

Lovász-Type Theorems and Game Comonads

Anuj Dawar and Tomáš Jakl

Department of Computer Science and Technology
University of Cambridge, UK
Email: anuj.dawar@cl.cam.ac.uk, tj330@cam.ac.uk

Luca Reggio

Department of Computer Science
University of Oxford, UK
Email: luca.reggio@cs.ox.ac.uk

Abstract—Lovász (1967) showed that two finite relational structures A and B are isomorphic if, and only if, the number of homomorphisms from C to A is the same as the number of homomorphisms from C to B for any finite structure C . Soon after, Pultr (1973) proved a categorical generalisation of this fact. We propose a new categorical formulation, which applies to any locally finite category with pushouts and a proper factorisation system. As special cases of this general theorem, we obtain two variants of Lovász’ theorem: the result by Dvořák (2010) that characterises equivalence of graphs in the k -dimensional Weisfeiler–Leman equivalence by homomorphism counts from graphs of tree-width at most k , and the result of Grohe (2020) characterising equivalence with respect to first-order logic with counting and quantifier depth k in terms of homomorphism counts from graphs of tree-depth at most k . The connection of our categorical formulation with these results is obtained by means of the game comonads of Abramsky et al. We also present a novel application to homomorphism counts in modal logic.

I. INTRODUCTION

Over fifty years ago, Lovász [20] proved that two finite graphs (or, more generally, any two finite relational structures) A and B are isomorphic if, and only if, for every finite graph C , the number of homomorphisms from C to A is the same as the number of homomorphisms from C to B . While one direction of this equivalence is obvious, the other direction establishes that the isomorphism type of a finite graph A is characterised by an infinite integer vector V indexed by \mathcal{F} , the collection of (isomorphism classes of) finite graphs, where for $C \in \mathcal{F}$, V_C is the number of homomorphisms from C to A . This seminal result has led to extensions and investigations in many different directions. One that concerns us is that by restricting the collection \mathcal{F} to a natural sub-collection, we can often obtain characterisations of natural coarsenings of the relation of isomorphism. For any subclass \mathcal{F}' of \mathcal{F} , say that a pair of graphs G and H is \mathcal{F}' -homomorphism equivalent if for any K in \mathcal{F}' the number of homomorphisms from K to G is the same as the number of homomorphisms from K to H .

For instance, Dvořák [11] shows that for any k , if we consider the collection \mathcal{T}_k of all graphs of tree-width at most k , then a pair of graphs are \mathcal{T}_k -homomorphism equivalent if, and only if, they are equivalent with respect to the k -dimensional Weisfeiler–Leman (k -WL) equivalence. The k -WL equivalence relations are a widely studied family of approximations

of the graph isomorphism relation (see [17] for a recent exposition) with many equivalent characterisations in terms of combinatorics, logic, algebra and linear optimisation. From our point of view, one important such characterisation is the fact that two graphs G and H are k -WL equivalent if, and only if, they are not distinguished by any sentence of \mathcal{C}^{k+1} —first-order logic with counting quantifiers restricted to no more than $k + 1$ distinct variables [9]. This logical formulation of the equivalence extends naturally to finite structures over any relational vocabulary.

Dvořák’s result offers yet another characterisation of k -WL equivalence, this time in terms of homomorphism counts from the equally widely studied class of graphs \mathcal{T}_k . An important special case is that of $k = 1$. Since \mathcal{T}_1 is the class of finite forests, and 1-WL equivalence is the same as *fractional isomorphism* (see [26]), the result also provides an elegant characterisation of equivalence of graphs with respect to the number of homomorphisms from trees. It is also known that two graphs are co-spectral if, and only if, they are Cyc-homomorphism equivalent, where Cyc is the class of simple cycles, and that two graphs are *quantum isomorphic* if, and only if, they are \mathcal{P} -homomorphism equivalent for the class \mathcal{P} of planar graphs [22]. Another recent addition to this collection of results is that of Grohe [14]. He shows that, if \mathcal{D}_n is the collection of graphs of *tree-depth* at most n , then two graphs are \mathcal{D}_n -homomorphism equivalent if, and only if, they are not distinguished by any sentence of \mathcal{C}_n —first-order logic with counting with quantifier depth at most n .

This variety of results relating homomorphism counts over restricted classes to approximations of isomorphism, often proved with very different methods, calls for the development of a more general theory. In the present paper we develop a general categorical result which yields, as example applications, the results of Lovász, Dvořák and Grohe mentioned above. The connection between our categorical result and the results of Dvořák and Grohe is established using the game comonads of Abramsky et al. [1], [3]. Specifically, in [1] the graded pebbling comonad \mathbb{P}_k on the category of σ -structures (for a finite relational signature σ) is introduced and it is shown that the coalgebras for this comonad correspond in a natural way with certain tree decompositions of width at most $k - 1$. In particular, a finite σ -structure admits a coalgebra structure for this comonad if, and only if, it has tree-width at most $k - 1$. At the same time, it is shown that isomorphism in the Kleisli category corresponding to the comonad (or equivalently in

Research supported in part by the EPSRC grant EP/T007257/1, and by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 837724.

the Eilenberg–Moore category) is exactly indistinguishability in \mathcal{C}^k . This brings together in one categorical construction the essential elements of Dvořák’s theorem. In an exactly analogous fashion, the Ehrenfeucht–Fraïssé comonad \mathbb{E}_n of [3] captures, at the level of coalgebras, the structures of tree-depth at most n and yields a notion of isomorphism corresponding to equivalence in \mathcal{C}_n . These categorical formulations of the combinatorial parameters tree-width and tree-depth on the one hand, and the logical equivalences with respect to \mathcal{C}^k and \mathcal{C}_n on the other, are what enable us to give a single result generalising the three theorems of Lovász, Dvořák and Grohe.

Our general result states that in any locally finite category with pushouts and a proper factorisation system, two objects m and n are isomorphic if, and only if, the number of morphisms from k to m is the same as that from k to n for any object k . This is proved in Section III. In Section IV, we use the game comonads to derive the theorems of Dvořák and Grohe from this general result and also a combination of the two characterising equivalence in \mathcal{C}_n^k —first-order logic with counting with quantifier depth n and k variables. In Section V, we apply the machinery we have developed to another game comonad: the graded modal comonad \mathbb{M}_k . This gives a new Lovász-style result for pointed Kripke structures, relating homomorphism counts from synchronization trees of bounded height to equivalence in a modal logic with counting. Finally, in Section VI we draw some conclusions from this about certain normal forms for first-order logic with counting.

Related work: Other categorical generalisations of Lovász’ theorem have been obtained by Pultr [25] and Isbell [15]. In modern terms, Pultr works with a finitely well-powered, locally finite category with (extremal epi, mono) factorisations and Isbell requires the category to be locally finite with a special type of factorisation system (which he calls “bicategory”). Our Theorem 5 is arguably better suited for applications to game comonads and, unlike the results in op. cit., does not hinge on a combinatorial counting argument but rather on an application of the inclusion-exclusion principle.

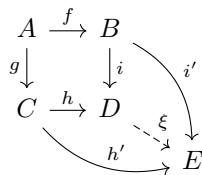
II. PRELIMINARIES ON GAME COMONADS

In this section, we recall the necessary material on (game) comonads and their coalgebras.

A. Categories

We briefly recall some basic notions of category theory. For a more thorough introduction, see e.g. [5] or [21].

Let \mathcal{A} be a category, and $f: A \rightarrow B$ and $g: A \rightarrow C$ two arrows in \mathcal{A} . The *pushout* of f and g , if it exists, consists of two arrows $h: C \rightarrow D$ and $i: B \rightarrow D$ such that $i \circ f = h \circ g$ and the following universal property is satisfied: For any two arrows $i': B \rightarrow E$ and $h': C \rightarrow E$ with $i' \circ f = h' \circ g$, there is a unique arrow $\xi: D \rightarrow E$ satisfying $\xi \circ i = i'$ and $\xi \circ h = h'$.



The *pullback* of two arrows $f: A \rightarrow C$ and $g: B \rightarrow C$ in \mathcal{A} , if it exists, is the pushout of f and g in the opposite category \mathcal{A}^{op} obtained by reversing the direction of the arrows in \mathcal{A} .

Pushouts and pullbacks are special instances of the concepts of colimits and limits of diagrams, respectively. For example, the pushout of $f: A \rightarrow B$ and $g: A \rightarrow C$ coincides with the colimit of the diagram

$$B \xleftarrow{f} A \xrightarrow{g} C.$$

For arbitrary colimits, we consider diagrams of any shape.

A functor $F: \mathcal{A} \rightarrow \mathcal{B}$ is *faithful* if, given any two parallel arrows $f, g: A \rightarrow A'$ in \mathcal{A} , $Ff = Fg$ implies $f = g$. Further, F is *full* if, for any two objects A and A' of \mathcal{A} , each arrow $FA \rightarrow FA'$ is of the form Ff for some $f: A \rightarrow A'$. A subcategory \mathcal{A} of a category \mathcal{B} is said to be *full* if the inclusion functor $\mathcal{A} \rightarrow \mathcal{B}$ is full.

A fundamental notion of category theory is that of an adjunction. Let $F: \mathcal{A} \rightarrow \mathcal{B}$ and $G: \mathcal{B} \rightarrow \mathcal{A}$ be any two functors between categories. We say that F is *left adjoint* to G (equivalently, G is *right adjoint* to F), and write $F \dashv G$, provided that for each object $B \in \mathcal{B}$ there is a morphism $\varepsilon_B: FGB \rightarrow B$ satisfying the following universal property: For every $A \in \mathcal{A}$ and morphism $f: FA \rightarrow B$ in \mathcal{B} , there exists a unique morphism $g: A \rightarrow GB$ in \mathcal{A} such that $\varepsilon_B \circ Fg = f$.

$$\begin{array}{ccc} FGB & \xrightarrow{\varepsilon_B} & B \\ Fg \uparrow & \nearrow f & \\ FA & & \end{array}$$

B. Comonads

A *comonad* (in Kleisli form) on a category \mathcal{A} is given by:

- an object map $\mathbb{C}: \mathcal{A} \rightarrow \mathcal{A}$;
- a morphism $\varepsilon_A: \mathbb{C}(A) \rightarrow A$ for every $A \in \mathcal{A}$;
- a *coextension operation* associating with any morphism $f: \mathbb{C}(A) \rightarrow B$ a morphism $f^*: \mathbb{C}(A) \rightarrow \mathbb{C}(B)$.

These data must satisfy the following equations for all morphisms $f: \mathbb{C}(A) \rightarrow B$ and $g: \mathbb{C}(B) \rightarrow C$:

$$\varepsilon_A^* = \text{id}_{\mathbb{C}(A)}, \quad \varepsilon_B \circ f^* = f, \quad (g \circ f^*)^* = g^* \circ f^*.$$

In particular, we can extend \mathbb{C} to a functor $\mathcal{A} \rightarrow \mathcal{A}$ by setting $\mathbb{C}(f) := (f \circ \varepsilon_A)^*$ for every morphism $f: A \rightarrow B$.¹

We recall next some examples of comonads that play an important role in this paper. Let σ be an arbitrary relational signature. The category Σ has as objects σ -structures (denoted A, B, \dots) and as morphisms σ -homomorphisms (sometimes simply called *homomorphisms*), i.e. functions $f: A \rightarrow B$ such that, for all relation symbols $R \in \sigma$, $f(R^A) \subseteq R^B$ where R^A and R^B are the interpretations of R in A and B , respectively.

For each positive integer n , the *Ehrenfeucht–Fraïssé* comonad \mathbb{E}_n on Σ (cf. [3]) is defined as follows. For any

¹It is easy to see that, setting $\delta_A := \text{id}_{\mathbb{C}(A)}^*$ for every $A \in \mathcal{A}$, the tuple $(\mathbb{C}, \varepsilon, \delta)$ is a comonad in the more traditional sense, where ε is the counit and δ the comultiplication. In fact, these two formulations are equivalent.

$A \in \Sigma$, the universe of $\mathbb{E}_n(A)$ is the set $A^{\leq n}$ of all non-empty sequences of length at most n . Before defining the interpretations of the relation symbols, let us define the map

$$\varepsilon_A: A^{\leq n} \rightarrow A, \quad \varepsilon_A[a_1, \dots, a_l] := a_l.$$

With this notation, for each $R \in \sigma$ of arity j , we define $R^{\mathbb{E}_n(A)}$ to be the set of those tuples (s_1, \dots, s_j) of sequences which are pairwise comparable in the prefix order, and such that $(\varepsilon_A(s_1), \dots, \varepsilon_A(s_j)) \in R^A$. The coextension operation sends a σ -homomorphism $f: \mathbb{E}_n(A) \rightarrow B$ to the σ -homomorphism

$$f^*: \mathbb{E}_n(A) \rightarrow \mathbb{E}_n(B), \quad f^*[a_1, \dots, a_j] := [b_1, \dots, b_j]$$

where $b_i := f[a_1, \dots, a_i]$ for all $1 \leq i \leq j$. It is not difficult to see that these data define a comonad on the category Σ .

From the viewpoint of model-comparison games, the elements of $A^{\leq n}$ represent the plays in the σ -structure A of length at most n , and a σ -homomorphism $\mathbb{E}_n(A) \rightarrow B$ is a strategy for Duplicator in the Ehrenfeucht–Fraïssé game with n rounds, where Spoiler plays always in A and Duplicator responds in B . For more details, cf. [3], [4]. Next we recall from [1] another comonad on Σ , which models pebble games.

For each positive integer k , set $\mathbf{k} := \{1, \dots, k\}$. Given a σ -structure A , we consider the set $(\mathbf{k} \times A)^+$ of all non-empty sequences of elements of $\mathbf{k} \times A$. We call a pair $(p, a) \in \mathbf{k} \times A$ a *move*. Whenever $[(p_1, a_1), \dots, (p_l, a_l)]$ is a sequence of moves, p_i is called the *pebble index* of the move (p_i, a_i) . As with \mathbb{E}_n , define the map

$$\varepsilon_A: (\mathbf{k} \times A)^+ \rightarrow A, \quad \varepsilon_A[(p_1, a_1), \dots, (p_l, a_l)] := a_l$$

sending a play to its last move. We let $\mathbb{P}_k(A)$ be the σ -structure with universe $(\mathbf{k} \times A)^+$ and such that, for every relation $R \in \sigma$ of arity j , its interpretation $R^{\mathbb{P}_k(A)}$ consists of those tuples of sequences (s_1, \dots, s_j) such that: (i) the s_i are pairwise comparable in the prefix order; (ii) whenever s_i is a prefix of $s_{i'}$, the pebble index of the last move in s_i does not appear in the suffix of s_i in $s_{i'}$; and (iii) $(\varepsilon_A(s_1), \dots, \varepsilon_A(s_j)) \in R^A$. Finally, the coextension operation is the same, mutatis mutandis, as with \mathbb{E}_n . Again, it is not difficult to see that these data define a comonad \mathbb{P}_k on Σ , called the *pebbling comonad*.

The Ehrenfeucht–Fraïssé comonad \mathbb{E}_n can be “combined” with the pebbling comonad \mathbb{P}_k to obtain a new comonad $\mathbb{P}_{k,n}$ on Σ , cf. [24]. For every $A \in \Sigma$, the universe of $\mathbb{P}_{k,n}(A)$ is the set $(\mathbf{k} \times A)^{\leq n}$ of all non-empty sequences of length at most n (this corresponds to bounding the length of plays in pebble games). The homomorphisms ε_A and the coextension operation for $\mathbb{P}_{k,n}$ are defined as the restrictions of the corresponding concepts for the pebbling comonad \mathbb{P}_k .

C. Coalgebras

Let \mathbb{C} be a comonad (in Kleisli form) on a category \mathcal{A} . A *coalgebra* for \mathbb{C} is a pair $(A, A \xrightarrow{\alpha} \mathbb{C}(A))$ such that $A \in \mathcal{A}$, α is a morphism in \mathcal{A} , and the following diagrams commute:

$$\begin{array}{ccc} A & \xrightarrow{\alpha} & \mathbb{C}(A) \\ \text{id}_A \searrow & & \downarrow \varepsilon_A \\ & & A \end{array} \quad \begin{array}{ccc} A & \xrightarrow{\alpha} & \mathbb{C}(A) \\ \alpha \downarrow & & \downarrow \text{id}_{\mathbb{C}(A)} \\ \mathbb{C}(A) & \xrightarrow{\mathbb{C}(\alpha)} & \mathbb{C}(\mathbb{C}(A)) \end{array}$$

We refer to α as a *coalgebra structure* for A . In the case of the game comonads $\mathbb{E}_n, \mathbb{P}_k$, and $\mathbb{P}_{k,n}$, the σ -structures admitting a coalgebra structure can be characterised in an elegant way in terms of key combinatorial parameters, as we now explain.

Let (F, \leq) be any poset. For all $x, y \in F$, we write $x \uparrow y$ whenever x and y are comparable, i.e. either $x \leq y$ or $y \leq x$, and let $\downarrow x := \{y \in F \mid y \leq x\}$. A *forest* is a poset (F, \leq) such that, for all $x \in F$, the set $\downarrow x$ is finite and totally ordered. The *height* of (F, \leq) is $\sup_{x \in F} |\downarrow x|$.

If $G = (V, \curvearrowright)$ is a graph (where V is the set of vertices and \curvearrowright the adjacency relation), a *forest cover* of G consists of a forest (F, \leq) and an injective function $f: V \rightarrow F$ such that $v \curvearrowright v'$ entails $f(v) \uparrow f(v')$ for all $v, v' \in V$.

Finally, recall that the *Gaifman graph* of a σ -structure A is $\mathcal{G}_A := (A, \curvearrowright)$ where $a \curvearrowright a'$ if $a \neq a'$ and there exist $R \in \sigma$ and a tuple $(a_1, \dots, a_l) \in R^A$ such that $a = a_i$ and $a' = a_j$ for some $i, j \in \{1, \dots, l\}$. A σ -structure A has *tree-depth at most n* if there is a forest cover (F, \leq) of \mathcal{G}_A of height $\leq n$.

The following is a direct consequence of [3, Theorem 17]:

Proposition 1. *A σ -structure A admits a coalgebra structure $\alpha: A \rightarrow \mathbb{E}_n(A)$ if, and only if, it has tree-depth at most n .*

In the same way as σ -structures that admit a coalgebra structure for \mathbb{E}_n can be characterised in terms of tree-depth, the existence of a coalgebra structure for the pebbling comonad \mathbb{P}_k corresponds to tree-width. We recall the relevant definitions.

Note that, if (F, \leq) is a forest cover of a graph $G = (V, \curvearrowright)$, with injective map $f: V \rightarrow F$, we can assume without loss of generality that $F = V$ and $f = \text{id}_V$. We shall assume this for the remainder of this section. A *k -pebble forest cover* of a graph $G = (V, \curvearrowright)$ consists of a forest cover (V, \leq) of G and a pebbling function $p: V \rightarrow \mathbf{k}$ such that, whenever $v \curvearrowright v'$ with $v \leq v'$, we have $p(v) \neq p(v')$ for all $v < v' \leq v'$.

A σ -structure A has *tree-width at most $k-1$* if its Gaifman graph \mathcal{G}_A admits a k -pebble forest cover.² The following result is a consequence of [1, Proposition 22].

Proposition 2. *A σ -structure A admits a coalgebra structure $\alpha: A \rightarrow \mathbb{P}_k(A)$ if, and only if, it has tree-width at most $k-1$.*

Remark 3. The tree-width of a graph is usually defined in terms of tree decompositions. The definition given above is an equivalent reformulation. For a proof of the fact that a finite graph admits a tree decomposition of width $< k$ if, and only if, it admits a k -pebble forest cover, see [3, Theorem 19].

In the same spirit, in the case of the comonad $\mathbb{P}_{k,n}$ we have the following characterisation, cf. [24, Theorem 2.14].

²It is customary in graph theory to use $k-1$, instead of k , so that trees have tree-width 1.

Proposition 4. A σ -structure A admits a coalgebra structure $\alpha: A \rightarrow \mathbb{P}_{k,n}(A)$ if, and only if, it has a k -pebble forest cover of height $\leq n$.

Finally, recall that coalgebras for a comonad \mathbb{C} on \mathcal{A} form themselves a category $\text{EM}(\mathbb{C})$, the *Eilenberg–Moore category* of \mathbb{C} . The objects of $\text{EM}(\mathbb{C})$ are the coalgebras (A, α) for \mathbb{C} , and morphisms $(A, \alpha) \rightarrow (B, \beta)$ in $\text{EM}(\mathbb{C})$ are morphisms $h: A \rightarrow B$ in \mathcal{A} such that $\mathbb{C}(h) \circ \alpha = \beta \circ h$. There is an obvious forgetful functor $U^{\mathbb{C}}: \text{EM}(\mathbb{C}) \rightarrow \mathcal{A}$ which sends a coalgebra (A, α) to A . This functor has a right adjoint $F^{\mathbb{C}}: \mathcal{A} \rightarrow \text{EM}(\mathbb{C})$ defined as follows: for any $A \in \mathcal{A}$, $F^{\mathbb{C}}(A) := (\mathbb{C}(A), \delta_A)$ where $\delta_A := \text{id}_{\mathbb{C}(A)}^*$. Further, if h is a morphism in \mathcal{A} , we set $F^{\mathbb{C}}(h) := \mathbb{C}(h)$.

$$\begin{array}{ccc} & F^{\mathbb{C}} & \\ \text{EM}(\mathbb{C}) & \xleftarrow{\quad} & \mathcal{A} \\ & U^{\mathbb{C}} & \end{array}$$

III. A CATEGORICAL LOVÁSZ-TYPE THEOREM

A category is *locally finite* if there are only finitely many morphisms between any two of its objects. A locally finite category \mathcal{A} is said to be *combinatorial* if, for all $m, n \in \mathcal{A}$,

$$m \cong n \iff |\text{hom}_{\mathcal{A}}(k, m)| = |\text{hom}_{\mathcal{A}}(k, n)| \quad \forall k \in \mathcal{A}.$$

Thus, Lovász’ theorem [20] states that, for any finite relational signature σ , the category Σ_f of finite σ -structures with homomorphisms is combinatorial. The aim of this section is to prove the following generalisation of Lovász’ result.

Theorem 5. Let \mathcal{A} be a locally finite category. If \mathcal{A} has pushouts and a proper factorisation system, then it is combinatorial.

A. Proper factorisation systems

We start by recalling the notion of weak factorisation system. Given morphisms e and m in \mathcal{A} , we say that e has the *left lifting property* with respect to m , or that m has the *right lifting property* with respect to e if, for every commutative square as on the left-hand side below,

$$\begin{array}{ccc} \bullet & \xrightarrow{e} & \bullet \\ \downarrow & & \downarrow \\ \bullet & \xrightarrow{m} & \bullet \end{array} \quad \begin{array}{ccc} \bullet & \xrightarrow{e} & \bullet \\ \downarrow & \searrow d & \downarrow \\ \bullet & \xrightarrow{m} & \bullet \end{array}$$

there is a (not necessarily unique) *diagonal filler*, i.e. an arrow d such that the right-hand diagram above commutes. If this is the case, we write $e \pitchfork m$. For any class \mathcal{F} of morphisms in \mathcal{A} , let ${}^{\pitchfork}\mathcal{F}$ (respectively \mathcal{F}^{\pitchfork}) be the class of morphisms having the left (respectively right) lifting property with respect to every morphism in \mathcal{F} .

Definition 6. A pair of classes of morphisms $(\mathcal{E}, \mathcal{M})$ in a category \mathcal{A} is a *weak factorisation system* provided it satisfies the following conditions:

- 1) every morphism f in \mathcal{A} can be written as $f = m \circ e$ with $e \in \mathcal{E}$ and $m \in \mathcal{M}$;
- 2) $\mathcal{E} = {}^{\pitchfork}\mathcal{M}$ and $\mathcal{M} = \mathcal{E}^{\pitchfork}$.

A *proper factorisation system* is a weak factorisation system $(\mathcal{E}, \mathcal{M})$ such that all morphisms in \mathcal{E} are epimorphisms and all morphisms in \mathcal{M} are monomorphisms.³

We describe two different proper factorisation systems in the category Σ of σ -structures, for a fixed relational signature σ . Each homomorphism of σ -structures $f: A \rightarrow B$ factors, as a function, through its set-theoretic image \tilde{A} :

$$A \rightarrow \tilde{A} \rightarrow B.$$

There are two natural ways to turn \tilde{A} into a σ -structure: we can equip it with the structure induced by either B , or A . In the first case, for every $R \in \sigma$ of arity n , we set $R^{\tilde{A}} := R^B \cap \tilde{A}^n$ and, in the second, we set $R^{\tilde{A}} := f(R^A) \cap \tilde{A}^n$.

These two ways of turning \tilde{A} into a σ -structure yield two different weak factorisation systems $(\mathcal{E}, \mathcal{M})$ in the category Σ .⁴ For example, the first one yields the weak factorisation system where $\mathcal{E} = \{\text{surjective homomorphisms}\}$ and $\mathcal{M} = \{\text{strong/induced embeddings}\}$. Both factorisation systems are proper. Moreover, since the image of a finite σ -structure under a homomorphism is also finite, these factorisation systems restrict to the locally finite category Σ_f of finite σ -structures.

The category Σ also has pushouts. Given homomorphisms of σ -structures $f: A \rightarrow B$ and $g: A \rightarrow C$, their pushout D is computed as the quotient of the disjoint sum of σ -structures $B + C$ by the least equivalence relation \sim such that $f(a) \sim g(a)$ for every $a \in A$. Equivalently, D is obtained by equipping D' , the pushout of the functions f and g in the category of sets, depicted in the following diagram,

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ g \downarrow & & \downarrow f' \\ C & \xrightarrow{g'} & D' \end{array}$$

with the smallest relational structure that turns f' and g' into homomorphisms. Using the same idea, it is not difficult to see that Σ is cocomplete, i.e. it has all colimits, cf. [6, p. 201].

By the previous description of pushouts, it is clear that the pushout in Σ of finite σ -structures is again a finite σ -structure. It follows that Σ_f has all pushouts, and thus it satisfies the assumptions of Theorem 5. This justifies our claim that the latter result is a generalisation of Lovász’ theorem.

We state next some well known properties of weak and proper factorisation systems (cf., e.g., [12] or [27]):

Lemma 7. Let $(\mathcal{E}, \mathcal{M})$ be a weak factorisation system in \mathcal{A} . The following hold:

- (a) \mathcal{E} and \mathcal{M} are closed under compositions;
- (b) $\mathcal{E} \cap \mathcal{M} = \{\text{isomorphisms}\}$;
- (c) the pushout in \mathcal{A} of an \mathcal{E} -morphism along any morphism, if it exists, is again in \mathcal{E} .

Moreover, if $(\mathcal{E}, \mathcal{M})$ is proper, the following hold:

³It is an easy observation that any proper factorisation system is an *orthogonal* factorisation system, meaning that the diagonal fillers are unique.

⁴These are the (epi, regular mono) and (regular epi, mono) factorisation systems, respectively. Cf. [6, pp. 200–201].

- (d) $g \circ f \in \mathcal{E}$ implies $g \in \mathcal{E}$;
(e) $g \circ f \in \mathcal{M}$ implies $f \in \mathcal{M}$.

B. Proof of Theorem 5

For the remainder of this section, we fix a category \mathcal{A} admitting a proper factorisation system $(\mathcal{E}, \mathcal{M})$. \mathcal{M} -morphisms will be denoted by \rightarrow , and \mathcal{E} -morphisms by \twoheadrightarrow .

An \mathcal{E} -pushout square in \mathcal{A} is a pushout square consisting entirely of \mathcal{E} -morphisms. We are interested in functors $\mathcal{A}^{\text{op}} \rightarrow \mathbf{FinSet}$, where \mathbf{FinSet} is the category of finite sets and functions between them, that turn \mathcal{E} -pushout squares in \mathcal{A} into pullbacks in \mathbf{FinSet} . Recall that pullbacks in \mathbf{FinSet} admit the following explicit description: a commutative square in \mathbf{FinSet} is a pullback if, and only if,

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ g \downarrow & & \downarrow i \\ C & \xrightarrow{h} & D \end{array} \quad \begin{array}{l} \forall b \in B, \forall c \in C, \text{ if } i(b) = h(c) \text{ then} \\ \exists ! a \in A. (f(a) = b \text{ and } g(a) = c). \end{array}$$

If h and i are injective then, by identifying B and C with subsets of D , the pullback A can be identified with $B \cap C$.

The following lemma shows that the hom-functors involved in the homomorphism counting theorem send pushout squares to pullbacks. (This is a consequence of the more general fact that representable functors preserve all limits that exist in their domain, cf. [21, pp. 116-117], but we provide here also a direct elementary proof.) Recall that, for any object n of a locally finite category \mathcal{B} , the functor $\text{hom}_{\mathcal{B}}(-, n): \mathcal{B}^{\text{op}} \rightarrow \mathbf{FinSet}$ sends a morphism $f: a \rightarrow b$ in \mathcal{B} to the function $- \circ f: \text{hom}_{\mathcal{B}}(b, n) \rightarrow \text{hom}_{\mathcal{B}}(a, n)$.

Lemma 8. *For any locally finite category \mathcal{B} and $n \in \mathcal{B}$, the functor $\text{hom}_{\mathcal{B}}(-, n): \mathcal{B}^{\text{op}} \rightarrow \mathbf{FinSet}$ sends all pushout squares that exist in \mathcal{B} to pullbacks in \mathbf{FinSet} .*

Proof. Just observe that, given a pushout square in \mathcal{B} and the corresponding diagram in \mathbf{FinSet} , as displayed below,

$$\begin{array}{ccc} a & \xrightarrow{f} & b \\ g \downarrow & & \downarrow i \\ c & \xrightarrow{h} & d \end{array} \quad \begin{array}{ccc} \text{hom}_{\mathcal{B}}(d, n) & \xrightarrow{- \circ i} & \text{hom}_{\mathcal{B}}(b, n) \\ - \circ h \downarrow & & \downarrow - \circ f \\ \text{hom}_{\mathcal{B}}(c, n) & \xrightarrow{- \circ g} & \text{hom}_{\mathcal{B}}(a, n) \end{array}$$

if $\alpha \in \text{hom}_{\mathcal{B}}(b, n)$ and $\beta \in \text{hom}_{\mathcal{B}}(c, n)$ are such that $\alpha \circ f = \beta \circ g$ then, by the universal property of the pushout, there is a unique $\gamma \in \text{hom}_{\mathcal{B}}(d, n)$ satisfying $\gamma \circ i = \alpha$ and $\gamma \circ h = \beta$. \square

Lemma 9. *If a functor $F: \mathcal{A}^{\text{op}} \rightarrow \mathbf{FinSet}$ sends \mathcal{E} -pushout squares in \mathcal{A} to pullbacks, then it sends \mathcal{E} -morphisms to injections.*

Proof. Let $e: n \twoheadrightarrow m$ be an \mathcal{E} -morphism in \mathcal{A} . Because e is an epimorphism, it follows directly that the square on the left-hand side below is a pushout square in \mathcal{A} .

$$\begin{array}{ccc} n & \xrightarrow{e} & m \\ e \downarrow & & \downarrow \text{id} \\ m & \xrightarrow{\text{id}} & m \end{array} \quad \begin{array}{ccc} F(m) & \xrightarrow{\text{id}} & F(m) \\ \text{id} \downarrow & & \downarrow F(e) \\ F(m) & \xrightarrow{F(e)} & F(n) \end{array}$$

Since identities are \mathcal{E} -morphisms, the square on the right-hand side above is a pullback in \mathbf{FinSet} . In turn, this is equivalent to $F(e)$ being an injection. \square

Generic and degenerate elements: The key step in the proof of the classical Lovász' theorem consists in showing that, if $|\text{hom}_{\Sigma_f}(C, A)| = |\text{hom}_{\Sigma_f}(C, B)|$ for all C in Σ_f , then also $|\text{inj}_{\Sigma_f}(C, A)| = |\text{inj}_{\Sigma_f}(C, B)|$ for all C in Σ_f , where $\text{inj}_{\Sigma_f}(C, A)$ is the set of all injective homomorphisms $C \hookrightarrow A$, and similarly for $\text{inj}_{\Sigma_f}(C, B)$. We aim to show the analogous fact in our setting, where injective homomorphisms are replaced by \mathcal{M} -morphisms.

The intuition behind the following definition is that a σ -homomorphism $f: C \rightarrow A$ is non-injective precisely when there exist an onto non-injective σ -homomorphism $e: C \twoheadrightarrow C'$ and a σ -homomorphism $g: C' \rightarrow A$ such that $f = g \circ e$. In terms of the hom-functor $E := \text{hom}_{\Sigma_f}(-, A)$, $f \in E(C)$ is non-injective if, and only if, $f = E(e)(g)$ for some onto non-injective σ -homomorphism $e: C \twoheadrightarrow C'$ and $g \in E(C')$.

Definition 10. A *strict quotient* in \mathcal{A} is an \mathcal{E} -morphism which is not an isomorphism. (Equivalently, by Lemma 7(b), a strict quotient is an arrow in $\mathcal{E} \setminus \mathcal{M}$.)

Further, consider a functor $E: \mathcal{A}^{\text{op}} \rightarrow \mathbf{FinSet}$. For every $k \in \mathcal{A}$, we say that $s \in E(k)$ is *degenerate* if there exist a strict quotient $f: k \twoheadrightarrow l$ and $t \in E(l)$ such that $E(f)(t) = s$. Otherwise, s is called *generic*. The subset of $E(k)$ consisting of the generic elements is denoted by $E[k]$.

The next lemma shows that this definition matches our intuition for hom-functors on \mathcal{A} .

Lemma 11. *Let $E = \text{hom}_{\mathcal{A}}(-, n)$ for some $n \in \mathcal{A}$. For any $k \in \mathcal{A}$, $E[k]$ is the set of all \mathcal{M} -morphisms $k \rightarrow n$.*

Proof. Let f be an arbitrary element of $E(k) = \text{hom}_{\mathcal{A}}(k, n)$. Assume f is generic, and take its $(\mathcal{E}, \mathcal{M})$ factorisation:

$$\begin{array}{ccccc} & & f & & \\ & & \downarrow & & \\ k & \xrightarrow{g} & l & \xrightarrow{h} & n \end{array}$$

Then $E(g)(h) = f$, and so g must be an \mathcal{M} -morphism. By Lemma 7(a), $f = h \circ g$ is also an \mathcal{M} -morphism.

Conversely, suppose f is an \mathcal{M} -morphism and pick an \mathcal{E} -morphism $g: k \twoheadrightarrow l$ and $h \in E(l)$ such that $E(g)(h) = f$, i.e. $h \circ g = f$. By Lemma 7(e), g is an \mathcal{M} -morphism, so it is not a strict quotient. That is, f is generic. \square

The following is the main technical lemma of this section, and it ultimately relies on an application of the inclusion-exclusion principle.

Lemma 12. *Assume \mathcal{A} has pushouts, and let E and F be functors $\mathcal{A}^{\text{op}} \rightarrow \mathbf{FinSet}$ sending \mathcal{E} -pushout squares in \mathcal{A} to pullbacks. If $E(k) \cong F(k)$ for all $k \in \mathcal{A}$, then $E[k] \cong F[k]$ for all $k \in \mathcal{A}$.*

Proof. Let E, F be as in the statement, and suppose that $E(k) \cong F(k)$ for all $k \in \mathcal{A}$. We must show that, for every $k \in \mathcal{A}$, $E[k] \cong F[k]$. Equivalently, we show that the set

$E(\llbracket k \rrbracket)$ of degenerate elements of $E(k)$ is in bijection with the set $F(\llbracket k \rrbracket)$ of degenerate elements of $F(k)$.

Observe that $E(\llbracket k \rrbracket)$ coincides with

$$\bigcup \{ \text{Im}(E(f)) \mid f: k \rightarrow l \text{ is a strict quotient in } \mathcal{A} \} \subseteq E(k),$$

and similarly for $F(\llbracket k \rrbracket)$. Since $E(k)$ and $F(k)$ are finite, there are finite sets $\mathcal{S}_1, \mathcal{S}_2$ of strict quotients of k such that

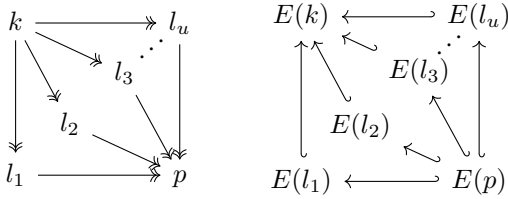
$$E(\llbracket k \rrbracket) = \bigcup \{ \text{Im}(E(f)) \mid f \in \mathcal{S}_1 \},$$

$$F(\llbracket k \rrbracket) = \bigcup \{ \text{Im}(F(f)) \mid f \in \mathcal{S}_2 \}.$$

Let $\mathcal{S} := \mathcal{S}_1 \cup \mathcal{S}_2$. By the inclusion-exclusion principle,

$$|E(\llbracket k \rrbracket)| = \sum_{J \subseteq \mathcal{S}, J \neq \emptyset} (-1)^{|J|+1} \left| \bigcap_{f \in J} \text{Im}(E(f)) \right|$$

and similarly for $|F(\llbracket k \rrbracket)|$. Therefore, to prove $E(\llbracket k \rrbracket) \cong F(\llbracket k \rrbracket)$ it suffices to show that $\left| \bigcap_{f \in J} \text{Im}(E(f)) \right| = \left| \bigcap_{f \in J} \text{Im}(F(f)) \right|$ for any non-empty $J \subseteq \mathcal{S}$. To this end, fix such a J and consider the *wide pushout* in \mathcal{A} of the strict quotients in J , i.e. the diagram obtained by taking consecutive pushouts of the elements of J , as shown in the diagram on the left-hand side below.



By Lemma 7(c), all arrows in this diagram are \mathcal{E} -morphisms. Hence, by Lemma 9, the diagram obtained by applying the functor E , depicted on the right-hand side above, consists of injections. Since E sends \mathcal{E} -pushouts to pullbacks, the diagram on the right is a wide pullback of injections in **FinSet**, and so $E(p) \cong \bigcap_{f \in J} \text{Im}(E(f))$. Similarly, $F(p) \cong \bigcap_{f \in J} \text{Im}(F(f))$. As $E(p) \cong F(p)$ by assumption, we get $\left| \bigcap_{f \in J} \text{Im}(E(f)) \right| = \left| \bigcap_{f \in J} \text{Im}(F(f)) \right|$, thus concluding the proof. \square

We are now in a position to prove the main result of this section:

Proof of Theorem 5. Fix arbitrary objects $m, n \in \mathcal{A}$. For the non-trivial direction, assume $|\text{hom}_{\mathcal{A}}(k, m)| = |\text{hom}_{\mathcal{A}}(k, n)|$ for all $k \in \mathcal{A}$. Let E and F denote, respectively, the functors $\text{hom}_{\mathcal{A}}(-, m): \mathcal{A}^{\text{op}} \rightarrow \mathbf{FinSet}$ and $\text{hom}_{\mathcal{A}}(-, n): \mathcal{A}^{\text{op}} \rightarrow \mathbf{FinSet}$. These functors send \mathcal{E} -pushout squares in \mathcal{A} to pullbacks by Lemma 8. Thus, by Lemma 12, $E(\llbracket k \rrbracket) \cong F(\llbracket k \rrbracket)$ for all $k \in \mathcal{A}$. According to Lemma 11, there is a bijection

$$\{\mathcal{M}\text{-morphisms } k \rightarrow m\} \cong \{\mathcal{M}\text{-morphisms } k \rightarrow n\}.$$

Setting $k = m$, the existence of an \mathcal{M} -morphism $m \rightarrow m$ (namely the identity) entails the existence of an \mathcal{M} -morphism $i: m \rightarrow n$. Similarly, there exists an \mathcal{M} -morphism $j: n \rightarrow m$. A standard argument then shows that $m \cong n$ (see e.g. [25]). We briefly sketch a proof. The set L of all \mathcal{M} -morphisms $m \rightarrow m$ is a monoid with respect to composition and contains

$j \circ i$. Since all \mathcal{M} -morphisms are monos, L satisfies the left cancellation law $ab = ac \Rightarrow b = c$. But every finite monoid satisfying the left cancellation law is a group, hence $j \circ i$ has an inverse. It follows from Lemma 7(b),(d) that j is an isomorphism. \square

Remark 13. Direct inspection of the preceding proofs shows that Theorem 5 can be slightly strengthened by requiring only the existence of pushouts of \mathcal{E} -morphisms along \mathcal{E} -morphisms.

C. Specialising to finite coalgebras

To conclude this section, we show that Theorem 5 can be used to lift Lovász' homomorphism counting result from Σ_f to categories of finite coalgebras for comonads on Σ . This result is then applied in Sections IV–V to obtain homomorphism counting results in finite model theory and modal logic, respectively.

Given a comonad \mathbb{C} on Σ , we say that a coalgebra (A, α) for \mathbb{C} is *finite* if A is a finite σ -structure. The full subcategory of $\text{EM}(\mathbb{C})$ defined by the finite coalgebras is denoted $\text{EM}_f(\mathbb{C})$.

Corollary 14. *Let \mathbb{C} be any comonad on Σ . Then $\text{EM}_f(\mathbb{C})$, the category of finite coalgebras for \mathbb{C} , is combinatorial.*

In order to prove the previous result, we need the following notions. Recall that, for any category \mathcal{A} and object $A \in \mathcal{A}$, two epimorphisms $f: A \rightarrow B$ and $g: A \rightarrow C$ are *equivalent*, written $f \sim g$, if there exists an isomorphism $h: B \rightarrow C$ such that $h \circ f = g$. It is not difficult to see that \sim is an equivalence relation on the collection of epimorphisms with domain A . The category \mathcal{A} is said to be *well-copowered* if, for each $A \in \mathcal{A}$, the collection of all \sim -equivalence classes of epimorphisms with domain A is a set (as opposed to a proper class). Any cocomplete category that is well-copowered admits a proper factorisation system, see e.g. [8, Propositions 4.4.2, 4.4.3] and [5, Proposition 14.7].

Moreover, recall that a functor $F: \mathcal{A} \rightarrow \mathcal{B}$ *creates colimits* if, for every diagram \mathcal{D} in \mathcal{A} such that $F(\mathcal{D})$ admits a colimit B in \mathcal{B} , there exists a colimit A of \mathcal{D} in \mathcal{A} with $F(A) \cong B$. In particular, if \mathcal{B} is cocomplete and F creates colimits, then \mathcal{A} is also cocomplete and the functor F preserves colimits. Further, F *creates isomorphisms* if, whenever $h: B \rightarrow F(A)$ is an isomorphism in \mathcal{B} , there exists an isomorphism $h': A' \rightarrow A$ in \mathcal{A} such that $F(h') = h$. It is a well known fact that, for any comonad \mathbb{C} on \mathcal{A} , the forgetful functor $U^{\mathbb{C}}: \text{EM}(\mathbb{C}) \rightarrow \mathcal{A}$ creates colimits and isomorphisms (see e.g. [5, Proposition 20.12] for a proof of the dual statement).

Proof of Corollary 14. We show that the category $\text{EM}_f(\mathbb{C})$ satisfies the hypotheses of Theorem 5. The forgetful functor $U^{\mathbb{C}}: \text{EM}(\mathbb{C}) \rightarrow \Sigma$ is faithful and restricts to a functor $\text{EM}_f(\mathbb{C}) \rightarrow \Sigma_f$. Thus, for all $(A, \alpha), (B, \beta) \in \text{EM}_f(\mathbb{C})$, there is an injection $\text{hom}_{\text{EM}_f(\mathbb{C})}((A, \alpha), (B, \beta)) \hookrightarrow \text{hom}_{\Sigma_f}(A, B)$. Because the latter set is finite, so is the former. Therefore, $\text{EM}_f(\mathbb{C})$ is locally finite.

We prove next that $\text{EM}_f(\mathbb{C})$ has all pushouts. Because the inclusion $\text{EM}_f(\mathbb{C}) \hookrightarrow \text{EM}(\mathbb{C})$ is full and faithful, it suffices to prove that the pushout in $\text{EM}(\mathbb{C})$ of finite coalgebras exists

and is a finite coalgebra. Since Σ is cocomplete (cf. the discussion following Definition 6) and $U^{\mathbb{C}}: \text{EM}(\mathbb{C}) \rightarrow \Sigma$ creates colimits, the category $\text{EM}(\mathbb{C})$ is cocomplete and the forgetful functor $U^{\mathbb{C}}: \text{EM}(\mathbb{C}) \rightarrow \Sigma$ preserves all colimits. In particular, the pushout of finite coalgebras exists in $\text{EM}(\mathbb{C})$ and is a finite coalgebra.

The functor $U^{\mathbb{C}}: \text{EM}(\mathbb{C}) \rightarrow \Sigma$ preserves colimits, and in particular epimorphisms. Using the fact that $U^{\mathbb{C}}$ is faithful and creates isomorphisms, it is not difficult to see that $\text{EM}(\mathbb{C})$ is well-copowered because so is Σ . Thus, $\text{EM}(\mathbb{C})$ admits a proper factorisation system $(\mathcal{E}, \mathcal{M})$. Now, consider a morphism $(A, \alpha) \rightarrow (B, \beta)$ in $\text{EM}_f(\mathbb{C})$ and its $(\mathcal{E}, \mathcal{M})$ -factorisation in $\text{EM}(\mathbb{C})$:

$$(A, \alpha) \xrightarrow{e} (\tilde{A}, \tilde{\alpha}) \xrightarrow{m} (B, \beta)$$

Since $U^{\mathbb{C}}$ preserves epimorphisms, $U(e): A \rightarrow \tilde{A}$ is an epimorphism in Σ , that is a surjective homomorphism. It follows that \tilde{A} is a finite σ -structure and so $(\tilde{A}, \tilde{\alpha}) \in \text{EM}_f(\mathbb{C})$.

The fact that $\text{EM}_f(\mathbb{C})$ admits a proper factorisation system is then a consequence of the following easy observation: Let \mathcal{A} be any category equipped with a proper factorisation system $(\mathcal{E}, \mathcal{M})$, and \mathcal{B} a full subcategory of \mathcal{A} . Assume that, whenever $A \rightarrow B$ is a morphism in \mathcal{B} and $A \rightarrow \tilde{A} \rightarrow B$ is its $(\mathcal{E}, \mathcal{M})$ -factorisation in \mathcal{A} , also $\tilde{A} \in \mathcal{B}$. Then $(\mathcal{E} \cap \mathcal{B}, \mathcal{M} \cap \mathcal{B})$ is a proper factorisation system in \mathcal{B} . \square

Remark 15. If \mathbb{C} is a comonad on Σ that restricts to a comonad \mathbb{C}' on Σ_f , then $\text{EM}(\mathbb{C}')$ is isomorphic to $\text{EM}_f(\mathbb{C})$ and thus $\text{EM}(\mathbb{C}')$ is combinatorial by Corollary 14. Also, note that the same proof as for Corollary 14 applies, mutatis mutandis, if Σ is replaced by the category of *pointed* σ -structures, having as objects the pairs (A, a) with $A \in \Sigma$ and $a \in A$, and as morphisms the σ -homomorphisms preserving the distinguished points. This will be needed in Section V for applications in the setting of modal logic.

IV. APPLICATIONS TO FINITE MODEL THEORY

In this section, we apply Corollary 14 to the game comonads introduced in Section II to give new proofs of results of Grohe and Dvořák connecting homomorphism counts and indistinguishability in appropriate logic fragments. We start by establishing a general result (Theorem 16) and then proceed to show how Grohe's and Dvořák's theorems follow from it.

Before stating this general result, we illustrate the underlying ideas for a generic game comonad \mathbb{C} . In Section II we have recalled how game comonads capture combinatorial parameters of σ -structures. On the other hand, they can also express equivalence with respect to appropriate logic fragments. For this, we need to extend the signature to account for equality in the logic. Let $\sigma^+ := \sigma \cup \{I\}$ be the relational signature obtained by adding a binary relation symbol I to σ , and let Σ^+ be the category of σ^+ -structures with homomorphisms. There is an *embedding*, i.e. a full and faithful functor, $J: \Sigma \rightarrow \Sigma^+$ which sends a σ -structure A to the σ^+ -structure $J(A)$ where I is interpreted as the identity.

As the comonad \mathbb{C} on Σ is defined for arbitrary relational signatures, we have a corresponding comonad \mathbb{C}^+ on Σ^+ . Typically, it turns out that, for an appropriate logic fragment \mathcal{L} ,

$$A \equiv_{\mathcal{L}} B \iff F^{\mathbb{C}^+}(J(A)) \cong F^{\mathbb{C}^+}(J(B))$$

for all $A, B \in \Sigma_f$, where $U^{\mathbb{C}^+}: \Sigma^+ \rightarrow \text{EM}(\mathbb{C}^+)$ is the adjunction associated with \mathbb{C}^+ , cf. Section II-C.

Provided that \mathbb{C}^+ sends finite structures to finite structures, we can apply Corollary 14 to translate the isomorphism $F^{\mathbb{C}^+}(J(A)) \cong F^{\mathbb{C}^+}(J(B))$ into a statement about homomorphism counts for $J(A)$ and $J(B)$ in Σ_f^+ . It then remains to go back to σ -structures to obtain a statement about homomorphism counts for A and B in Σ_f . To this end, note that the functor $J: \Sigma \rightarrow \Sigma^+$ has a left adjoint

$$H: \Sigma^+ \rightarrow \Sigma$$

sending a σ^+ -structure D to D^-/\sim , where D^- is the σ -reduct of D and \sim is the equivalence relation generated by I^D (for a proof of this fact, see Lemma 25).

We now state the general result (whose proof is deferred to Section IV-C) from which Grohe's and Dvořák's theorems will be derived in Sections IV-A and IV-B, respectively. Let us write $\text{im}(U^{\mathbb{C}})$ for the full subcategory of Σ consisting of the objects of the form $U^{\mathbb{C}}(A, \alpha) = A$ for $(A, \alpha) \in \text{EM}(\mathbb{C})$, and similarly for $\text{im}(U^{\mathbb{C}^+})$.

Theorem 16. *Let \mathbb{C} and \mathbb{C}^+ be comonads on Σ and Σ^+ , respectively, and \mathcal{L} a logic fragment. Suppose the following conditions are satisfied:*

- 1) *for all $A, B \in \Sigma_f$, $A \equiv_{\mathcal{L}} B$ if, and only if, $F^{\mathbb{C}^+}(J(A)) \cong F^{\mathbb{C}^+}(J(B))$;*
- 2) *\mathbb{C}^+ sends finite σ^+ -structures to finite σ^+ -structures;*
- 3) *the embedding $J: \Sigma \rightarrow \Sigma^+$ and its left adjoint H restrict to $\Sigma_f \cap \text{im}(U^{\mathbb{C}})$ and $\Sigma_f^+ \cap \text{im}(U^{\mathbb{C}^+})$.*

Then, for any finite σ -structures A and B ,

$$A \equiv_{\mathcal{L}} B \text{ if, and only if, } |\text{hom}_{\Sigma_f}(C, A)| = |\text{hom}_{\Sigma_f}(C, B)|$$

for every finite σ -structure C in $\text{im}(U^{\mathbb{C}})$.

Remark 17. The proofs of Grohe's and Dvořák's theorems essentially reduce to showing that the assumptions of Theorem 16 are satisfied for the appropriate comonads. The “combinatorial core” of these results, which requires a specific argument for each comonad, corresponds to verifying that the functor H restricts to $\Sigma_f^+ \cap \text{im}(U^{\mathbb{C}^+}) \rightarrow \Sigma_f \cap \text{im}(U^{\mathbb{C}})$. This amounts to saying that the operation $D \mapsto D^-/\sim$ does not increase the tree-depth or tree-width of D , and can be understood as an equality elimination result (cf. Section VI).

A. Bounded tree-depth

Let \mathcal{C} be the extension of first-order logic obtained by adding, for each natural number i , a counting quantifier $\exists_{\geq i}$. The semantics of these quantifiers is the following: for any structure A , we have $A \models \exists_{\geq i} x. \varphi$ if, and only if, $A \models \varphi[a/x]$ holds for at least i distinct elements $a \in A$. Let \mathcal{C}_n be the fragment of \mathcal{C} consisting of formulas of quantifier depth $\leq n$.

The aim of this section is to use Theorem 16 to give a new proof of a recent result of Grohe [14]:⁵

Theorem 18. *For any A, B finite σ -structures, $A \equiv_{C_n} B$ if, and only if, $|\text{hom}_{\Sigma_f}(C, A)| = |\text{hom}_{\Sigma_f}(C, B)|$ for every finite σ -structure C with tree-depth at most n .*

Let \mathbb{E}_n^+ be the Ehrenfeucht–Fraïssé comonad on Σ^+ , and $U^{\mathbb{E}_n^+} \dashv F^{\mathbb{E}_n^+} : \Sigma^+ \rightarrow \text{EM}(\mathbb{E}_n^+)$ the associated adjunction (cf. Section II-C). By Proposition 1, a σ -structure has tree-depth at most n precisely when it belongs to $\text{im}(U^{\mathbb{E}_n^+})$. Thus, to prove Theorem 18, it suffices to show that the assumptions of Theorem 16 are satisfied for $\mathbb{C} = \mathbb{E}_n$, $\mathbb{C}^+ = \mathbb{E}_n^+$, and $\mathcal{L} = \mathcal{C}_n$.

Item 1 in Theorem 16 is a consequence of [3, Theorem 12]:

Proposition 19. *For any two σ -structures A, B , we have $A \equiv_{C_n} B$ if, and only if, $F^{\mathbb{E}_n^+}(J(A)) \cong F^{\mathbb{E}_n^+}(J(B))$.*

Item 2 holds because the comonad \mathbb{E}_n^+ restricts to finite σ^+ -structures, and item 3 is a consequence of the following proposition—which concludes the proof of Grohe’s theorem.

Proposition 20. *The embedding $J : \Sigma \rightarrow \Sigma^+$ and its left adjoint H restrict to the full subcategories consisting of finite structures with tree-depth at most n .*

Proof. Clearly, if A is a (finite) σ -structure with tree-depth at most n , then $J(A)$ has tree-depth at most n . It remains to prove that, for any $D \in \Sigma_f^+$, if D has tree-depth at most n then so does $H(D)$. To improve readability, set $C := H(D)$.

Suppose (F, \leq) is a forest of height $\leq n$ and $f : D \rightarrow F$ an injective map such that $f(d_1) \uparrow f(d_2)$ whenever $d_1 \sim d_2$ in the Gaifman graph \mathcal{G}_D . Recall from the definition of H that, at the level of sets, $C = D/\sim$ where \sim is the equivalence relation generated by I^D . If $d \in D$, let $[d]$ denote its \sim -equivalence class. For each $d \in D$, set

$$\xi(d) := \min \{f(d') \mid d' \in [d]\}.$$

This minimum exists because $\{f(d') \mid d' \in [d]\}$ is a connected subset of F . (Alternatively, note that $\xi(d)$ coincides with the minimum of $\{f(d') \mid d' \in [d]\} \cap \downarrow f(d)$, which is a finite non-empty totally ordered set.) Define

$$g : C \rightarrow F, \quad g([d]) := \xi(d).$$

Clearly, g is well-defined and it is injective because so is f . It remains to prove that, for all $d_1, d_2 \in D$, $[d_1] \sim [d_2]$ in \mathcal{G}_C entails $g([d_1]) \uparrow g([d_2])$, for then it follows that (F, \leq) , along with the map g , is a forest cover of \mathcal{G}_C .

If $[d_1] \sim [d_2]$, there are $d'_1 \in [d_1]$ and $d'_2 \in [d_2]$ satisfying $d'_1 \sim d'_2$, and so $f(d'_1) \uparrow f(d'_2)$. Assume $f(d'_1) \leq f(d'_2)$. Since $g([d_1]) \leq f(d'_1)$ and $g([d_2]) \leq f(d'_2)$, it follows that both $g([d_1])$ and $g([d_2])$ belong to $\downarrow f(d'_2)$, which is totally ordered. Therefore, $g([d_1]) \uparrow g([d_2])$. Reasoning in a similar manner, we see that $f(d'_2) \leq f(d'_1)$ implies $g([d_1]) \uparrow g([d_2])$. \square

⁵Grohe proved this result for undirected vertex-coloured graphs, but pointed out that his proof can be extended to arbitrary σ -structures.

B. Bounded tree-width

Let \mathcal{C}^k be the k -variable fragment of \mathcal{C} —first-order logic with counting quantifiers. In this section, we show how to derive from Theorem 16 the following variant of Dvořák’s theorem [11]:⁶

Theorem 21. *For any A, B finite σ -structures, $A \equiv_{C^k} B$ if, and only if, $|\text{hom}_{\Sigma_f}(C, A)| = |\text{hom}_{\Sigma_f}(C, B)|$ for every finite σ -structure C with tree-width at most $k - 1$.*

The k -variable counting logic \mathcal{C}^k appearing in Theorem 21 is captured by the pebbling comonad \mathbb{P}_k^+ on Σ^+ . In fact, it follows directly from [1, Theorem 18] that, whenever A, B are finite σ -structures, $F^{\mathbb{P}_k^+}(J(A)) \cong F^{\mathbb{P}_k^+}(J(B))$ if, and only if, $A \equiv_{C^k} B$. However, the comonad \mathbb{P}_k^+ does not satisfy item 2 in Theorem 16, because $\mathbb{P}_k^+(A)$ is infinite even when $A \in \Sigma_f^+$.

To circumvent this problem, we consider the comonads $\mathbb{P}_{k,n}^+$, which do restrict to finite σ^+ -structures. Again by (the proof of) [1, Theorem 18], for any finite σ -structures A, B ,

$$A \equiv_{C_n^k} B \iff F^{\mathbb{P}_{k,n}^+}(J(A)) \cong F^{\mathbb{P}_{k,n}^+}(J(B))$$

where C_n^k consists of the formulas that are simultaneously in \mathcal{C}_n and \mathcal{C}^k . Thus, items 1 and 2 in Theorem 16 are satisfied for $\mathbb{C} = \mathbb{P}_{k,n}$, $\mathbb{C}^+ = \mathbb{P}_{k,n}^+$, and $\mathcal{L} = \mathcal{C}_n^k$.

Assume for a moment that item 3 is also satisfied. Then, combining Theorem 16 with Proposition 4, we obtain that

$$A \equiv_{C_n^k} B \text{ if, and only if, } |\text{hom}_{\Sigma_f}(C, A)| = |\text{hom}_{\Sigma_f}(C, B)|$$

for all finite σ -structures C whose Gaifman graph \mathcal{G}_C admits a k -pebble forest cover of height $\leq n$. Observe next that (i) $A \equiv_{C^k} B$ precisely when $A \equiv_{C_n^k} B$ for all n , and (ii) a finite σ -structure C has tree-width at most $k - 1$ if, and only if, \mathcal{G}_C admits a k -pebble forest cover of height $\leq n$ for some n . Therefore, $A \equiv_{C^k} B$ if, and only if,

$$|\text{hom}_{\Sigma_f}(C, A)| = |\text{hom}_{\Sigma_f}(C, B)|$$

for all finite σ -structures C with tree-width at most $k - 1$, thus settling Theorem 21.

In the remainder of this section, we prove that item 3 in Theorem 16 is indeed satisfied. As mentioned already, $\Sigma_f \cap \text{im}(U^{\mathbb{P}_{k,n}})$ consists of the finite σ -structures that admit a k -pebble forest cover of height $\leq n$, and similarly for $\Sigma_f^+ \cap \text{im}(U^{\mathbb{P}_{k,n}^+})$. To improve readability, let us denote these categories by $\mathcal{T}_{k,n}$ and $\mathcal{T}_{k,n}^+$, respectively. Note that $J : \Sigma_f \rightarrow \Sigma_f^+$ restricts to a functor $\mathcal{T}_{k,n} \rightarrow \mathcal{T}_{k,n}^+$. It remains to show that its left adjoint H restricts to $\mathcal{T}_{k,n}^+ \rightarrow \mathcal{T}_{k,n}$.

Recall from Section II-C that a forest cover of a σ -structure A can be identified with a forest order \leq on A that is *compatible*, in the sense that $a \uparrow a'$ whenever $a \sim a'$ in \mathcal{G}_A . Suppose that A is equipped with a forest order \leq and let $p : A \rightarrow \mathbf{k}$ be any function. We define a relation \bowtie on A as follows: $a \bowtie a'$ (read as *a sees a'*) if $a \uparrow a'$ and

$$\min(a, a') < z \leq \max(a, a') \implies p(z) \neq p(\min(a, a'))$$

⁶Dvořák’s result is for undirected graphs without loops. For a discussion of how his result can be recovered in our framework, see Remark 24.

for all $z \in A$. With this notation, a k -pebble forest cover of A can be identified with a triple (A, \leq, p) where \leq is a forest order on A and $p: A \rightarrow \mathbf{k}$ is such that $a \frown a'$ implies $a \bowtie a'$ for all $a, a' \in A$. Just observe that \leq is a compatible forest order since $a \bowtie a'$ entails $a \uparrow a'$.

Saying that H restricts to $\mathcal{T}_{k,n}^+ \rightarrow \mathcal{T}_{k,n}$ means that, whenever a finite σ^+ -structure A admits a k -pebble forest cover of height $\leq n$, then so does the quotient A^-/\sim where A^- is the σ -reduct of A and \sim is the least equivalence relation containing I^A . The difficulty in showing this consists in defining the pebbling function on A^-/\sim . We do this by considering consecutive one-step quotients (inspired by equality elimination from Section VI), where at each step a new pair of distinct elements in I^A is identified. As A is finite, this construction terminates after finitely many steps and yields the quotient A^-/\sim .

Let $\mathbb{A} := (A, \leq, p)$ be a finite σ^+ -structure together with a k -pebble forest cover of height $\leq n$. Suppose that there is a pair $(u, v) \in I^A$ such that $u \neq v$. Then $u \uparrow v$, and we can assume without loss of generality that $u < v$. The one-step quotient $\mathbb{A}/_{u \sim v} := (A', \leq', p')$ of \mathbb{A} is defined as follows:

- the carrier of $\mathbb{A}/_{u \sim v}$ is the set $A' := A \setminus \{v\}$;
- the forest order \leq' is the restriction of \leq to A' ;
- for each $R \in \sigma \cup \{I\}$, the relation $R^{A'}$ is obtained by replacing all occurrences of v by u in each tuple of R^A .
- the pebbling function $p': A' \rightarrow \mathbf{k}$ is defined by

$$w \mapsto \begin{cases} p(u) & \text{if } v \leq w, p(w) = p(v), \text{ and } u \not\bowtie w_{\min} \\ p(v) & \text{if } v \leq w, p(w) = p(u), \text{ and } v \bowtie w_{\min} \\ p(w) & \text{otherwise} \end{cases}$$

where $w_{\min} := \min\{w' \in A \mid v < w' \leq w \text{ and } p(w') = p(w)\}$.

Clearly, \leq' is a forest order of height $\leq n$ on A' . Further, $a \frown a'$ implies $a \bowtie a'$ for all $a, a' \in A'$. This follows from the next lemma, which is proved by a careful case analysis.

Lemma 22. *Let $w, w' \in A \setminus \{v\}$. Then the following hold:*

- 1) $v \bowtie w$ in \mathbb{A} implies $u \bowtie w$ in $\mathbb{A}/_{u \sim v}$;
- 2) $w \bowtie w'$ in \mathbb{A} implies $w \bowtie w'$ in $\mathbb{A}/_{u \sim v}$.

Proof. (1) Suppose $w \in A \setminus \{v\}$ and $v \bowtie w$ in \mathbb{A} . There are two possibilities, either $v < w$ or $w < v$. In both cases we have that $w \uparrow u$ and so we only need to check the condition on the pebbling function p' . Assume that $v < w$. As $(u, v) \in I^A$, we know that $u \frown v$ and so $u \bowtie v$ in \mathbb{A} . Because $p'(z) = p(z)$ whenever $v \not\leq z$, to show that $u \bowtie w$ in $\mathbb{A}/_{u \sim v}$ we only need to check that

$$\forall z \in A, \quad v < z \leq w \implies p'(z) \neq p'(u). \quad (1)$$

Observe that $p'(u) = p(u)$. Also, by definition of p' , for any z such that $v < z \leq w$ the equality $p'(z) = p(u)$ is satisfied if, and only if, either $p(z) = p(v)$ and $u \not\bowtie z_{\min}$, or $p(z) = p(u)$ and $v \bowtie z_{\min}$. However, none of these two scenarios can occur because $v \bowtie w$ implies $p(z) \neq p(v)$ and $v \bowtie z_{\min}$. Hence, the condition in (1) is satisfied.

On the other hand, assume that $w < v$. Then $u \uparrow w$. If $u < w$ then $u \bowtie v$ implies $u \bowtie w$ in \mathbb{A} . Similarly, if $w \leq u$ then $v \bowtie w$ entails $u \bowtie w$ in \mathbb{A} . Therefore, $u \bowtie w$ in $\mathbb{A}/_{u \sim v}$ because $p(z) = p'(z)$ for every z such that $v \not\leq z$.

(2) Let $w, w' \in A \setminus \{v\}$ be such that $w \bowtie w'$ in \mathbb{A} . We can assume without loss of generality that $w \leq w'$. If $v \not\leq w'$ then $p(z) = p'(z)$ for every z such that $w \leq z \leq w'$, and so $w \bowtie w'$ in $\mathbb{A}/_{u \sim v}$. Thus, suppose $v \leq w'$. It follows that $v \uparrow w$ and so either $v < w$ or $w < v$.

First we assume $v < w \leq w'$. Note that, for any z such that $w < z \leq w'$, if $p'(z) = p'(w)$ and $p(z) \notin \{p(u), p(v)\}$, then by definition of p' we must have $p(z) = p(w)$, contradicting the fact that $w \bowtie w'$ in \mathbb{A} . So, we only have to take care of the case where either $p(z) = p(u)$ or $p(z) = p(v)$. Set

$$U := \{x \in A \mid v < x \leq w' \text{ and } p(x) = p(u)\} \\ V := \{x \in A \mid v < x \leq w' \text{ and } p(x) = p(v)\}.$$

Suppose U and V are non-empty and so they admit minimal elements \bar{u} and \bar{v} , respectively. Then, either $\bar{u} < \bar{v}$ or $\bar{v} < \bar{u}$. If $\bar{u} < \bar{v}$ then any $z \geq \bar{u}$ such that $p(z) = p(u)$ satisfies $v \bowtie z_{\min}$ (because $z_{\min} = \bar{u}$) and so $p'(z) = p(v)$, and similarly any $z \geq \bar{v}$ such that $p(z) = p(v)$ satisfies $u \not\bowtie z_{\min}$ (because $z_{\min} = \bar{v}$) and so $p'(z) = p(u)$. On the other hand, if $\bar{v} < \bar{u}$ then $p'(z) = p(z)$ for any $z \geq \bar{v}$. Either way, for any z between w and w' , $p'(z) = p'(w)$ implies $p(z) = p(w)$. Hence, $w \bowtie w'$ in \mathbb{A} entails $w \bowtie w'$ in $\mathbb{A}/_{u \sim v}$. A similar reasoning applies, mutatis mutandis, when $U = \emptyset$ or $V = \emptyset$.

Suppose next that $w < v < w'$. Then $p'(w) = p(w)$ and, as $w \bowtie w'$ in \mathbb{A} , it follows that $p(w) \neq p(v)$. If $p(w) \neq p(u)$ then we conclude that $w \bowtie w'$ in $\mathbb{A}/_{u \sim v}$ because $p'(z) = p'(w)$ entails $p(z) = p(w)$ for any element z between w and w' . To conclude the proof, suppose towards a contradiction that $p(w) = p(u)$ and there exists a $z \in A \setminus \{v\}$ such that $w < z \leq w'$ and $p'(z) = p'(w)$. Then, $p'(z) = p(u)$. Since $w \bowtie w'$ in \mathbb{A} , we get $p(z) \neq p(u)$ and so $v < z$, $p(z) = p(v)$ and $u \not\bowtie z_{\min}$ in \mathbb{A} . As $u \bowtie v$ in \mathbb{A} , there must exist an x such that

$$w < v < x \leq z_{\min} \leq z \leq w'$$

and $p(x) = p(u)$, contradicting the fact that $w \bowtie w'$ in \mathbb{A} . \square

Applying this construction iteratively, we obtain the following analogue of Proposition 20, which concludes the proof of Dvořák's theorem.

Proposition 23. *The functor H restricts to $\mathcal{T}_{k,n}^+ \rightarrow \mathcal{T}_{k,n}$.*

Proof. Let $\mathbb{A}_0 = (A, \leq, p)$ be a finite σ^+ -structure together with a k -pebble forest cover of height $\leq n$. Then we have a sequence $\mathbb{A}_0, \mathbb{A}_1, \mathbb{A}_2, \dots$ such that \mathbb{A}_{i+1} is defined as $\mathbb{A}_i /_{u \sim v}$ for some pair of distinct elements $(u, v) \in I^{A_i}$, where A_i is the σ^+ -structure underlying \mathbb{A}_i . This is iterated as long as such a pair (u, v) can be found. Let \mathbb{A}_n be the last element of the sequence. Note that, by construction, the σ -reduct of A_n is isomorphic to $H(A)$. Because A_n admits a k -pebble forest cover of height $\leq n$ witnessed by \mathbb{A}_n —by repeatedly applying Lemma 22—so does $H(A)$. \square

Remark 24. An advantage of the categorical approach to homomorphism counting is that it specialises to any full subcategory \mathcal{A} of Σ , provided the game comonad in question restricts to \mathcal{A} . For example, let σ consist of a single binary relation symbol, so that Σ is the category of (directed) graphs and graph homomorphisms. As the pebbling comonad \mathbb{P}_k restricts to the full subcategory of Σ defined by undirected graphs without loops, we obtain an analogue of Theorem 21 where A, B, C range over the class of finite undirected graphs without loops. A similar observation applies to the Ehrenfeucht–Fraïssé comonad \mathbb{E}_n , yielding a variant of Theorem 18 for undirected graphs without loops.

C. A proof of Theorem 16

We start by showing that the functor $H: \Sigma^+ \rightarrow \Sigma$ sending D to D^-/\sim , where D^- is the σ -reduct of D and \sim is the equivalence relation generated by I^D , is left adjoint to J .

Lemma 25. $H: \Sigma^+ \rightarrow \Sigma$ is left adjoint to the embedding J .

Proof. It suffices to show that for every $D \in \Sigma^+$, the unique σ^+ -homomorphism $e_D: D \rightarrow JH(D)$ extending the quotient map $D^- \rightarrow H(D)$ satisfies the following universal property: For every σ^+ -homomorphism $j: D \rightarrow J(C)$, there exists a unique σ -homomorphism $h: H(D) \rightarrow C$ making the following diagram commute.

$$\begin{array}{ccc} D & \xrightarrow{e_D} & JH(D) \\ & \searrow j & \downarrow J(h) \\ & & J(C) \end{array}$$

Pick an arbitrary σ^+ -homomorphism $j: D \rightarrow J(C)$. Since j preserves the relation I , $d_1 \sim d_2$ entails $j(d_1) = j(d_2)$ for all $d_1, d_2 \in D$. Hence, the following map is well defined

$$h: H(D) \rightarrow C, \quad h([d]) := j(d)$$

where $[d]$ denotes the \sim -equivalence class of d . Moreover, h is a σ -homomorphism. Just observe that, for any relation symbol $R \in \sigma$ of arity n , if $([d_1], \dots, [d_n]) \in R^{H(D)}$ then there are $d'_1 \sim d_1, \dots, d'_n \sim d_n$ such that $(d'_1, \dots, d'_n) \in R^D$. Thus, $(h([d_1]), \dots, h([d_n])) = (j(d'_1), \dots, j(d'_n)) \in R^C$. Further, $j = J(h) \circ e_D$ by construction. To conclude, note that h is unique with this property because e_D is surjective. \square

Remark 26. Since both J and H send finite structures to finite structures, they restrict to an adjunction $H \dashv J: \Sigma_f \rightarrow \Sigma_f^+$.

Before proceeding to the proof of Theorem 16, we make the following easy observation.

Lemma 27. Assume that J and H restrict to full subcategories $\mathcal{A} \subseteq \Sigma_f$ and $\mathcal{A}^+ \subseteq \Sigma_f^+$. Then, for any $A, B \in \Sigma_f$, the following are equivalent:

- 1) $\text{hom}_{\Sigma_f}(C, A) \cong \text{hom}_{\Sigma_f}(C, B) \quad \forall C \in \mathcal{A}$;
- 2) $\text{hom}_{\Sigma_f^+}(C', J(A)) \cong \text{hom}_{\Sigma_f^+}(C', J(B)) \quad \forall C' \in \mathcal{A}^+$.

Proof. Assume item 1 holds. For any $C' \in \mathcal{A}^+$,

$$\begin{aligned} \text{hom}_{\Sigma_f^+}(C', J(A)) &\cong \text{hom}_{\Sigma_f}(H(C'), A) && (H \dashv J) \\ &\cong \text{hom}_{\Sigma_f}(H(C'), B) && (H(C') \in \mathcal{A}) \\ &\cong \text{hom}_{\Sigma_f^+}(C', J(B)). && (H \dashv J) \end{aligned}$$

Conversely, suppose item 2 holds. Then, for any $C \in \mathcal{A}$,

$$\begin{aligned} \text{hom}_{\Sigma_f}(C, A) &\cong \text{hom}_{\Sigma_f^+}(J(C), J(A)) && (J \text{ embedding}) \\ &\cong \text{hom}_{\Sigma_f^+}(J(C), J(B)) && (J(C) \in \mathcal{A}^+) \\ &\cong \text{hom}_{\Sigma_f}(C, B). && (J \text{ embedding}) \end{aligned}$$

This concludes the proof. \square

Proof of Theorem 16. For convenience of notation, let $\text{EM}_f^+ := \text{EM}_f(\mathbb{C}^+)$. Combining item 1 with Corollary 14 applied to the comonad \mathbb{C}^+ , we see that, for all $A, B \in \Sigma_f$, $A \leq_{\mathbb{C}} B$ if, and only if,

$$|\text{hom}_{\text{EM}_f^+}(D, F^{\mathbb{C}^+}(J(A)))| = |\text{hom}_{\text{EM}_f^+}(D, F^{\mathbb{C}^+}(J(B)))| \quad (2)$$

for all finite coalgebras $D \in \text{EM}_f^+$. Note that here we used the fact that, by item 2, $F^{\mathbb{C}^+}(D')$ is a finite coalgebra whenever D' is a finite σ^+ -structure. In view of the adjunction $U^{\mathbb{C}^+} \dashv F^{\mathbb{C}^+}$, the condition in equation (2) is equivalent to

$$|\text{hom}_{\Sigma_f^+}(U^{\mathbb{C}^+}(D), J(A))| = |\text{hom}_{\Sigma_f^+}(U^{\mathbb{C}^+}(D), J(B))|$$

for all $D \in \text{EM}_f^+$. Finally, by item 3, an application of Lemma 27 with $\mathcal{A} := \Sigma_f \cap \text{im}(U^{\mathbb{C}})$ and $\mathcal{A}^+ := \Sigma_f^+ \cap \text{im}(U^{\mathbb{C}^+})$ yields the desired statement. \square

V. HOMOMORPHISM COUNTING IN MODAL LOGIC

In this section we prove a new homomorphism counting result for (multi-)modal logic. This is derived by applying the categorical framework from Section III to a game comonad for modal logic defined on the category of pointed Kripke structures. Since there is no equality in the logic, this time we can dispense with the extra relation I . This considerably simplifies the proofs, compared to the previous section.

Let σ be a signature consisting of relation symbols of arity at most 2. Each unary relation symbol P yields a propositional variable p , and each binary relation symbol R_α yields modalities \Box_α and \Diamond_α . We can think of a σ -structure A as a Kripke structure for this multi-modal logic, where P^A gives the valuation for the propositional variable p , and R_α^A gives the accessibility relation for the modalities \Box_α and \Diamond_α .

Each formula φ in this modal logic admits a translation into a first-order formula $\llbracket \varphi \rrbracket_x$ in one free variable x ; this is known as the *standard translation*, see e.g. [7, §2.4]. We let $\llbracket p \rrbracket_x := P(x)$ and let $\llbracket \cdot \rrbracket_x$ commute with Boolean connectives. Further, set $\llbracket \Box_\alpha \varphi \rrbracket_x := \forall y. R_\alpha(x, y) \rightarrow \llbracket \varphi \rrbracket_y$ and $\llbracket \Diamond_\alpha \varphi \rrbracket_x := \exists y. R_\alpha(x, y) \wedge \llbracket \varphi \rrbracket_y$ where y is a fresh variable. Consider a *pointed Kripke structure*, that is a pair (A, a) where A is a σ -structure and $a \in A$. Then $A, a \models \varphi$ according to Kripke semantics if, and only if, $A \models \llbracket \varphi \rrbracket_x[a/x]$ in the standard model-theoretic sense.

Henceforth, we assume all Kripke structures under consideration are pointed. Let \mathbf{Kri} be the category of Kripke structures and σ -homomorphisms preserving the distinguished element. For every natural number k , there is a *modal comonad* \mathbb{M}_k on \mathbf{Kri} , see [3]. For any Kripke structure (A, a) , the carrier of $\mathbb{M}_k(A, a)$ is the set of all paths of length $\leq k$ starting from a :

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

where R_1, \dots, R_n are binary relations from σ . The distinguished element of $\mathbb{M}_k(A, a)$ is the trivial path (a) of length 0. If $P \in \sigma$ is unary then $P^{\mathbb{M}_k(A, a)}$ is defined as the set of paths p such that the last element a_n of p belongs to P^A . For a binary relation symbol $R \in \sigma$, let $R^{\mathbb{M}_k(A, a)}$ be the set of pairs of paths (p, p') such that p' is obtained by extending p by one step along R . The morphism $\varepsilon_{(A, a)}: \mathbb{M}_k(A, a) \rightarrow (A, a)$ sends a path to its last element; for the definition of the coextension operation, see [3, §2.3].

The comonad \mathbb{M}_k captures a well known combinatorial parameter of Kripke structures, as we now recall. Let us say that a Kripke structure (A, a) is *rooted* if, for any $a' \in A$, there is a path from a to a' . If we further require that this path be unique, then (A, a) is a *synchronization tree*. A synchronization tree (A, a) has *height* $\leq k$ if, for each $a' \in A$, the length of the unique path from a to a' is at most k . For a proof of the following result, see [4, Proposition 6.6].

Proposition 28. *A rooted Kripke structure (A, a) admits a coalgebra structure $(A, a) \rightarrow \mathbb{M}_k(A, a)$ if, and only if, (A, a) is a synchronization tree of height $\leq k$.*

On the other hand, the modal comonad \mathbb{M}_k characterises logical equivalence with respect to an appropriate modal logic. Let us extend the multi-modal logic introduced above by adding, for every positive integer n and binary relation symbol $R_\alpha \in \sigma$, *graded modalities* \Box_α^n and \Diamond_α^n defining $A, a \models \Diamond_\alpha^n \varphi$ if there are at least n distinct R_α -successors of a that satisfy φ (and $\Box_\alpha^n \varphi = \neg \Diamond_\alpha^n \neg \varphi$). A formula in this graded modal logic has *modal depth* $\leq k$ if it contains at most k nested modalities. Let $\mathcal{C}_k^{\text{ML}}$ be the image of the standard translation of formulas with graded modalities that have modal depth $\leq k$. The following result is a special case of [3, Theorem 16].

Proposition 29. *For finite Kripke structures (A, a) and (B, b) ,*

$$(A, a) \equiv_{\mathcal{C}_k^{\text{ML}}} (B, b) \quad \text{iff} \quad F^{\mathbb{M}_k}(A, a) \cong F^{\mathbb{M}_k}(B, b).$$

We aim to exploit the two preceding propositions to obtain a characterisation of the equivalence relation $\equiv_{\mathcal{C}_k^{\text{ML}}}$ over finite Kripke structures in terms of homomorphism counts from finite synchronization trees of height $\leq k$ (Theorem 31 below).

To this end, let \mathbf{Kri}_f and \mathbf{RKri}_f be the full subcategories of \mathbf{Kri} consisting, respectively, of the finite Kripke structures, and finite rooted Kripke structures. Note that \mathbb{M}_k restricts to a comonad on \mathbf{Kri}_f , that we denote again by \mathbb{M}_k . Further, as $\mathbb{M}_k(A, a)$ is a finite synchronization tree whenever (A, a) is a finite Kripke structure, \mathbb{M}_k restricts to a comonad \mathbb{M}_k^* on the category \mathbf{RKri}_f .

Lemma 30. *The following statements hold:*

- 1) *the inclusion $\mathbf{RKri}_f \hookrightarrow \mathbf{Kri}_f$ has a right adjoint R ;*
- 2) *R lifts to a functor $\bar{R}: \mathbf{EM}(\mathbb{M}_k) \rightarrow \mathbf{EM}(\mathbb{M}_k^*)$ making the following diagram commute.*

$$\begin{array}{ccc} \mathbf{EM}(\mathbb{M}_k) & \xrightarrow{\bar{R}} & \mathbf{EM}(\mathbb{M}_k^*) \\ \uparrow F^{\mathbb{M}_k} & & \uparrow F^{\mathbb{M}_k^*} \\ \mathbf{Kri}_f & \xrightarrow{R} & \mathbf{RKri}_f \end{array}$$

Proof. (1) For any finite Kripke structure (A, a) , let $R(A, a)$ be the substructure of (A, a) consisting of the elements accessible from a . To show that this assignment extends to a functor R that is a right adjoint to the inclusion $\mathbf{RKri}_f \hookrightarrow \mathbf{Kri}_f$, it is enough to observe that, for any $(A, a) \in \mathbf{Kri}_f$, $(B, b) \in \mathbf{RKri}_f$, and homomorphism $h: (B, b) \rightarrow (A, a)$, there exists a unique homomorphism $h': (B, b) \rightarrow R(A, a)$ such that $h = e \circ h'$, where $e: R(A, a) \rightarrow (A, a)$ is the substructure embedding.

(2) Let \bar{R} be the functor sending a coalgebra $\alpha: (A, a) \rightarrow \mathbb{M}_k(A, a)$ to its restriction $\alpha|_{R(A, a)}: R(A, a) \rightarrow \mathbb{M}_k^* R(A, a)$. Then $\bar{R}F^{\mathbb{M}_k}(A, a)$ coincides with $F^{\mathbb{M}_k^*} R(A, a)$, for any (A, a) in \mathbf{Kri}_f , because the carrier of $F^{\mathbb{M}_k}(A, a) = (\mathbb{M}_k(A, a), \delta_A)$ is rooted and consists only of paths of length $\leq k$ that start at a . The equality on morphisms holds for the same reason. \square

We are now ready to prove a Lovász-type theorem for graded modal logic.

Theorem 31. *Let (A, a) and (B, b) be finite (pointed) Kripke structures. Then $(A, a) \equiv_{\mathcal{C}_k^{\text{ML}}} (B, b)$ if, and only if,*

$$|\text{hom}_{\mathbf{Kri}_f}((C, c), (A, a))| = |\text{hom}_{\mathbf{Kri}_f}((C, c), (B, b))|$$

for every finite synchronization tree (C, c) of height $\leq k$.

Proof. Let $\mathbb{A} := (A, a)$ and $\mathbb{B} := (B, b)$. For the purpose of this proof, write $V: \mathbf{RKri}_f \hookrightarrow \mathbf{Kri}_f$ for the inclusion functor and let $\mathbf{EM}^* := \mathbf{EM}(\mathbb{M}_k^*)$. By Proposition 28, the statement about homomorphism counts is equivalent to

$$|\text{hom}_{\mathbf{Kri}_f}(V(U^{\mathbb{M}_k^*}(\mathbb{D})), \mathbb{A})| = |\text{hom}_{\mathbf{Kri}_f}(V(U^{\mathbb{M}_k^*}(\mathbb{D})), \mathbb{B})|$$

for all coalgebras \mathbb{D} in \mathbf{EM}^* . In turn, by virtue of the adjunctions $V \dashv R$ and $U^{\mathbb{M}_k^*} \dashv F^{\mathbb{M}_k^*}$, this is equivalent to

$$|\text{hom}_{\mathbf{EM}^*}(\mathbb{D}, F^{\mathbb{M}_k^*}(R(\mathbb{A})))| = |\text{hom}_{\mathbf{EM}^*}(\mathbb{D}, F^{\mathbb{M}_k^*}(R(\mathbb{B})))|$$

for all $\mathbb{D} \in \mathbf{EM}^*$. Since \mathbf{EM}^* is combinatorial by Corollary 14 and Remark 15, the last statement is equivalent to $F^{\mathbb{M}_k^*}(R(\mathbb{A})) \cong F^{\mathbb{M}_k^*}(R(\mathbb{B}))$ and thus, by Lemma 30, to $\bar{R}(F^{\mathbb{M}_k}(\mathbb{A})) \cong \bar{R}(F^{\mathbb{M}_k}(\mathbb{B}))$. Note that the restriction of \bar{R} to the image of $F^{\mathbb{M}_k}$ is full and faithful. Hence, $\bar{R}(F^{\mathbb{M}_k}(\mathbb{A})) \cong \bar{R}(F^{\mathbb{M}_k}(\mathbb{B}))$ if, and only if, $F^{\mathbb{M}_k}(\mathbb{A}) \cong F^{\mathbb{M}_k}(\mathbb{B})$ which, by Proposition 29, is equivalent to $\mathbb{A} \equiv_{\mathcal{C}_k^{\text{ML}}} \mathbb{B}$. \square

VI. LOGICAL NORMAL FORMS

Let \mathcal{C}_∞^k denote the closure of \mathcal{C}^k under infinitary conjunctions and disjunctions. This infinitary logic has been much studied in finite model theory as a means of delimiting the expressive power of fixed-point logics with counting. In this

logic we can express all, and only, the properties of finite structures that are invariant under the equivalence relation $\equiv_{\mathcal{C}^k}$. In fact, since it is known that each finite structure is characterised up to $\equiv_{\mathcal{C}^k}$ by a single sentence of \mathcal{C}^k , it follows that every class of finite structures that is invariant under $\equiv_{\mathcal{C}^k}$ is defined by a single infinitary disjunction of \mathcal{C}^k sentences (see [23] for details). An entirely analogous normal form holds for $\mathcal{C}_{n,\infty}$ —the closure of \mathcal{C}_n under infinitary conjunctions and disjunctions (see [19, Chapter 8]).

The theorems of Dvořák and Grohe provide a route to alternative normal forms for these infinitary logics. To see this, note that for any finite structure A , we can write a primitive positive sentence γ_A of first-order logic such that for any B , $B \models \gamma_A$ if, and only if, there is a homomorphism from A to B . Here, *primitive positive* means that the sentence is built up from atomic formulas using only conjunctions and existential quantification. The sentence γ_A is called the *canonical conjunctive query of A* (see [13, Chapter 6]). It is known that if A has tree-width strictly less than k , then γ_A can be written using no more than k variables [18, Lemma 5.2]. Similarly, if A has tree-depth at most n , then γ_A can be chosen to have quantifier depth at most n [28, Lemma 2.14].

Given a positive integer t , we can transform γ_A by standard means into a sentence γ_A^t of \mathcal{C} which asserts that there are at least t distinct homomorphisms from A . That is, for any B , $B \models \gamma_A^t$ if, and only if, $|\text{hom}(A, B)| \geq t$. To show this, we establish something more general.

Let $\gamma(\bar{x})$ be a primitive positive formula with free variables among \bar{x} . For any structure B and any interpretation ι taking the variables in \bar{x} to elements of B , we write $B \models \gamma[\iota]$ to indicate that the formula γ is satisfied in B when the free variables \bar{x} are interpreted according to ι . Since γ is primitive positive, we have $B \models \gamma[\iota]$ precisely if there is a function κ mapping the existential quantifiers in γ to elements of B so that every atomic formula $R(\bar{y})$ occurring in γ is satisfied in B when each variable y in \bar{y} is interpreted by $\iota(y)$, if this occurrence of y is free in γ , and by κ applied to the quantifier binding y otherwise. We write $|B \models \gamma[\iota]|$ to denote the number of distinct functions κ witnessing $B \models \gamma[\iota]$ and observe that $|B \models \gamma_A| = |\text{hom}(A, B)|$.

We now define by induction on the structure of γ a formula $\gamma^t(\bar{x})$ of \mathcal{C} with the property that for any structure B and interpretation ι of \bar{x} in B , we have $B \models \gamma^t[\iota]$ if, and only if, $|B \models \gamma[\iota]| \geq t$. For atomic γ , γ^1 is just γ and γ^t is *false* if $t > 1$. For γ a conjunction $\gamma_1 \wedge \gamma_2$, let $T := \{(t_1, t_2) \in \mathbb{N}^2 \mid t_1 \cdot t_2 \geq t\}$. Then

$$\gamma^t := \bigvee_{(t_1, t_2) \in T} (\gamma_1^{t_1} \wedge \gamma_2^{t_2}).$$

When γ is $\exists x. \gamma_1$, let F be the collection of all finite partial functions f on \mathbb{N} such that $\sum_{s \in \text{dom}(f)} sf(s) \geq t$. Then

$$\gamma^t := \bigvee_{f \in F} \bigwedge_{s \in \text{dom}(f)} \exists_{\geq f(s)} x. \gamma_1^s.$$

Note that both the quantifier rank and the total number of variables that appear in γ_A^t are the same as for γ_A , and γ_A^t

still has no negation symbols. It then follows from Dvořák's theorem that any sentence φ of \mathcal{C}_{∞}^k is equivalent in the finite (i.e., over finite structures) to an infinite Boolean combination of sentences of \mathcal{C}^k of the form γ_A^t . Indeed, φ is equivalent to

$$\bigvee_{B \in \mathcal{M}} \bigwedge_{A \in \mathcal{T}} \left(\gamma_A^{|\text{hom}(A, B)|} \wedge \neg \gamma_A^{|\text{hom}(A, B)|+1} \right)$$

where \mathcal{M} consists of one representative for each $\equiv_{\mathcal{C}^k}$ -class of structures in $\{B \mid B \models \varphi\}$ and \mathcal{T} consists of one representative for each isomorphism class of finite structures that have tree-width at most $k-1$.

Likewise, it is a consequence of Grohe's theorem that any sentence of $\mathcal{C}_{n,\infty}$ is equivalent (in the finite) to an infinite Boolean combination of sentences of \mathcal{C}_n of the form γ_A^t ; a similar normal form was also exhibited in [14, §5]. This yields an interesting normal form for these counting logics. In particular, there is no need of universal quantifiers or of equality, as neither of these is used in the sentences γ_A^t . Further, if we allow dual counting quantifiers $\exists_{\leq i}$ with the semantics that $A \models \exists_{\leq i} x. \varphi$ if, and only if, $A \models \varphi[a/x]$ holds for at most i distinct elements $a \in A$, then we do not need negation at all in our formulas.

From our point of view, the most interesting aspect of this is the elimination of equality symbols. Consider for example the simple first-order sentence in the language of graphs: $\exists x \exists y. (x \neq y \wedge E(x, y))$ asserting the existence of an edge between two distinct vertices. At first sight, it does not seem possible to express this property without using the equality symbol. However, it is possible to do so with infinitary Boolean connectives and counting quantifiers as we illustrate.

First, consider the formula $\exists_{\geq t} xy. E(x, y)$ asserting that there exist at least t pairs x, y for which $E(x, y)$. Note this is just γ^t when γ is the sentence $\exists x \exists y. E(x, y)$. Now, the sentence $\exists x \exists y. (x \neq y \wedge E(x, y))$ is equivalent to

$$\bigvee_{i \in \mathbb{N}} (\exists_{\geq i} xy. E(x, y) \wedge \exists_{\leq i-1} x. E(x, x)).$$

We now give a direct proof, in our categorical framework, that we can eliminate equality from \mathcal{C}_{∞}^k and $\mathcal{C}_{n,\infty}$.

Theorem 32. *Every sentence of \mathcal{C}_{∞}^k is equivalent in the finite to one without equality, and every sentence of $\mathcal{C}_{n,\infty}$ is equivalent in the finite to one without equality.*

Proof. We give the proof only for $\mathcal{C}_{n,\infty}$ as the other is entirely analogous. It suffices to prove that if two finite structures A, B are not distinguished by any sentence of \mathcal{C}_n without equality, then they are not distinguished by any sentence of \mathcal{C}_n .

Suppose then that A and B are not distinguished by any sentence of \mathcal{C}_n without equality. We thus have $F^{\mathbb{E}_n}(A) \cong F^{\mathbb{E}_n}(B)$ in $\text{EM}(\mathbb{E}_n)$ and so $\text{hom}(D, F^{\mathbb{E}_n}(A)) \cong \text{hom}(D, F^{\mathbb{E}_n}(B))$ for every D in $\text{EM}_f(\mathbb{E}_n)$. By the adjunction $U^{\mathbb{E}_n} \dashv F^{\mathbb{E}_n}$, we have $\text{hom}(U^{\mathbb{E}_n}(D), A) \cong \text{hom}(U^{\mathbb{E}_n}(D), B)$ for every D in $\text{EM}_f(\mathbb{E}_n)$. Since a finite structure has tree-depth at most n if, and only if, it admits an \mathbb{E}_n -coalgebra structure, this implies $\text{hom}(C, A) \cong \text{hom}(C, B)$ for every finite structure C with tree-depth at most n and hence, by Theorem 18, $A \equiv_{\mathcal{C}_n} B$. \square

VII. CONCLUSION

The ideas developed in this paper connect homomorphism counting results in finite model theory with the framework of game comonads recently put forward by Abramsky et al.

Our categorical generalisation of Lovász’ theorem allows us to give uniform proofs of results by Dvořák and Grohe. The only part of the proof specific to each of these results is the combinatorial argument underlying the fact that the adjunction between σ -structures and σ^+ -structures restricts to structures of bounded tree-width and bounded tree-depth, respectively. We also establish a new homomorphism counting result for modal logic which characterises the equivalence relation $\equiv_{C_k^{\text{ML}}}$ over finite Kripke structures in terms of homomorphism counts.

As a by-product of the modularity of this general framework, we are immediately able to deduce: (i) the characterisation of the equivalence relation $\equiv_{C_n^k}$ in terms of homomorphism counts from structures admitting a k -pebble forest cover of height $\leq n$, (ii) the specialisation of our results to undirected graphs without loops (cf. Remark 24), and (iii) equality elimination results for counting logics (Theorem 32).

We believe that our method lays a pathway to discovering more Lovász-type theorems. In particular, any comonad on the category of σ -structures that satisfies the conditions of Theorem 16 will yield a Lovász-type theorem. The natural next step to test this theory is to apply our results to the game comonads introduced in [2] and [10].

On the other hand, there are many Lovász-type theorems that do not immediately fit into our framework. For example, quantum isomorphism [22] and co-spectrality are characterised by restricting the homomorphism counts to planar graphs and cycles, respectively. Are there suitable comonads (or monads) that could allow us to recover these characterisations in a uniform manner? Another line for future investigation consists in fine-tuning the general theory introduced in this paper. For example, can we refine Theorem 31 to yield a homomorphism counting result for p-morphisms, the natural concept of morphism between Kripke structures?

The ideas that we have described in this paper contribute to advancing the theory of game comonads, which connects two distinct areas of logic in computer science: on the one hand finite and algorithmic model theory, and on the other hand categorical and structural methods in semantics. Game comonads have been successfully applied, e.g., to investigate equirank-variable homomorphism preservation theorems in finite model theory [24]. We broaden the applicability of this theory by bringing these techniques to the area of Lovász-type results in finite model theory.

ACKNOWLEDGMENT

We acknowledge useful discussions in the Cambridge-Oxford Comonads seminar. The last-named author is grateful to André Joyal for explaining to him (private communication) the construction $E(-) \mapsto E[-]$ for polyadic spaces [16], which features in Lemma 12.

REFERENCES

- [1] S. Abramsky, A. Dawar, and P. Wang, “The pebbling comonad in finite model theory”, in *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS*, 2017, pp. 1–12.
- [2] S. Abramsky and D. Marsden, “Comonadic semantics for guarded fragments”, Preprint available at arXiv:2008.11094, 2020.
- [3] S. Abramsky and N. Shah, “Relating Structure and Power: Comonadic semantics for computational resources”, in *27th EACSL Annual Conference on Computer Science Logic, CSL*, 2018, pp. 2:1–2:17.
- [4] —, “Relating Structure and Power: Extended version”, To appear in *Journal of Logic and Computation*, preprint available at arXiv:2010.06496, 2020.
- [5] J. Adámek, H. Herrlich, and G. E. Strecker, *Abstract and concrete categories*. John Wiley & Sons, Inc., New York, 1990.
- [6] J. Adámek and J. Rosický, *Locally presentable and accessible categories*, ser. London Mathematical Society Lecture Note Series. Cambridge University Press, Cambridge, 1994, vol. 189.
- [7] P. Blackburn, M. de Rijke, and Y. Venema, *Modal logic*, ser. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 2001, vol. 53.
- [8] F. Borceux, *Handbook of categorical algebra. 1*, ser. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, 1994, vol. 50.
- [9] J.-Y. Cai, M. Fürer, and N. Immerman, “An optimal lower bound on the number of variables for graph identification”, *Combinatorica*, vol. 12, no. 4, pp. 389–410, 1992.
- [10] A. Ó Conghaile and A. Dawar, “Game comonads & generalised quantifiers”, in *29th EACSL Annual Conference on Computer Science Logic, CSL*, 2021, pp. 16:1–16:17.
- [11] Z. Dvořák, “On recognizing graphs by numbers of homomorphisms”, *Journal of Graph Theory*, vol. 64, no. 4, pp. 330–342, 2010.
- [12] P. J. Freyd and G. M. Kelly, “Categories of continuous functors, I”, *Journal of Pure and Applied Algebra*, vol. 2, no. 3, pp. 169–191, 1972.
- [13] E. Grädel, P. Kolaitis, L. Libkin, M. Marx, J. Spencer, M. Vardi, Y. Venema, and S. Weinstein, *Finite Model Theory and Its Applications*. Springer, 2007.
- [14] M. Grohe, “Counting bounded tree depth homomorphisms”, in *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science*, 2020, pp. 507–520.
- [15] J. Isbell, “Some inequalities in hom sets”, *Journal of Pure and Applied Algebra*, vol. 76, no. 1, pp. 87–110, 1991.
- [16] A. Joyal, “Polyadic spaces and elementary theories”, *Notices of the AMS*, vol. 18, no. 3, p. 563, 1971, preliminary report.
- [17] S. Kiefer, “The Weisfeiler-Leman algorithm: an exploration of its power”, *ACM SIGLOG News*, vol. 7, no. 3, pp. 5–27, 2020.
- [18] P. G. Kolaitis and M. Y. Vardi, “Conjunctive query containment and constraint satisfaction”, *Journal of Computer and System Sciences*, vol. 61, pp. 302–332, 2000.
- [19] L. Libkin, *Elements of Finite Model Theory*. Springer, 2004.
- [20] L. Lovász, “Operations with structures”, *Acta Math. Acad. Sci. Hungar.*, vol. 18, pp. 321–328, 1967.
- [21] S. Mac Lane, *Categories for the working mathematician*, 2nd ed., ser. Graduate Texts in Mathematics. Springer-Verlag, New York, 1998, vol. 5.
- [22] L. Mančinská and D. E. Roberson, “Quantum isomorphism is equivalent to equality of homomorphism counts from planar graphs”, in *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, 2020, pp. 661–672.
- [23] M. Otto, *Bounded Variable Logics and Counting — A Study in Finite Models*, ser. Lecture Notes in Logic. Springer-Verlag, 1997, vol. 9.
- [24] T. 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*, vol. 352, pp. 191–209, 2020.
- [25] A. Pultr, “Isomorphism types of objects in categories determined by numbers of morphisms”, *Acta Scientiarum Mathematicarum*, vol. 35, pp. 155–160, 1973.
- [26] M. V. Ramana, E. R. Scheinerman, and D. Ullman, “Fractional isomorphism of graphs”, *Discrete Mathematics*, vol. 132, pp. 247–265, 1994.
- [27] E. Riehl, “Factorization systems”, notes available at <http://www.math.jhu.edu/~eriehl/factorization.pdf>.
- [28] B. Rossman, “Homomorphism preservation theorems”, *Journal of the ACM*, vol. 55, 2008.