal{A})}$. We split into three cases: For $Q = P_j$, we have that:

$$\begin{aligned} w \in C_j^{\mathbb{E}_k\rho(\mathcal{A})} &\Leftrightarrow \varepsilon_{\rho(\mathcal{A})}(w) \in C_j^{\rho(\mathcal{A})} \\ &\Leftrightarrow \varepsilon_{\rho(\mathcal{A})}(w) \in C_i^{\mathcal{A}} \cup C_j^{\mathcal{A}} \\ &\Leftrightarrow \varepsilon_{\mathcal{A}}(w) \in C_i^{\mathcal{A}} \cup C_j^{\mathcal{A}} \\ &\Leftrightarrow \varepsilon_{\mathcal{A}}(w) \in C_i^{\mathcal{A}} \text{ or } \varepsilon_{\mathcal{A}}(w) \in C_j^{\mathcal{A}} \\ &\Leftrightarrow w \in C_i^{\mathbb{E}_k\mathcal{A}} \text{ or } w \in C_j^{\mathbb{E}_k\mathcal{A}} \\ &\Leftrightarrow w \in C_i^{\mathbb{E}_k\mathcal{A}} \cup C_j^{\mathbb{E}_k\mathcal{A}} \\ &\Leftrightarrow w \in C_j^{\rho(\mathbb{E}_k\mathcal{A})} \end{aligned}$$

The other two cases, $Q = C_i$ or $Q \in \sigma^m \backslash \{C_i, C_j\}$, are easy to check. □

Thus as a corollary of this proposition and Theorem 3.2.19, we recover semantically the following FVM theorems about the recolouring operation $\rho_{i\to j}$:

Corollary. *For all $k > 0$, $i, j \in [m]$, and σ-structures $\mathcal{A}, \mathcal{B}$,*

- *if $\mathcal{A} \Rrightarrow^{\exists^+\mathcal{L}^k} \mathcal{B}$, then $\rho_{i\to j}(\mathcal{A}) \equiv^{\exists^+\mathbf{FO}_k} \rho_{i\to j}(\mathcal{B})$.*
- *if $\mathcal{A} \equiv^{\#\mathcal{L}^k} \mathcal{B}$, then $\rho_{i\to j}(\mathcal{A}) \equiv^{\#\mathbf{FO}_k} \rho_{i\to j}(\mathcal{B})$.*

We now proceed with proving FVM theorems for the R-creation operation and $\exists^+\mathbf{FO}$ and $\#\mathbf{FO}$. For every $R \in \sigma$ with arity r and $\vec{i} = (i_1, \ldots, i_r) \in [m]^r$, we can formulate the $\eta_{\vec{i}}^R$ operation as a functor $\eta_{\vec{i}}^R \colon \mathsf{Struct}(\sigma^m) \to \mathsf{Struct}(\sigma^m)$ where

- On objects, for $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma^m))$ with universe A, $\eta_{\vec{i}}^R(\mathcal{A})$ has the same universe $\mathcal{A}$, and relational interpretations defined as:

$$\begin{aligned} R^{\eta_{\vec{i}}^R(\mathcal{A})} &= R^{\mathcal{A}} \cup C_{i_1} \times \cdots \times C_{i_r} \\ Q^{\eta_{\vec{i}}^R(\mathcal{A})} &= Q^{\mathcal{A}} \text{ for all } Q \in \sigma^m \backslash \{R\} \end{aligned}$$

- On morphisms, note that a morphism $f \colon \mathcal{A} \to \mathcal{B} \in \mathbf{Ob}(\mathsf{Struct}(\sigma^m))$ induces a morphism $\eta_{\vec{i}}^R(f) \in \mathbf{Mor}(\mathsf{Struct}(\sigma^m))$ with the same underlying set function.

For every $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma^m))$, we have a σ^m-morphism

$$\kappa_{\mathcal{A}} \colon \mathbb{E}_k(\rho_{i\to j}(\mathcal{A})) \to \rho_{i\to j}(\mathbb{E}_k(\mathcal{A}))$$

which as a set function is the identity on the universe $\mathbb{E}_k(A)$ of $\mathbb{E}_k(\mathcal{A})$. These family of maps specify a coKleisli law $\kappa \colon \mathbb{E}_k \circ \rho_{i\to j} \to \rho_{i\to j} \circ \mathbb{E}_k$

Proposition 5.3.6. *$\kappa_{\mathcal{A}} \colon \mathbb{E}_k \circ \eta_{\vec{i}}^R \to \eta_{\vec{i}}^R \circ \mathbb{E}_k$ is a coKleisli law.*

Proof. Since the underlying set function is the identity, verifying that κ is natural and satisfies diagrams (5.3) and (5.4) is trivial. All that remains to show is that for every $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}(\sigma^m))$, $\kappa_{\mathcal{A}}$ is a σ^m-morphism. For clarity of notation, we use $\eta = \eta^R_{\vec{i}}$. Suppose $Q \in \sigma^m$, then we must show that $Q^{\mathbb{E}_k \eta(\mathcal{A})} \subseteq Q^{\eta(\mathbb{E}_k \mathcal{A})}$. We split into two cases:

For $Q = R$, suppose that $(w_1, \ldots, w_r) \in R^{\mathbb{E}_k \eta(\mathcal{A})}$, then the $w_1, \ldots, w_r$ are pairwise comparable and either $(\varepsilon_{\eta(\mathcal{A})}(w_1), \ldots, \varepsilon_{\eta(\mathcal{A})}(w_r)) \in R^{\mathcal{A}}$ or $\varepsilon_{\eta(\mathcal{A})}(w_j) \in P^{\mathcal{A}}_{i_j}$ for all $j \leq r$. If $(\varepsilon_{\eta(\mathcal{A})}(w_1), \ldots, \varepsilon_{\eta(\mathcal{A})}(w_r)) \in R^{\eta(\mathcal{A})}$, then by $\eta(\mathbb{E}_k(\mathcal{A}))$ and $\mathbb{E}_k(\mathcal{A})$ having the same universe and $\varepsilon_{\eta(\mathcal{A})} = \varepsilon_{\mathcal{A}}$, $(\varepsilon_{\mathcal{A}}(w_1), \ldots, \varepsilon_{\mathcal{A}}(w_r)) \in R^{\mathcal{A}}$. By the $w_1, \ldots, w_r$ being pairwise comparable, $(w_1, \ldots, w_r) \in R^{\mathbb{E}_k \mathcal{A}}$. If $\varepsilon_{\eta(\mathcal{A})}(w_j) \in P^{\mathcal{A}}_{i_j}$ for all $j \leq r$, then $\varepsilon_{\mathcal{A}}(w_j) \in P^{\mathcal{A}}_{i_j}$ and therefore $w_j \in P^{\mathbb{E}_k \mathcal{A}}_{i_j}$ for all $j \leq r$. Hence, $(w_1, \ldots, w_r) \in R^{\mathbb{E}_k \mathcal{A}}$ or $(w_1, \ldots, w_r) \in P^{\mathcal{A}}_{i_1} \times \cdots \times P^{\mathcal{A}}_{i_r}$. Therefore, by definition of $\eta^R_{\vec{i}}(\mathbb{E}_k \mathcal{A})$, $(w_1, \ldots, w_r) \in R^{\eta(\mathbb{E}_k \mathcal{A})}$ and $R^{\mathbb{E}_k \eta(\mathcal{A})} \subseteq R^{\eta(\mathbb{E}_k \mathcal{A})}$.

For $Q \in \sigma^m \backslash \{R\}$, the verification is trivial as $\eta^R_{\vec{i}}$ does not change the interpretations of these relations. ☐

Thus as a corollary of this proposition and Theorem 3.2.19, we recover semantically the following FVM theorems about the R-creation operations $\eta^R_{\vec{i}}$:

Corollary. *For all $k > 0$, $R \in \sigma$ of arity r, $\vec{i} \in [m]^r$ and σ-structures $\mathcal{A}, \mathcal{B}$,*

- *if $\mathcal{A} \Rrightarrow^{\exists^+ \mathcal{L}^k} \mathcal{B}$, then $\eta^R_{\vec{i}}(\mathcal{A}) \equiv^{\exists^+ \mathbf{FO}_k} \eta^R_{\vec{i}}(\mathcal{B})$.*
- *if $\mathcal{A} \equiv^{\# \mathcal{L}^k} \mathcal{B}$, then $\eta^R_{\vec{i}}(\mathcal{A}) \equiv^{\# \mathbf{FO}_k} \eta^R_{\vec{i}}(\mathcal{B})$.*

Chapter 6

Mixed distributive laws

In the previous chapter, we investigated coKleisli laws which were specified by families of morphisms of type $\kappa_X \colon \mathbb{T}(F(X)) \to F(\mathbb{S}(X))$ where F is a functorial operation and $\mathbb{S}, \mathbb{T}$ are (products of) game comonads arising from arboreal covers of $\mathsf{Struct}(\sigma)$ or $\mathsf{Struct}_\star(\sigma)$. We demonstrated how such laws reproduced Feferman-Vaught-Mostowski theorems in logic.

In this chapter, we investigate special case of these laws where $\mathbb{S} = \mathbb{T}$ and $F = \mathbf{M}$ is the underlying functor of a monad $(\mathbf{M}, \eta, \mu)$. Monads are dual to comonads and are ubiquitous in mathematics and theoretical computer science. In this case, the coKleisli law κ of $\mathbb{T}$ over $\mathbf{M}$ must also satisfy additional coherence axioms with the unit η and multiplication μ of M so that κ is also Kleisli law of $\mathbb{T}$ over $\mathbf{M}$.

The original motivation for this line of research arose from the search for a graded distributive law $\kappa^d \colon \mathbb{P}_k\mathbf{Q}_d \to \mathbf{Q}_d\mathbb{P}_k$ of the k-pebble game comonad $\mathbb{P}_k$ over the graded quantum monad $\mathbf{Q}_d$ on relational structures from [4]. Just as the existence of a coKleisli morphism $\mathbb{P}_k\mathcal{A} \to \mathcal{B}$ corresponds to the k-consistency approximation $\mathcal{A} \Rrightarrow^{\exists^+\mathcal{L}^k} \mathcal{B}$ to $\mathcal{A} \to \mathcal{B}$, a Kleisli morphism $\mathcal{A} \to \mathbf{Q}_d\mathcal{B}$ corresponds the ‘quantum homomorphism’ approximation to $\mathcal{A} \to \mathcal{B}$, originally defined in [65] via a non-local game. We could combine these two approximations to $\mathcal{A} \to \mathcal{B}$ by investigating morphisms of type $\mathbb{P}_k\mathcal{A} \to \mathbf{Q}_d\mathcal{B}$. A distributive law would give a mechanism for composing morphisms $\mathbb{P}_k\mathcal{A} \to \mathbf{Q}_d\mathcal{B}$ potentially illuminating a categorical setting for quantum finite model theory.

As a side effect of our initial search for the distributive law $\kappa^d \colon \mathbb{P}_k\mathbf{Q}_d \to \mathbf{Q}_d\mathbb{P}_k$, we ended up proving a no-go theorem demonstrating the non-existence of comonad-monad distributive laws for (co)monads sharing some similar features as $\mathbb{P}_k$ and $\mathbf{Q}_d$. Our journey to this no-go theorem was by first making a series of three simplifications to the search for a law $\kappa^d \colon \mathbb{P}_k\mathbf{Q}_d \to \mathbf{Q}_d\mathbb{P}_k$. The first step was to ignore pebbles and the grading of $\mathbb{P}_k$

simplifying the question to existence of a distributive law $\kappa\colon \mathbb{E}\mathbf{Q}_d \to \mathbf{Q}_d\mathbb{E}$. The second step was to try to define κ on universes–ignoring the role of relations–simplifying the question to the existence of a distributive law $\kappa\colon N^+\mathbf{Q}_d \to \mathbf{Q}_d N^+$ were N^+ is the prefix list comonad [72] and $\mathbf{Q}_d$ is the quantum monad on the empty signature. The final step was to view the quantum monad 'at the level of possibilities' simplifying the question to the existence of a distributive law $\kappa\colon N^+\mathcal{P} \to \mathcal{P}N^+$ where $\mathcal{P}$ is the powerset monad. It was this final step which led us to prove, using a Plotkin-style argument, that there is in fact no distributive law of type $N^+\mathcal{P} \to \mathcal{P}N^+$. Our final no-go theorem generalises this argument to a much wider class of (co)monads showing that there is no distributive law $\kappa\colon WM \to MW$ where the comonad W is in a large class of directed containers $\mathcal{C}_W$ and M is in a large class of monads equipped with a notion of 'uniform sampling' $\mathcal{C}_M$. Despite this no-go theorem, we will conclude that for structures in signatures σ which have a distinguished binary reflexive, symmetric 'commeasurability' relation, there is a graded distributive law $\kappa^d\colon \mathbb{E}_k^{\odot}\mathbf{Q}_d^{\odot} \to \mathbf{Q}_d^{\odot}\mathbb{E}_k^{\odot}$ where $\mathbb{E}_k^{\odot}, \mathbf{Q}_d^{\odot}$ are tailored versions of $\mathbb{E}_k, \mathbf{Q}_d$ that fit these commeasurable σ-structures.

6.1 Monads and mixed distributive laws

We recall the definition of monad, dual to the definition of comonad defined in Section 2.3.

Definition 6.1.1. The triple $(\mathbf{M}\colon \mathcal{E} \to \mathcal{E}, \eta, \mu)$ is a *monad (in standard form) over* $\mathcal{E}$ where $\mathbf{M}\colon \mathcal{E} \to \mathcal{E}$ is functor, $\eta\colon \mathsf{Id}_{\mathcal{E}} \to \mathbf{M}$ is the *unit* natural transformation, and $\mu\colon \mathbf{MM} \to \mathbf{M}$ satisfying the following equations for every $X \in \mathbf{Ob}(\mathcal{E})$

$$\mu_X \circ \mathbf{M}\mu_X = \mu_X \circ \mu_{\mathbf{M}X}; \quad \mu_X \circ \mathbf{M}\eta_X = \mu_X \circ \eta_{\mathbf{M}X} = \mathsf{id}_{\mathbf{X}}$$

Two common examples of monads over **Set** are the powerset and distribution monads.

Example 6.1.2. The powerset monad $(\mathcal{P}, \eta, \mu)$ has underlying functor $\mathcal{P}\colon \mathbf{Set} \to \mathbf{Set}$ which maps an object X to the set of subsets $\mathcal{P}(X)$ and maps a function $f\colon X \to Y$ to the direct image function $\mathcal{P}(f)\colon \mathcal{P}(X) \to \mathcal{P}(Y)$. The unit $\eta\colon \mathsf{Id}_{\mathbf{Set}} \to \mathcal{P}$ has components η_X defined as $\eta_X(x) = \{x\}$. The multiplication $\mu\colon \mathcal{P}\mathcal{P} \to \mathcal{P}$ has components μ_X defined as $\mu_X(U) = \bigcup U$. We can also define the finite powerset monad $(\mathcal{P}_f, \eta, \mu)$, non-empty powerset monad $(\mathcal{P}^+, \eta, \mu)$, and finite non-empty powerset monad $(\mathcal{P}_f^+, \eta, \mu)$. In each of these cases, the unit and multiplications are just appropriate restrictions of the unit and multiplication of the (full) powerset monad. In order to refer to these variants on the powerset monad uniformly, we will use $\wp$ to denote one of $\mathcal{P}, \mathcal{P}^+, \mathcal{P}_f^+, \mathcal{P}_f$.

Example 6.1.3. The discrete probability distribution monad $(\mathbf{D}_{\mathbb{R}^+}, \eta, \mu)$ has underlying functor $\mathbf{D}_{\mathbb{R}^+}\colon \mathbf{Set} \to \mathbf{Set}$ which maps a set X to the set of probability distributions $\varphi\colon X \to \mathbb{R}^+$, i.e. functions such that for all but finitely many $x \in X$, $\varphi(x) = 0$ and $\sum_{x\in X} \varphi(x) = 1$, and maps a function $f\colon X \to Y$ to the function $\mathbf{D}_{\mathbb{R}^+}(f)\colon \mathbf{D}_{\mathbb{R}^+}(X) \to \mathbf{D}_{\mathbb{R}^+}(Y)$ defined as

$$\mathbf{D}_{\mathbb{R}^+}(f)(\varphi) = \lambda y. \sum_{x\in f^{-1}(y)} \varphi(x).$$

The unit $\eta\colon \mathsf{Id}_{\mathbf{Set}} \to \mathbf{D}_{\mathbb{R}^+}$ has components η_X where $\eta_X(x)$ is the Dirac delta distribution at point $x \in X$. The multiplication $\mu\colon \mathbf{D}_{\mathbb{R}^+}\mathbf{D}_{\mathbb{R}^+} \to \mathbf{D}_{\mathbb{R}^+}$ has components μ_X defined as

$$\mu_X(\Psi) = \lambda x. \sum_{\psi\in \mathbf{D}_{\mathbb{R}^+}(X)} \Psi(\psi)\psi(x).$$

This monad is an instance of a wide class of distribution monads on **Set** that can be defined over any semiring $\mathcal{S}$. These semiring monads are defined and investigated in Section 6.2.3.

We conclude the background with the definition of comonad-monad mixed distributive law.

Definition 6.1.4. A *mixed distributive law of comonad* (W, ε, δ) *over monad* (M, η, μ) is a natural transformation $\kappa\colon WM \to MW$ satisfying the following diagrams:

$$\begin{tikzcd} & W \arrow[dl, "W\eta"'] \arrow[dr, "\eta W"] & \\ WM \arrow[rr, "\kappa"'] & & MW \end{tikzcd} \tag{6.1}$$

$$\begin{tikzcd} WM \arrow[rr, "\kappa"'] \arrow[dr, "\varepsilon M"'] & & MW \arrow[dl, "M\varepsilon"] \\ & M & \end{tikzcd} \tag{6.2}$$

$$\begin{tikzcd} WMM \arrow[r, "\kappa M"] \arrow[d, "W\mu"'] & MWM \arrow[r, "M\kappa"] & MMW \arrow[d, "\mu W"] \\ WM \arrow[rr, "\kappa"'] & & MW \end{tikzcd} \tag{6.3}$$

$$\begin{tikzcd} WM \arrow[rr, "\kappa"] \arrow[d, "\delta M"'] & & MW \arrow[d, "M\delta"] \\ WWM \arrow[r, "W\kappa"] & WMW \arrow[r, "\kappa W"] & MWW \end{tikzcd} \tag{6.4}$$

With a mixed distributive law $\kappa\colon WM \to MW$ for (co)monads W, M on category $\mathcal{E}$, we can define a *biKleisli composition* to compose morphisms of type $WA \to MB$. Given morphisms $f\colon WA \to MB \in \mathbf{Mor}(\mathcal{E})$ and $g\colon WB \to MC \in \mathbf{Mor}(\mathcal{E})$, we can produce a morphism

$$WA \xrightarrow{\delta} WWA \xrightarrow{W(f)} WMB \xrightarrow{\kappa} MWB \xrightarrow{M(g)} MMC \xrightarrow{\mu} MC \tag{6.5}$$

of type $WA \to MC \in \mathbf{Mor}(\mathcal{E})$.

6.2 No-go theorems

In this section, we build up to proving a general no-go theorem demonstrating that there is no mixed distributive law of the form in Definition 6.1.4 for a wide class of comonads $\mathcal{C}_W$ and monads $\mathcal{C}_M$ over **Set**.

6.2.1 Prefix list comonad over powerset monad

We warm-up to the general no-go theorem by demonstrating that there is no natural transformation of type $\kappa\colon N^+\wp \to \wp N^+$ which satisfies diagram (6.1) and diagram (6.4). Thus, there is no comonad-monad distributive law of the prefix list comonad over the powerset monad. This work was motivated by the intuition that the unbounded EF comonad $\mathbb{E}$ on the empty signature is N^+ and that the quantum monad 'at the level of possibilities' is the finite non-empty powerset monad $\mathcal{P}_f^+$.

Theorem 6.2.1. *If $(\wp, \eta)$ is the powerset endofunctor, and $\kappa\colon N^+\wp \to \wp N^+$ is pointed law, then for all $l = [Y_1, \ldots, Y_n] \in N^+\wp^+(X) \subseteq N^+\wp(X)$,*

$$\kappa([Y_1, \ldots, Y_n]) = \{[y_1, \ldots, y_n] \mid \forall i \in [n], y_i \in Y_i\}. \tag{6.6}$$

The proof of this theorem involves using as three step 'Plotkin'-style argument where we chase naturality squares. For the first two steps, we restrict to proving Equation (6.6) holds for sequences $l = [Y_1, \ldots, Y_n] \in N^+\wp^+(X)$ which satisfy the following pairwise disjoint condition:

(PD) For all $i, j \in [n]$, if $Y_i \cap Y_j \neq \varnothing$, then $i = j$.

The first step in the proof of Theorem 6.2.1 is to prove the $\subseteq$ inclusion in Equation (6.6) using the unit axiom and chasing the naturality square for a 'collapse' function c_l for every $l \in N^+(\wp^+(X))$:

Lemma 6.2.2 (Collapse)**.** *If κ satisfies the hypotheses of Theorem 6.2.1 and $l = [Y_1, \ldots, Y_n] \in N^+\wp^+(X) \subseteq N^+\wp(X)$ satisfies condition (PD), then*

$$\kappa([Y_1, \ldots, Y_n]) \subseteq \{[y_1, \ldots, y_n] \mid \forall i \in [n], y_i \in Y_i\}.$$

Proof. The argument proceeds in through two cases.

(1) Unpacking the unit axiom, we show that equation (6.6) holds if for all $i \in [n]$, $Y_i = \{y_i\}$. Consider the singleton $\{[y_1, \dots, y_n]\} \in \wp^+(N^+(X))$

$$\begin{aligned}
\{[y_1, \dots, y_n]\} &= \eta_{N^+(X)}([y_1, \dots, y_n]) && \text{definition of } \eta \\
&= \kappa_X \circ N^+(\eta_X)([y_1, \dots, y_n]) && \text{unit axiom} \\
&= \kappa_X([\eta_X(y_1), \dots, \eta_X(y_n)]) && N^+ \text{ on morphisms} \\
&= \kappa_X([\{y_1\}, \dots, \{y_n\}]) && \text{definition of } \eta \\
&= \kappa_X([Y_1, \dots, Y_n]) && \text{case supposition}
\end{aligned}$$

(2) For all $l = [Y_1, \dots, Y_n] \in N^+(\wp^+(X))$ and $i \in [n]$, since $Y_i \neq \varnothing$, there exists a $y_i \in Y_i$. We construct a 'collapse the Y_i' function $c_l \colon X \to X$. The definition of $c_l \colon X \to X$ is

$$c_l(x) = \begin{cases} y_i & \text{if } x \in Y_i \\ x & \text{otherwise} \end{cases}.$$

By the (PD) assumption, the function c is well-defined. Observe that $\wp^+(c_l)[Y_1, \dots, Y_n] = [\{y_1\}, \dots, \{y_n\}]$. Chasing l along the κ-naturality square of c_l, we obtain

$$\begin{tikzcd}
{[Y_1, \dots, Y_n]} \arrow[r, mapsto, "\kappa_X"] \arrow[d, mapsto, "N^+(\wp(c_l))"'] & \kappa_X([Y_1, \dots, Y_n]) \arrow[d, mapsto, "\wp(N^+(c_l))"] \\
{[\{y_1\}, \dots, \{y_n\}]} \arrow[r, mapsto, "\kappa_X"] & \{[y_1, \dots, y_n]\}
\end{tikzcd}$$

where the bottom arrow follows from the first step. Since $\wp^+(N^+(c_l))$ maps the set $\kappa_X(l)$ to the singleton $\{[y_1, \dots, y_n]\}$, we can make two observations. Observe that for every $[y'_1, \dots, y'_n] \in \kappa_X(l)$, $N^+(c_l)([y'_1, \dots, y'_n]) = [y_1, \dots, y_n]$. From the definition of $\wp^+(N^+(c_l))$, we conclude that

$$\kappa_X(l) \subseteq \{[y'_1, \dots, y'_n] \mid \forall i \in [n], y'_i \in c^{-1}(y_i) = Y_i\}.$$

□

The second step in the proof of Theorem 6.2.1 is to prove the $\supseteq$ inclusion in Equation (6.6) using Lemma 6.2.2 and chasing the naturality square for a 'swap' bijection b_l for every $l \in N^+(\wp^+(X))$.

Lemma 6.2.3 (Swap). *If κ satisfies the hypotheses of Theorem 6.2.1 and $l = [Y_1, \dots, Y_n] \in N^+\wp^+(X) \subseteq N^+\wp(X)$ satisfies condition (PD), then*

$$\kappa([Y_1, \dots, Y_n]) \supseteq \{[y_1, \dots, y_n] \mid \forall i \in [n], y_i \in Y_i\}.$$

Proof. For all $l = [Y_1, \ldots, Y_n] \in N^+(\wp^+(X))$ satisfying condition (PD), we construct a 'swap' bijection $b_l \colon X \to X$. Suppose $[y_1, \ldots, y_n] \in N^+(X)$ such that $y_i \in Y_i$. We need to show that $[y_1, \ldots, y_n] \in \kappa_X(l)$. By the previous case and the fact that $\kappa_X(l) \in \wp^+(N^+(X))$ is a non-empty subset of $N^+(X)$, we know that there exists at least one $[y'_1, \ldots, y'_n] \in \kappa_X(l)$ such that for all $i \in [n]$, $y'_i \in Y_i$. For every $i \in [n]$, let $\pi_i \colon X \to X$ be the permutation which swaps $y'_i \in Y_i$ and $y_i \in Y_i$ while fixing every other element of X. As π_i fixes the set Y_i and, by the (PD) condition, leaves Y_j unchanged for all $j \neq i \in [n]$, we can conclude that $l = N^+(\wp^+(\pi_i))(l)$. Hence, by the κ-naturality square of π_i, we have that $\wp^+(N^+(\pi_i))(\kappa_X(l)) = \kappa_X(l)$. Therefore, we can compose all these swapping bijections (in any order) $\{\pi_i\}_{i \in [n]}$ to obtain a bijection $b_l \colon X \to X$ such that $\wp^+(N^+(b))(\kappa_X(l)) = \kappa_X(l)$. This means that:

$$\begin{aligned}
[y'_1, \ldots, y'_n] \in \kappa_X(l) &\Rightarrow N^+(b)[y'_1, \ldots, y'_n] \in \kappa_X(l) \\
&\Rightarrow [b(y'_1), \ldots, b(y'_n)] \in \kappa_X(l) \\
&\Rightarrow [\pi_1(y'_1), \ldots, \pi_n(y'_n)] \in \kappa_X(l) \\
&\Rightarrow [y_1, \ldots, y_n] \in \kappa_X(l)
\end{aligned}$$

Since $[y'_1, \ldots, y'_n] \in \kappa_X(l)$, we can conclude that $[y_1, \ldots, y_n] \in \kappa_X(l)$. □

In the final step for the proof of Theorem 6.2.1, we show that the pairwise-disjoint condition (PD) does not constitute a loss of generality.

Proof of Theorem 6.2.1. Consider $l = [Y_1, \ldots, Y_n] \in N^+(\wp^+(X))$ (the Y_i are not necessarily pairwise-disjoint), we first consider the set $Z = [n] \times X$. There is a list $l' = [Y'_1, \ldots, Y'_n] \in N^+(\wp(Z))$ where $Y'_i = \{(i, y_i) \mid \forall y_i \in Y_i\} \in \wp(Z)$ satisfying (PD). By construction, $N^+(\wp^+(t))(l') = l$ where $t \colon Z \to X$ is the projection onto the second component. Since l' satisfies (PD), we can use Lemma 6.2.2, Lemma 6.2.3 and the naturality square of t to compute that

$$\begin{aligned}
\kappa_X(l) &= \kappa_X(N^+\wp^+(t)(l')) \\
&= \wp^+(N^+(t)(\kappa_Y(l'))) \\
&= \wp^+(N^+(t)(\{[(1, y_1), \ldots, (n, y_n)] \mid \forall i \in [n], y_i \in Y_i\})) \\
&= \{[t(1, y_1), \ldots, t(n, y_n)] \mid \forall i \in [n], y_i \in Y_i\} \\
&= \{[y_1, \ldots, y_n] \mid \forall i \in [n], y_i \in Y_i\}.
\end{aligned}$$

□

Since every distributive law $\kappa\colon N^+\wp \to \wp N^+$ must satisfy the unit axiom (6.1), then by Theorem 6.2.1, for all lists with non-empty sets, κ must satisfy Equation (6.6). However, using this definition of κ and chasing a simple element in $N^+\wp^+(\{a,b\})$ reveals that it cannot satisfy the comultiplication diagram (6.4). Therefore, there can be no distributive law of the prefix list comonad N^+ over the powerset pointed endofunctor $(\wp, \eta)$.

Theorem 6.2.4. *There is no distributive law $\kappa\colon N^+\wp \to \wp N^+$ of $(N^+, \varepsilon, \delta)$ over $(\wp, \eta, \mu)$*

Proof. Suppose for contradiction there exists a distributive law $\kappa\colon N^+\wp \to \wp N^+$. As κ must satisfy the unit axiom (6.1), κ is a pointed law of N^+ over $(\wp, \eta)$. Theorem 6.2.1 implies that for lists which contain only non-empty sets, the components of κ satisfy Equation (6.6). From equation (6.6), the definition of δ, and considering the list $l = [\{a,b\},\{b\}] \in N^+(\wp(X))$ for $X = \{a,b\}$, we obtain the following inequality contradicting the comultiplication axiom 6.4:

$$\begin{aligned}
\wp\delta_X \circ \kappa_X([\{a,b\},\{b\}]) &= \wp\delta_X(\{[a,b],[b,b]\}) \\
&= \{[[a],[a,b]],[[b],[b,b]]\} \\
&\neq \{[[a],[a,b]],[[b],[a,b]],[[a],[b,b]],[[b],[b,b]]\} \\
&= \kappa_{N^+X}([\{[a],[b]\},\{[a,b],[b,b]\}]) \\
&= \kappa_{N^+X} \circ N^+\kappa_X([[\{a,b\}],[\{a,b\},\{b\}]]) \\
&= \kappa_{N^+X} \circ N^+\kappa_X \circ \delta_{\wp X}([\{a,b\},\{b\}])
\end{aligned}$$

□

Interestingly, for the list $l' = [\{b\},\{a,b\}]$, $\wp(\delta_X) \circ \kappa(l') = \kappa_{N^+X} \circ N^+\kappa_X \circ \delta_{\wp X}(l')$. The last element is the returned by ε since that is the root position of N^+. The fact that the set $\{a,b\}$ is in the root position of l' whereas $\{a,b\}$ is not in the root position of l leads to the difference in outcomes when chasing diagram (6.4). Indeed, we can generalise this counterexample to directed containers since these are comonads that are equipped with a notion of root position and subshape. This use of a two element set in the non-root position is the key to generalising Theorem 6.2.1 to a class of directed containers.

6.2.2 Directed containers over uniform choice

We generalise the argument from Section 6.2.1 from the prefix list comonad N^+ to non-linear directed containers W and from the powerset monad to any monad (in fact, any pointed endofunctor) M which has a meaningful notion of 'uniform sampling'. Directed containers are class of comonads of **Set** whose underlying endofunctor specified by set of shapes and positions within those shapes. This is the notion of container over **Set** [1] as in the following definition.

Definition 6.2.5. A *container* $S \lhd P$, is a set S and a functor $P \colon S \to \mathbf{Set}$ where we consider S as a discrete category, i.e. P defines an S-indexed family of sets. We consider S to be a set of shapes, and for each $s \in S$, $P(s)$ is the set of positions associated with a shape $s \in S$. The *induced endofunctor on* $S \lhd P$ is $[S \lhd P] \colon \mathbf{Set} \to \mathbf{Set}$ where

- for a set X, $[S \lhd P]X = \{(s, l) \mid s \in S, l \colon P(s) \to X\}$; and
- for a function $g \colon X \to Y$, $[S \lhd P]g \colon [S \lhd P]X \to [S \lhd P]Y$ is defined as $(s, l) \mapsto (s, g \circ l)$.

We abuse terminology and say an endofunctor $F \colon \mathbf{Set} \to \mathbf{Set}$ is a container if $F \cong [S \lhd P]$ for some set S and functor $P \colon S \to \mathbf{Set}$. Thus, a container $F \cong [S \lhd P]$ is exactly a polynomial endofunctor $F \colon \mathbf{Set} \to \mathbf{Set}$ given by the expression $F = \sum_{s \in S} y^{P(s)}$.

Ahman, Chapman, and Uustalu in [16], showed that comonads whose underlying endofunctor is a container are equivalent to containers which are equipped with additional 'directed' structure. Intuitively, directed containers are containers where each position has an associated subshape and every such subshape has a root.

Definition 6.2.6. A *directed container* $D = (F, \mathsf{o}, \downarrow, \oplus)$ consists of

- a container $F = [S \lhd P]$;
- for each shape $s \in S$, a root position $\mathsf{o}_s \in P(s)$;
- for each position $p \in P(s)$, a subshape $s \downarrow p \in S$;
- for each position $p' \in P(s \downarrow p)$ in the subshape $s \downarrow p$, a translation into a position $p \oplus_s p' \in P(s)$ in the global shape $s \in S$;

satisfying the equations

$$s \downarrow \mathsf{o}_s = s \quad s \downarrow (p \oplus_s p') = (s \downarrow p) \downarrow p' \tag{6.7}$$

$$p \oplus_s \mathsf{o}_{s \downarrow p} = p = \mathsf{o}_s \oplus_s p \tag{6.8}$$

$$(p \oplus_s p') \oplus_s p'' = p \oplus_s (p' \oplus_{s \downarrow p} p''). \tag{6.9}$$

A directed container $D = ([S \lhd P], \mathsf{o}, \downarrow, \oplus)$, has an *induced comonad* $([S \lhd P], \varepsilon, \delta)$ *of* D where:

- The counit has components defined as

$$\varepsilon_X(s, l \colon P(s) \to X) = l(\mathsf{o}_s)$$

- The comultiplication has components defined as

$$\delta_X(s, l\colon P(s) \to X) = (s, \lambda p.(s \downarrow p, \lambda q.l(p \oplus_s q)))$$

We will abuse terminology and say a comonad (W, ε, δ) is a directed container if it is isomorphic to the induced comonad of a directed container.

In order to carve out this class of monads which have a meaningful notion of 'uniform sampling', we take inspiration from the algebraic presentations of finitary monads which arise from universal algebra. In particular, this perspective formalises a term as a natural transformation into the free monad M associated to an algebraic theory. Thus, we define a *n-ary open term for an endofunctor* $M\colon \mathcal{C} \to \mathcal{C}$ on a category $\mathcal{C}$ with finite products, as a natural transformation $\beta\colon \mathsf{Id}_{\mathcal{C}} \times \cdots \times \mathsf{Id}_{\mathcal{C}} \to M$ where the domain of β is endofunctor sending an object X to its n-th Cartesian power X^n. Beyond the algebraic portion of this definition, we will also need to restrict to monads which have a meaningful notion of support. Formally, working within the category of **Set**, we say that a pair (M, supp) is a *supported endofunctor* if $M\colon \mathbf{Set} \to \mathbf{Set}$ is an endofunctor and $\mathsf{supp}\colon M \to \mathcal{P}$ is a natural transformation from M to the covariant powerset functor. With these notions in place, we arrive at our formalisation of a pointed endofunctor which has 'uniform sampling'.

Definition 6.2.7. Given a supported pointed endofunctor (M, η, supp), a n-ary open term $\beta\colon \mathsf{Id}_{\mathcal{C}} \times \cdots \times \mathsf{Id}_{\mathcal{C}} \to M$ for M is a *n-uniform choice term* if

(1) *β is idempotent:* For all $X \in \mathbf{Ob}(\mathbf{Set})$ and $x \in X$,

$$\beta(x, \ldots, x) = \eta(x);$$

(2) *β is commutative:* For all $X \in \mathbf{Ob}(\mathbf{Set})$, $x_1, \ldots, x_n \in X$, and permutations $\pi\colon [n] \to [n]$,

$$\beta(x_1, \ldots, x_n) = \beta(x_{\pi(1)}, \ldots, x_{\pi(n)});$$

(3) supp *preserves β:* For all $X \in \mathbf{Ob}(\mathbf{Set})$ and $x_1, \ldots, x_n \in X$,

$$\mathsf{supp}(\beta(x_1, \ldots, x_n)) = \{x_1, \ldots, x_n\}.$$

We will say M is a *n-uniform choice pointed endofunctor* if there exists a natural transformation $\mathsf{supp}\colon M \to \mathcal{P}$ and n-uniform choice term β for the supported pointed endofunctor (M, η, supp). We will say a monad (M, η, μ) is a *n-uniform choice monad* if (M, η) is *n-uniform choice pointed endofunctor.* Since the powerset supported endofunctor $(\mathcal{P}, \eta^{\mathcal{P}}, \mathsf{id}_{\mathcal{P}})$ itself has n-uniform choice term $\beta^{\mathcal{P}}(x_1, \ldots, x_n) = \{x_1, \ldots, x_n\}$ for every

$n > 0$, the terminology for condition (3) in Definition 6.2.7 is justified. In fact, for every n-uniform choice endofunctor M, it follows that $\mathsf{supp}\colon M \to \mathcal{P}$ is a pointed endofunctor morphism:

$$\mathsf{supp}(\eta(x)) = \mathsf{supp}(\beta(x, \ldots, x)) = \{x\} = \eta^{\mathcal{P}}(x). \tag{6.10}$$

Moreover, we have a partial converse: every supported pointed endofunctor (M, η, supp) such that $\mathsf{supp}\colon M \to \mathcal{P}$ is pointed endofunctor morphism is a 1-uniform choice endofunctor where $\beta = \eta$.

For a n-uniform choice pointed endofunctor (M, η) with n-uniform choice term β and $X \in \mathbf{Ob}(\mathbf{Set})$, we define the set of terms as

$$U(X) := \{\beta(x_1, \ldots, x_n) \mid x_1, \ldots, x_n \in X\} \subseteq M(X).$$

To illustrate, take M to be the discrete probability distribution monad and define β as

$$\beta(x_1, \ldots, x_n) = \sum_{i \in [n]} \frac{1}{n} x_i.$$

In this case, the set $U(X)$ contains uniformly samplings from multisets of X. That is, $\beta(x_1, \ldots, x_n)$ is the probability distribution resulting from uniformly sampling from the multiset of the input list $[x_1, \ldots, x_n]$. In particular, if each of the $x_1, \ldots, x_n$ are distinct, then $\beta(x_1, \ldots, x_n)$ is the uniform distribution on the subset $\{x_1, \ldots, x_n\}$ of X of cardinality n.

Theorem 6.2.8. *If $F = [S \triangleleft P]$ is a container, (M, η) is a n-uniform choice endofunctor with n-uniform choice term β, and $\rho\colon FM \to MF$ is pointed law, then for all $(s, l) \in FU(X) \subseteq FM(X)$,*

$$\mathsf{supp}(\rho(s, l)) = \{(s, j) \mid \forall p \in P(s), j(p) \in \mathsf{supp}(l(p))\}. \tag{6.11}$$

The proof mirrors collapse-swap-relabel argument of the uniqueness proof of Theorem 6.2.1. Thus, we prove analogues of Lemmas 6.2.2 and 6.2.3. Condition (1) that β is idempotent is key to demonstrating the 'collapse' argument in Lemma 6.2.9. Condition (2) that β is commutative is key to demonstrating that the 'swap' argument in Lemma 6.2.10.

Similar to Theorem 6.2.1, we first consider elements $(s, l) \in FM(X)$ where the following pairwise disjoint condition holds:

(PDS) For all $p, q \in P(s)$, if $\mathsf{supp}(l(p)) \cap \mathsf{supp}(l(q)) \neq \varnothing$, then $p = q$.

Lemma 6.2.9 (Collapse)**.** *If F, M, and ρ satisfy the hypotheses of Theorem 6.2.8 and $(s, l) \in FU(X) \subseteq FM(X)$ satisfies condition (PDS), then*

$$\mathsf{supp}(\rho(s, l)) \subseteq \{(s, j) \mid \forall p \in P(s) \in \mathsf{supp}(l(p))\}.$$

Proof. The argument proceeds in two steps.

(1) Unpacking the unit axiom, we will show that equation (6.11) holds if (s, l) is such that for all $p \in P(s)$, there exists a $x_p \in X$, $l(p) = \eta(x_p)$. There is function $j = \lambda p.x_p \colon P(s) \to X$. Consider the singleton $\{(s, \lambda p.x_p)\} \in \mathcal{P}(F(X))$. For clarity of notation, we use η for the unit of M and $\eta^{\mathcal{P}}$ for the unit of $\mathcal{P}$.

$$\begin{aligned}
\{(s, \lambda p.x_p)\} &= \eta^{\mathcal{P}}(s, \lambda p.x_p) && \text{definition of } \eta^{\mathcal{P}} \\
&= \mathsf{supp}(\eta(s, \lambda p.x_p)) && \text{equation (6.10)} \\
&= \mathsf{supp}(\rho(F(\eta)(s, \lambda p.x_p))) && \text{unit axiom} \\
&= \mathsf{supp}(\rho(s, \lambda p.\eta(x_p))) && F \text{ functorial action} \\
&= \mathsf{supp}(\rho(s, l)) && \text{definition of } l
\end{aligned}$$

From the above equation and $\{(s, \lambda p.x_p)\} = \{(s, j) \mid \forall p \in P(s), j(p) \in \mathsf{supp}(l(p))\}$, we obtain that the desired result.

(2) For all $(s, l) \in F(U(X))$ and $p \in P(s)$, by the definition of $U(X)$, $l(p) = \beta(x_1^p, \ldots, x_n^p)$ for some $x_1^p, \ldots, x_n^p \in X$, and $\mathsf{supp}(l(p)) = \{x_1^p, \ldots, x_n^p\} \neq \varnothing$. Hence, for every $p \in P(s)$, we can choose a $y^p = x_1^p \in \mathsf{supp}(l(p))$ in order to construct a 'collapse the $\mathsf{supp}(l(p))$' function $c \colon X \to X$. The definition of $c \colon X \to X$ is

$$c(x) = \begin{cases} y^p & \text{if } x \in \mathsf{supp}(l(p)) \\ x & \text{otherwise} \end{cases}.$$

By the (PDS) assumption, the function c is well-defined. Next, we prove that $M(c) \circ l = \lambda p.\eta(y^p)$. Suppose $p \in P(s)$:

$$\begin{aligned}
M(c) \circ l(p) &= M(c)(\beta(x_1^p, \ldots, x_n^p)) && l(p) \in U(X) \\
&= \beta(c(x_1^p), \ldots, c(x_n^p)) && \beta\text{-naturality square of } c \\
&= \beta(y^p, \ldots, y^p) && \text{definition of } c \\
&= \eta(y^p) && \beta \text{ idempotent}
\end{aligned}$$

Chasing (s, l) along the ρ-naturality square of c composed with the supp-naturality square of $F(c)$, we obtain

$$\begin{tikzcd}
(s, l) \arrow[r, mapsto, "\rho"] \arrow[d, mapsto, "F(M(c))"'] & \rho(s, l) \arrow[r, mapsto, "\mathsf{supp}"] \arrow[d, mapsto, "M(F(c))"] & \mathsf{supp}(\rho(s, l)) \arrow[d, mapsto, "\wp(F(c))"] \\
(s, \lambda p.\eta(y^p)) \arrow[r, mapsto, "\rho"] & \eta(s, \lambda p.y^p) \arrow[r, mapsto, "\mathsf{supp}"] & \{(s, \lambda p.y^p)\}
\end{tikzcd}$$

where the bottom left arrow follows from the first step. Since $\wp(F(c))$ maps the set $\mathsf{supp}(\rho(s, l))$ to the singleton $\{(s, \lambda p.y_p)\}$, we can conclude that for every $(t, j) \in$

$\mathsf{supp}(\rho(s,l))$, $F(c)(t,j) = (s, \lambda p.y^p)$. Since the function $F(c)$ preserves shape, for every $(t,j) \in \mathsf{supp}(\rho(s,l))$, $t = s$. Moreover, from the definition of $\wp(F(c))$, we conclude that

$$\mathsf{supp}(\rho(s,l)) \subseteq \{(s,j) \mid \forall p \in P(s), j(p) \in c^{-1}(y_p) = \mathsf{supp}(l(p))\}.$$

$\square$

Lemma 6.2.10 (Swap)**.** *If F, M, and ρ satisfy the hypotheses of Theorem 6.2.8 and $(s,l) \in FU(X) \subseteq FM(X)$ satisfies condition (PDS), then*

$$\mathsf{supp}(\rho(s,l)) \supseteq \{(s,j) \mid \forall p \in P(s) \in \mathsf{supp}(l(p))\}.$$

Proof. For all $(s,l) \in F(U(X))$ satisfying condition (PDS), we construct a 'swap' bijection $b\colon X \to X$. By the assumption that $(s,l) \in F(U(X))$, for all $p \in P(s)$, there exists $x_1^p, \dots, x_n^p \in X$ such that $l(p) = \beta(x_1^p, \dots, x_n^p)$ and $\mathsf{supp}(l(p)) = \{x_1^p, \dots, x_n^p\}$. Suppose $(s,j) \in F(X)$ such that $j(p) = x_z^p \in \mathsf{supp}(l(p))$, we need to show that $(s,j) \in \mathsf{supp}(\rho(s,l))$. By the previous case and the fact that $\mathsf{supp}(\rho(s,l)) \in \wp(F(X))$ is a non-empty subset of $F(X)$, we know there exists at least one $(s,j') \in \mathsf{supp}(\rho(s,l))$ such that $j'(p) = x_{z'}^p \in \mathsf{supp}(l(p))$. For every $p \in P(s)$, let $s_p\colon X \to X$ be the permutation which swaps $x_z^p, x_{z'}^p \in \mathsf{supp}(l(p))$ and $(z\ z')\colon [n] \to [n]$ be the permutation which swaps $z, z' \in [n]$. By construction, for all $j \in [n]$, $s_p(x_j^p) = x_{(z\ z')(j)}^p$. Next, we prove that $M(s_p) \circ l = l$. Suppose $p \in P(s)$:

$$\begin{aligned}
M(s_p) \circ l(p) &= M(s_p)(\beta(x_1^p, \dots, x_n^p)) && l(p) \in U(X) \\
&= \beta(s_p(x_1^p), \dots, s_p(x_n^p)) && \beta\text{-naturality square of } s_p \\
&= \beta(x_{(z\ z')(1)}^p, \dots, x_{(z\ z')(n)}^p) && \text{definition of } s_p \\
&= \beta(x_1^p, \dots, x_n^p) && \beta \text{ commutative} \\
&= l(p) && l(p) \in U(X)
\end{aligned}$$

Thus, the function s_p fixes $l(p)$ and by the (PDS) condition leaves $l(q)$ unchanged for all $q \neq p$, so $F(M(s_p))(s,l) = (s,l)$. Hence, we can conclude from the ρ-naturality square of s_p:

$$\begin{aligned}
\rho(F(M(s_p))(s,l)) &= M(F(s_p))(\rho(s,l)) \\
\rho(s,l) &= M(F(s_p))(\rho(s,l))
\end{aligned}$$

Therefore, we compose all these swapping bijections $\{s_p\}_{p\in P(s)}$ (in any order) to obtain a bijection $b\colon X \to X$ such that $M(F(b))(\rho(s,l)) = \rho(s,l)$ and $b \circ j' = j$. By the supp-naturality square of $F(b)$, we obtain that:

$$\begin{aligned}\mathsf{supp}(\rho(s,l)) &= \mathsf{supp}(M(F(b))(\rho(s,l))) \\ &= \wp(F(b))(\mathsf{supp}(\rho(s,l)))\end{aligned}$$

Unpacking the definition of $\wp(F(b))$ and using the above equality, we obtain that

$$\begin{aligned}(s,j') \in \mathsf{supp}(\rho(s,l)) &\Rightarrow F(b)(s,j') \in \mathsf{supp}(\rho(s,l)) \\ &\Rightarrow (s, b\circ j') \in \mathsf{supp}(\rho(s,l)) \\ &\Rightarrow (s,j) \in \mathsf{supp}(\rho(s,l))\end{aligned}$$

Since $(s,j') \in \mathsf{supp}(\rho(s,l))$, we can conclude that $(s,j) \in \mathsf{supp}(\rho(s,l))$ as desired. □

Finally, we complete the proof of Theorem 6.2.8 by repeating a version of the 'relabel' argument of Theorem 6.2.1 which demonstrates that assuming (PDS) does not constitute a loss of generality.

Proof of Theorem 6.2.8. Suppose $(s,l) \in FU(X)$ where for all $p \in P(s)$, there exists $x^p_1,\dots,x^p_n \in X$ such that $l(p) = \beta(x^p_1,\dots,x^p_n)$. We factor $l = M(t)\circ z$ where $z\colon P(s) \to M(P(s)\times X)$ is defined as

$$z(p) = \beta((p,x^p_1),\dots,(p,x^p_n))$$

and $t\colon P(s)\times X \to X$ is the projection onto the second component. By construction, $(s,z) \in FU(P(s)\times X)$ satisfies condition (PDS), so applying Lemmas 6.2.9 and 6.2.10, we obtain the equation

$$\mathsf{supp}(\rho(s,z)) = \{(s,m) \mid \forall p \in P(s), m(p) \in \mathsf{supp}(z(p))\} \tag{6.12}$$

Through diagram chasing, we obtain the desired equality:

$$\begin{aligned}&\mathsf{supp}(\rho(s,l)) \\ =\ &\{\ l = M(t)\circ z\ \} \\ &\mathsf{supp}(\rho(s, M(t)\circ z)) \\ =\ &\{\ \text{definition of } F\ \} \\ &\mathsf{supp}(\rho(F(M(t))(s,z))) \\ =\ &\{\ \rho\text{-naturality square of } t\ \} \\ &\mathsf{supp}(M(F(t))(\rho(s,z)))\end{aligned}$$

$$
\begin{aligned}
&= \quad \{\ \mathsf{supp}\text{-naturality square of } F(t)\ \} \\
&\quad \wp(F(t))(\mathsf{supp}(\rho(s,z))) \\
&= \quad \{\ \text{equation (6.12)}\ \} \\
&\quad \wp(F(t))(\{(s,m) \mid \forall p \in P(s), m(p) \in \mathsf{supp}(z(p))\}) \\
&= \quad \{\ \text{definition of } \wp\ \} \\
&\quad \{F(t)(s,m) \mid \forall p \in P(s), m(p) \in \mathsf{supp}(z(p))\} \\
&= \quad \{\ \text{definition of } F\ \} \\
&\quad \{(s, t \circ m) \mid \forall p \in P(s), m(p) \in \mathsf{supp}(z(p))\} \\
&= \quad \{\ z(p) = \beta((p, x_1^p), \ldots, (p, x_n^p))\ \} \\
&\quad \{(s, t \circ m) \mid \forall p \in P(s), m(p) \in \{(p, x_1^p), \ldots, (p, x_n^p)\}\} \\
&= \quad \{\ j = t \circ m\ \} \\
&\quad \{(s, j) \mid \forall p \in P(s), j(p) \in \{t(p, x_1^p), \ldots, t(p, x_n^p)\}\} \\
&= \quad \{\ \text{definition of } t\ \} \\
&\quad \{(s, j) \mid \forall p \in P(s), j(p) \in \{x_1^p, \ldots, x_n^p\}\} \\
&= \quad \{\ \mathsf{supp} \text{ preserves } \beta\ \} \\
&\quad \{(s, j) \mid \forall p \in P(s), j(p) \in \mathsf{supp}(\beta(x_1^p, \ldots, x_n^p))\} \\
&= \quad \{\ \text{definition of } l(p)\ \} \\
&\quad \{(s, j) \mid \forall p \in P(s), j(p) \in \mathsf{supp}(l(p))\}
\end{aligned}
$$

□

The equation (6.11) we proved in Theorem 6.2.8 allows us to compute via a simple counting argument the cardinality of the support of $\rho(s,l)$ for $(s,l) \in FU(X)$.

Corollary 6.2.11. *If $F = [S \lhd P]$ is a container, and $\rho\colon FM \to MF$, then for all $(s,l) \in FU(X) \subseteq FM(X)$*

$$|\mathsf{supp}(\rho(s,l))| = \prod_{p \in P(s)} |\mathsf{supp}(l(p))| \tag{6.13}$$

Proof. Equation (6.11) states that $\mathsf{supp}(\rho(s,l))$ is equal to the set of all pairs (s,j) where for each position $p \in P(s)$, $j(p)$ is a sample from the set $\mathsf{supp}(l(p))$. For each position $p \in P(s)$, sample an element from $\mathsf{supp}(l(p))$. Each sampling is independent, thus we obtain the desired product. □

We use $\mathcal{C}_M$ to denote the class of monads (M, η, μ) such that M is a n-uniform choice endofunctor for $n \geq 2$.

Theorem 6.2.12. *If* $(W, \varepsilon, \delta) \in \mathcal{C}_W$ *and* $(M, \eta, \mu) \in \mathcal{C}_M$*, then there is no distributive law* $\rho\colon WM \to MW$ *of* (W, ε, δ) *over* (M, η, μ)*.*

Proof. Suppose for contradiction, there exists a distributive law $\rho\colon WM \to MW$. By Theorem 6.2.8, ρ must satisfy equation (6.11) for every $(s, l) \in W(U(X))$. We will construct an $(s, l) \in W(U(X))$ which does not satisfy the comultiplication axiom (6.4).

As $W \in \mathcal{C}_W$, we have that $|P(s)| \geq 2$. Hence, there exists a position $v \in P(s)$ such that $v \neq \mathsf{o}_s$. As $M \in \mathcal{C}_M$, there exists a n-uniform choice term β for some $n \geq 2$. Assume without loss of generality that $X = \{x_1, \ldots, x_n, x\}$. Let $l(v) = \beta(x_1, \ldots, x_n)$ and for all $q \neq v \in P(s)$, $l(q) = \eta(x)$. By construction, $|\mathsf{supp}(l(v))| = n$ and for all $q \neq v$, $|\mathsf{supp}(l(q))| = |\{x\}| = 1$. From this construction of (s, l) and equation (6.13), we calculate:

$$\begin{aligned}
|(\mathsf{supp} \circ M(\delta_X) \circ \rho_X)(s, l)| &= |(\mathcal{P}(\delta_X) \circ \mathsf{supp} \circ \rho_X)(s, l)| && \mathsf{supp}\text{-naturality of } \delta_X \\
&= |\mathsf{supp}(\rho_X(s, l))| && \delta_X \text{ injective} \\
&= \prod_{p \in P(s)} |\mathsf{supp}(l(p))| && \text{equation (6.13)} \\
&= |\mathsf{supp}(l(v))| && |\mathsf{supp}(l(q))| = 1 \text{ for } q \neq v \\
&= n && |\mathsf{supp}(l(v)| = 1
\end{aligned}$$

By contrast, consider $(s, m) = W(\rho_X) \circ \delta_{M(X)}(s, l)$ and note by the functorial action of W,

$$(s, \mathsf{supp} \circ m) = W(\mathsf{supp} \circ \rho) \circ \delta_{W(X)}(s, l).$$

From the definition of δ and Corollary 6.2.11, we have that

$$|\mathsf{supp}(m(p))| = \prod_{q \in P(s{\downarrow}p)} |\mathsf{supp}(l(p \oplus_s q))|. \tag{6.14}$$

Applying Corollary 6.2.11 again, we obtain the inequality:

$$\begin{aligned}
|\mathsf{supp} \circ \rho_{M(X)} \circ W(\rho_X) \circ \delta_{M(X)}(s, l)| &= \prod_{p \in P(s)} |\mathsf{supp}(m(p))| && \text{equation (6.13)} \\
&= \prod_{p \in P(s)} \prod_{q \in P(s{\downarrow}p)} |\mathsf{supp}(l(p \oplus_s q))| && \text{equation (6.14)} \\
&\geq |\mathsf{supp}(l(v \oplus_s \mathsf{o}_{s\downarrow v}))| |\mathsf{supp}(l(\mathsf{o}_s \oplus_s v))| \\
&= |\mathsf{supp}(l(v))| |\mathsf{supp}(l(v))| && \text{equation (6.8)} \\
&= n^2 && |\mathsf{supp}(l(v))| = n
\end{aligned}$$

By our supposition, ρ satisfies the comultiplication axiom, so

$$\begin{aligned}(M(\delta_X) \circ \rho_X)(s, l) &= (\rho_{W(X)} \circ W(\rho_X) \circ \delta_{M(X)})(s, l) \\ |(\mathsf{supp} \circ M(\delta_X) \circ \rho_X)(s, l)| &= |(\mathsf{supp} \circ \rho_{W(X)} \circ W(\rho_X) \circ \delta_{M(X)})(s, l)| \\ n &\geq n^2\end{aligned}$$

Contradiction, $n \geq 2$ implies that $n \not\geq n^2$. $\square$

Example 6.2.13. The discrete probability distribution monad $(\mathbf{D}_{\mathbb{R}^+}, \eta, \mu)$ from Example 6.1.3 is in $\mathcal{C}_M$. This monad has a natural transformation $\mathsf{supp}\colon \mathbf{D}_{\mathbb{R}^+} \to \mathcal{P}$ which maps a probability distribution to its underlying support, i.e. $\mathsf{supp}(\varphi) = \{x \mid \varphi(x) \neq 0\}$, and for every $n > 0$, a n-uniform choice term $\beta\colon \mathsf{Id}_{\mathbf{Set}} \times \cdots \times \mathsf{Id}_{\mathbf{Set}} \to \mathcal{D}$ defined as

$$\beta(x_1, \ldots, x_n) = \sum_{i \in [n]} \frac{1}{n} x_i.$$

Example 6.2.14. Every variation of the powerset $\wp$ monad, e.g. full, finite, non-empty, from Example 6.1.2 has a transformation $\mathsf{supp}\colon \wp \to \mathcal{P}$ given by the inclusion into the full powerset functor. Moreover, for every $n > 0$, $\wp$ has a n-uniform choice term $\beta\colon \mathsf{Id}_{\mathbf{Set}} \times \cdots \times \mathsf{Id}_{\mathbf{Set}} \to \wp$ defined as

$$\beta(x_1, \ldots, x_n) = \{x_1, \ldots, x_n\}.$$

Both these examples are instances of monads $M\colon \mathbf{Set} \to \mathbf{Set}$ where elements of $M(X)$ could be considered $\mathcal{S}$-weighted formal sums of elements in X for some semiring $\mathcal{S}$, e.g. semiring of non-negative reals for Example 6.2.13 and Boolean semiring for Example 6.2.14. In the following section, we determine necessary and sufficient conditions for these monads to be in the class $\mathcal{C}_M$.

6.2.3 Distribution and multiset monads of semirings

We begin by recalling the definition of semiring. A semiring is given by the data $\mathcal{S} = (S, 0_{\mathcal{S}}, 1_{\mathcal{S}}, +, \cdot)$ where $(S, 0_{\mathcal{S}}, +)$ is an 'additive' commutative monoid and $(S, 1_{\mathcal{S}}, \cdot)$ is a 'multiplicative' monoid where multiplication $\cdot$ distributes over addition $+$, i.e. for all x, y, z $\;x \cdot (y + z) = x \cdot y + x \cdot z$, $(y + z) \cdot x = y \cdot x + z \cdot x$, and $0_{\mathcal{S}} \cdot x = 0_{\mathcal{S}} = x \cdot 0_{\mathcal{S}}$. A semiring morphism is a set function preserving the addition, multiplication, additive unit, and multiplicative unit. Since the semiring of natural numbers $(\mathbb{N}, 0, 1, +, *)$ is the initial object in the category of semirings and semiring homomorphisms, for every semiring $\mathcal{S}$, there is a unique semiring morphism $\top^{\mathcal{S}}\colon \mathbb{N} \to \mathcal{S}$.

The *multiset monad over* $\mathcal{S}$ is $(\mathbf{M}_{\mathcal{S}}, \eta, \mu)$ where $\mathbf{M}_{\mathcal{S}}\colon \mathbf{Set} \to \mathbf{Set}$ is the endofunctor such that:

- for a set X, $\mathbf{M}_{\mathcal{S}}(X)$ is the set of all functions of type $\varphi\colon X \to S$ where φ has *finite support*, i.e. for all but finitely many $x \in X$, $\varphi(x) = 0_{\mathcal{S}}$. Due to the finite support restriction, elements $\varphi \in \mathbf{M}_{\mathcal{S}}(X)$ can be written as formal sums $\varphi = \Sigma_{x \in X} \varphi(x).x$.
- for a function, $f\colon X \to Y$, $\mathbf{M}_{\mathcal{S}}(f)\colon \mathbf{M}_{\mathcal{S}}(X) \to \mathbf{M}_{\mathcal{S}}(Y)$ maps φ to $\lambda y.\Sigma_{x \in f^{-1}(y)} \varphi(x)$.
- The unit has components defined as $\eta_X(x) = \Delta_x$ where $\Delta_x(x) = 1_{\mathcal{S}}$ and $\Delta_x(y) = 0_{\mathcal{S}}$ for $y \neq x$.
- The multiplication has components defined as $\mu_X(\varphi) = \lambda x.\Sigma_{\psi \in \mathbf{M}_{\mathcal{S}}(X)} \varphi(\psi) \cdot \psi(x)$.

We define the support of a multiset $\varphi \in \mathbf{M}_{\mathcal{S}}(X)$ as the set $\mathsf{supp}(\varphi) = \{x \in X \mid \varphi(x) \neq 0_{\mathcal{S}}\}$.

The distribution monad $(\mathbf{D}_{\mathcal{S}}, \eta, \mu)$ over semiring $\mathcal{S}$ underlying functor $\mathbf{D}_{\mathcal{S}}\colon \mathbf{Set} \to \mathbf{Set}$ where $\mathbf{D}_{\mathcal{S}}(X)$ is the subset of $\mathbf{M}_{\mathcal{S}}(X)$ such that elements $\varphi \in \mathbf{M}_{\mathcal{S}}(X)$ satisfy the normalisation condition $\sum_{x \in X} s_i = 1_{\mathcal{S}}$. The unit and multiplication of $(\mathbf{D}_{\mathcal{S}}, \eta, \mu)$ are obtained by restricting the corresponding maps of the multiset monad to $\mathbf{D}_{\mathcal{S}}(X)$.

In the case where the addition operation $+$ of $\mathcal{S}$ can be extended to an arbitrary sum operation which also distributes over multiplication, $\mathcal{S}$ is a complete semiring. As this sum operation in complete semirings is well-defined for infinite subsets, we can remove the finite support restriction on the functions in $\mathbf{M}_{\mathcal{S}}(X)$ and $\mathbf{D}_{\mathcal{S}}(X)$ to obtain analogous monads $\mathbf{M}_{\mathcal{S}}^{\infty}$ and $\mathbf{D}_{\mathcal{S}}^{\infty}$.

We enumerate sufficient conditions for when a distribution $\mathbf{D}_{\mathcal{S}}$ and multiset monad $\mathbf{M}_{\mathcal{S}}$ over a semiring $\mathcal{S}$ is in the class of monads $\mathcal{C}_M$ in order to understand the scope of Theorem 6.2.12.

(US1) $\mathcal{S}$ *is zero-sumfree*: If $a + b = 0_{\mathcal{S}}$, then $a = 0_{\mathcal{S}}$ and $b = 0_{\mathcal{S}}$.

(US2) $\mathcal{S}$ *has a natural non-trivial unit* $n_{\mathcal{S}}$*:* There exists some $n \geq 2$ such that $n_{\mathcal{S}} = \top^{\mathcal{S}}(n)$ is a unit. i.e. there exists $t \in \mathcal{S}$ such that $n_{\mathcal{S}} t = t n_{\mathcal{S}} = 1_{\mathcal{S}}$. If such a $t \in \mathcal{S}$ exists, then t is unique and so we can denote t as $\frac{1}{n_{\mathcal{S}}}$.

Lemma 6.2.15. *Let* $M = \mathbf{D}_{\mathcal{S}}$ *or* $M = \mathbf{M}_{\mathcal{S}}$ *for some semiring* $\mathcal{S}$*.* $\mathcal{S}$ *is zero-sumfree if and only if* (M, supp^M) *is a supported endofunctor.*

Proof. By contrapositive, assume that $\mathcal{S}$ is not zero-sumfree, then there exists $r, t \in \mathcal{S}$ such that $r + t = 0_{\mathcal{S}}$, but $r \neq 0_{\mathcal{S}}$ or $t \neq 0_{\mathcal{S}}$. Note that it follows from the semiring axiom $0_{\mathcal{S}} + a = a = a + 0_{\mathcal{S}}$, that $r \neq 0_{\mathcal{S}} \Leftrightarrow t \neq 0_{\mathcal{S}}$. Thus, both $r \neq 0_{\mathcal{S}}$ and $t \neq 0_{\mathcal{S}}$. Let $X = \{x, y, z\}$, $A = \{a, b\}$ and $f\colon X \to A$ defined as $x, y \mapsto a$ and $z \mapsto b$. Consider the distribution $\varphi = r.x + t.y + 1_{\mathcal{S}}.z \in \mathbf{D}_{\mathcal{S}}(X)$. As $r + t = 0_{\mathcal{S}}$, φ satisfies the normalisation condition and is indeed a distribution in $\mathbf{D}_{\mathcal{S}}(X)$. We obtain the following inequality:

$$\begin{aligned}
&(\mathsf{supp}_A \circ \mathbf{D}_{\mathcal{S}}(f))(r.x + t.y + 1_{\mathcal{S}}.z) \\
&\qquad = \mathsf{supp}_A((r+t).a + 1_{\mathcal{S}}.b) && \text{definition of } \mathbf{D}_{\mathcal{S}}(f) \\
&\qquad = \mathsf{supp}_A(1_{\mathcal{S}}.b) && r + t = 0_{\mathcal{S}} \\
&\qquad = \{b\} && \text{definition of } \mathsf{supp} \\
&\qquad \neq \{a, b\} && a \neq b \\
&\qquad = \mathcal{P}(f)(\{x, y, z\}) && \text{definition of } \mathcal{P}(f) \\
&\qquad = (\mathcal{P}(f) \circ \mathsf{supp}_X)(r.x + t.y + 1_{\mathcal{S}}.z) && \text{definition of } \mathsf{supp},\ r \neq 0_{\mathcal{S}}
\end{aligned}$$

Therefore, $\mathsf{supp}_A \circ \mathbf{D}_{\mathcal{S}}(f) \neq \mathcal{P}(f) \circ \mathsf{supp}_X$, the supp-naturality square of f does not commute. Hence, the family of maps supp_X is not a natural transformation.

Conversely, suppose $\mathcal{S}$ is zero-sumfree and that $f\colon X \to Y \in \mathbf{Mor}(\mathbf{Set})$. We must show that $\mathsf{supp}_Y(\mathbf{D}_{\mathcal{S}}(f)(\varphi)) = \mathcal{P}(f)(\mathsf{supp}_X(\varphi))$ for all $\varphi \in \mathbf{D}_{\mathcal{S}}(X)$. Suppose $\varphi = \sum_{i \in I} s_i.x_i \in \mathbf{D}_{\mathcal{S}}(X)$ and for all $i \in I$, $s_i \neq 0_{\mathcal{S}}$. By definition of $\mathbf{D}_{\mathcal{S}}(f)$, $\mathbf{D}_{\mathcal{S}}(f)(\varphi) = \sum_{y \in f(X)} (\sum_{j \in J_y} s_j) y$ where $J_y = \{i \in I \mid x_i \in f^{-1}(y)\} \subseteq I$ and $f(X)$ is the image of f. Since $\mathcal{S}$ is zero-sumfree, the sums $\sum_{j \in J_y} s_j$ are non-zero whenever $y = f(x_i)$. Hence, we can compute

$$\begin{aligned}
\mathsf{supp}_Y(\mathbf{D}_{\mathcal{S}}(f)(\varphi)) &= \{f(x_i) \mid i \in I\} \\
&= \mathcal{P}(f)(\{x_i \mid i \in I\}) && \mathcal{P} \text{ on morphisms} \\
&= \mathcal{P}(f)(\mathsf{supp}_X(\varphi))
\end{aligned}$$

□

Lemma 6.2.16. *Let $M = \mathbf{D}_{\mathcal{S}}$ or $M = \mathbf{M}_{\mathcal{S}}$. $\mathcal{S}$ has a natural non-trivial unit $n_{\mathcal{S}}$ if and only if there exists a n-ary idempotent and commutative term $\beta\colon \mathsf{Id}_{\mathbf{Set}} \times \cdots \times \mathsf{Id}_{\mathbf{Set}} \to M$ for $n \geq 2$ defined as:*

$$\beta(x_1, \dots, x_n) = \sum_{i \in [n]} \frac{1}{n_{\mathcal{S}}} x_i$$

In particular, if $\mathcal{S}$ is zero-sumfree, then β is a n-uniform choice term.

Proof. Suppose $\mathcal{S}$ has a natural non-trivial unit $n_{\mathcal{S}}$ and β is defined, using $n_{\mathcal{S}}$, as in the statement. To verify that β is idempotent, suppose $X \in \mathbf{Ob}(\mathbf{Set})$ and $x \in X$, then:

$$\begin{aligned}
\beta(x, \dots, x) &= \sum_{i\in[n]} \frac{1}{n_{\mathbb{S}}}.x && \text{definition of } \beta \\
&= \frac{1}{n_{\mathbb{S}}}(\sum_{i\in[n]} 1_{\mathbb{S}}).x && \text{distribution axiom of semirings} \\
&= \frac{1}{n_{\mathbb{S}}} n_{\mathbb{S}}.x && \text{definition of } n_{\mathbb{S}} \\
&= 1_{\mathbb{S}}.x && n_{\mathbb{S}} \text{ is a unit} \\
&= \eta(x) && \text{definition of } \eta
\end{aligned}$$

To verify β is commutative, suppose $X \in \mathbf{Ob}(\mathbf{Set})$ and $x_1, \dots, x_n \in X$, then:

$$\beta(x_{\pi(1)}, \dots, x_{\pi(n)}) = \sum_{i\in[n]} \frac{1}{n_{\mathbb{S}}}.x_{\pi(i)} \overset{*}{=} \sum_{i\in[n]} \frac{1}{n_{\mathbb{S}}}.x_i = \beta(x_1, \dots, x_n)$$

where the starred equality $\overset{*}{=}$ follows from the fact that rearranging terms in the formal sum yields the same distribution.

Conversely, suppose β is idempotent and commutative n-ary open term for $n \geq 2$, then we want to show that $n_{\mathbb{S}}$ is a natural non-trivial unit. Consider a set of n distinct elements $X = \{x_1, \dots, x_n\}$ and $\beta(x_1, \dots, x_n) = \varphi \in M(X)$. We want to show that $\varphi(x_1)$ is the inverse to $n_{\mathbb{S}}$. We first show that for every $i \in [n]$, $\varphi(x_1) = \varphi(x_i)$. To accomplish this, for every $i \in [n]$, consider the permutation $\tau_i \colon X \to X$ which swaps x_i and x_1. By construction, $\tau_i(x_j) = x_{(1\,i)(j)}$ where $(1\,i) \colon [n] \to [n]$ is the transposition swapping $1 \in [n]$ and $i \in [n]$.

$$\begin{aligned}
\varphi(x_i) &= \beta(x_1, \dots, x_n)(x_i) && \varphi = \beta(x_1, \dots, x_n) \\
&= \beta(x_{(1\,i)(1)}, \dots, x_{(1\,i)(n)})(x_i) && \beta \text{ commutative} \\
&= \beta(\tau_i(x_1), \dots, \tau_i(x_n))(x_i) && \text{definition of } \tau_i \\
&= M(\tau_i)(\beta(x_1, \dots, x_n))(x_i) && \beta\text{-naturality square of } \tau_i \\
&= M(\tau_i)(\varphi)(x_i) && \varphi = \beta(x_1, \dots, x_n) \\
&= \sum_{x_j \in \tau_i^{-1}(x_i)} \varphi(x_j) && M \text{ on morphisms} \\
&= \varphi(x_1) && \tau_i^{-1}(x_i) = \{x_1\}
\end{aligned}$$

We now show that $\varphi(x_1) = \varphi(x_2) = \cdots = \varphi(x_n)$ is the inverse of $n_{\mathbb{S}}$. To accomplish this, we consider the 'constant x_1 function' $c\colon X \to X$ which sends every element in X to x_1.

$$
\begin{aligned}
1_{\mathcal{S}} &= \eta(x_1)(x_1) && \text{definition of } \eta \\
&= \beta(x_1, \ldots, x_1)(x_1) && \beta \text{ idempotent} \\
&= \beta(c(x_1), \ldots, c(x_n))(x_1) && \text{definition of } c \\
&= M(c)(\beta(x_1, \ldots, x_n))(x_1) && \beta\text{-naturality square of } c \\
&= M(c)(\varphi)(x_1) && \varphi = \beta(x_1, \ldots, x_n) \\
&= \sum_{x_j \in c^{-1}(x_1)} \varphi(x_j) && M \text{ on morphisms} \\
&= \sum_{i \in [n]} \varphi(x_i) && c^{-1}(x_1) = \{x_1, \ldots, x_n\} = X \\
&= \sum_{i \in [n]} \varphi(x_1) && \forall i \in [n], \varphi(x_i) = \varphi(x_1)
\end{aligned}
$$

From the identity axiom of monoids, i.e. $\varphi(x_1) = \varphi(x_1)1_{\mathcal{S}} = 1_{\mathcal{S}}\varphi(x_1)$, and the two distribution axioms of semirings, we obtain that:

$$
\begin{aligned}
1_{\mathcal{S}} &= \varphi(x_1)(\sum_{i \in [n]} 1_{\mathcal{S}}) = \varphi(x_1)n_{\mathcal{S}}; \text{ and} \\
1_{\mathcal{S}} &= (\sum_{i \in [n]} 1_{\mathcal{S}})\varphi(x_1) = n_{\mathcal{S}}\varphi(x_1).
\end{aligned}
$$

Thus, $\varphi(x_1)$ is the inverse of $n_{\mathcal{S}}$, i.e. $\varphi(x_1) = \frac{1}{n_{\mathcal{S}}}$. By $n \geq 2$, $n_{\mathcal{S}}$ is a non-trivial unit. Moreover, by $\beta(x_1, \ldots, x_n) = \varphi$ and the fact that $\varphi(x_i) = \varphi(x_1) = \frac{1}{n_{\mathcal{S}}}$ for all $i \in [n]$, β is defined as in the statement of the theorem.

Finally, if we assume that $\mathcal{S}$ is zero-sumfree, then $n_{\mathcal{S}} \neq 0_{\mathcal{S}}$ and $\frac{1}{n_{\mathcal{S}}} \neq 0_{\mathcal{S}}$. Thus,

$$
\mathsf{supp}(\beta(x_1, \ldots, x_n)) = \mathsf{supp}(\sum_{i \in [n]} \frac{1}{n_{\mathcal{S}}} x_i) = \{x_1, \ldots, x_n\}.
$$

Hence, in the case where $\mathcal{S}$ is zero-sumfree, β is a n-ary uniform choice term. ☐

Theorem 6.2.17. *Let $M = \mathbf{D}_{\mathcal{S}}$ or $M = \mathbf{M}_{\mathcal{S}}$ for semiring $\mathcal{S}$. $\mathcal{S}$ satisfies conditions (US1)-(US2) if and only if $\mathbf{D}_{\mathcal{S}} \in \mathcal{C}_M$ and $\mathbf{M}_{\mathcal{S}} \in \mathcal{C}_M$. If $\mathcal{S}$ is also complete, then $\mathbf{D}^{\infty}_{\mathcal{S}} \in \mathcal{C}_M$ and $\mathbf{M}^{\infty}_{\mathcal{S}} \in \mathcal{C}_M$.*

Proof. Follows from combining the statements of Lemma 6.2.15 and Lemma 6.2.16. For the case of $\mathbf{D}^{\infty}_{\mathcal{S}}$ and $\mathbf{M}^{\infty}_{\mathcal{S}}$ when $\mathcal{S}$ is complete, we note that the proofs of these lemmas do not make essential use of the finite support restriction. ☐

As a corollary, we obtain an instance of Theorem 6.2.12 for multiset and distribution monads.

Corollary 6.2.18. *If $W \in \mathcal{C}_W$ and $\mathcal{S}$ satisfies conditions (US1)- (US2), then there is no distributive law $\rho\colon WM \to MW$ where $M = \mathbf{D}_{\mathcal{S}}$ or $M = \mathbf{M}_{\mathcal{S}}$. If $\mathcal{S}$ is also complete, then there is no distributive law $\rho\colon WM \to MW$ where $M = \mathbf{D}^{\infty}_{\mathcal{S}}$ or $M = \mathbf{M}^{\infty}_{\mathcal{S}}$.*

Example 6.2.19. The monads in Example 6.2.14 and Example 6.2.13 have a n-uniform choice term *for every* $n > 0$. This example illustrates why the weaker condition of having a uniform distribution of size n *for some* $n \geq 2$ genuinely widens the scope of Theorem 6.2.12. The multiset $\mathbf{M}_{\mathcal{S}}$ and distribution $\mathbf{D}_{\mathcal{S}}$ monads over the (sub)-semiring of non-negative rationals of the form $\frac{n}{3^k}$ for $n, k \in \mathbb{N}$ are in $\mathcal{C}_M$. In this case, since 2 is not a unit, there are no uniform distributions $\varphi \in \mathbf{D}_{\mathcal{S}}(X)$ such that $|\mathsf{supp}(\varphi)| = 2$ However, for every $k \in \mathbb{N}$, 3^k is a unit and so $\mathbf{D}_{\mathcal{S}}(X)$ has uniform distributions φ such that $|\mathsf{supp}(\varphi)| = 3^k$.

As another application of Corollary 6.2.18, we show the non-existence of a distributive law $\rho\colon \mathbb{E} \circ \widehat{\mathbf{D}_{\mathbb{R}+}} \to \widehat{\mathbf{D}_{\mathbb{R}+}} \circ \mathbb{E}$ where $k > 1$ and $\widehat{\mathbf{D}_{\mathbb{R}+}}$ is a lifting of the ordinary probability distribution monad $\mathbf{D}_{\mathbb{R}+}$. This monad was defined in Adam Connolly's Thesis [26] and its Kleisli morphisms capture the basic linear programming relaxations (BLP) of the constraint satisfaction problem.

Example 6.2.20. The BLP monad $\widehat{\mathbf{D}_{\mathbb{R}+}}$ has functor $\widehat{\mathbf{D}_{\mathbb{R}+}}\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma)$ which sends a relational structure $\mathcal{A}$ to a new structure $\widehat{\mathbf{D}_{\mathbb{R}+}}(\mathcal{A})$ with universe $\mathbf{D}_{\mathbb{R}+}(A)$. For every relation symbol $R \in \sigma$, we define

$$R^{\widehat{\mathbf{D}_{\mathbb{R}+}}(\mathcal{A})} = \{(\sum_{\vec{a} \in R^{\mathcal{A}}} \gamma(\vec{a}).\vec{a}[1] \ldots, \sum_{\vec{a} \in R^{\mathcal{A}}} \gamma(\vec{a}).\vec{a}[m]) \mid \gamma\colon R^{\mathcal{A}} \to [0,1], \sum_{\vec{a} \in R^{\mathcal{A}}} \gamma(\vec{a}) = 1\}$$

Unit and multiplication have the same element-wise definition as the distribution monad $\mathbf{D}_{\mathbb{R}+}$ on **Set**.

Corollary 6.2.21. *There is no distributive law $\rho\colon \mathbb{E}_k\widehat{\mathbf{D}_{\mathbb{R}+}} \to \widehat{\mathbf{D}_{\mathbb{R}+}}\mathbb{E}_k$*

Proof. We can view **Set** as a full subcategory $\mathbf{Set}(\sigma)$ of $\mathsf{Struct}(\sigma)$ where a set X is identified with the object of $\mathsf{Struct}(\sigma)$ with universe X and empty relational interpretations, i.e. for every $R \in \sigma$, $R^X = \varnothing$. Observe that the (co)monads $\mathbb{E}$ and $\widehat{\mathbf{D}_{\mathbb{R}+}}$ restrict to the subcategory $\mathbf{Set}(\sigma)$ and these restrictions can be identified with the (co)monads N^+_k and $\mathbf{D}_{\mathbb{R}+}$ over **Set**, respectively. Thus, if we suppose for contradiction there exists a distributive law $\rho\colon \mathbb{E}\widehat{\mathbf{D}_{\mathbb{R}+}} \to \widehat{\mathbf{D}_{\mathbb{R}+}}\mathbb{E}$ over $\mathsf{Struct}(\sigma)$, this would yield a distributive law $\rho'\colon N^+\mathbf{D}_{\mathbb{R}+} \to \mathbf{D}_{\mathbb{R}+}N^+$ over **Set** contradicting Corollary 6.2.18. A more formal categorical proof of this argument may be found in [55]. □

To understand the limitations of Theorem 6.2.12 and Corollary 6.2.18, we exhibit some examples of semirings which do not satisfy Conditions (US1)-(US2).

Example 6.2.22 (Non-Example)**.** Every ring R is not zero-sum-free, so the R-module monad $\mathcal{M}_R$ and distribution monad $\mathcal{D}_R$, by Lemma 6.2.15, are not in $\mathcal{C}_M$.

Example 6.2.23 (Non-Example)**.** There are semirings which are zero-sum-free, but fail to satisfy the condition (US2). For instance, the semiring of naturals $(\mathbb{N}, 0, 1, +, *)$ does not satisfy condition (US2). Therefore, although $\mathcal{C}_M$ has many multiset monads over other semirings $\mathcal{S}$, $\mathcal{C}_M$ does not contain the ordinary multiset/bag monad $\mathcal{M} = \mathcal{M}_{\mathbb{N}}$. Moreover, $\mathcal{D}_{\mathbb{N}}$ over the semiring $(\mathbb{N}, 0, 1, +, *)$ of natural numbers only has distributions which are singletons and is, in fact, isomorphic to the identity monad.

Example 6.2.24 (Non-Example)**.** For a example of a zero-sumfree semiring $\mathcal{S}$ which does not satisfy condition (US2), but where $\mathbf{D}_{\mathcal{S}}$ is not the identity monad, consider $\mathcal{S} = \mathbb{N}[x, y]/(x + y = 1)$. This semiring is the quotient of the free commutative semiring on the set $\{x, y\}$ by the equation $x + y = 1_{\mathcal{S}}$. The additional equation $x + y = 1_{\mathcal{S}}$ ensures that $\mathbf{D}_{\mathcal{S}}$ is not the identity monad by allowing for non-singleton distributions, i.e. $x.a + y.b \in \mathbf{D}_{\mathcal{S}}(\{a, b\})$. However, this is the only non-singleton distribution in $\mathbf{D}_{\mathcal{S}}(\{a, b\})$ and neither x nor y are inverses to $\top^{\mathcal{S}}(n)$ for some $n \geq 2$. Thus, $\mathcal{S}$ does not satisfy condition (US2), and $\mathbf{D}_{\mathcal{S}} \notin \mathcal{C}_M$. We adapt this example in Section 6.3.2 as a step towards demonstrating the existence of a distributive law of the EF game comonad over the quantum monad on relational structures.

6.2.4 Filter monad

In this section, through the example of the filter monad, we demonstrate that the class of monads under the scope of no-go Theorem 6.2.12 is not limited to multiset and distribution monads. The filter monad $(\mathcal{F}, \eta, \mu)$ over **Set**, appearing in [31], maps a set X to the families of subsets $U \in \mathcal{F}(X)$ which satisfy the following closure conditions:

(F1) If $S \in U$ and $S \subseteq S'$, then $S' \in U$.

(F2) If $S \in U$ and $S' \in U$, then $S \cap S' \in U$.

For a function $f \colon X \to Y$, $\mathcal{F}(f) \colon \mathcal{F}(X) \to \mathcal{F}(Y)$ is the double inverse image function, i.e.

$$\mathcal{F}(f)(U) = \{T \subseteq Y \mid f^{-1}(T) \in U\}.$$

The unit $\eta \colon \mathsf{Id}_{\mathbf{Set}} \to \mathcal{F}$ maps $x \in X$ to the family of subsets $S \subseteq X$ such that $x \in S$. The multiplication of $\mathcal{F}$ is cumbersome to state formally, but as the proof of Theorem 6.2.12 demonstrates, we will not need it for this application.

We will now show that the filter monad is a n-uniform choice monad for every $n > 1$.

To describe elements of $\mathcal{F}(X)$ we use the notation $\langle S\rangle$ for the filter generated by the subset S of X. Given any filter $U \in \mathcal{F}(X)$, it is uniquely determined by its minimal element $\min(U)$ with respect to $\subseteq$. We define $\min(U)$ as the $S \in U$ such that if $S' \in U$ and $S' \subseteq S$, then $S = S'$. For the notation to be well-defined $\min(U) = S$ has to be the unique element satisfying this minimality condition. To verify this, suppose S' is also minimal, then $S \cap S' \subseteq S$ and $S \cap S' \in U$ by (F2), by S being minimal $S \cap S' = S$. Similarly, $S \cap S' \subseteq S'$, so $S \cap S' = S'$. Thus, by transitivity of equality, $S = S \cap S' = S'$. Hence, every filter $U \in \mathcal{F}(X)$ is of the form $\langle \min(U)\rangle$.

Lemma 6.2.25. *The family of maps* $\mathsf{supp}_X \colon \mathcal{F}(X) \to \mathcal{P}(X)$ *indexed by* $X \in \mathbf{Ob}(\mathbf{Set})$

$$\mathsf{supp}_X(U) = \min(U)$$

form a natural transformation. In particular, $(\mathcal{F}, \eta, \mathsf{supp})$ *is a supported pointed endofunctor.*

Proof. Suppose $f \colon X \to Y$, then we need to show that $\mathcal{P}(f) \circ \mathsf{supp}_X = \mathsf{supp}_Y \circ \mathcal{F}(f)$. Unpacking this expression, we need to show that for every $U \in \mathcal{F}(X)$, the following equality between sets holds:

$$\min(\mathcal{F}(f)(U)) = \mathcal{P}(f)(\min(U)) \tag{6.15}$$

The following fact relating the preimage $f^{-1} \colon \mathcal{P}(Y) \to \mathcal{P}(X)$ and direct image $\mathcal{P}(f) \colon \mathcal{P}(X) \to \mathcal{P}(Y)$ will be useful for our proof:

(IM) For all $S \subseteq X$, $T \subseteq Y$, $\mathcal{P}(f)(S) \subseteq T$ if and only if $S \subseteq f^{-1}(T)$.

Suppose $T = \min(\mathcal{F}(f)(U))$. By definition of $\mathcal{F}(f)(U)$, $f^{-1}(T) \in U$. Thus, for $S = \min(U)$, by (F2), $f^{-1}(T) \cap S \in U$.

$$\begin{aligned} f^{-1}(T) \cap S &\subseteq f^{-1}(T) && \text{intersection} \\ S &\subseteq f^{-1}(T) && S \text{ minimal, } f^{-1}(T) \cap S = S \\ \mathcal{P}(f)(S) &\subseteq T && (IM) \end{aligned}$$

Hence, by T being the minimal element of $\mathcal{F}(f)(U)$, $\mathcal{P}(f)(S) = T$. Therefore, $T = \mathcal{P}(f)(\min(U))$. □

Lemma 6.2.26. *For every* $n > 0$, *there is a* n*-uniform choice term* $\beta \colon \mathsf{Id}_{\mathbf{Set}} \times \cdots \times \mathsf{Id}_{\mathbf{Set}} \to \mathcal{F}$ *for* $\mathcal{F}$ *with components* β_X *for* $X \in \mathbf{Ob}(\mathbf{Set})$ *defined as:*

$$\beta_X(x_1, \ldots, x_n) = \langle \{x_1, \ldots, x_n\}\rangle.$$

Proof. To verify idempotence, $\beta_X(x, \ldots, x) = \langle\{x\}\rangle$ which is the family of subsets which contain $x \in X$, so $\beta_X(x, \ldots, x) = \eta_X(x)$. To verify commutativity, we note that a permutation of the arguments of β just fixes the set $\{x_1, \ldots, x_n\}$ and thus fixes its closure. To verify that supp preserves β, we note that $\min(\langle\{x_1, \ldots, x_n\}\rangle)$ is the set $\{x_1, \ldots, x_n\}$, and so $\mathsf{supp}(\langle\{x_1, \ldots, x_n\}\rangle) = \{x_1, \ldots, x_n\}$. □

From these two lemmas, we can conclude that the filter monad $\mathcal{F} \in \mathcal{C}_M$ and obtain the following corollary of Theorem 6.2.12:

Corollary 6.2.27. *If $W \in \mathcal{C}_W$, then there is no distributive law $\rho\colon W\mathcal{F} \to \mathcal{F}W$.*

6.3 Examples of mixed distributive laws

In this section, we exhibit and verify a few example of mixed distributive laws that involving (co)monads just outside the boundary of the general no-go Theorem 6.2.12.

6.3.1 Coreader comonad over powerset monad

The definition of the class of comonads $\mathcal{C}_W$ for which the no-go theorem applies is all directed containers where there exists an $s \in S$, $|P(s)| \geq 2$. Since directed containers are equivalent to polynomial comonads, the class $\mathcal{C}_W$ can also be described as the class of non-linear polynomial comonads over **Set**. From the definition of directed container, the linear polynomial comonads are the 'product by S', also known as, the S-coreader comonad for any set $S \in \mathbf{Ob}(\mathbf{Set})$. For a fixed set S, the underlying functor $S \times (\cdot)$ of the coreader comonad $(S \times (\cdot), \varepsilon, \delta)$ maps a set X to $S \times X$ and a function f to $\mathsf{id}_S \times f$. The counit has components defined as $\varepsilon(s, x) = x$ and comultiplication has components defined as $\delta(s, x) = (s, (s, x))$. The identity comonad is recovered as the case where S is a singleton. The coreader comonad on set S is isomorphic to the induced comonad of a directed container on $[S \triangleleft P]$ where for every $s \in S$, $P(s)$ is the singleton set $\{\top\}$. As $P(s)$ is trivial for every $s \in S$, the directed container structure is such that $\mathsf{o}_s = \top$, $s \downarrow \top = s$, and $\top \oplus_s \top = \top$. Using the no-go Theorem 6.2.12, we can give a characterisation of coreader comonads in terms of mixed distributive laws over the powerset monad.

Theorem 6.3.1. *Let (W, ε, δ) be a directed container with $W = [S \triangleleft P]$. $W = S \times (\cdot)$ is the coreader comonad on S if and only if there exists a distributive law $\kappa\colon W\mathcal{P} \to \mathcal{P}W$.*

Proof. $\Rightarrow$ Suppose $W = S \times (\cdot)$ is the coreader comonad. Let $\kappa\colon W\mathcal{P} \to \mathcal{P}W$ be the natural transformation with components defined as

$$\kappa_X(s, Y) := \{(s, t) \mid t \in Y\}$$

for every $Y \in \mathcal{P}(X)$ and $s \in S$. It is easy to check that κ satisfies the diagrams (6.1),(6.2),(6.3),(6.4).

$\Leftarrow$ Conversely, suppose that there exists a distributive law $\kappa \colon W\mathcal{P} \to \mathcal{P}W$. By the contrapositive of Theorem 6.2.12 and $\mathcal{P} \in \mathcal{C}_M$, we must conclude that $W \notin \mathcal{C}_W$. Negating the definition of the class $\mathcal{C}_W$, for all $s \in S$, $|P(s)| < 2$. By W being a directed container, for every $s \in S$, there is root position $\mathsf{o}_s \in P(s)$, so $P(s)$ is non-empty. Therefore, $|P(s)| = 1$, so we can take $P(s) = \{\top\}$. This is precisely the definition of the coreader comonad. □

6.3.2 Prefix list comonad over single measurement monad

In this section, we explore a partial converse to Corollary 6.2.18, by constructing semirings $\mathcal{E}$ such that there is a distributive law $\rho \colon N^+\mathbf{D}_\mathcal{E} \to \mathbf{D}_\mathcal{E}N^+$. For every $n \geq 2$, we consider a generating set $S = \{e_1, \dots, e_n\}$ to produce the free semiring $\mathbb{N}[e_1, \dots, e_n]$ on S, and a set of equations $\Delta(n)$:

$$e_i^2 = e_i \text{ for every } i \in [n] \tag{6.16}$$

$$e_i e_j = 0 \text{ for every } i \neq j \in [n] \tag{6.17}$$

$$\sum_{i \in [n]} e_i = 1 \tag{6.18}$$

These equations generate an ideal $\langle \Delta(n) \rangle$ in $\mathbb{N}[e_1, \dots, e_n]$. Thus, we can define the quotient semiring $\mathcal{E}^n = {}^{\mathbb{N}[e_1, \dots, e_n]}/_{\langle \Delta(n) \rangle}$. Throughout the rest of the section, we will assume that $\mathcal{E} = \mathcal{E}^n$ for some $n \geq 2$.

Remark 6.3.2. The equations (6.16)-(6.18) were deliberately chosen to mimic the properties of projectors on a Hilbert space which form a single projector-valued measurement (PVM). These PVMs are convex sums of projectors arising in the context of quantum information. We will introduce PVMs in the next example section, Section 6.3.3. Accordingly, we refer to the semirings $\mathcal{E}^n$ as single measurement semirings and the monad $\mathbf{D}_\mathcal{E}$ as a single measurement monad. The construction of a distributive law $\rho \colon N^+\mathbf{D}_\mathcal{E} \to \mathbf{D}_\mathcal{E}N^+$ is a warm-up to constructing a distributive law of the k-round EF comonad over the quantum monad.

We now construct a distributive law $\rho \colon N^+\mathbf{D}_\mathcal{E} \to \mathbf{D}_\mathcal{E}N^+$ of the prefix list comonad N^+ over $\mathbf{D}_\mathcal{E}$. For every $Y \in \mathbf{Ob}(\mathbf{Set})$, the component ρ_Y is defined as

$$\rho_Y[\varphi_1, \dots, \varphi_n] = \lambda[y_1, \dots, y_m]. \begin{cases} \prod_{i \in [n]} \varphi_i(y_i) & \text{if } m = n \\ 0 & \text{otherwise} \end{cases}$$

Which we can express as the formal sum:

$$\rho_Y[\varphi_1, \dots, \varphi_n] = \sum_{[y_1,\dots,y_n]\in} \left(\prod_{i\in[n]} \varphi_i(y_i) \right) .[y_1, \dots, y_n]. \tag{6.19}$$

Proposition 6.3.3. *ρ is natural transformation.*

Proof. Suppose $f \colon X \to Y$ is a function. We define for every $i \in [n]$,

$$\psi_i(y) = \mathbf{D}_{\mathcal{E}}(f)(\varphi_i)(y) = \sum_{a\in f^{-1}(y)} \varphi_i(x). \tag{6.20}$$

To verify the ρ-naturality square of f, consider $[\varphi_1, \dots, \varphi_n] \in N^+(\mathbf{D}_{\mathcal{E}}(X))$:

$$\begin{aligned}
& \mathbf{D}_{\mathcal{E}}(N^+(f)) \circ \rho(s) \\
= \;& \{ \text{ definition of } \rho \text{ } \} \\
& N^+(\mathbf{D}_{\mathcal{E}}(f)) \left(\sum_{[x_1,\dots,x_n]\in N^+(X)} (\prod_{i\in[n]} \varphi_i(X_i)).[x_1, \dots, x_n] \right) \\
= \;& \{ \text{ } \mathbf{D}_{\mathcal{E}} \text{ on morphisms } \} \\
& \sum_{[x_1,\dots,x_n]\in N^+(X)} (\prod_{i\in[n]} \varphi_i(x_i)).N^+(f)([x_1, \dots, x_n]) \\
\\
= \;& \{ \text{ } N^+ \text{ on morphisms } \} \\
& \sum_{[x_1,\dots,x_n]\in N^+(X)} (\prod_{i\in[n]} \varphi_i(x_i)).([f(x_1), \dots, f(x_n)])
\end{aligned}$$

Using the distribution axiom of semirings, for every index $i \in [n]$, we can inductively 'pull-out all terms in the formal sum whose i-th element starts with $y_i = f(x_i)$.

$$\begin{aligned}
& \mathbf{D}_{\mathcal{E}}(N^+(f)) \circ \rho_X(s) \\
= \;& \{ \text{ distribution axiom of semirings } \} \\
& \sum_{[x_1,\dots,x_n]\in N^+(X)} \sum_{x_1\in f^{-1}(y_1)} \varphi_1(x_1) \left(\prod_{2\leq j\leq n} \varphi_j(x_j).[f(x_1), \dots, f(x_n)] \right) \\
= \;& \{ \text{ inductively apply distribution axiom } \} \\
& \sum_{[x_1,\dots,x_n]\in N^+(X)} (\sum_{x_1\in f^{-1}(y_1)} \varphi_1(x_1))(\sum_{x_2\in f^{-1}(y_2)} \varphi_2(x_2)) \dots (\sum_{x_n\in f^{-1}(y_n)} \varphi_n(x_n)).[f(x_1), \dots, f(x_n)] \\
= \;& \{ \text{ equation (6.20) } \} \\
& \sum_{[x_1,\dots,x_n]\in N^+(X)} \psi_1(y_1) \dots \psi_n(y_n).[f(x_1), \dots, f(x_n)] \\
= \;& \{ \text{ support in image of } N^+(f) \text{ } \} \\
& \sum_{[y_1,\dots,y_n]\in \mathbf{Im}(N^+(f))} \psi_1(y_1) \dots \psi_n(y_n).[y_1, \dots, y_n]
\end{aligned}$$

$$
\begin{aligned}
= \;& \{ \text{ product notation and } 0_{\mathcal{S}} \text{ terms for } N^+(Y)\backslash\mathbf{Im}(N^+(f)) \ \} \\
& \sum_{[y_1,\ldots,y_n]\in N^+(Y)} (\prod_{i\in[n]} \psi_i(y_i)).[y_1,\ldots,y_n] \\
= \;& \{ \text{ definition of } \rho \text{ (6.19) } \} \\
& \rho_Y([\psi_1,\ldots,\psi_n]) \\
= \;& \{ \text{ equation (6.20) } \} \\
& \rho_Y([\mathbf{D}_{\mathcal{E}}(f)(\varphi_1),\ldots,\mathbf{D}_{\mathcal{E}}(f)(\varphi_n)]) \\
= \;& \{ \ N^+ \text{ on morphisms } \} \\
& \rho_Y \circ N^+(\mathbf{D}_{\mathcal{E}}(f))(s)
\end{aligned}
$$

□

Before we verify that this definition of ρ indeed yields a distributive law, we first explain why the counterexample, demonstrating that the failure of the comultiplication axiom at the core of the no-go theorems, Theorem 6.2.4 and Theorem 6.2.12, can not be carried out in the case of $\mathcal{E}$. Recall that the counterexample in Theorem 6.2.4 traced the list $[\{a,b\},\{b\}] \in N^+(\wp(\{a,b\}))$ around the comultiplication diagram 6.2. One chase around the diagram yielded the set

$$\{[[a],[a,b]],[[b],[b,b]]\}.$$

Whereas the other direction yielded the set

$$\{[[a],[a,b]],[[b],[a,b]],[[a],[b,b]],[[b],[b,b]]\}$$

with the additional 'covariant' lists $[[b],[a,b]]$ and $[[a],[b,b]]$. If we attempt the same chase for our definition of ρ, in the case where $\mathcal{E} = \mathcal{E}^2$, we see that equation (6.17) eliminates these covariant terms and equation (6.16) reduces the degree in the other 'independent' terms as in the boldfaced expressions below:

$$
\begin{aligned}
& \rho_{N^+X} \circ N^+\rho_X \circ \delta_{\mathbf{D}_{\mathcal{E}}X}([e_1.a + e_2.b, 1.b]) \\
= \;& \{ \text{ definition of } \delta \ \} \\
& \rho_{N^+X} \circ N^+\rho_X([[e_1.a + e_2.b],[e_1.a + e_2.b, 1.b]]) \\
= \;& \{ \ N^+ \text{ on functions } \} \\
& \rho_{N^+X}([\rho_X([e_1.a + e_2.b]),\rho_X([e_1.a + e_2.b, 1.b])]) \\
= \;& \{ \text{ definition of } \rho \ \} \\
& \rho_{N^+X}([e_1.[a] + e_2.[b], e_1.[a,b] + e_2.[b,b]]) \\
= \;& \{ \text{ definition of } \rho \ \} \\
& e_1^2.[[a],[a,b]] + \mathbf{e_2e_1}.[[b],[a,b]] + \mathbf{e_1e_2}.[[a],[b,b]] + e_2^2.[[b],[b,b]]
\end{aligned}
$$

$$
\begin{aligned}
= \;& \{ \text{ equation (6.17) } \} \\
& e_1^2.[[a],[a,b]] + \mathbf{0}.[[b],[a,b]] + \mathbf{0}.[[a],[b,b]] + e_2^2.[[b],[b,b]] \\
= \;& \{ \text{ remove 0 weighted terms } \} \\
& \mathbf{e_1^2}.[[a],[a,b]] + \mathbf{e_2^2}.[[b],[b,b]] \\
= \;& \{ \text{ equation (6.16) } \} \\
& \mathbf{e_1}.[[a],[a,b]] + \mathbf{e_2}.[[b],[b,b]] \\
= \;& \{ \text{ definition of } \delta \ \} \\
& e_1.\delta_X([a,b]) + e_2.\delta_X([b,b]) \\
= \;& \{ \ \mathbf{D}_{\mathcal{E}} \text{ on functions } \} \\
& \mathbf{D}_{\mathcal{E}}(\delta_X)(\{e_1.[a,b] + e_2.[b,b]\}) \\
= \;& \{ \text{ definition of } \rho_X \ \} \\
& \mathbf{D}_{\mathcal{E}}(\delta_X) \circ \rho_X([e_1.a + e_2.b, 1.b])
\end{aligned}
$$

Generalising this diagram chase to arbitrary elements allows us to verify the comultiplication axiom (6.4) of Definition 6.1.4. We now proceed with verifying that ρ is indeed a mixed distributive law.

Proposition 6.3.4. *$\rho \colon N^+\mathbf{D}_{\mathcal{E}} \to \mathbf{D}_{\mathcal{E}}N^+$ is a mixed distributive law.*

Proof. Unit: To verify the unit axiom (6.1), consider $[x_1, \dots, x_n] \in N^+(X)$:

$$
\begin{aligned}
& \rho \circ N^+(\eta)([x_1, \dots, x_n]) \\
= \;& \{ \ N^+ \text{ on morphisms } \} \\
& \rho([\eta(x_1), \dots, \eta(x_n)]) \\
= \;& \{ \text{ definition of } \eta \ \} \\
& \rho([1.x_1, \dots, 1.x_n]) \\
= \;& \{ \text{ definition of } \rho \text{ and } \varphi_i = 1.x_i \ \} \\
& (\prod_{i \in [n]} 1).[x_1, \dots, x_n] \\
= \;& \{ \text{ identity axiom } \mathcal{E} \ \} \\
& 1.[x_1, \dots, x_n] \\
= \;& \{ \text{ definition of } \eta \ \} \\
& \eta([x_1, \dots, x_n])
\end{aligned}
$$

Counit: To verify the counit axiom (6.22), consider $[\varphi_1, \dots, \varphi_n] \in N^+(\mathbf{D}_{\mathcal{E}}(X))$:

$$
\mathbf{D}_{\mathcal{E}}(\varepsilon) \circ \rho([\varphi_1, \dots, \varphi_n])
$$

$$
\begin{aligned}
= &\quad \{ \text{ definition of } \rho \ \} \\
& \mathbf{D}_{\mathcal{E}}(\varepsilon)\left(\sum_{[x_1,\ldots,x_n]\in N^{+}(X)} (\prod_{i\in[n]} \varphi_i(x_i)).[x_1,\ldots,x_n]\right) \\
= &\quad \{ \ \mathbf{D}_{\mathcal{E}} \text{ on morphisms } \} \\
& \sum_{[x_1,\ldots,x_n]\in N^{+}(X)} (\prod_{i\in[n]} \varphi_i(x_i)).\varepsilon([x_1,\ldots,x_n]) \\
= &\quad \{ \text{ definition of } \varepsilon \ \} \\
& \sum_{[x_1,\ldots,x_n]\in N^{+}(X)} (\prod_{i\in[n]} \varphi_i(x_i)).x_n \\
= &\quad \{ \text{ distribution axiom } \mathsf{Proj}(d) \ \} \\
& \sum_{x_n\in X}\left(\sum_{[x_1,\ldots,x_{n-1}]\in N^{+}(X)} \prod_{i\in[\mathsf{n-1}]} \varphi_i(x_i)\right)\varphi_n(x_n).x_n \\
= &\quad \{ \text{ normalisation } \} \\
& \sum_{x_n\in X} 1\varphi_n(x_n).x_n \\
= &\quad \{ \text{ identity } 1\in\mathcal{E} \ \} \\
& \sum_{x_n\in X} \varphi_n(x_n).x_n \\
= &\quad \{ \ \varphi_n \text{ as a formal sum } \} \\
& \varphi_n \\
= &\quad \{ \text{ definition of } \varepsilon \ \} \\
& \varepsilon[\varphi_1,\ldots,\varphi_n]
\end{aligned}
$$

Multiplication: To verify the multiplication axiom (6.3), consider $[\Psi_1,\ldots,\Psi_n]\in N^{+}(\mathbf{D}_{\mathcal{E}}(\mathbf{D}_{\mathcal{E}}(X))$:

$$
\begin{aligned}
& \mu_{N^{+}(X)}\circ\mathbf{D}_{\mathcal{E}}(\rho)\circ\rho_{\mathbf{D}_{\mathcal{E}}(X)}([\Psi_1,\ldots,\Psi_n]) \\
= &\quad \{ \text{ definition of } \rho \ \} \\
& \mu_{N^{+}(X)}\circ\mathbf{D}_{\mathcal{E}}(\rho)\left(\sum_{[\varphi_1,\ldots,\varphi_n]\in N^{+}(\mathbf{D}_{\mathcal{E}}(X))} \prod_{i\in[n]} \Psi_i(\varphi_i)).[\varphi_1,\ldots,\varphi_n]\right) \\
= &\quad \{ \ \mathbf{D}_{\mathcal{E}} \text{ on morphisms } \} \\
& \mu_{N^{+}(X)}\left(\sum_{[\varphi_1,\ldots,\varphi_n]\in N^{+}(\mathbf{D}_{\mathcal{E}}(X))} \prod_{i\in[n]} \Psi_i(\varphi_i)).\rho([\varphi_1,\ldots,\varphi_n])\right) \\
= &\quad \{ \text{ definition of } \rho \ \} \\
& \mu_{N^{+}(X)}\left(\sum_{[\varphi_1,\ldots,\varphi_n]\in N^{+}(\mathbf{D}_{\mathcal{E}}(X))} \prod_{i\in[n]} \Psi_i(\varphi_i)).(\sum_{[x_1,\ldots,x_n]\in N^{+}(X)} (\prod_{i\in[n]} \varphi_i(x_i)).[x_1,\ldots,x_n])\right)
\end{aligned}
$$

$$
\begin{aligned}
= \quad & \{ \text{ definition of } \mu \ \} \\
& \sum_{[x_1,\ldots,x_n]\in N^+(X)} \ \sum_{[\varphi_1,\ldots,\varphi_n]\in N^+(\mathbf{D}_{\mathcal{E}}(X))} \left(\prod_{i\in[n]} \Psi_i(\varphi_i)\varphi_i(x_i) \right).[x_1,\ldots,x_n] \\
= \quad & \{ \text{ definition of } \rho \ \} \\
& \rho_X([\sum_{\varphi_1\in\mathbf{D}_{\mathcal{E}}(X)} \sum_{x_1\in X} \Psi_1(\varphi_1)(\varphi_1(x_1)).x_1, \ldots, \sum_{\varphi_n\in\mathbf{D}_{\mathcal{E}}(X)} \sum_{x_n\in X} \Psi_n(\varphi_n)(\varphi_n(x_n)).x_n] \\
= \quad & \{ \text{ definition of } \mu_X \ \} \\
& \rho_X([\mu_X(\Psi_1),\ldots,\mu_X(\Psi_n)]) \\
= \quad & \{ \ N^+ \text{ on morphisms } \} \\
& \rho_X \circ N^+(\mu_X)([\Psi_1,\ldots,\Psi_n])
\end{aligned}
$$

Comultiplication: To verify the comultiplication axiom (6.4), we need to be explicit about the number of monomials in $\mathcal{E}$, so we assume that $\mathcal{E} = \mathcal{E}^n$. Consider $\boldsymbol{\varphi} = [\varphi_1,\ldots,\varphi_m] \in N^+(\mathbf{D}_{\mathcal{E}}(X))$:

$$
\begin{aligned}
\rho_{\mathbf{D}_{\mathcal{E}}} \circ N^+(\rho) \circ \delta(\boldsymbol{\varphi}) &= \rho_{\mathbf{D}_{\mathcal{E}}} \circ N^+(\rho)([[\varphi_1],\ldots,[\varphi_1,\ldots,\varphi_m]]) \\
&= \rho_{\mathbf{D}_{\mathcal{E}}}([\rho([\varphi_1]),\ldots,\rho([\varphi_1,\ldots,\varphi_m])]) \\
&= \rho_{\mathbf{D}_{\mathcal{E}}} \left(\sum_{[x_1]} \varphi(x_1).[x_1], \cdots \sum_{[x_1,\ldots,x_m]} (\prod_{i\in[m]} \varphi_i(x_i)).[x_1,\ldots,x_m] \right) \\
&= \sum_{[s_1,\ldots,s_m]\in T} \ \prod_{j\in[m]} \left(\sum_{s_j=[x_1,\ldots,x_j]} \ \prod_{i\in[j]} \varphi_j(x_j) \right).[[x_1],\ldots,[x_1,\ldots,x_m]]
\end{aligned}
$$

where T is the image of the $\delta_{\mathbf{D}_{\mathcal{E}}}$. Using the +-commutative and distributive axioms of semirings we can rearrange the terms in the above expression to a more convenient form. To attain this rearrangement, we first note that by the normalisation condition, the weights in each of the φ_i must sum to 1. By the construction of $\mathcal{E}$, the only such sums are 1 itself and the sum of the monomials $e_i \in \mathcal{E}$ as in equation (6.18). Thus, for every $j \in [m]$, either $\varphi_j = 1.x_j$ for some $x_j \in X$ or $\varphi_j = e_1.x_j^1 + \cdots + e_n.x_j^n$ for some $x_j^1,\ldots,x_j^n \in X$. By equation (6.18), we can view the former case as instance of the latter where for all $z \in [n]$, $x_j^z = x_j$. Therefore, we can assume that for all $j \in [m]$, $\varphi_j = e_1.x_j^1 + \cdots + e_n.x_j^n$. Under this assumption, the only terms that are potentially non-zero correspond to the cases where $[s_1,\ldots,s_m] \in T$ is such that the $s_j = [x_1^{z_1},\ldots,x_j^{z_j}]$ for some $z_1,\ldots,z_j \in [n]$; so we define $T_{\boldsymbol{\varphi}}$ be the collection of all $[s_1,\ldots,s_m] \in T$ of this form. With this definition and noting that for all $j \in [m]$ and $i \in [n]$, $\varphi_j(x^{z_i}) = e_{z_i}$ we can write the above sum as:

$$
\begin{aligned}
& \rho_{\mathbf{D}_{\mathcal{E}}} \circ N^+(\rho) \circ \delta(\boldsymbol{\varphi}) \\
& \qquad = \sum_{[s_1,\ldots,s_m]\in T_{\boldsymbol{\varphi}}} \left(\sum_{s_m=[x_1^{z_1},\ldots,x_m^{z_m}]} e_{z_1} \cdot e_{z_2} \cdot \cdots \cdot e_{z_m} \right).[[x_1^{z_1}],\ldots,[x_1^{z_1},\ldots,x_m^{z_m}]]
\end{aligned}
$$

However, note that because of equation (6.17), all of the terms of the form $e_{z_1} \cdot \dots \cdot e_{z_m}$ where at least two of the $z_1, \dots, z_m$ are not equal become 0. Intuitively, these terms are the weights of the 'covariant' lists of lists and the terms that remain are for the 'independent' lists of lists: $[[x_1^z], \dots, [x_1^z, \dots, x_m^z]]$ for every $z \in [n]$. More precisely, the terms that remain are:

$$
\begin{aligned}
\rho_{\mathbf{D}_{\mathcal{E}}} \circ N^{+}(\rho) \circ \delta(\boldsymbol{\varphi}) &= \sum_{z\in[n]} \left(\sum_{s=[x_1^z,\dots,x_m^z]} e_z^m \right) .[[x_1^z], \dots, [x_1^z, \dots, x_m^z]] \\
&= \sum_{z\in[n]} \left(\sum_{s=[x_1^z,\dots,x_m^z]} e_z \right) .[[x_1^z], \dots, [x_1^z, \dots, x_m^z]] \\
&= \sum_{z\in[n]} \left(\sum_{s=[x_1^z,\dots,x_m^z]} e_z \right) .\delta([x_1^z, \dots, x_m^z]) \\
&= \mathbf{D}_{\mathcal{E}}(\delta)(\sum_{z\in[n]} e_z.[x_1^z, \dots, x_m^z]) \\
&= \mathbf{D}_{\mathcal{E}}(\delta) \circ \rho(\boldsymbol{\varphi})
\end{aligned}
$$

□

6.3.3 Ehrenfeucht-Fraïssé comonad over quantum monad

In this section, we exhibit a distributive law of a version of the k-round EF comonad $\mathbb{E}_k^{\odot}$ over a version of the quantum monad $\mathbf{Q}_d^{\odot}$ from [4] which are tailored to a category of 'commeasurable' σ-structures $\mathsf{Struct}^{\odot}(\sigma)$ defined below. We start by introducing the quantum monad over the category of σ-structures $\mathsf{Struct}(\sigma)$ from [4]. The quantum monad is not in fact a monad, but a *graded monad.* The definition of graded monad is rather general and in fact, parametric over monoidal categories [40, 67]. However for the purpose of introducing the quantum monad, we only need to define $(\mathbb{N}^{+}, 1, *)$-graded monads.

Definition 6.3.5. A $\mathbb{N}^{+}$-graded monad $(\{M_n\}, \eta, \{\mu^{n,m}\})$ on $\mathcal{E}$ is a $\mathbb{N}^{+}$ indexed family of endofunctors $M_n \colon \mathcal{E} \to \mathcal{E}$, a unit natural transformation $\eta \colon \mathsf{Id}_{\mathcal{E}} \to M_1$, and graded multiplication given by a family of natural transformations $\{\mu^{n,m} \colon M_n M_m \to M_{nm}\}_{n,m\in\mathbb{N}^{+}}$ satisfying the following diagrams for every $i, j, k \in \mathbb{N}^{+}$:

$$\begin{tikzcd}
M_i M_j M_k \arrow[r, "M_i \mu^{j,k}"] \arrow[d, "\mu^{i,j} M_k"'] & M_i M_{jk} \arrow[d, "\mu^{i,jk}"] \\
M_{ij} M_k \arrow[r, "\mu^{ij,k}"'] & M_{ijk}
\end{tikzcd}
\qquad
\begin{tikzcd}
M_i M_j M_k \arrow[r, "M_i \mu^{j,k}"] \arrow[d, "\mu^{i,j} M_k"'] & M_i M_{j*k} \arrow[d, "\mu^{i,jk}"] \\
M_{ij} M_k \arrow[r, "\mu^{ij,k}"'] & M_{ijk}
\end{tikzcd}$$

Since our aim is to exhibit a distributive law of the k-round EF comonad over the $\mathbb{N}^+$-graded monad $\mathbf{Q}_d$, we must also adapt Definition 6.1.4 of a comonad-monad distributive law to take into allow for the $\mathbb{N}^+$-grading on the monad.

Definition 6.3.6. A *mixed distributive law of a comonad* (W, η, δ) *over a* $\mathbb{N}^+$*-graded monad* $(\{M_d\}, \eta, \{\mu^{d,d'}\})$ is a family of natural transformations $\kappa^d \colon WM_d \to M_dW$ satisfying the following commutative diagrams for all $d, d' \in \mathbb{N}^+$:

$$\begin{tikzcd}
 & W \arrow[dl, "W\eta"'] \arrow[dr, "\eta W"] & \\
WM_1 \arrow[rr, "\kappa^1"'] & & M_1W
\end{tikzcd} \tag{6.21}$$

$$\begin{tikzcd}
WM_d \arrow[rr, "\kappa^d"] \arrow[dr, "\varepsilon M_d"'] & & M_dW \arrow[dl, "M_d\varepsilon"] \\
 & M &
\end{tikzcd} \tag{6.22}$$

$$\begin{tikzcd}
WM_dM_{d'} \arrow[r, "\kappa^d M_{d'}"] \arrow[d, "W\mu^{d,d'}"'] & M_dWM_{d'} \arrow[r, "M_d\kappa^{d'}"] & M_dM_{d'}W \arrow[d, "\mu^{d,d'}W"] \\
WM_{dd'} \arrow[rr, "\kappa^{dd'}"'] & & M_{dd'}W
\end{tikzcd} \tag{6.23}$$

$$\begin{tikzcd}
WM_d \arrow[rr, "\kappa^d"] \arrow[d, "\delta M_d"'] & & M_dW \arrow[d, "M_d\delta"] \\
WWM_d \arrow[r, "W\kappa^d"] & WM_dW \arrow[r, "\kappa^d W"] & M_dWW
\end{tikzcd} \tag{6.24}$$

Distributive laws of graded comonads over graded monads are defined in [41]. Definition 6.3.6 above is an instance of this definition which the grading on the comonad is trivial and the grading on the monad is given by the multiplicative monoid $(\mathbb{N}^+, 1, *)$.

We now proceed with defining the $\mathbb{N}^+$-graded quantum monad $\mathbf{Q}_d$ on $\mathsf{Struct}(\sigma)$. In quantum mechanics, pure states of a quantum system are vectors in a complex Hilbert space. For quantum information, it is common to restrict to finite dimensional Hilbert spaces. Let $\mathcal{H}_d$ denote the d-dimensional Hilbert space consisting of d-length tuples of complex numbers. Classical information about a quantum system is obtained by performing measurements on these states. Measurements are represented in a phase-invariant way via *projector-valued measurements (PVMs).* To define PVMs we first define (orthogonal) projectors.

Definition 6.3.7. Given a d-dimensional complex Hilbert space $\mathcal{H}_d$, a *projector* is a matrix E representing a linear map of type $\mathcal{H}_d \to \mathcal{H}_d$ satisfying the equations $E = E^2 = E^\dagger$ where $E^\dagger$ denotes the conjugate transpose of E.

The d-dimensional identity matrix I_d is a projector on the whole space $\mathcal{H}_d$ and the d-dimensional zero matrix $\mathbf{0}_d$ is a projector on the zero subspace of $\mathcal{H}_d$. Let $\mathsf{Proj}(d)$ denote set of projectors on $\mathcal{H}_d$. The product EF of two projectors E, F is a projector if and only if they commute $EF - FE = \mathbf{0}_d$. The projector E is said to be *orthogonal to F* if $EF = \mathbf{0}_d$. We are now ready to define PVMs.

Definition 6.3.8. Given a set A, *a d-dimensional projector-valued measurement over A* or *d-dim PVM over A* is a function $\vartheta\colon A \to \mathsf{Proj}(d)$ satisfying the following conditions:

(1) *ϑ has finite support:* For all but finitely many $a \in A$, $\vartheta(a) = \mathbf{0}_d$

(2) *ϑ is pairwise orthogonal:* If $a \neq a'$, then $\vartheta(a)\vartheta(a') = \mathbf{0}_d$.

(3) *normalisation:* The sum $\sum_{a \in A} \vartheta(a) = I_d$

As with the probability distributions, we can write PVMs as a formal sums $\vartheta = \sum_{a \in A} \vartheta(a).a$.

Given a d-dim PVM ϑ over A, the probability of measuring outcome $a \in A$ on a quantum system in pure quantum state, represented as vector $v \in \mathcal{H}_d$, is $v^{\dagger}Pv \geq 0$ where $P = \vartheta(a) \in \mathsf{Proj}(d)$. Two d-dim *PVMs* ϑ, ϑ' are jointly measurable if for all a, a', $\vartheta(a)$ commutes with $\vartheta'(a')$.

We are now ready to define the $\mathbb{N}^+$-graded quantum monad $(\{\mathbf{Q}_d\}, \eta, \{\mu^{d,d'}\})$ on $\mathsf{Struct}(\sigma)$.

- For every $d \in \mathbb{N}^+$, the object mapping of the endofunctor $\mathbf{Q}_d\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}(\sigma)$ maps a σ-structure $\mathcal{A}$ with universe A to a σ-structure $\mathbf{Q}_d(\mathcal{A})$ with universe

  $$\mathbf{Q}_d(A) = \{\vartheta\colon A \to \mathsf{Proj}(d) \mid \vartheta \text{ is a } d\text{-dim PVM.}\};$$

  and for every m-relation $R \in \sigma$, $(\vartheta_1, \dots, \vartheta_r) \in R^{\mathbf{Q}_d(\mathcal{A})}$ if and only if

  (1) for all $i, j \in [m]$, ϑ_i is jointly measurable with ϑ_j.

  (2) If $(a_1, \dots, a_m) \notin R^{\mathcal{A}}$, then $\vartheta_1(a_1) \dots \vartheta_m(a_m) = \prod_{i \in [m]} \vartheta_i(a_i) = \mathbf{0}_d$.

- For every $d \in \mathbb{N}^+$ and every σ-morphism $f\colon \mathcal{A} \to \mathcal{B} \in \mathbf{Mor}(\mathsf{Struct}(\sigma))$, the functorial action produces a σ-morphism $\mathbf{Q}_d(f)\colon \mathbf{Q}_d(\mathcal{A}) \to \mathbf{Q}_d(\mathcal{B})$ mapping A-valued PVM ϑ to B-valued PVM $\lambda b. \sum_{a \in f^{-1}(b)} \vartheta(a)$.

- The unit natural transformation $\eta\colon \mathsf{Id}_{\mathsf{Struct}(\sigma)} \to \mathbf{Q}_1$ has components $\eta_{\mathcal{A}}\colon \mathcal{A} \to \mathbf{Q}_1(\mathcal{A})$ for every $\mathcal{A} \in \mathsf{Struct}(\sigma)$ where $\eta_{\mathcal{A}}(a) = \Delta_a$ is the 'Dirac delta PVM':

  $$\Delta_a(a') = \begin{cases} I_1 & \text{if } a = a' \\ \mathbf{0}_1 & \text{if } a \neq a' \end{cases}$$

- To define the multiplication of the quantum monad, we recall that given $E \in \mathsf{Proj}(d)$ and $F \in \mathsf{Proj}(d')$, we can form the tensor product to obtain a projector $E \otimes F \in \mathsf{Proj}(dd')$. Thus, for every d, d', the transformation $\mu^{d,d'} \colon \mathbf{Q}_d\mathbf{Q}_{d'} \to \mathbf{Q}_{dd'}$ has components $\mu^{d,d'}_{\mathcal{A}}$ for every $\mathcal{A} \in \mathsf{Struct}(\sigma)$ where
$$\mu^{d,d'}_{\mathcal{A}}(\Theta) = \lambda a. \sum_{\vartheta \in \mathbf{Q}_{d'}(\mathcal{A})} \Theta(\vartheta) \otimes \vartheta(a).$$

Evidently, the definition of the quantum monad resembles the definition of the discrete probability distribution monad where $\mathbb{R}^+$-weighted convex sums are replaced with $\mathsf{Proj}(d)$-weighted convex sums, i.e. d-dimensional PVMs. In fact, taking into account that each PVM consists of weights which satisfy idempotence $E^2 = E$ and pairwise orthogonality $EF = \mathbf{0}_d$ for $F \neq E$, the quantum monad most closely resembles the definition of the single measurement monad from Section 6.3.2. In fact, $\mathbf{Q}_d$ can almost be see as a distribution monad over the 'semiring' of d-dimensional projectors $\mathsf{Proj}(d)$. However, this is not accurate as $\mathsf{Proj}(d)$ is not a semiring because the multiplication of projectors is a partial operation, i.e. $EF \in \mathsf{Proj}(d)$ is only a projector if E and F commute. Another salient difference between the quantum monad and distribution monads, is that the graded multiplication of the quantum monad uses the tensor product, rather then the multiplication of $\mathsf{Proj}(d)$. The use of tensor product, instead of projector multiplication, is important for modelling quantum homomorphisms, but also manages to avoid the issue of $\mathsf{Proj}(d)$'s multiplication being partial.

Despite this, given the resemblance of $\mathbf{Q}_d$ to $\mathbf{D}_{\mathcal{S}}$ and $\mathbb{E}_k$ to N^+, we can attempt to mimic the definition of $\rho \colon N^+\mathbf{D}_{\mathcal{S}} \to \mathbf{D}_{\mathcal{S}}N^+$ by defining a graded distributive law $\kappa^d \colon \mathbb{E}_k\mathbf{Q}_d \to \mathbf{Q}_d\mathbb{E}_k$ whose components κ_d would send $[\varphi_1, \ldots, \varphi_n]$ to the PVM where $[a_1, \ldots, a_n]$ has the weight $\prod_{i \in I} \varphi_i(a_i)$. However, the partiality of projector multiplication in $\mathsf{Proj}(d)$ means that this would not be well-defined. We may avoid this partiality issue in defining κ^d by moving from $\mathsf{Struct}(\sigma)$ to the category of commeasurable σ-structures $\mathsf{Struct}^{\odot}(\sigma)$. To define $\mathsf{Struct}^{\odot}(\sigma)$, we must first define the category of reflexive graphs **RGraph**. The objects of **RGraph** are undirected graphs $G = (V, E)$ such that each vertex $v \in V$ has a distinguished self-loop from $(v, v) \in E$. Morphisms in **RGraph** are graph homomorphisms which send self-loops to self-loops. We write $\odot$ for the commeasurability relation. The category of $\mathsf{Struct}^{\odot}(\sigma)$ has objects which are σ-structures $\mathcal{A}$ such that the universe $A \in \mathbf{Ob}(\mathbf{RGraph})$ and that two elements $a, a' \in \mathcal{A}$ can be appear in related tuple if $a \odot a'$. Commeasurable σ-structures form the category $\mathsf{Struct}^{\odot}(\sigma)$ where morphisms are morphisms in **RGraph** in addition to the interpretation of relation symbols $R \in \sigma$.

We now adapt the definition of the k-round EF comonad produce a comonad $(\mathbb{E}^{\odot}_k, \varepsilon, \delta)$ on commeasurable σ-structures $\mathsf{Struct}^{\odot}(\sigma)$. The key difference between $\mathbb{E}_k$ and $\mathbb{E}^{\odot}_k$ is that $\mathbb{E}_k^{\odot}$ only consists of sequences of elements which are pairwise commeasurable. For a commeasurable σ-structure $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}^{\odot}(\sigma))$ with universe $A \in \mathbf{Ob}(\mathbf{RGraph})$, the universe $\mathbb{E}_k^{\odot}(A)$ of $\mathbb{E}_k^{\odot}(\mathcal{A})$ has underlying set

$$\mathbb{E}_k^{\odot}(A) = \{[a_1, \ldots, a_n] \mid n \leq k, \forall i \in [n], a_i \odot a_j\}.$$

and commeasurability relation $\odot$ being prefix comparable, i.e. $s \odot s'$ if $s \sqsubseteq s'$ or $s' \sqsubseteq s$. Just as with $\mathbb{E}_k$, the counit $\varepsilon\colon \mathsf{Id}_{\mathsf{Struct}^{\odot}(\sigma)} \to \mathbb{E}_k^{\odot}$ returns the last element. For every m-ary relation $R \in \sigma$, the define the interpretation

$$\begin{aligned} R^{\mathbb{E}_k^{\odot}(\mathcal{A})}(s_1, \ldots, s_m) \Leftrightarrow{} & \forall i, j \in [m], s_i \odot s_j \\ & \text{and } R^{\mathcal{A}}(\varepsilon_{\mathcal{A}}(s_1), \ldots, \varepsilon_{\mathcal{A}}(s_m)) \end{aligned}$$

Notice that since $\odot$ for $\mathbb{E}_k^{\odot}(A)$, by definition, is prefix comparability this coincides with the definition of relational interpretation $R^{\mathbb{E}_k(\mathcal{A})}$. The action of $\mathbb{E}_k^{\odot}$ on $\mathsf{Struct}^{\odot}(\sigma)$-morphisms and the definition δ are defined as for $\mathbb{E}_k$.

Remark 6.3.9. Ostensibly, this move to the setting of commeasurable σ-structures $\mathsf{Struct}^{\odot}(\sigma)$ to fix the partiality issue in the definition of κ^d may seem unmotivated. However, this fix was inspired by a similar move in the study of logic for (quantum) contextuality where total Boolean algebras are replaced with partial Boolean algebras [3, 57]. In that setting, similar to our case, partial Boolean algebras are equipped with a commeasurability relation such that the meet and join operations are only well-defined for commeasurable elements. Moreover, we can recover ordinary σ-structures $\mathsf{Struct}(\sigma)$ from $\mathsf{Struct}^{\odot}(\sigma)$ since there is a reflection $U \dashv G$ where $U\colon \mathsf{Struct}^{\odot}(\sigma) \to \mathsf{Struct}(\sigma)$ is the forgetful functor removing the commeasurability relation and $G\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}^{\odot}(\sigma)$ sends a σ-structure $\mathcal{A}$ to the the commeasurable σ-structure where the commeasurability relation is the full relation, i.e. $\odot = A^2$. Thus, if we take the relative-lifting of $\mathbb{E}_k$ by $G\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}^{\odot}(\sigma)$, then the coKleisli category of the relative comonad $\mathbb{E}_k^{\odot} \circ G$ over G is essentially the same as the coKleisli category of $\mathbb{E}_k$. Adding the full relation, to the signature, as we did via the functor $G\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}^{\odot}(\sigma)$, has precedence in the study of quantum constraint satisfaction. Namely, adding the full relation was required to ensure closure under primitive positive reductions when investigating the quantum analogue of Schaefer's dichotomy theorem on Boolean-valued constraint satisfaction problems [18].

We now adapt the definition of the graded quantum monad to produce a $\mathbb{N}^{+}$-graded monad $(\{\mathbf{Q}_d^{\odot}\}, \eta, \{\mu^{d,d'}\})$ on commeasurable σ-structures $\mathsf{Struct}^{\odot}(\sigma)$. For a commeasurable σ-structure $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}^{\odot}(\sigma))$ with universe $A \in \mathbf{Ob}(\mathbf{RGraph})$, the universe

$\mathbf{Q}^{\odot}_d(A)$ of $\mathbf{Q}^{\odot}_d(\mathcal{A})$ has the same underlying set as $\mathbf{Q}_d(A)$, i.e. all d-dimensional PVMs over A, and commeasurability relation defined as

$$\begin{aligned} &\vartheta \odot \vartheta' \Leftrightarrow \vartheta \text{ and } \vartheta' \text{ are jointly measurable.} \\ &\qquad\qquad \text{and if } a \odot a' \text{ does not hold, then } \vartheta(a)\vartheta(a') = \mathbf{0}_d. \end{aligned}$$

For the other relations $R \in \sigma$, the interpretation $R^{\mathbf{Q}^{\odot}_d(\mathcal{A})}$ is the same set as $R^{\mathbf{Q}_d(\mathcal{A})}$.

With these definitions of $\mathbb{E}^{\odot}_k$ and $\mathbf{Q}^{\odot}_d$ in place, we are now able to officially define the graded mixed distributive law of $\mathbb{E}^{\odot}_k$ over $\mathbf{Q}^{\odot}_d$ as the family of natural transformations $\kappa^d \colon \mathbb{E}^{\odot}_k\mathbf{Q}^{\odot}_d \to \mathbf{Q}^{\odot}_d\mathbb{E}^{\odot}_k$ with components defined as

$$\kappa^d_{\mathcal{A}}([\vartheta_1, \ldots, \vartheta_n]) = \lambda[a_1, \ldots, a_m]. \begin{cases} \prod \vartheta_i(a_i) & \text{if } m = n \\ \mathbf{0}_d & \text{otherwise} \end{cases}$$

Which we can express as the formal sum:

$$\kappa^d_{\mathcal{A}}([\vartheta_1, \ldots, \vartheta_n]) = \sum_{[a_1,\ldots,a_n] \in \mathbb{E}^{\odot}_k(\mathcal{A})} \left(\prod_{i \in [n]} \vartheta_i(a_i) \right) .[a_1, \ldots, a_n]. \tag{6.25}$$

Note the product in this definition of κ is well-defined precisely because $\mathbb{E}^{\odot}_k$ was tailored to only consists of lists which are pairwise commeasurable.

We first verify that, for each $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}^{\odot}(\sigma))$ and $d > 0$, $\kappa^d_{\mathcal{A}}$ defined by equation (6.25) is indeed a morphism in $\mathsf{Struct}^{\odot}(\sigma)$.

Proposition 6.3.10. *For every* $\mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}^{\odot}(\sigma))$ *and* $d > 0$*, the function* $\kappa^{\mathcal{A}}_d$ *is a* $\mathsf{Struct}^{\odot}(\sigma)$*-morphism.*

Proof. In the following, given a $\boldsymbol{\vartheta} = [\vartheta_1, \ldots, \vartheta_n] \in \mathbb{E}^{\odot}_k(\mathbf{Q}^{\odot}_d(\mathcal{A})$ and $s = [a_1, \ldots, a_n] \in \mathbb{E}^{\odot}_k(\mathcal{A})$, we use the shorthand notation $\boldsymbol{\vartheta(s)} = \prod_{i\in[n]} \vartheta_i(a_i)$.

Suppose $R \in \sigma$ is a m-ary relation and $(\boldsymbol{\vartheta^1}, \ldots, \boldsymbol{\vartheta^m}) \in R^{\mathbb{E}^{\odot}_k(\mathbf{Q}^{\odot}_d(\mathcal{A}))}$ where assume for all $j \in [m]$, $\boldsymbol{\vartheta^j} = [\vartheta^j_1, \ldots, \vartheta^j_{n(j)}]$ for some $n(j) \leq k$. We need to show that $(\kappa^d(\boldsymbol{\vartheta^1}), \ldots, \kappa^d(\boldsymbol{\vartheta^m})) \in R^{\mathbf{Q}^{\odot}_d(\mathbb{E}^{\odot}_k(\mathcal{A}))}$.

To verify the joint measurability condition in the definition of $R^{\mathbf{Q}^{\odot}_d(\mathbb{E}^{\odot}_k(\mathcal{A}))}$, we first observe that all of the elements $\boldsymbol{\vartheta^1}, \ldots, \boldsymbol{\vartheta^m}$ are pairwise commeasurable, i.e. prefix comparable, for all $i, j \in [m]$, $\boldsymbol{\vartheta^j} \sqsubseteq \boldsymbol{\vartheta^i}$ or $\boldsymbol{\vartheta^i} \sqsubseteq \boldsymbol{\vartheta^j}$. Without loss of generality, assume $\boldsymbol{\vartheta^i} \sqsubseteq \boldsymbol{\vartheta^j}$. Since $\boldsymbol{\vartheta^j}$ is a list of pairwise commeasurable elements in $\mathbf{Q}^{\odot}_d(\mathcal{A})$ and $\boldsymbol{\vartheta^i} \sqsubseteq \boldsymbol{\vartheta^j}$, each of the PVMs $\vartheta^i_w \in [\boldsymbol{\vartheta^i}]$ is jointly measurable with every PVM $\vartheta^j_z \in [\boldsymbol{\vartheta^j}]$. By the definition of jointly measurable, for all $a, a' \in A$, $\vartheta^i_w \in [\boldsymbol{\vartheta^i}]$, and $\vartheta^j_z \in [\boldsymbol{\vartheta^j}]$, $\vartheta^i_w(a)$ commutes with $\vartheta^j_z(a')$.

It follows that for every $s, s' \in \mathbb{E}_k(\mathcal{A})$, $\boldsymbol{\vartheta^i(s)}$ commutes with $\boldsymbol{\vartheta^j(s')}$, so $\kappa^d(\boldsymbol{\vartheta^i})$ is jointly measurable with $\kappa^d(\boldsymbol{\vartheta^j})$.

To verify the compatibility condition in the definition of $R^{\mathbf{Q}_d^{\odot}(\mathbb{E}_k^{\odot}(\mathcal{A}))}$, suppose $\vec{s} = (s_1, \dots, s_m) \notin \mathbb{E}_k^{\odot}(\mathcal{A})$. We have to show that $K(\vec{s}) = \prod_{i \in [m]} \kappa^d(\boldsymbol{\vartheta^i})(s_i) = \mathbf{0}$. In this verification, as with the joint measurability verification, we assume without loss of generality $\boldsymbol{\vartheta^i} \sqsubseteq \boldsymbol{\vartheta^j}$. There are three cases we consider:

(1) There exists $i \in [m]$, $|s_i| \neq |\boldsymbol{\vartheta^i}|$, then by definition of κ^d, $\kappa^d(\boldsymbol{\vartheta^i})(s_i) = \mathbf{0}_d$. Thus, in this case, the whole product $K(\vec{s}) = \mathbf{0}_d$.

Thus, we can assume the next two cases that for all $i \in [m]$, $|s_i| = |\boldsymbol{\vartheta^i}|$.

(2) There exists $i, j \in [m]$ where $s_i \not\sqsubseteq s_j$. It follows that there exists some least index c where $a_c \in [s_i]$ is the c-th element of s_i and $a_c \notin [s_j]$. Let a'_c be the c-th element of s_j which is not equal to a_c. By the assumption that $\boldsymbol{\vartheta^i} \sqsubseteq \boldsymbol{\vartheta^j}$, for all $z \leq n(i)$, $\vartheta^i_z = \vartheta^j_z$. In particular, $\vartheta^i_c = \vartheta^j_c$. Since the PVMs $\boldsymbol{\vartheta^1}, \dots, \boldsymbol{\vartheta^m}$ are pairwise jointly measurable, we can commute the terms of the product $K(\vec{s})$ so that $\vartheta^i_c(a_c)\vartheta^j_c(a'_c)$ are adjacent. Since $a_c \neq a'_c$ and projectors in PVMs are orthogonal, $\vartheta^i_c(a_c)\vartheta^j_c(a'_c) = \vartheta^i_c(a_c)\vartheta^i_c(a'_c) = \mathbf{0}_d$. Thus, the whole product $K(\vec{s}) = \mathbf{0}_d$.

(3) $(\varepsilon_{\mathcal{A}}(s_1), \dots, \varepsilon_{\mathcal{A}}(s_m)) \notin R^{\mathcal{A}}$. Since $(\boldsymbol{\vartheta^1}, \dots, \boldsymbol{\vartheta^m}) \in R^{\mathbb{E}_k^{\odot}(\mathbf{Q}_d^{\odot}(\mathcal{A}))}$, by the compatibility condition in the definition of $R^{\mathbb{E}_k^{\odot}(\mathbf{Q}_d^{\odot}(\mathcal{A}))}$, $(\vartheta^1_{n(1)}, \dots, \vartheta^m_{n(m)}) \in R^{\mathbf{Q}_d^{\odot}(\mathcal{A})}$. By the compatibility condition of $R^{\mathbf{Q}_d^{\odot}(\mathcal{A})}$ and $(\varepsilon_{\mathcal{A}}(s_1), \dots, \varepsilon_{\mathcal{A}}(s_m)) \notin R^{\mathcal{A}}$, we have that

$$\vartheta^1_{n(1)}(\varepsilon_{\mathcal{A}}(s_1)) \dots \vartheta^m_{n(m)}(\varepsilon_{\mathcal{A}}(s_m)) = \prod_{i \in [m]} \vartheta^i_{n(i)}(\varepsilon_{\mathcal{A}}(s_i)) = \mathbf{0}_d.$$

As with the previous case, by pairwise joint measurability of $\boldsymbol{\vartheta^1}, \dots, \boldsymbol{\vartheta^m}$, we can rearrange the product so that the above sequence appears and conclude that $K(\vec{s}) = \mathbf{0}_d$.

Thus, in all three cases, $K(\vec{s}) = \prod_{i \in [m]} \kappa^d(\boldsymbol{\vartheta^i})(s_i) = \mathbf{0}$ as desired. Since the joint measurability and compatibility conditions are satisfied, $(\kappa^d(\boldsymbol{\vartheta^1}), \dots, \kappa^d(\boldsymbol{\vartheta^m})) \in R^{\mathbf{Q}_d^{\odot}(\mathbb{E}_k^{\odot}(\mathcal{A}))}$ and κ^d is a σ-morphism.

Observe that the definition of $\odot$ for $\mathbf{Q}_d^{\odot}(\mathbb{E}_k(\mathcal{A}))$ is the same as any other relational interpretation. Thus, we can repeat the proof above in the case where $R = \odot$, to show that κ^d preserves the commeasurability relation. Thus, we can conclude that κ^d defines a morphism in $\mathsf{Struct}^{\odot}(\sigma)$. $\square$

We next verify that for every $d > 0$, the family of morphisms $\{\kappa^d_{\mathcal{A}} \mid \mathcal{A} \in \mathbf{Ob}(\mathsf{Struct}^{\odot}(\sigma))\}$ forms a natural transformation.

Proposition 6.3.11. *For every* $d > 0$, $\kappa^d \colon \mathbb{E}_k^{\odot} \circ \mathbf{Q}_d^{\odot} \to \mathbf{Q}_d^{\odot} \circ \mathbb{E}_k^{\odot}$ *is a natural transformation.*

Proof. The proof follows essentially the same lines as Proposition 6.3.3. Suppose $f \colon A \to B$ is a morphism in $\mathsf{Struct}^{\odot}(\sigma)$. We define for every $i \in [n]$,

$$\psi_i(b) = \mathbf{Q}_d^{\odot}(f)(\varphi_i)(b) \sum_{a \in f^{-1}(b)} \varphi_i(a) \in \mathbf{Q}_d^{\odot}(B). \tag{6.26}$$

To verify the κ^d naturality square of f, consider $s = [\varphi_1, \ldots, \varphi_n] \in \mathbb{E}_k^{\odot}\mathbf{Q}_d^{\odot}(A)$:

$$\begin{aligned}
& \mathbf{Q}_d^{\odot}(\mathbb{E}_k^{\odot}(f)) \circ \kappa_A^d(s) \\
= \; & \{ \text{ definition of } \kappa^d \text{ } \} \\
& \mathbf{Q}_d^{\odot}(\mathbb{E}_k^{\odot}(f)) \left(\sum_{[a_1,\ldots,a_n] \in \mathbb{E}_k^{\odot}(A)} (\prod_{i \in [n]} \varphi_i(a_i)).[a_1, \ldots, a_n] \right) \\
= \; & \{ \text{ } \mathbf{Q}_d^{\odot} \text{ on morphisms } \} \\
& \sum_{[a_1,\ldots,a_n] \in \mathbb{E}_k^{\odot}(A)} (\prod_{i \in [n]} \varphi_i(a_i)).\mathbb{E}_k^{\odot}(f)([a_1, \ldots, a_n]) \\
= \; & \{ \text{ } \mathbb{E}_k^{\odot} \text{ on morphisms } \} \\
& \sum_{[a_1,\ldots,a_n] \in \mathbb{E}_k^{\odot}(A)} (\prod_{i \in [n]} \varphi_i(a_i)).([f(a_1), \ldots, f(a_n)])
\end{aligned}$$

Using the distributive property of projectors in $\mathsf{Proj}(d)$, for every $i \in [n]$, we can inductively 'pull-out' all terms $[f(a_1), \ldots, f(a_i)]$ in the formal sum whose i-th element starts with $b_i = f(a_i)$.

$$\begin{aligned}
& \mathbf{Q}_d^{\odot}(\mathbb{E}_k^{\odot}(f)) \circ \kappa_A(s) \\
= \; & \{ \text{ distribution property of projectors } \} \\
& \sum_{[a_1,\ldots,a_n] \in \mathbb{E}_k^{\odot}(A)} \sum_{a_1 \in f^{-1}(b_1)} \varphi_1(a_1) \left(\prod_{2 \le j \le n} \varphi_j(a_j).[f(a_1), \ldots, f(a_n)] \right) \\
= \; & \{ \text{ inductively apply distribution property } \} \\
& \sum_{[a_1,\ldots,a_n] \in \mathbb{E}_k^{\odot}(A)} (\sum_{a_1 \in f^{-1}(b_1)} \varphi_1(a_1))(\sum_{a_2 \in f^{-1}(b_2)} \varphi_2(a_2)) \ldots (\sum_{a_n \in f^{-1}(b_n)} \varphi_n(a_n)).[f(a_1), \ldots, f(a_n)] \\
= \; & \{ \text{ equation (6.26) } \} \\
& \sum_{[a_1,\ldots,a_n] \in \mathbb{E}_k^{\odot}(A)} \psi_1(b_1) \ldots \psi_n(b_n).[f(a_1), \ldots, f(a_n)] \\
= \; & \{ \text{ support in image of } \mathbb{E}_k^{\odot}(f) \text{ } \} \\
& \sum_{[b_1,\ldots,b_n] \in \mathbf{Im}(\mathbb{E}_k^{\odot}(f))} \psi_1(b_1) \ldots \psi_n(b_n).[b_1, \ldots, b_n]
\end{aligned}$$

$$\begin{aligned}
&= \quad \{ \text{ product notation and } \mathbf{0} \text{ terms for } \mathbb{E}_k^{\odot}(B)\backslash\mathbf{Im}(\mathbb{E}_k^{\odot}(f)) \} \\
&\qquad \sum_{[b_1,\ldots,b_n]\in\mathbb{E}_k^{\odot}(B)} (\prod_{i\in[n]} \psi_i(b_i)).[b_1,\ldots,b_n] \\
&= \quad \{ \text{ definition of } \kappa_d \text{ (6.25) } \} \\
&\qquad \kappa_B^d([\psi_1,\ldots,\psi_n]) \\
&= \quad \{ \text{ equation (6.26) } \} \\
&\qquad \kappa_B^d([\mathbf{Q}_d^{\odot}(f)(\varphi_1),\ldots,\mathbf{Q}_d^{\odot}(f)(\varphi_n)]) \\
&= \quad \{ \ \mathbb{E}_k^{\odot} \text{ on morphisms } \} \\
&\qquad \kappa_B^d \circ \mathbb{E}_k^{\odot}(\mathbf{Q}_d^{\odot}(f))(s)
\end{aligned}$$

□

Finally, we check that the collection of natural transformations κ^d for every $d > 0$ satisfies the four diagrams of Definition 6.3.6 verifying that it is indeed a graded distributive law of comonad $\mathbb{E}_k^{\odot}$ over graded monad $\mathbf{Q}_d^{\odot}$.

Proposition 6.3.12. $\{\kappa^d \colon \mathbb{E}_k^{\odot}\mathbf{Q}_d^{\odot} \to \mathbf{Q}_d^{\odot}\mathbb{E}_k^{\odot}\}_{d\in\mathbb{N}^+}$ *is a graded mixed distributive law.*

Proof. We must show that for diagrams in for the unit axiom (6.21), counit axiom (6.22), multiplication axiom (6.23), and comultiplication axiom (6.24) in Definition 6.3.6 are satisfied.

Unit: To verify the unit axiom (6.21), consider $[a_1,\ldots,a_n] \in \mathbb{E}_k^{\odot}(\mathcal{A})$:

$$\begin{aligned}
&\kappa^1 \circ \mathbb{E}_k^{\odot}(\eta)([a_1,\ldots,a_n]) \\
= &\quad \{ \ \mathbb{E}_k^{\odot} \text{ on morphisms } \} \\
&\kappa^1([\eta(a_1),\ldots,\eta(a_n)]) \\
= &\quad \{ \text{ definition of } \eta \ \} \\
&\kappa^1([I_1.a_1,\ldots,I_1.a_n]) \\
= &\quad \{ \text{ definition of } \kappa \text{ and } \varphi_i = I_1 a_i \ \} \\
&(\prod_{i\in[n]} I_1).[a_1,\ldots,a_n] \\
= &\quad \{ \text{ identity axiom } \mathsf{Proj}(1) \ \} \\
&I_1.[a_1,\ldots,a_n] \\
= &\quad \{ \text{ definition of } \eta \ \} \\
&\eta([a_1,\ldots,a_n])
\end{aligned}$$

Counit: To verify the counit axiom (6.22), consider $[\varphi_1,\ldots,\varphi_n] \in \mathbb{E}_k^{\odot}(\mathbf{Q}_d^{\odot}(\mathcal{A}))$:

$$\mathbf{Q}_d^{\odot}(\varepsilon) \circ \kappa^d([\varphi_1,\ldots,\varphi_n])$$

$$
\begin{aligned}
&= \quad \{ \text{ definition of } \kappa^d \ \} \\
&\quad \mathbf{Q}_d^{\odot}(\varepsilon)\left(\sum_{[a_1,\ldots,a_n]\in\mathbb{E}_k^{\odot}(A)} (\prod_{i\in[n]} \varphi_i(a_i)).[a_1,\ldots,a_n] \right) \\
&= \quad \{ \ \mathbf{Q}_d^{\odot} \text{ on morphisms } \} \\
&\quad \sum_{[a_1,\ldots,a_n]\in\mathbb{E}_k^{\odot}(A)} (\prod_{i\in[n]} \varphi_i(a_i)).\varepsilon([a_1,\ldots,a_n]) \\
&= \quad \{ \text{ definition of } \varepsilon \ \} \\
&\quad \sum_{[a_1,\ldots,a_n]\in\mathbb{E}_k^{\odot}(A)} (\prod_{i\in[n]} \varphi_i(a_i)).a_n \\
&= \quad \{ \text{ distribution axiom } \mathsf{Proj}(d) \ \} \\
&\quad \sum_{a_n\in A}\left(\sum_{[a_1,\ldots,a_{n-1}]\in\mathbb{E}_k^{\odot}(A)} \prod_{i\in[\mathsf{n}-1]} \varphi_i(a_i)\right)\varphi_n(a_n).a_n \\
&= \quad \{ \text{ normalisation } \} \\
&\quad \sum_{a_n\in A} I_d\varphi_n(a_n).a_n \\
&= \quad \{ \text{ identity } I_d\in\mathsf{Proj}(d) \ \} \\
&\quad \sum_{a_n\in A} \varphi_n(a_n).a_n \\
&= \quad \{ \ \varphi_n \text{ as a formal sum } \} \\
&\quad \varphi_n \\
&= \quad \{ \text{ definition of } \varepsilon \ \} \\
&\quad \varepsilon[\varphi_1,\ldots,\varphi_n]
\end{aligned}
$$

Multiplication: For this verification, we will need to use the following the mixed product property of the tensor product $\otimes$:

$$(E\otimes F)(G\otimes H) = EG\otimes FH \tag{6.27}$$

To verify the multiplication axiom (6.23), consider $[\Psi_1,\ldots,\Psi_n]\in\mathbb{E}_k(\mathbf{Q}_d^{\odot}(\mathbf{Q}_{d'}^{\odot}(A))$:

$$
\begin{aligned}
&\mu_{\mathbb{E}_k^{\odot}(\mathcal{A})}\circ\mathbf{Q}_d^{\odot}\kappa^{d'}\circ\kappa^d_{\mathbf{Q}_{d'}^{\odot}}([\Psi_1,\ldots,\Psi_n]) \\
&= \quad \{ \quad \text{definition of } \kappa^d \quad \} \\
&\quad \mu_{\mathbb{E}_k^{\odot}(\mathcal{A})}\circ\mathbf{Q}_d^{\odot}\kappa^{d'}\left(\sum_{[\varphi_1,\ldots,\varphi_n]\in\mathbb{E}_k^{\odot}(\mathbf{Q}_{d'}^{\odot}(A))} (\prod_{i\in[n]} \Psi_i(\varphi_i)).[\varphi_1,\ldots,\varphi_n] \right) \\
&= \quad \{ \quad \mathbf{Q}_d^{\odot} \text{ on morphisms } \quad \} \\
&\quad \mu_{\mathbb{E}_k^{\odot}(\mathcal{A})}\left(\sum_{[\varphi_1,\ldots,\varphi_n]\in\mathbb{E}_k^{\odot}(\mathbf{Q}_{d'}^{\odot}(A))} (\prod_{i\in[n]} \Psi_i(\varphi_i)).\kappa^{d'}([\varphi_1,\ldots,\varphi_n]) \right)
\end{aligned}
$$

$$= \quad \{ \text{ definition of } \kappa^{d'} \ \}$$

$$\mu_{\mathbb{E}_k^{\odot}(\mathcal{A})}\left(\sum_{[\varphi_1,\ldots,\varphi_n]\in\mathbb{E}_k^{\odot}(\mathbf{Q}_{d'}^{\odot}(A))} (\prod_{i\in[n]} \Psi_i(\varphi_i)).(\sum_{[a_1,\ldots,a_n]\in\mathbb{E}_k^{\odot}(A)} (\prod_{i\in[n]} \varphi_i(a_i)).[a_1,\ldots,a_n]) \right)$$

$$= \quad \{ \text{ definition of } \mu \ \}$$

$$\sum_{[a_1,\ldots,a_n]\in\mathbb{E}_k^{\odot}(A)} \ \sum_{[\varphi_1,\ldots,\varphi_i]\in\mathbb{E}_k^{\odot}(\mathbf{Q}_d^{\odot}(A))} (\prod_{i\in[n]} \Psi_i(\varphi_i) \otimes \prod_{i\in[n]} \varphi_i(a_i)).[a_1,\ldots,a_n]$$

$$= \quad \{ \text{ equation (6.27) } \}$$

$$\sum_{[a_1,\ldots,a_n]\in\mathbb{E}_k^{\odot}(A)} \ \sum_{[\varphi_1,\ldots,\varphi_i]\in\mathbb{E}_k^{\odot}(\mathbf{Q}_d^{\odot}(A))} (\prod_{i\in[n]} \Psi_i(\varphi_i) \otimes \varphi_i(a_i)).[a_1,\ldots,a_n]$$

$$= \quad \{ \text{ definition of } \kappa^{dd'} \ \}$$

$$\kappa^{dd'}[\sum_{\varphi_1\in\mathbf{Q}_{d'}^{\odot}(A)} \ \sum_{a_1\in A} (\Psi_1 \otimes \varphi_1(a_1)).a_1, \ldots, \sum_{\varphi_n\in\mathbf{Q}_{d'}^{\odot}(A)} \ \sum_{a_n\in A} (\Psi_n \otimes \varphi_n(a_n)).a_n]$$

$$= \quad \{ \text{ definition of } \mu_{\mathcal{A}} \ \}$$

$$\kappa^{dd'}[\mu(\Psi_1),\ldots,\mu(\Psi_n)]$$

$$= \quad \{ \ \mathbb{E}_k^{\odot} \text{ on morphisms } \}$$

$$\kappa^{dd'} \circ \mathbb{E}_k^{\odot}(\mu)[\Psi_1,\ldots,\Psi_n]$$

Comultiplication: To verify the comultiplication axiom (6.24), consider $[\varphi_1,\ldots,\varphi_n] \in \mathbb{E}_k^{\odot}(\mathbf{Q}_d^{\odot}(A))$:

$$\begin{aligned} \kappa^d_{\mathbf{Q}_d^{\odot}} \circ \mathbb{E}_k^{\odot}(\kappa^d) \circ \delta_{\mathbf{Q}_d^{\odot}}([\varphi_1,\ldots,\varphi_n]) &= \kappa^d_{\mathbf{Q}_d^{\odot}} \circ \mathbb{E}_k^{\odot}(\kappa^d)([[\varphi_1],[\varphi_1,\varphi_2],\ldots,[\varphi_1,\ldots,\varphi_n]]) \\ &= \kappa^d_{\mathbf{Q}_d^{\odot}}([\kappa^d([\varphi_1]),\ldots,\kappa^d([\varphi_1,\ldots,\varphi_n])]) \end{aligned}$$

Using the notation $\boldsymbol{\varphi_i} = [\varphi_1,\ldots,\varphi_i]$ and $\boldsymbol{\varphi_j}([a_1,\ldots,a_j]) = \prod_{i\in[j]} \varphi_i(a_i)$, we obtain the following from the definition of κ^d:

$$\begin{aligned} \kappa^d_{\mathbf{Q}_d^{\odot}} \circ \mathbb{E}_k^{\odot}\kappa^d \circ \delta_{\mathbf{Q}_d^{\odot}}([\varphi_1,\ldots,\varphi_n]) &= \kappa^d_{\mathbf{Q}_d^{\odot}}([\sum_{s=[a_1]} \boldsymbol{\varphi_1}(s).s,\ldots, \sum_{s=[a_1,\ldots,a_n]} \boldsymbol{\varphi_n}(s).s]) \\ &= \sum_{[s_1,\ldots,s_n]\in\mathbb{E}_k^{\odot}(\mathbb{E}_k^{\odot}(\mathcal{A}))} \prod(\boldsymbol{\varphi_i}(s_i)).[s_1,\ldots,s_n] \end{aligned}$$

Let $S_i = \{s_i = [a_1,\ldots,a_i] \in \mathbb{E}_k^{\odot}(\mathcal{A}) \mid \boldsymbol{\varphi_i}(s_i) \neq 0\}$ and let $T = \{[s_1,\ldots,s_n] \mid s_i \in S_i\}$. From the definition of $\boldsymbol{\varphi_i}$, every term where $[s_1,\ldots,s_n] \notin T$ is 0. Thus, we obtain that:

$$\kappa^d_{\mathbf{Q}_d^{\odot}} \circ \mathbb{E}_k^{\odot}(\kappa^d) \circ \delta_{\mathbf{Q}_d^{\odot}}([\varphi_1,\ldots,\varphi_n]) = \sum_{[s_1,\ldots,s_n]\in T} \prod(\boldsymbol{\varphi_i}(s_i)).[s_1,\ldots,s_n]$$

In T, there are two types of terms, *independent terms* $[s_1, \ldots, s_n]$ where for all $i \in [n] \backslash n$, $s_i \sqsubseteq s_{i+1}$, and *covariant terms* $[s_1, \ldots, s_n]$ where for some $i \in [n]$, $s_i \not\sqsubseteq s_{i+1}$. Let $T_x \subseteq T$ and $T_c \subseteq T$ denote the set of independent and covariant terms respectively. Using this notation, we obtain that:

$$\begin{aligned}\kappa^d_{\mathbf{Q}^{\odot}_d} \circ \mathbb{E}^{\odot}_k(\kappa^d) \circ \delta_{\mathbf{Q}^{\odot}_d}([\varphi_1, \ldots, \varphi_n]) = &\sum_{[s_1,\ldots,s_n] \in T_x} \prod (\boldsymbol{\varphi_i}(s_i)).[s_1, \ldots, s_n] \\ &+ \sum_{[s_1,\ldots,s_n] \in T_c} \prod (\boldsymbol{\varphi_i}(s_i)).[s_1, \ldots, s_n]\end{aligned}$$

For each element in $c = [s_1, \ldots, s_n] \in T_c$, let i_c be the least index such that $s_{i_c} \not\sqsubseteq s_{i_c+1}$. Using the fact that projectors in $\varphi_1, \ldots, \varphi_n$ are commeasurable, we can rearrange the terms in product $\prod \boldsymbol{\varphi_i}(s_i)$, such that the term $\varphi_{i_c}(a_{i_c})\varphi_{i_c+1}(a'_{i_c})$, where a_{i_c} is the last element of s_{i_c} and a'_{i_c} is the penultimate element of s_{i_c+1}, appears in the resulting product. Since $a_{i_c} \neq a'_{i_c}$ and $\varphi_{i_c}(a_{i_c})$, $\varphi_{i_c+1}(a'_{i_c})$ are different projectors of the same PVM φ_{i_c}, we have that, $\varphi_{i_c}(a_{i_c})\varphi_{i_c+1}(a'_{i_c}) = \mathbf{0}$, and thus the product $\prod \boldsymbol{\varphi_i}(s_i)$ is 0 for every $[s_1, \ldots, s_n] \in T_c$. Thus, we obtain that:

$$\kappa^d_{\mathbf{Q}^{\odot}_d} \circ \mathbb{E}^{\odot}_k(\kappa^d) \circ \delta_{\mathbf{Q}^{\odot}_d}([\varphi_1, \ldots, \varphi_n]) = \sum_{[s_1,\ldots,s_n] \in T_x} \prod (\boldsymbol{\varphi_i}(s_i)).[s_1, \ldots, s_n]$$

Using the fact that the projectors in the PVMs $\varphi_1, \ldots, \varphi_n$ are commeasurable and that each $t \in T_x$ is of the form $t = \delta([a_1, \ldots, a_n])$ where for all $i \in [n]$, $\varphi_i(a_i) \neq 0$, we can rearrange the terms in the sum over T_x and obtain that:

$$\begin{aligned}&\kappa^d_{\mathbf{Q}^{\odot}_d} \circ \mathbb{E}^{\odot}_k(\kappa^d) \circ \delta_{\mathbf{Q}^{\odot}_d}([\varphi_1, \ldots, \varphi_n]) \\ &\quad = \sum_{[s_1,\ldots,s_n] \in T_x} \prod_{i \in [n]} (\varphi_i(a_i))^2.[s_1, \ldots, s_n] \\ &\quad = \sum_{[a_1,\ldots,a_n] \in S_n} \prod_{i \in [n]} (\varphi_i(a_i))^2.\delta([a_1, \ldots, a_n]) \\ &\quad = \sum_{[a_1,\ldots,a_n] \in S_n} \prod_{i \in [n]} \varphi_i(a_i).\delta([a_1, \ldots, a_n]) \\ &\quad = \mathbf{Q}^{\odot}_d(\delta) \left(\sum_{[a_1,\ldots,a_n] \in S_n} \prod_{i \in [n]} \varphi_i(a_i).[a_1, \ldots, a_n] \right) \\ &\quad = \mathbf{Q}^{\odot}_d(\delta) \circ \kappa^d([\varphi_1, \ldots, \varphi_n])\end{aligned}$$

□

Remark 6.3.13. The verifications that κ^d satisfies the counit and comultiplication diagrams in Proposition 6.3.12, in particular, establish that for every $d > 0$, there is a coKleisli law of $\mathbb{E}^{\odot}_k$ over $\mathbf{Q}^{\odot}_d$. Thus, we may apply the abstract morphism (Theorem 5.2.2) and isomorphism FVM theorem (Theorem 5.2.2) to demonstrate that $\mathbf{Q}_d$ preserves the $\Rrightarrow^{\mathbf{FO}_k}$ and $\equiv^{\mathbf{FO}_k}$ relations, respectively. This is not a direct application, but involves using the the relative-liftings of $\mathbb{E}_k^{\odot}, \mathbf{Q}_d^{\odot}$ along the functor $R\colon \mathsf{Struct}(\sigma) \to \mathsf{Struct}^{\odot}(\sigma)$.

With this distributive law in place, we can compose morphisms of type $\mathbb{E}_k\mathcal{A} \to \mathbf{Q}_d\mathcal{B}$ using the expression from biKleisli composition in equation (6.5). This is rather interesting since a biKleisli morphism $\mathbb{E}_k\mathcal{A} \to \mathbf{Q}_d\mathcal{B}$ can either be interpreted as a 'quantum strategy for Duplicator in ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$' or as 'quantum perfect strategy in a homomorphism game with Spoiler's contexts in ${}^{\exists^+}\mathbf{EF}_k(\mathcal{A}, \mathcal{B})$'. Finding a recipe for composing strategies across these two types of games, given their non-categorical definitions would likely have been difficult task without the use of (co)monads. Furthermore, as the no-go theorem from the previous section indicates, these similar compositions are by no means guaranteed. Exploring how the existence of this distributive law could be used to address questions in quantum finite model theory is outside the scope of this thesis, but will be explored in the forthcoming work [56] with my co-author Amin Karamlou.

Chapter 7

Conclusion

The paper *Relating Structure and Power* [13] demonstrated that the encoding of pebble games in finite model theory, as comonad over relational structures, first discovered by Abramsky, Dawar, and Wang in [5], was not simply coincidence, but one instance of the wider phenomena of model comparison game comonads. However, this paper did more than just provide an impetus for the game comonads research program, it also articulated and cast the program as a part of a wider vision of unifying two research strands within Logic in Computer Science:

Power: This strand uses logic to resolve questions of expressive power and algorithmic complexity in databases, constraint satisfaction, artificial intelligence, and computer-assisted verification.

Structure: This strand studies semantics and compositionality in logic to facilitate the creation of bug-free, correct, and modular software systems.

The goal of unification is to reap benefits to both strands.

For the Power strand, structural tools provide a perspective that organises, generalises, and reformulates common arguments employed across different concrete problems. This birds-eye view opens avenues to discover new variations of old results and new instances of common constructions.

For the Structure strand, tailoring structural abstractions to these concrete problems provide new challenges and motivate questions about these abstractions that drive their further development.

Model comparison game comonads are instance of this vision by drawing a connection between (finite) model theory, exemplifying power, and categorical semantics, exemplifying structure. It is my belief that the contributions documented in this thesis provide evidence that both these benefits have emerged from the game comonads program.

7.1 Summary of work

Second-order logics via relative-liftings In Section 3.3 of Chapter 3, we demonstrated how a relative-lifting of the Ehrenfeucht-Fraïssé comonad could be used to capture equivalences in logics with equality. This technique of relative-lifting is not limited to providing an interpretation for equality but for any interpretation that is functorial. In particular, we showed how use relative-lifting to provide interpretations for first-order logic with a connectivity predicate. By upgrading the Ehrenfeucht-Fraïssé comonad to the multi-sorted setting, we used this technique to obtain a relative comonad which captured equivalence in monadic second-order logic.

Linear arboreal covers In Chapter 4, we strengthened the axioms of arboreal categories to exclude 'branching' behaviours. We also showed that the any arboreal category $\mathcal{C}$ satisfying a strong connectedness condition has a linear arboreal subcategory $\mathcal{C}^L$. The linear arboreal subcategory $\mathcal{C}^L$ is related to $\mathcal{C}$ via right adjoint $T\colon \mathcal{C}^L \to \mathcal{C}$. Utilising this right adjoint, we demonstrated how to derive a linear arboreal cover from an arboreal cover, and thus obtain linear variants of game comonads. These linear variants of game comonads 'linearise' the branching behavioural relations native to $\mathcal{C}$. This axiomatic work motivated by generalising the passage from the pebbling comonad $\mathbb{P}_k$ of to its linear variant $\mathbb{P}_k^L$, i.e. the pebble-relation comonad from [68]. Through this example of the linear pebble comonad $\mathbb{P}_k^L$, we showed how these linear behavioural relations are related to all-in-one variants of the model comparison games. These all-in-one games capture restricted conjunction fragments of $\exists^+\mathcal{L}^k$ and $\exists\mathcal{L}^k$, and the restriction conjunction with global counting fragment of $\#\mathcal{L}^k$. Moreover, we showed that just as $\mathbb{P}_k$-coalgebras encode the data of a tree decomposition of width $< k$, $\mathbb{P}_k^L$-coalgebras encode the data of a path decomposition of width $< k$. Thus, yielding a coalgebraic definition of pathwidth. Through the example of the linear modal comonad $\mathbb{M}_k^L$, we showed how the 'branching' relations of simulation, property-preserving simulation, bisimulation, and graded counting bisimulation can be uniformly 'linearised' into trace inclusion, labelled trace inclusion, labelled trace equivalence, and bijective trace equivalence.

FVM theorems via coKleisli laws In Chapter 5, we investigated given (products of) game comonads $\mathbb{S}, \mathbb{T}$, how distributive laws of type $\kappa\colon \mathbb{S}F \to F\mathbb{T}$ reproduce FVM theorems for functorial operations F. First, we proved that merely having a 'discrete' natural transformation of type $\kappa\colon \mathbb{T}F \to F\mathbb{S}$ is sufficient to obtain a FVM theorem for the positive-existential fragments $\exists^{+}\mathcal{J}$ of logics $\mathcal{J}$ captured by arboreal covers associated to game comonads in the product $\mathbb{S} = \prod_{i \in I} \mathbb{S}_i$ and $\mathbb{T}$. Second, we proved how if κ satisfied the axioms of coKleisli law then this was sufficient to obtain a FVM theorem for extensions of $\mathcal{J}$ by counting quantifiers $\#\mathcal{J}$. We discussed the conditions in [52] that κ and F must satisfy in order to obtain a FVM theorem for the existential $\exists\mathcal{J}$ and full fragments $\mathcal{J}$. These abstract FVM theorems were illustrated by recovering classical FVM theorems for coproduct operations. We also applied these theorems to the algebraic operations used to build structures of bounded clique-width. In the case where F was the identity functor on either $\mathsf{Struct}(\sigma)$ or $\mathsf{Struct}_{\star}(\sigma)$, we semantically recovered translations between the logics associated with game comonad $\mathbb{T}$ to the logic associated with game comonad $\mathbb{S}$.

Comonad-monad distributive laws, negative and positive In Chapter 6, we continued our investigation of distributive laws by focusing on the case where the comonads $\mathbb{C} = \mathbb{D} = W$ were the same and F was the underlying functor M of a monad (M, η, μ) satisfying additional coherence axioms with the unit η and multiplication μ of M. This investigation was motivated by the question of whether there existed a graded comonad-monad distributive law $\kappa\colon \mathbb{E}_k\mathbf{Q}_d \to \mathbf{Q}_d\mathbb{E}_k$ where $\mathbf{Q}_d$ is the quantum monad over $\mathsf{Struct}(\sigma)$ from [4]. The search for such a distributive law lead us to prove a general no-go theorem demonstrating that there cannot exist any comonad-monad distributive law $\rho\colon WM \to MW$ for wide class of comonads $W \in \mathcal{C}_W$ and monads $M \in \mathcal{C}_M$ over $\mathbf{Set}$. The class of $\mathcal{C}_W$ is all directed containers that are not coreader comonads, or equivalently, all polynomial comonads which are not linear. The class of $\mathcal{C}_M$ is all monads which contain a 'uniform distribution' of size ≥ 2. As a corollary, we showed there are no distributive law of the k-round EF or k-pebble comonads over the extension of the discrete probability distribution monad to $\mathsf{Struct}(\sigma)$ from Adam Connolly's PhD thesis [26]. By contrast, we exhibited a distributive law $\kappa\colon \mathbb{E}_k^{\odot}\mathbf{Q}_d^{\odot} \to \mathbf{Q}_d^{\odot}\mathbb{E}_k^{\odot}$ for versions of $\mathbb{E}_k^{\odot}$ and $\mathbf{Q}_d^{\odot}$ tailored to σ-structures with a reflexive, symmetric commeasurability relation. These distributive laws were based on a distributive law of the prefix list comonad $\kappa\colon N^{+}\mathcal{D}_{\mathcal{E}} \to \mathcal{D}_{\mathcal{E}}N^{+}$ over the distribution monad of the 'single measurement semiring' $\mathcal{E}$. This illustrated that the one of the key properties which enables a distributive law of $\mathbb{E}_k^{\odot}$ over $\mathbf{Q}_d^{\odot}$, in contrast to its classical cousin $\mathcal{D}_{\mathbb{R}_{\geq 0}}$, is the arithmetic of projector-valued measurements involved in the definition of $\mathbf{Q}_d^{\odot}$.

7.2 Further directions

We conclude this thesis by mentioning important questions that arose from this research and suggest directions of further inquiry.

Categorical definition of branch bijective, injective, surjective games In Section 3.1.2 about the abstract games from [10], we exhibited the games ${}^{\exists^+}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, ${}^{\exists}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, and $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ as instances of abstract path matching games ${}^{\exists^+}\mathcal{G}(X, Y)$, ${}^{\exists}\mathcal{G}(X, Y)$, and $\mathcal{G}(X, Y)$ which were defined for any pair of objects X, Y in an arboreal category $\mathcal{C}$. We also showed that ${}^{\#}\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is equivalent to isomorphism between free coalgebras $R_k(\mathcal{A}), R_k(\mathcal{B})$ in $\mathcal{R}^P_k(\sigma)$. Can we define an abstract bijective game ${}^{\#}\mathcal{G}(X, Y)$ which provides an 'operational' definition of isomorphism in any arboreal category $\mathcal{C}$ similar to how ${}^{\exists^+}\mathcal{G}(X, Y)$ provides a definition of morphism? One key restriction of bijective games is that only coincide with equivalence in a logic extended with counting quantifiers if the structures are finite. Does this mean that ${}^{\#}\mathcal{G}(X, Y)$ can only be defined on locally-finitely presentable (lfp) arboreal categories as in [76]? Isomorphism in locally-finitely presentable categories, where each object is a filtered colimit of compact/finite objects, are characterised in terms of so-called 'confluent sequences' between the filtered cocones that present an object [66]. The definition of confluent sequence was inspired by back-and-forth model comparison games. Could a Duplicator winning strategy in an abstract bijective game ${}^{\#}\mathcal{G}(X, Y)$ for an lfp arboreal category $\mathcal{C}$ be described as a confluent sequence between the path embedding cocones that generate X and Y in $\mathcal{C}$? Furthermore, in the paper [25], and in Adam Connolly's PhD thesis [26], for every $n, k > 0$, the 'Hella' game comonad $\mathbb{H}_{k,n}$ was defined so that isomorphism in the Kleisli category characterised equivalence in k-variable logic extended with all n-ary generalised quantifiers $\mathcal{L}^k(\mathcal{Q}_n)$. The isomorphism power theorem for $\mathbb{H}_{n,k}$ was proved via looking at a n-ary extended version of the bijective k-pebble game. As an instance of the investigation in [25], variants of this game, where f was weakened to be surjection or injection, were shown to characterise the at least counting quantifier $\exists_{\geq n}$ fragment of $\#\mathcal{L}^k$ and universal fragment of $\mathcal{L}^k$. Can we similarly define the abstract branch-injective and branch-surjective games in $\mathcal{C}$? What are the morphisms or constructions, e.g. pathwise embeddings, spans, etc. in an arboreal category $\mathcal{C}$ which correspond to these games?

Relative comonads which are not relative-liftings In Section 3.3 of Chapter 3, we demonstrated how lifting a game comonad $\mathbb{C}$ to a relative comonad $\mathbb{C} \circ J$ for a functor J enables the study of equivalence in logics where symbols carry with them the intended interpretation picked out by the functor J. However, not every relative comonad defined as in Definition 2.3.1 is obtained by lifting of an ordinary comonad. Are there relative comonads which are *not* obtained through a relative-lifting of game comonad that could be used to characterise equivalence in other logics?

Other (weak) adjoint constructions to capture different equivalences In Section 4.1.1, we constructed an adjunction $I \dashv T$ from an arboreal category $\mathcal{C}$ whose objects are strongly connected to linear arboreal subcategory $\mathcal{C}^L$ whose objects and morphisms exclude 'branching behaviour'. We saw how that generalised forms of simulation, property-preserving simulation, bisimulation and counting bisimulation native to $\mathcal{C}$ were transformed into their linear variants trace inclusion, labelled trace inclusion, labelled trace equivalence, and bijective trace equivalence in $\mathcal{C}^L$. Further, in the case where only restricted to objects without branching behaviour, we obtained a weak adjoint $I^q \dashv T^q$ from an arboreal category to a quasi-linear arboreal category $\mathcal{C}^{qL}$. Illustrated with the example 4.2.3, we also obtained a categorical characterisation of complete trace equivalence. This reveals that by composing an arboreal cover $\mathcal{E}$ by $\mathcal{C}$ with (weak) adjoint functors on an arboreal category $\mathcal{C}$, we can obtain additional behavioural equivalences to the ones in Definition 3.1.19. What other (weak) adjoint functors on $\mathcal{C}$ exist that capture additional equivalences in fragments of the logic $\mathcal{J}$ associated to $\mathcal{C}$?

FVM theorems for operations with structures on indexing sets In Chapter 5, we showed how considering coKleisli laws of game comonads over functors could recover FVM theorems and logical translations. FVM theorems demonstrate when a operation combining multiple structures preserves logical equivalence or maps equivalence in some logic to equivalence in another logic. However, the original results of Feferman and Vaught in [35] were more general by involving operations that had structure on the indices of the input to an operation. Chapter 5 only recovered the case where the indices are a set, i.e. have no structure. How do we recover these more general Feferman-Vaught-Mostowski theorems?

No-go theorems for probability monads In Section 6.2.2, we proved that for any non-linear polynomial **Set**-comonad, i.e. directed container which is not a coreader comonad, does not have a distributive law over any **Set**-monad with a meaningful notion of 'uniform sampling'. The primary example of such a monad is the discrete probability distribution monad $\mathcal{D}_{\mathbb{R}_{\geq 0}}$. However, $\mathcal{D}_{\mathbb{R}_{\geq 0}}$ is only the most rudimentary example of a 'probability monad' investigated in probabilistic programming semantics [39, 46]. Typically, probability monads are defined over categories where objects have some amount of topological structure. For instance, the Giry monad over measure spaces [42], the

Radon monad over compact Hausdorff spaces, and the Borel measure monad on topological spaces [39]. On the other hand, directed containers can be defined over a wide class of categories [1]. Is there way to extend our no-go theorem to categories equipped with a probability monad and have sufficient structure to define directed containers? The Kleisli category for each of these probability monads are Markov categories [38]. In Markov categories, objects are equipped with a comonoidal structure. In particular, this means that each object X has an associated $\mathbf{copy}_X \colon X \to X \otimes X$ morphism. Deterministic morphisms $f \colon X \to Y$, i.e. morphisms that do not exhibit any probabilistic behaviours, are then axiomatised as morphisms which commute with the copy morphism, i.e. $\mathbf{copy}_Y \circ f = (f \otimes f) \circ \mathbf{copy}_X$. In our proof of the no-go theorem, the failure of the comultiplication axiom for a distributive law of N^+ over $\mathcal{D}_{\mathbb{R}_{\geq 0}}$ amounts to a failure of the uniform distribution on two elements to commute with the comultiplication for any polynomial comonad which involves 'copying', i.e. is non-linear. Can the characterisation of non-deterministic morphisms in Markov categories as not commuting with copy maps be used to generalise our no-go theorem to other probability monads?

Decidable first-order logic via structure theory of arboreal covers One of the first theorems during the emergence of mathematical logic and computer science in the early 20th century proved that the satisfiability problem for first-order logic is undecidable. By contrast, the satisfiability problem for the much less expressive modal logic, which we can consider as a fragment of first-order logic via a standard translation, is decidable. Thus, since the field's inception, practitioners have been searching for fragments of first-order logic which are decidable, but more expressive than the modal fragment. In fact, the holy grail of this research program is to characterise the most expressive fragment of first-order logic with a decidable satisfiability problem. Much of this work has been guided by proving decidable satisfiability problem for fragments which generalise features of the modal fragment, for instance the two-variable fragment, guarded fragments, and variable ordered fragments. These fragments seem to correlate with arboreal covers which exhibit stronger structural properties, e.g. the arboreal cover capturing modal logic is a coreflection and the arboreal cover capturing guarded logic satisfies the bisimilar companion property articulated in [11]. By contrast, arboreal covers which do not satisfy this property seem to correlate to expressive, but undecidable, fragments. Indeed, as Section 6 of the survey paper [2] formulates, we should be able to use the structural properties of arboreal covers to classify the expressive power of logical fragments. Could a "tame-vs-wild" structure theory of arboreal covers shed light on characterising the most expressive fragment of first-order logic which has decidable satisfiability problem? As first step towards this connection, we note that one of the techniques to show that a first-order fragment is decidable is to prove that the fragment has a so-called ‘tree-model property’. Can the tree-model property be formalised abstractly in terms of arboreal categories and covers?

Appendix A

Multi-sorted model theory

Definition A.0.1. A *S-sorted binary relation* γ from A to B is an S-indexed family of subsets $\gamma_s \subseteq A_s \times B_s$

Definition A.0.2. Let $I \subseteq S$. A S-sorted binary relation γ from A to B is a *partial function on* I if for all $s \in I$ and $a \in A_s$, $(a, b_1), (a, b_2) \in \gamma_s \Rightarrow b_1 = b_2$.

Definition A.0.3. Given a S-sorted binary relation from A to B, the *converse relation* γ^c is a S-sorted binary relation from B to A such that:

$$(a, b) \in \gamma_s \Leftrightarrow (b, a) \in \gamma_s^c \subseteq B_s \times A_s$$

Definition A.0.4. An element $\vec{a} = (a_1{:}s_1, \dots, a_n{:}s_n) \in A_{s_1} \times \dots \times A_{s_n}$ is said to be *A-vector of sort type* $\langle s_1, \dots, s_n\rangle$.

Definition A.0.5. A S-sorted signature σ is a set of relational symbols $R \in \sigma$, each with an associated arity given by a finite tuple $\langle s_1, \dots, s_m\rangle$ of elements in S.

Definition A.0.6. An *S-sorted σ-structure* $\mathcal{A}$ has universe given by an S-sorted set A, and interpretations

$$R^{\mathcal{A}} \subseteq A_{s_1} \times \dots \times A_{s_m}$$

for every relational symbol $R \in \sigma$ with arity $\langle s_1, \dots, s_m\rangle$.

Similarly to the single-sorted case, a σ-homomorphism is given by relation preserving S-sorted function.

Definition A.0.7. An *S-sorted σ-morphism* $h\colon \mathcal{A} \to \mathcal{B}$ is a S-sorted function $h\colon A \to B$ such that for every relational symbol $R \in \sigma$ of arity $\langle s_1, \dots, s_m\rangle$,

$$R^{\mathcal{A}}(a_1{:}s_1, \dots, a_m{:}s_m) \Rightarrow R^{\mathcal{B}}(h(a_1{:}s_1), \dots, h(a_m{:}s_m))$$

for all $a_1{:}s_1, \dots, a_m{:}s_m \in A$.

We will write $\mathsf{Struct}^S(\sigma)$ for the category S-sorted σ-structures and σ-morphisms.

The syntax of S-sorted first-order logic with equality on I in relational signature σ, $\mathbf{FO}^{I,S}(\sigma)$ can be defined recursively. We present this syntax in pseudo-BNF form. We assume, for every sort $s \in S$, we have an infinite supply of variables $x_1{:}s, x_2{:}s, \ldots, y_1{:}s, y_2{:}s, \ldots$, etc. The assumption is that when instantiated over S-sorted universe A, $x{:}s$ varies over the set A_s. For every vector of variables $\vec{x} = (x_1{:}s_1, \ldots, x_n{:}s_n)$, we define formulas in $\mathcal{J}(\vec{x})$ recursively:

$$\begin{aligned}
\phi(\vec{x}) ::= {} & t_1{:}s = t_2{:}s && (s \in I) \\
& \mid R(t_1{:}s_1, \ldots, t_r{:}s_r) && (R \text{ of arity } \langle s_1, \ldots, s_r\rangle) \\
& \mid \neg\phi(\vec{x}) \\
& \mid \phi_1(\vec{x}) \wedge \phi_2(\vec{x}) \;\mid\; \phi_1(\vec{x}) \vee \phi_2(\vec{x}) \\
& \mid \forall y{:}s\phi(y{:}s, \vec{x}) \;\mid\; \exists y{:}s\phi(y{:}s, \vec{x})
\end{aligned}$$

Finally, in order to define the winning conditions for the model comparison games in this multi-sorted setting, we will need the following definitions.

Definition A.0.8. Given S-sorted σ-structures $\mathcal{A}, \mathcal{B}$ with universe A, B (respectively), we say that S-sorted binary relation γ from A to B is a *σ-correspondence* if, for all $R \in \sigma$ with arity $\langle s_1, \ldots, s_n\rangle$,

$$R^{\mathcal{A}}(a_1, \ldots, a_n) \Rightarrow R^{\mathcal{B}}(b_1, \ldots, b_n)$$

whenever $(a_1, b_1) \in \gamma_{s_1}, \ldots, (a_n, b_n) \in \gamma_{s_n}$.

Definition A.0.9. If γ is a σ-correspondence, from $\mathcal{A}$ to $\mathcal{B}$, and γ^c is a σ-correspondence from $\mathcal{B}$ to $\mathcal{A}$, then γ is a *tight σ-correspondence from $\mathcal{A}$ to $\mathcal{B}$*

Definition A.0.10. Given a subset $I \subseteq S$ of sorts, a σ-correspondence γ from $\mathcal{A}$ to $\mathcal{B}$ is also a partial function on I is a *partial σ-morphism on I.*

Definition A.0.11. If γ and γ^c are partial σ-morphisms on I, then γ is a *partial σ-isomorphism on I.*

Equivalently, γ is a partial σ-isomorphism on I if it is a tight σ-correspondence and γ_s is a partial injective function for all $s \in I \subseteq S$. We now prove the Ehrenfeucht-Fraïssé theorem for S-sorted first-order logic. We extending the notation from single-sorted σ-structures and let $(\mathcal{A}, \vec{a})$ denote a S-sorted σ-structure with a distinguished choice of parameters $\vec{a} = (a^1{:}s^1, \ldots, a^n{:}s^n)$ of sort type $\mathfrak{s} = \langle s^1, \ldots, s^n\rangle$. The structure $(\mathcal{A}, \vec{a})$ can be equivalently thought of as the σ-structure $\mathcal{A}$ in an extended signature with constants of sort type $\mathfrak{s}$. In order to obtain an easy inductive proof, we will prove the Ehrenfeucht-Fraïssé theorem for these σ-structures choice of parameters $\vec{a}$ which provide interpretations for free variables in a S-sorted formula . Given two vectors $\vec{a} = (a^1{:}s^1, \ldots, a^n{:}s^n)$ of $\vec{b} = (b^1{:}s^1, \ldots, b^n{:}s^n)$ of the same sort type in A and B, respectively, we use $\gamma(\vec{a}, \vec{b})$ to denote the S-sorted partial function such that $a^i \mapsto b^i$.

Definition A.0.12. Given S-sorted σ structures with distinguished vectors of the same sort type $(\mathcal{A}, \vec{a}), (\mathcal{B}, \vec{b})$, the *$k$-round S-sorted Ehrenfeucht-Fraïssé game* $\mathbf{EF}^I_k((\mathcal{A}, \vec{a}), (\mathcal{B}, \vec{b}))$ *with equality on I sorts* is played between two players, Spoiler and Duplicator. In each round $i \leq k$,

- Spoiler chooses a sort $s_i \in S$ and element $a_i \in A_{s_i}$ or $b_i \in B_{s_i}$.
- Duplicator responds with an element $b_i \in B_{s_i}$ or $a_i \in A_{s_i}$ in the other structure's s_i sort.

Duplicator wins round i, if the following S-relation γ_i, defined component-wise,

$$\gamma_{i,s} = \gamma(\vec{a}, \vec{b})_s \cup \{(a_j, b_j) \mid \text{ for } s_j = s \text{ and } j \leq i\}$$

is a partial isomorphism on I of S-sorted σ-structures from $\mathcal{A}$ to $\mathcal{B}$.

Theorem A.0.13. *The following are equivalent for all σ-structures $\mathcal{A}, \mathcal{B}$ with parameters $\vec{a}, \vec{b}$ of the same sort type $\mathfrak{s}$:*

(1) Duplicator has a winning strategy in $\mathbf{EF}_k((\mathcal{A}, \vec{a}), (\mathcal{B}, \vec{b}))$

(2) $(\mathcal{A}, \vec{a}) \equiv^{\mathbf{FO}_k} (\mathcal{B}, \vec{b})$

Proof. (1) $\Rightarrow$ (2) We prove the statement by induction on k. Suppose $k = 0$ and that Duplicator has a winning strategy in $\mathbf{EF}_0((\mathcal{A}, \vec{a}), (\mathcal{B}, \vec{b}))$. By the winning condition, the relation $\gamma = \gamma(\vec{a}, \vec{b})$ is a partial isomorphism. Hence, from the definition of partial isomorphism, $\vec{a}$ and $\vec{b}$ satisfy the same quantifier-free formulas in $\mathcal{A}$ and $\mathcal{B}$, respectively.

For the inductive step, suppose Duplicator has a winning strategy in $\mathbf{EF}_{k+1}((\mathcal{A}, \vec{a}), (\mathcal{B}, \vec{b}))$ and consider the sentence ϕ of rank $k+1$. It follows from the definition of quantifier rank, that every sentence ϕ of rank $k+1$ is a Boolean combination of sentences of the form $\exists y{:}s\psi(y{:}s)$ for sort $s \in S$ where ψ is a formula of rank $\leq k$. Hence, by recursion on the construction of formulas $\phi(\vec{x})$, it suffices to only prove that $(\mathcal{A}, \vec{a}) \vDash \phi(\vec{x}) \Leftrightarrow (\mathcal{B}, \vec{b}) \vDash \phi(\vec{x})$ for $\phi = \exists y{:}s\psi(y{:}s, \vec{x})$.

Suppose $(\mathcal{A}, \vec{a}) \vDash \exists y{:}s\psi(y{:}s, \vec{x})$, then there exists some $a \in A_s$, such that $(\mathcal{A}, a\vec{a}) \vDash \psi(y{:}s, \vec{x})$. Consider the $\mathbf{EF}_{k+1}((\mathcal{A}, \vec{a}), (\mathcal{B}, \vec{b}))$ game where Spoiler plays $s \in S$ and $a \in A_s$ in the first round. Let Duplicator's response be $b \in B_s$ according to her winning strategy. Duplicator then has a winning strategy in the $\mathbf{EF}_k((\mathcal{A}, a\vec{a}), (\mathcal{B}, b\vec{b}))$ game by playing according the subsequent k moves of the $\mathbf{EF}_{k+1}((\mathcal{A},\vec{a}),(\mathcal{B},\vec{b}))$ game. Hence, by the inductive hypothesis, we have that for all σ formulas $\phi'(y{:}s,\vec{x})$ of rank $\leq k$, $(\mathcal{A},a)\vDash\phi'(y{:}s,\vec{x})\Leftrightarrow(\mathcal{B},b)\vDash\phi'(y{:}s,\vec{x})$. In particular, $(\mathcal{B},b\vec{b})\vDash\psi(y{:}s,\vec{x})$. Hence, $\mathcal{B}\vDash\exists y{:}s\psi(y{:}s,\vec{x})$. A symmetric argument establishes that $\mathcal{B}\vDash\exists y{:}s\psi(y{:}s)\Rightarrow\mathcal{A}\vDash\exists y{:}s\psi(y{:}s)$.

(1) $\Rightarrow$ (2) We prove the statement by induction on k. For $k=0$, $\vec{a}$ and $\vec{b}$ satisfy the same atomic and negated atomic formulas in $\mathcal{A}$ and $\mathcal{B}$ proving that γ is a partial isomorphism.

For the inductive step, assume $(\mathcal{A},\vec{a})$ and $(\mathcal{B},\vec{b})$ satisfy the same formulas of rank $k+1$. We need to show that Duplicator has a winning strategy in the $\mathbf{EF}_{k+1}((\mathcal{A},\vec{a}),(\mathcal{B},\vec{b}))$ game. Suppose, in the first round of $k+1$, Spoiler chooses $s\in S$ and $a\in A_s$. By σ having finitely many relational symbols and k finite, there is finite collection of in-equivalent formulas $\varphi_1,\ldots,\varphi_M$ of rank k with variables $(y{:}s,x{:}s_1,\ldots,x{:}s_n)$ satisfied by $a\vec{a}$, i.e. $(\mathcal{A},a\vec{a})\vDash\varphi_i(y{:}s,\vec{x})$ where $\vec{x}=(x{:}s_1,\ldots,x{:}s_n)$. Hence, $(\mathcal{A},\vec{a})\vDash\exists y{:}s(\bigwedge_{i\leq M}\varphi_i(y{:}s,\vec{x}))$. The sentence $\exists y{:}s(\bigwedge_{i\leq M}\varphi_i(y{:}s,\vec{x}))$ has rank $k+1$, so by the supposition, $\mathcal{B}\vDash\exists y{:}s(\bigwedge_{i\leq M}\varphi_i(y{:}s,\vec{x}))$. Therefore, there exists a $b\in B_s$ witnessing the existential quantifier $\exists x{:}s$; and so Duplicator responds with this $b\in B_s$. By construction, $(\mathcal{A},a\vec{a})\vDash\varphi(y{:}s,\vec{x})\Leftrightarrow(\mathcal{B},b\vec{b})\vDash\varphi(y{:}s,\vec{x})$ for all σ formulas of rank $\leq k$. By the inductive hypothesis, Duplicator has a winning strategy in the $\mathbf{EF}_k((\mathcal{A},a\vec{a}),(\mathcal{B},b\vec{b}))$ game. Therefore, Duplicator proceeding in the subsequent k moves of $\mathbf{EF}_{k+1}((\mathcal{A},\vec{a}),(\mathcal{B},\vec{b}))$ according to the moves of $\mathbf{EF}_k((\mathcal{A},a),(\mathcal{B},b))$ yields a winning strategy for $\mathbf{EF}_{k+1}((\mathcal{A},\vec{a}),(\mathcal{B},\vec{b}))$. ☐

Proposition A.0.14. *For all σ-structures $\mathcal{A},\mathcal{B}$,*

$$\mathcal{A}\equiv^{\mathbf{MSO}_k(\sigma)}\mathcal{B}\text{ if and only if }J_{\in}(\mathcal{A})\equiv^{\mathbf{FO}^{I,S}_k}J_{\in}(\mathcal{B}).$$

Proof. To simplify our reasoning we fix an enumeration of the the four types of variables we have encountered thus far. We enumerate the ordinary and set variables of $\mathbf{MSO}(\sigma)$, respectively, as

$$x_1,x_2,x_3,\ldots\text{ and }X_1,X_2,X_3,\ldots$$

Similarly, we enumerate the f-sorted and s-sorted variables of $\mathbf{FO}^{I,S}(\sigma)$, respectively, as

$$y_1{:}\mathbf{f},y_2{:}\mathbf{f},y_3{:}\mathbf{f},\ldots\text{ and }z_1{:}\mathbf{s},z_2{:}\mathbf{s},z_3{:}\mathbf{s},\ldots$$

This allows us to write translations between the two types of formulas as follows. Let

$F\colon \mathbf{MSO}(\sigma) \to \mathbf{FO}^{I,S}(\sigma^{\mathbf{e}})$ be defined inductively as follows

$$\begin{aligned}
F(x_i \in X_j) &= e(y_i, z_j)\\
F(x_i = x_j) &= y_i{:}\mathbf{f} = y_j{:}\mathbf{f}\\
F(R(x_{i_1}, x_{i_2}, \dots, x_{i_n})) &= R(y_{i_1}, y_{i_2}, \dots, y_{i_n})\\
F(\varphi \wedge \psi) &= F(\varphi) \wedge F(\psi)\\
F(\neg\varphi) &= \neg F(\varphi)\\
F(\exists x_i\, \varphi) &= \exists\, y_i{:}\mathbf{f}\, F(\varphi)\\
F(\exists X_i\, \varphi) &= \exists\, z_i{:}\mathbf{f}\, F(\varphi)
\end{aligned}$$

Similarly, we obtain a translation in the other direction $M\colon \mathbf{FO}^{I,S}(\sigma^{\mathbf{e}}) \to \mathbf{MSO}(\sigma)$ defined by

$$\begin{aligned}
M(\mathbf{e}(y_i, z_j)) &= x_i \in X_j\\
M(y_i{:}\mathbf{f} = y_j{:}\mathbf{f}) &= x_i = x_j\\
M(R(y_{i_1}{:}\mathbf{f}, y_{i_2}{:}\mathbf{f}, \dots, y_{i_n}{:}\mathbf{f})) &= R(x_{i_1}, x_{i_2}, \dots, x_{i_n})\\
M(\varphi \wedge \psi) &= S(\varphi) \wedge S(\psi)\\
M(\neg\varphi) &= \neg S(\varphi)\\
M(\exists\, y_i{:}\mathbf{f}\, \varphi) &= \exists\, x_i\, S(\varphi)\\
M(\exists\, z_i{:}\mathbf{s}\, \varphi) &= \exists\, X_i\, S(\varphi)
\end{aligned}$$

By induction on the complexity of formulas, using the standard semantics, we can show that for all $\varphi \in \mathbf{MSO}(\sigma)$, $\mathcal{A} \vDash \varphi \Leftrightarrow J(\mathcal{A}) \vDash F(\varphi)$ and conversely, for all $\varphi \in \mathbf{FO}_k^{I,S}(\sigma^{\mathbf{e}})$, $\mathcal{A} \vDash M(\varphi) \Leftrightarrow J(\mathcal{A}) \vDash \varphi$. $\square$

Bibliography

[1] Michael Abbott, Thorsten Altenkirch, and Neil Ghani. Categories of containers. In *International Conference on Foundations of Software Science and Computation Structures*, pages 23–38. Springer, 2003.

[2] Samson Abramsky. Structure and power: An emerging landscape. *Fundamenta Informaticae*, 186(1-4):1–26, August 2022.

[3] Samson Abramsky and Rui Soares Barbosa. The Logic of Contextuality. In Christel Baier and Jean Goubault-Larrecq, editors, *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*, volume 183 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:18, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[4] Samson Abramsky, Rui Soares Barbosa, Nadish de Silva, and Octavio Zapata. The quantum monad on relational structures, 2017.

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

[6] Samson Abramsky, Thomas Laure, and Luca Reggio. Existential and positive games: a comonadic and axiomatic view, 2024.

[7] Samson Abramsky and Dan Marsden. Comonadic semantics for guarded fragments. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–13. IEEE, 2021.

[8] Samson Abramsky and Dan Marsden. Comonadic semantics for hybrid logic. In Stefan Szeider, Robert Ganian, and Alexandra Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022, August 22-26, 2022, Vienna, Austria*, volume 241 of *LIPIcs*, pages 7:1–7:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[9] Samson Abramsky and Luca Reggio. Arboreal categories and resources. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 115:1–115:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[10] Samson Abramsky and Luca Reggio. Arboreal categories: An axiomatic theory of resources. *Logical Methods in Computer Science*, 19(3), 2023.

[11] Samson Abramsky and Luca Reggio. Arboreal categories and homomorphism preservation theorems, 2023.

[12] Samson Abramsky and Nihil Shah. Relating structure and power: Comonadic semantics for computational resources. In *27th EACSL Annual Conference on Computer Science Logic, CSL 2018, September 4-7, 2018, Birmingham, UK*, pages 2:1–2:17, 2018.

[13] Samson Abramsky and Nihil Shah. Relating structure and power: Comonadic semantics for computational resources. *Journal of Logic and Computation*, 31(6):1390–1428, 2021.

[14] Samson Abramsky and Nikos Tzevelekos. Introduction to categories and categorical logic. *New structures for physics*, pages 3–94, 2011.

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

[16] Danel Ahman, James Chapman, and Tarmo Uustalu. When is a container a comonad? *Logical Methods in Computer Science*, 10(3), September 2014.

[17] Thosten Altenkirch, James Chapman, and Tarmo Uustalu. Monads need not be endofunctors. *Logical Methods in Computer Science*, 11(1), March 2015.

[18] Albert Atserias, Phokion G. Kolaitis, and Simone Severini. Generalized satisfiability problems via operator assignments. *Journal of Computer and System Sciences*, 105:171–198, 2019.

[19] Jon Beck. Distributive laws. In *Seminar on triples and categorical homology theory*, pages 119–140. Springer, 1969.

[20] Bartosz Bednarczyk and Mateusz Urbanczyk. Comonadic semantics for description logics games. In Ofer Arieli, Martin Homola, Jean Christoph Jung, and Marie-Laure Mugnier, editors, *Proceedings of the 35th International Workshop on Description*

Logics (DL 2022) co-located with Federated Logic Conference (FLoC 2022), Haifa, Israel, August 7th to 10th, 2022, volume 3263 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2022.

[21] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2002.

[22] Achim Blumensath. A model-theoretic characterisation of clique width. *Annals of Pure and Applied Logic*, 142(1):321–350, 2006.

[23] E. Casanovas, P. Dellunde, and R. Jansana. On elementary equivalence for equality-free logic. *Notre Dame Journal of Formal Logic*, 37(3):506–522, 1996.

[24] Adam Ó Conghaile. First year report: Game comonads, descriptive complexity & finite model theory, 2019.

[25] Adam Ó Conghaile and Anuj Dawar. Game comonads & generalised quantifiers. In Christel Baier and Jean Goubault-Larrecq, editors, *29th EACSL Annual Conference on Computer Science Logic, CSL 2021, January 25-28, 2021, Ljubljana, Slovenia (Virtual Conference)*, volume 183 of *LIPIcs*, pages 16:1–16:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[26] Adam Peter Connolly. *Game comonads and beyond: compositional constructions for logic and algorithms*. PhD thesis, University of Cambridge, 2022.

[27] B. Courcelle, J. A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computing Systems*, 33(2):125–150, March 2000.

[28] Bruno Courcelle and Joost Engelfriet. *Relational structures*, page 621–685. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2012.

[29] Victor Dalmau. Linear Datalog and Bounded Path Duality of Relational Structures. *Logical Methods in Computer Science*, 1(1):5, April 2005. arXiv: cs/0504027.

[30] Anuj Dawar, Tomás Jakl, and Luca Reggio. Lovász-type theorems and game comonads. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–13. IEEE, 2021.

[31] Alan Day. Filter monads, continuous lattices and closure systems. *Canadian Journal of Mathematics*, 27(1):50–59, 1975.

[32] Maarten de Rijke. A note on graded modal logic. *Studia Logica: An International Journal for Symbolic Logic*, 64(2):271–283, 2000.

[33] Zdeněk Dvořák. On recognizing graphs by numbers of homomorphisms. *Journal of Graph Theory*, 64(4):330–342, November 2009.

[34] Andrzej Ehrenfeucht. An application of games to the completeness problem for formalized theories. *Fund. Math*, 49(129-141):13, 1961.

[35] Solomon Feferman and Robert L. Vaught. The first-order properties of algebraic systems. *Fundamenta Mathematicae*, 47, 1959.

[36] Roland Fraïssé. On the extension to relations of some properties of orders. In *Scientific Annals of the École Normale Supérieure*, volume 71, pages 363–388, 1954.

[37] P.J. Freyd and G.M. Kelly. Categories of continuous functors, i. *Journal of pure and applied algebra*, 2(3):169–191, 1972.

[38] Tobias Fritz, Tomáš Gonda, Paolo Perrone, and Eigil Fjeldgren Rischel. Representable markov categories and comparison of statistical experiments in categorical probability. *Theoretical Computer Science*, 961:113896, 2023.

[39] Tobias Fritz, Paolo Perrone, and Sharwin Rezagholi. Probability, valuations, hyperspace: Three monads on top and the support as a morphism. *Mathematical Structures in Computer Science*, 31(8):850–897, 2021.

[40] Soichiro Fujii, Shin-ya Katsumata, and Paul-André Melliès. Towards a formal theory of graded monads. In Bart Jacobs and Christof Löding, editors, *Foundations of Software Science and Computation Structures*, pages 513–530, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

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

[42] Michèle Giry. A categorical approach to probability theory. In *Categorical aspects of topology and analysis (Ottawa, Ont., 1980)*, volume 915 of *Lecture Notes in Mathematics*, pages 68–85. Springer, Berlin, 1982.

[43] Martin Grohe. Counting bounded tree depth homomorphisms. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. ACM, July 2020.

[44] Lauri Hella. Logical hierarchies in ptime. *Information and Computation*, 129(1):1–19, 1996.

[45] Matthew Hennessy and Robin Milner. On observing nondeterminism and concurrency. In *Proceedings of the 7th Colloquium on Automata, Languages and Programming*, page 299–309, Berlin, Heidelberg, 1980. Springer-Verlag.

[46] Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. A convenient category for higher-order probability theory. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, jun 2017.

[47] Neil Immerman. Upper and lower bounds for first order expressibility. *Journal of Computer and System Sciences*, 28:76–98, 1982.

[48] Neil Immerman and Eric Lander. *Describing graphs: A first-order approach to graph canonization.* Springer, 1990.

[49] Bart Jacobs. Semantics of weakening and contraction. *Annals of Pure and Applied Logic*, 69(1):73–106, 1994.

[50] Tomáš Jakl, Dan Marsden, and Nihil Shah. A game comonadic account of courcelle and feferman-vaught-mostowski theorems, 2022.

[51] Tomáš Jakl, Dan Marsden, and Nihil Shah. Generalizations of bilinear maps – technical report, 2022.

[52] Tomáš Jakl, Dan Marsden, and Nihil Shah. A categorical account of composition methods in logic. In *2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–14, Los Alamitos, CA, USA, jun 2023. IEEE Computer Society.

[53] André Joyal, Mogens Nielsen, and Glynn Winskel. Bisimulation from open maps. *Information and Computation*, 127(2):164–185, June 1996.

[54] Paul C. Kainen. Weak adjoint functors. *Mathematische Zeitschrift*, 122:1–9, 1971.

[55] Amin Karamlou and Nihil Shah. No go theorems: Directed containers that do not distribute over distribution monads. In *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science*, 2024.

[56] Amin Karamlou and Nihil Shah. Pebble games, logic, and quantum homomorphisms, 2024.

[57] Simon Kochen and Ernst Paul Specker. The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17(1):59–87, 1967.

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

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

[60] Moritz Lichter, Ilia Ponomarenko, and Pascal Schweitzer. Walk refinement, walk logic, and the iteration number of the weisfeiler-leman algorithm. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13, 2019.

[61] László Lovász. Operations with structures. *Acta Mathematica Academiae Scientiarum Hungaricae*, 18(3-4):321–328, September 1967.

[62] Roger Lyndon. Properties preserved under homomorphism. *Pacific journal of mathematics*, 9(1):143–154, 1959.

[63] Ernest Manes. *Algebraic Theories.* Springer New York, 1976.

[64] Ernest Manes and Philip Mulry. Monad compositions i: general constructions and recursive distributive laws. *Theory and Applications of Categories*, 18(7):172–208, 2007.

[65] Laura Mančinska and David E. Roberson. Quantum homomorphisms. *Journal of Combinatorial Theory, Series B*, 118:228–267, 2016.

[66] Vincenzo Marra and Luca Spada. Two isomorphism criteria for directed colimits. *arXiv preprint arXiv:1312.0432*, 2013.

[67] Stefan Milius, Dirk Pattinson, and Lutz Schröder. Generic Trace Semantics and Graded Monads. In Lawrence S. Moss and Pawel Sobocinski, editors, *6th Conference on Algebra and Coalgebra in Computer Science (CALCO 2015)*, volume 35 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 253–269, Dagstuhl, Germany, 2015. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[68] Yoàv Montacute and Nihil Shah. The pebble-relation comonad in finite model theory. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '22, New York, NY, USA, 2022. Association for Computing Machinery.

[69] Yoàv Montacute and Nihil Shah. The pebble-relation comonad in finite model theory. *Logical Methods in Computer Science*, Volume 20, Issue 2, May 2024.

[70] Yoàv Montacute and Nihil Shah. The pebble-relation comonad in finite model theory. *Logical Methods in Computer Science*, Volume 20, Issue 2, May 2024.

[71] Andrzej Mostowski. On direct products of theories. *The Journal of Symbolic Logic*, 17(1):1–31, 1952.

[72] Dominic Orchard. *Programming contextual computations.* PhD thesis, University of Cambridge, Computer Laboratory, 2014.

[73] Thomas Paine. A pebbling comonad for finite rank and variable logic, and an application to the equirank-variable homomorphism preservation theorem. *Electronic Notes in Theoretical Computer Science*, 352:191–209, 2020. The 36th Mathematical Foundations of Programming Semantics Conference, 2020.

[74] John Power and Hiroshi Watanabe. Combining a monad and a comonad. *Theoretical Computer Science*, 280(1-2):137–162, May 2002.

[75] Michaël Rao. Msol partitioning problems on graphs of bounded treewidth and clique-width. *Theoretical Computer Science*, 377(1):260–267, 2007.

[76] Luca Reggio and Colin Riba. Finitely accessible arboreal adjunctions and hintikka formulae, 2023.

[77] Emily Riehl. Factorization systems, 2008.

[78] Emily Riehl. *Category Theory in Context.* Aurora: Dover Modern Math Originals. Dover Publications, 2017.

[79] Benjamin Rossman. Homomorphism preservation theorems. *Journal of the ACM*, 55(3):1–53, 2008.

[80] Nicole Schirrmacher, Sebastian Siebertz, and Alexandre Vigny. First-order logic with connectivity operators. *ACM Trans. Comput. Logic*, 24(4), jul 2023.

[81] Nihil Shah. Game comonads in finite model theory. Master's thesis, University of Oxford, 2017.

[82] Alfred Tarski. Contributions to the theory of models. iii. *Indagationes Mathematicae (Proceedings)*, 58:56–64, 1955.

[83] Robert J. van Glabbeek. The linear time - branching time spectrum. In J. C. M. Baeten and J. W. Klop, editors, *CONCUR '90 Theories of Concurrency: Unification and Extension*, pages 278–297, Berlin, Heidelberg, 1990. Springer Berlin Heidelberg.

[84] Robert J. van Glabbeek. The linear time – branching time spectrum ii. In Eike Best, editor, *CONCUR'93*, pages 66–81, Berlin, Heidelberg, 1993. Springer Berlin Heidelberg.

[85] Jerzy Łoś. On the extending of models (i). *Fundamenta mathematicae*, 42(1):38–54, 1955.