Homomorphism Preservation Theorems

BENJAMIN ROSSMAN
Massachusetts Institute of Technology

The homomorphism preservation theorem (h.p.t.), a result in classical model theory, states that a first-order formula is preserved under homomorphisms on all structures (finite and infinite) if and only if it is equivalent to an existential-positive formula. Answering a long-standing question in finite model theory, we prove that the h.p.t. remains valid when restricted to finite structures (unlike many other classical preservation theorems, including the Łoś-Tarski theorem and Lyndon's positivity theorem). Applications of this result extend to constraint satisfaction problems and to database theory via a correspondence between existential-positive formulas and unions of conjunctive queries. A further result of this article strengthens the classical h.p.t.: we show that a first-order formula is preserved under homomorphisms on all structures if and only if it is equivalent to an existential-positive formula *of equal quantifier-rank.*

Categories and Subject Descriptors: F.4.1 [**Mathematical Logic and Formal Languages**]: Mathematical Logic—*Model Theory*; H.2.3 [**Database Management**]: Languages—*Query languages*

General Terms: Theory

Additional Key Words and Phrases: Finite model theory, first-order logic, homomorphisms, preservation theorems, tree-depth, quantifier-rank, conjunctive queries

1. INTRODUCTION

1.1 Preservation Theorems in Classical and Finite Model Theory

Classical model theory studies the general structures of mathematics through the lens of first-order logic. *Preservation theorems* are a group of results in classical model theory that describe relationships between syntactic and semantic properties of first-order formulas. The Łoś-Tarski theorem, Lyndon's positivity theorem and the homomorphism preservation theorem (h.p.t., for short) are three fundamental preservation theorems dating from the 1950s. Each of these theorems states that a certain syntactic class of formulas (respectively: existential, positive, existential-positive) contains, up to logical equivalence, all first-order formulas preserved under a certain algebraic relationship between structures (the existence of a particular kind of homomorphism).

THEOREM 1.1 (ŁOŚ-TARSKI THEOREM). *A first-order formula is preserved under embeddings on all structures if, and only if, it is logically equivalent to an exis-*

This work was supported by an MIT Akamai Presidential Fellowship and a National Defense Science and Engineering Graduate Fellowship and partially written during an internship at IBM's Almaden Research Center.
An extended abstract of this article appeared as [Rossman 2005].
Author's address: Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory, 32 Vassar St., Cambridge, MA, e-mail: brossman@theory.csail.mit.edu

tential formula.[1]

THEOREM 1.2 (LYNDON'S POSITIVITY THEOREM). *A first-order formula is preserved under surjective homomorphisms on all structures if, and only if, it is logically equivalent to a positive formula.*

THEOREM 1.3 (HOMOMORPHISM PRESERVATION THEOREM). *A first-order formula is preserved under homomorphisms on all structures if, and only if, it is logically equivalent to an existential-positive formula.*

In these statements and throughout this article, the class of "all structures" includes both finite and infinite structures. (We assume no background in logic; all relevant definitions are presented in §2.1. For more on classical model theory, see [Hodges 1993].)

Finite model theory is in large part the study of first-order logic on finite structures [Ebbinghaus and Flum 1996; Libkin 2004]. At first glance, finite model theory appears to be a subfield of classical model theory; after all, finite structures are a subclass of "all structures". In practice, however, there is a great divide between the finite and classical worlds. From the perspective of classical model theory, the intrinsically interesting structures are infinite. Many key classical theorems and techniques break down when restricted to finite structures. The compactness theorem (a supremely important tool in classical model theory) is the most conspicuous result to fail on finite structures. Other classical theorems become meaningless or irrelevant in the finite context (for instance, the Lowenheim-Skolem theorem). Others still are curiously inverted: whereas the set of valid first-order formulas is r.e. (recursively enumerable) [Gödel's completeness theorem] but not co-r.e. [Church's theorem], the set of first-order formulas that are valid on finite structures is co-r.e. but not r.e. [Trakhtenbrot's theorem]. A handful of classical theorems and techniques survive the passage to the finite setting (such as Ehrenfeucht-Fraïssé games), although not enough to bridge the gulf between the finite and classical worlds. And, of course, a great many results are native to finite model theory, having no classical counterpart. (For more on this taxonomy, see [Rosen 2002; Baldwin 2000; Kolaitis 1993].)

The project to classify the status of classical theorems when restricted to finite structures has been an active line of research in finite model theory beginning with [Gurevich 1984]. Empirically, the classical theorems which remain true on finite structures are those whose proofs in the classical setting work just as well when one considers only finite structures. On the other hand, the failure of compactness theorem harbingers the collapse of its many corollaries (including the classical preservation theorems, whose original proofs rely on compactness arguments). In a survey on classical and finite model theory, Rosen [2002] wrote that "there seems to be no example of a theorem [of classical model theory] that remains true when relativized to finite structures but for which there are entirely different proofs for the two cases. It would be interesting to find a theorem proved using the compactness theorem that can be established using a new method over finite structures. It

[1]An *embedding* of **A** into **B** is an isomorphism from **A** onto an induced substructure of **B**. The Łoś-Tarski Theorem is sometimes stated in its dual form: *a first-order formula is preserved under induced substructures if, and only if, it is logically equivalent to a universal formula.*

is known that many of the candidates for such a theorem, such as preservation and interpolation theorems, fail in the finite." Indeed, both the Łoś-Tarski theorem and Lyndon's positivity theorem are known to fail on finite structures.

Theorem 1.4 (Łoś-Tarski Fails on Finite Structures). *There exists a first-order formula that is preserved under embeddings on finite structures, but is not equivalent in the finite to an existential formula.*

Theorem 1.5 (Lyndon Fails on Finite Structures). *There exists a first-order formula that is preserved under surjective homomorphisms on finite structures, but is not equivalent in the finite to a positive formula.*

Theorem 1.4 was first proved by Tait [1959] and rediscovered by Gurevich and Shelah (see [Gurevich 1984]). Theorem 1.5 is due to Ajtai and Gurevich [1987]. In neither case is the failure on finite structures as blatant as the failure of the compactness theorem. The counterexample of Theorem 1.5 in particular is highly nontrivial; a simplified counterexample was later obtained by Stolboushkin [1995]. (For more on the failure of classical preservation theorems on finite structures, see [Alechina and Gurevich 1997; Rosen and Weinstein 1995].)

Meanwhile, the status of the h.p.t. on finite structures remained an intriguing open problem despite a number of partial solutions [Ajtai and Gurevich 1994; Atserias 2005; Grädel and Rosen 1999; Rosen 1995] and an incorrect claim in [Gurevich 1990].

1.2 Main Results

We finally resolve the status of the h.p.t., showing that it remains valid when restricted to finite structures. This is surprising since the h.p.t. seems to reside at the intersection of the Łoś-Tarski and Lyndon preservation theorems, which both fail on finite structures. A further result of this article, which we call the *equirank h.p.t.*, improves the classical h.p.t. in the general setting.[2]

Theorem 1.6 (Equirank Homomorphism Preservation Theorem). *A first-order sentence is preserved under homomorphisms on all structures if, and only if, it is equivalent to an existential-positive sentence of equal quantifier-rank.*

Theorem 1.7 (Finite Homomorphism Preservation Theorem). *A first-order sentence of quantifier-rank n is preserved under homomorphisms on finite structures if, and only if, it is equivalent in the finite to an existential-positive sentence of quantifier-rank $\rho(n)$ (for some explicit function $\rho : \mathbb{N} \longrightarrow \mathbb{N}$).*

These results are later restated as Theorems 4.12 and 5.16. In fact, we obtain slightly sharper results (Theorems 4.11 and 5.15) stated more generally as interpolation theorems. We remark that in both of the above theorems, one direction is straightforward: it is easy to see that every existential-positive sentence is preserved under homomorphisms (on finite structures). All of our effort thus goes into proving the converse direction in these theorems.

[2]As a matter of convenience, we now switch from speaking about *formulas* to instead speaking about *sentences* (i.e., formulas without free variables). There is no loss of generality; both theorems are valid when stated more generally for formulas instead of sentences (see the discussion in §7.1).

By the equirank h.p.t., every homomorphism-preserved first-order sentence Φ is equivalent to an existential-positive sentence Ψ with the same quantifier-rank. That is, there is no blow-up in quantifier-rank from Φ to Ψ. One might instead ask, what is the *length* of the shortest existential-positive sentence Ψ equivalent to Φ? A previously unpublished result of Gurevich and Shelah (which we include here as Theorem 6.1) states that, despite the non-increase in quantifier-rank, there is a potentially non-elementary blow-up in length from Φ to the shortest equivalent Ψ. The same counterexample implies that the best quantifier-rank bound $\rho(n)$ we manage to achieve in the finite h.p.t. (Theorem 1.7) is non-elementary.

1.3 Combinatorial Perspective

Both our main theorems have purely combinatorial interpretations. In fact, pretty much the entire technical development of the article takes place within a combinatorial framework (without reference to logic). From the combinatorial perspective, we are interested in three equivalence relations on structures (see §2 for precise definitions).

homomorphic equivalence ($\mathbf{A} \rightleftarrows \mathbf{B}$)

There exist homomorphisms $\mathbf{A} \to \mathbf{B}$ and $\mathbf{B} \to \mathbf{A}$.

n***-homomorphic equivalence*** ($\mathbf{A} \rightleftarrows^n \mathbf{B}$)

(combinatorial) $\mathbf{C} \to \mathbf{A} \iff \mathbf{C} \to \mathbf{B}$ for all finite structures $\mathbf{C}$ of tree-depth n.

(logical) $\mathbf{A}$ and $\mathbf{B}$ satisfy the same existential-positive sentences of quantifier-rank n.

n***-back-and-forth equivalence*** ($\mathbf{A} \equiv^n \mathbf{B}$)

(combinatorial) There exists an n-back-and-forth system of partial isomorphisms between $\mathbf{A}$ and $\mathbf{B}$.

(logical) $\mathbf{A}$ and $\mathbf{B}$ satisfy the same first-order sentences of quantifier-rank n.

While the notions of homomorphic equivalence and n-back-and-forth equivalence have a long history, n-homomorphic equivalence appears to be a new concept. Intuitively, n-homomorphic equivalence approximates homomorphic equivalence "up to tree-depth n". In a similar fashion, n-back-and-forth equivalence approximates isomorphism. Indeed, on finite structures, equivalences $\rightleftarrows^0, \rightleftarrows^1, \rightleftarrows^2, \dots$ converge to $\rightleftarrows$ (in the sense that $\mathbf{A} \rightleftarrows^n \mathbf{B}$ for all n if and only if $\mathbf{A} \rightleftarrows \mathbf{B}$), while $\equiv^0, \equiv^1, \equiv^2, \dots$ converge to isomorphism. By contrast, on infinite structures, $\bigcap_{n \in \mathbb{N}} \rightleftarrows_n$ is coarser than $\rightleftarrows$, while $\bigcap_{n \in \mathbb{N}} \equiv_n$ (called *elementary equivalence* in model theory) is coarser than isomorphism.

We now state the combinatorial versions of our two main results. These are the theorems we directly prove; their logical counterparts, Theorems 1.6 and 1.7, are in fact obtained as corollaries.

Theorem 1.8 (Equirank H.P.T., Combinatorial Version). *For all structures* $\mathbf{A}$ *and* $\mathbf{B}$ *such that* $\mathbf{A} \rightleftarrows^n \mathbf{B}$*, there exist structures* $\mathbf{A}'$ *and* $\mathbf{B}'$ *such that* $\mathbf{A} \rightleftarrows \mathbf{A}' \equiv^n \mathbf{B}' \rightleftarrows \mathbf{B}$.

THEOREM 1.9 (FINITE H.P.T., COMBINATORIAL VERSION). *For all finite structures* $\mathbf{A}$ *and* $\mathbf{B}$ *such that* $\mathbf{A} \rightleftarrows^{\rho(n)} \mathbf{B}$*, there exist finite structures* $\mathbf{A}'$ *and* $\mathbf{B}'$ *such that* $\mathbf{A} \rightleftarrows \mathbf{A}' \equiv^n \mathbf{B}' \rightleftarrows \mathbf{B}$.

1.4 Related Work and Connections to Computer Science

Prior to this article, the h.p.t. was known to hold on finite structures in various special cases. Rosen [1995] proved that a first-order formula in the $\forall^*\exists^*\forall^*$ prefix class is preserved under homomorphisms on finite structures if, and only if, it is equivalent in the finite to an existential-positive formula. Grädel and Rosen [1999] proved the finite h.p.t. for first-order formulas involving only two distinct variables. Ajtai and Gurevich [1994] showed that if a class of finite structures is definable both by a first-order sentence and by a DATALOG expression, then it is also definable by an existential-positive sentence. Although originally proved by very different means, this result follows immediately from the finite h.p.t., since all DATALOG expressions are preserved under homomorphisms.

The field of constraint satisfaction problems provides another example of a consequence of the finite h.p.t. that was originally proved by a different technique. For a finite structure $\mathbf{B}$, let CSP($\mathbf{B}$) denote the class of finite structures $\mathbf{A}$ that have a homomorphism to $\mathbf{B}$. The membership problem for the class CSP($\mathbf{B}$) is known as the *constraint satisfaction problem with template* $\mathbf{B}$. Atserias [2005] proved that if CSP($\mathbf{B}$) is definable by a first-order sentence, then the complementary class co-CSP($\mathbf{B}$) (of finite structures with no homomorphism to $\mathbf{B}$) is definable by an existential-positive sentence. This result follows immediately from the finite h.p.t., since for every template $\mathbf{B}$, the class co-CSP($\mathbf{B}$) is closed under homomorphisms. A nice corollary of this result is that a constraint satisfaction problem CSP($\mathbf{B}$) is first-order definable if, and only if, it has *finite duality*, meaning there exist finitely many "forbidden" structures $\mathbf{F}_1, \dots, \mathbf{F}_m$ such that $\mathbf{A} \to \mathbf{B} \iff \bigwedge_{i=1}^{m} \mathbf{F}_i \not\to \mathbf{A}$ for all $\mathbf{A}$.

We now mention a few results related to the finite h.p.t., but which do not directly follow from it. Feder and Vardi [2003] proved homomorphism preservation theorems on finite structures for various non-first-order logics including $\exists L^\omega$, SNP and 2SNP. In a different vein, Atserias, Dawar and Kolaitis [2006] showed that the h.p.t. holds on certain restricted classes of finite structures. Specifically, if $\mathcal{K}$ is a class of finite structures whose cores either have bounded degree or exclude a minor, then the h.p.t. holds when restricted to $\mathcal{K}$. That is, a first-order formula is preserved under homomorphisms on $\mathcal{K}$ if and only if it is equivalent on $\mathcal{K}$ to an existential-positive formula. The basic approach of Atserias, Dawar and Kolaitis, building on [Ajtai and Gurevich 1994], involves a completely different combinatorial framework. (To summarize the difference: they study the *minimal models* of an existential-positive sentence, while we study the *hom-minimal models*; see Definition 2.15.)

Numerous applications of finite model theory are found in computer science [Grädel et al. 2007]. A primary motivation for studying the validity of the h.p.t. on finite structures comes from the theory of relational databases. Indeed, many questions arising in the database context have been fruitfully studied using techniques from finite model theory. Finite relational structures have been a popular model for databases ever since [Codd 1970]. Several popular query languages correspond to sublogics or extensions of first-order logic. To take a pertinent example, *unions of conjunctive queries* (also called *select-project-join-union queries* and said to be the most common database queries in practice [Abiteboul et al. 1995]) are semantically equivalent to existential-positive formulas. The status of the h.p.t. on finite structures may thus be cast as a question about database query languages (cp. [Atserias et al. 2006]).

1.5 Overview of the Article

Section 2 covers the basics of structures and homomorphisms, including the important notion of *tree-depth*. This section also includes a concise review of first-order logic with emphasis on the properties of primitive-positive and existential-positive formulas. In particular, we describe a correspondence, originally due to Chandra and Merlin [1977], between primitive-positive sentences (up to logical equivalence) and finite structures (up to homomorphic equivalence).

In section 3, we introduce notions of n*-homomorphism* and n*-cores*, which approximate the familiar notions of homomorphism and cores "up to tree-depth n". We prove a number of lemmas about the lattice of n-homomorphic equivalence classes. (An equivalent logical theory of n*-existential-positive types* was developed in [Rossman 2005], a preliminary version of this article.)

Our main results, the equirank and finite homomorphism preservation theorems, are proved in sections 4 and 5. For the proof of the equirank h.p.t., we define a structural property of n*-extendability* (called n*-existential-positive saturation* in [Rossman 2005]) and a technique for constructing infinite n-extendable co-retracts of any structure. This proof serves as a blueprint and a warm-up for the more intricate proof of the finite h.p.t., in which we relax the n-extendability property in order to "finitize" the previous construction.

Section 6 addresses the question: given a homomorphism-preserved first-order sentence Φ of length n, how long is the shortest existential-positive sentence Ψ equivalent to Φ? We present a previously unpublished result of Gurevich and Shelah that for certain Φ, the length of Ψ is a non-elementary function of n. This result was announced in [Gurevich 1990], but a proof has before never appeared in print. The same example of Gurevich and Shelah also establishes that the quantifier-rank bound $\rho(n)$ in the finite h.p.t. is non-elementary.

We conclude in section 7 by stating some corollaries of our main results, documenting the failure on finite structures of the homomorphism interpolation theorem (a generalization of the h.p.t.) and raising a few open questions.

2. PRELIMINARIES

In §2.1 we present the basic definitions of *structures* and *homomorphisms* and *first-order logic* that are needed to understand the statements of our main results (and the other theorems mentioned in the introduction). Following these basic definitions, we consider in §2.2 the category of structures and homomorphisms over a set X. In §2.3 we define the *Gaifman graph* of a structure and introduce the key notion of *tree-depth*. We then review in §2.4 the definitions and basic properties of *retractions* and *cores*. This leads into a brief discussion in §2.5 of the *homomorphism lattice*. We then describe in §2.6 a fundamental correspondence between existential-positive sentences and antichains in the homomorphism lattice. We conclude this section in by defining n*-back-and-forth equivalence* of structures and giving a characterization in terms of systems partial isomorphisms §2.7.

Most lemmas in this section are stated without proof and should be considered easy exercises. References are given where results are not well-known folklore.

2.1 Basic Definitions

2.1.1 *Structures.* A *(relational) structure* is an object $\mathbf{A} = \langle A, R_1^{\mathbf{A}}, \dots, R_m^{\mathbf{A}} \rangle$ where A is a nonempty set, m is a natural number, $R_1, \dots, R_m$ are abstract *relation symbols* with associated *arities* $k_1, \dots, k_m$ (nonnegative integers), and each $R_i^{\mathbf{A}}$ is a k_i-ary relation on A (i.e., a subset of A^{k_i}). The set A is called the *universe* of $\mathbf{A}$ and may in general be infinite. The sequence of relation symbols $R_1, \dots, R_m$ together with corresponding arities $k_1, \dots, k_m$ comprise the *vocabulary* of $\mathbf{A}$. Relation $R_i^{\mathbf{A}}$ is called the *interpretation* of relation symbol R_i in $\mathbf{A}$. The *size* of $\mathbf{A}$ refers to the cardinality of A and is sometimes denoted by $|\mathbf{A}|$. We will generally consider structures with a common vocabulary, which we denote by σ. We emphasize that σ consists of finitely many relation symbols.[3] We consistently employ boldface letters $\mathbf{A}, \mathbf{B}, \mathbf{C}, \dots$ for structures and italic letters $A, B, C, \dots$ for the corresponding universes.

A structure $\mathbf{B}$ is a *substructure* of $\mathbf{A}$ (in symbols: $\mathbf{B} \subseteq \mathbf{A}$) if $B \subseteq A$ and $R^{\mathbf{B}} \subseteq R^{\mathbf{A}}$ for every relation symbol R in σ. It is an *induced substructure* of $\mathbf{A}$ if $R^{\mathbf{B}}$ is the restriction of $R^{\mathbf{A}}$ to B (i.e., $R^{\mathbf{B}} = R^{\mathbf{A}} \cap B^k$ where k is the arity of R) for every R in σ. For a subset $X \subseteq A$, we denote by $\mathbf{A}|_X$ the (unique) induced substructure of $\mathbf{A}$ with universe X.

2.1.2 *Homomorphisms.* Let $\mathbf{A} = \langle A, R_1^{\mathbf{A}}, \dots, R_m^{\mathbf{A}} \rangle$ and $\mathbf{B} = \langle B, R_1^{\mathbf{B}}, \dots, R_m^{\mathbf{B}} \rangle$ be structures (in the same vocabulary). A *homomorphism from* $\mathbf{A}$ *to* $\mathbf{B}$ is a function $h : A \longrightarrow B$ such that $h(R_i^{\mathbf{A}}) \subseteq R_i^{\mathbf{B}}$ (i.e., if $(a_1, \dots, a_{k_i}) \in R_i^{\mathbf{A}}$ then $(h(a_1), \dots, h(a_{k_i})) \in R_i^{\mathbf{B}}$) for every R_i.

2.1.3 *First-order logic.* We quickly review the basics of first-order logic. *(First-order) formulas* (in the vocabulary σ) are built up from atomic formulas using negation, conjunction, disjunction and existential and universal quantification:

$x = y, \quad Rx_1 \dots x_k$	atomic formulas,
$\neg\phi, \quad \phi \wedge \psi, \quad \phi \vee \psi$	negation, conjunction, disjunction,
$\exists x\ \phi, \quad \forall x\ \phi$	existential and universal quantification.

Here x, y and $x_1, \dots, x_k$ are variables, R is a k-ary relation symbol in σ, and ϕ and ψ are formulas. We assume familiarity with the concept of free and bound variables, as well as the semantics of first-order logic (i.e., what it means for a formula to be true in (satisfied by) a structure for a given assignment of free variables). As a matter of notation, formulas are often written out followed by an ordered list of free variables, in the style of $\phi(x_1, \dots, x_k)$. For a structure $\mathbf{A}$ and tuple $\vec{a} \in A^k$, the notation $\mathbf{A} \models \phi(\vec{a})$ asserts that the formula $\phi(\vec{x})$ is true in $\mathbf{A}$ with variables $x_1, \dots, x_k$ taking values $a_1, \dots, a_k$. Formulas with no free variables are called *sentences*. Structures which satisfy a given sentence are called *models* of that sentence. If ϕ is a sentence,

[3]In the usual parlance, σ is a finite relational vocabulary. In general, vocabularies may contain infinitely many symbols, including constant symbols and function symbols in addition to relation symbols. In §7.1, we discuss how our results extend to structures in more general vocabularies.

we respectively denote by $\mathrm{Mod}(\phi)$ and $\mathrm{Mod}_{\mathrm{fin}}(\phi)$ the classes of models and finite models of ϕ.

The *quantifier-count* of a formula ϕ is the total number of quantifiers in ϕ. The *quantifier-rank* of ϕ is the maximum nesting depth of quantifiers in ϕ. These quantities are denoted by $\mathrm{qcount}(\phi)$ and $\mathrm{qrank}(\phi)$, respectively. Quantifier-rank is obviously at most quantifier-count and often strictly less. For example, the formula $(\forall x\ (\exists y\ Rxy) \wedge (\exists z\ Rzx)) \vee (\exists x\ Rxx)$ has quantifier-rank 2 and quantifier-count 4.

Remark 2.1. It should be obvious that we do not insist that formulas be expressed in *prenex form*, that is, with all quantifiers up front followed by a propositional (quantifier-free) formula. This is an important qualification, as the notion of quantifier-rank is sensitive to the fact that quantifiers may be interlaced with conjunctions and disjunction. The standard procedure for transforming a first-order formula into an equivalent prenex formula preserves quantifier-count but potentially increases quantifier-rank.

Two formulas $\phi(x_1, \dots, x_k)$ and $\psi(x_1, \dots, x_k)$ with the same free variables are said to be *logically equivalent* [*in the finite*] if for all [finite] structures $\mathbf{A}$ and $\mathbf{B}$ and tuples $\vec{a} \in A^k$ and $\vec{b} \in B^k$, it holds that $\mathbf{A} \models \phi(a_1, \dots, a_k)$ if and only if $\mathbf{B} \models \phi(b_1, \dots, b_k)$. If ϕ and ψ are sentences, this condition is equivalent to the statement that $\mathrm{Mod}_{[\mathrm{fin}]}(\phi) = \mathrm{Mod}_{[\mathrm{fin}]}(\psi)$.

A formula $\phi(x_1, \dots, x_k)$ is *preserved under homomorphisms* [*on finite structures*] if for all [finite] structures $\mathbf{A}$ and $\mathbf{B}$ and tuples $\vec{a} \in A^k$ such that $\mathbf{A} \models \phi(a_1, \dots, a_k)$ and $h : \mathbf{A} \to \mathbf{B}$ is a homomorphism, it holds that $\mathbf{B} \models \phi(h(a_1), \dots, h(a_k))$.

Remark 2.2. Logically equivalent formulas are clearly also logically equivalent in the finite. Similarly, a formula which is preserved under homomorphisms is also preserved under homomorphisms on finite structures. In neither case, however, is the converse true. (One can cook up counterexamples by tinkering with *axioms of infinity*, that is, first-order sentences with infinite models but without any finite model.)

Four classes of formulas important to this article are primitive-positive, existential-positive, positive, and existential formulas. *Primitive-positive formulas* are built out of atomic formulas using only conjunction and existential quantification. (In the context of relational databases, these formulas define conjunctive queries.) *Existential-positive formulas* are built up from atomic formulas using disjunction in addition to conjunction and existential quantification. (The database analogue are unions of conjunctive queries, also known as select-project-join-union queries.) *Positive formulas* are first-order formulas without negations. *Existential formulas* are formulas in which every existential quantifier falls inside the scope of an even number of negations, while every universal quantifier falls inside the scope of an odd number of negations. Equivalently, a formula is existential if no universal quantifiers remain once all negations are pushed down to the level of atomic formulas (via rules such as $\neg\exists\phi \rightsquigarrow \forall\neg\phi$ and $\neg(\phi \vee \psi) \rightsquigarrow \neg\phi \wedge \neg\psi$ and $\neg\neg\phi \rightsquigarrow \phi$).

We have now defined all of the concepts appearing in the statements of our two homomorphism preservation theorems (as well as the Łoś-Tarski theorem and Lyndon's positivity theorem). In the remainder of this section, we present further background definitions and lemmas pertaining to the proofs of these theorems.

But first, we remark that one direction of the if-and-only-if statement is trivial in both the equirank and finite homomorphism preservation theorems (as well as the classical h.p.t.).

LEMMA 2.3. *Every existential-positive formula is preserved under homomorphisms [on finite structures].* □

2.2 Structures and Homomorphisms over a Set X

Let X be an arbitrary set. We call a structure $\mathbf{A}$ whose universe contains X (i.e., $X \subseteq A$) a *structure over* X. For structures $\mathbf{A}$ and $\mathbf{B}$ over X, a *homomorphism from* $\mathbf{A}$ *to* $\mathbf{B}$ *over* X is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ which fixes X pointwise. We write $\mathbf{A} \to_X \mathbf{B}$ if there exists a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ over X. We say that $\mathbf{A}$ and $\mathbf{B}$ are *homomorphically equivalent over* X and we write $\mathbf{A} \rightleftarrows_X \mathbf{B}$ if both $\mathbf{A} \to_X \mathbf{B}$ and $\mathbf{B} \to_X \mathbf{A}$. We say that $\mathbf{A}$ and $\mathbf{B}$ are *isomorphic over* X and we write $\mathbf{A} \cong_X \mathbf{B}$ if there exist homomorphisms $f : \mathbf{A} \to_X \mathbf{B}$ and $g : \mathbf{B} \to_X \mathbf{A}$ such that $g \circ f = \mathrm{id}_A$ and $f \circ g = \mathrm{id}_B$; in this case, we call f and g *isomorphisms over* X.

The coproduct of $\mathbf{A}$ and $\mathbf{B}$ in the category of structures and homomorphisms over X is called the X*-sum* of $\mathbf{A}$ and $\mathbf{B}$ and denoted by $\mathbf{A} \oplus_X \mathbf{B}$. Formally, this is the structure with universe $X \sqcup (A \setminus X) \sqcup (B \setminus X)$ whose relations are inherited from $\mathbf{A}$ and $\mathbf{B}$ via the natural inclusion maps $A \hookrightarrow X \sqcup (A \setminus X) \sqcup (B \setminus X) \hookleftarrow B$. We view $\oplus_X$ as an associative and commutative operation and we view $\mathbf{A}$ and $\mathbf{B}$ as substructures of $\mathbf{A} \oplus_X \mathbf{B}$, even if this is a slight fiction.[4] The coproduct of an indexed family $(\mathbf{A}_i)_{i \in I}$ of structures over X and the coproduct of a set $\mathscr{A}$ of structures over X are defined in the obvious way by extension and are respectively denoted by $\bigoplus_{i \in I} {}_X \mathbf{A}_i$ and $\bigoplus_X \mathscr{A}$.

The product of $\mathbf{A}$ and $\mathbf{B}$ in the category of structures and homomorphisms over X is denoted by $\mathbf{A} \otimes_X \mathbf{B}$. Informally, this is the structure with universe $A \times B$ in which the set $\Delta_X = \{(x,x) : x \in X\}$ (the "diagonal over X") is identified with X itself; formally, $\mathbf{A} \otimes_X \mathbf{B}$ has universe $X \sqcup ((A \times B) \setminus \Delta_X)$. Relations in $\mathbf{A} \otimes_X \mathbf{B}$ are defined by

$$R^{\mathbf{A} \otimes_X \mathbf{B}} = \{((a_1,b_1),\dots,(a_k,b_k)) \in (A \times B)^k : (a_1,\dots,a_k) \in R^{\mathbf{A}}, (b_1,\dots,b_k) \in R^{\mathbf{B}}\}.$$

In the special case where $X = \emptyset$, we simply write $\mathbf{A} \to \mathbf{B}$ and $\mathbf{A} \oplus \mathbf{B}$ and $\mathbf{A} \otimes \mathbf{B}$ instead of $\mathbf{A} \to_\emptyset \mathbf{B}$ and $\mathbf{A} \oplus_\emptyset \mathbf{B}$ and $\mathbf{A} \otimes_\emptyset \mathbf{B}$. In particular, $\mathbf{A} \oplus \mathbf{B}$ is the familiar disjoint union of structures.

Every structure over X is clearly also a structure over W for every $W \subseteq X$. Parts (1) and (2) of the next lemma extends this observation.

LEMMA 2.4. *Let $\mathbf{A}, \mathbf{B}, \mathbf{C}$ be structures over a set X and let $W \subseteq X$.*

(1) $\mathbf{A} \to_X \mathbf{B} \Longrightarrow \mathbf{A} \to_W \mathbf{B}$

[4] The situation is similar with the set-theoretic disjoint union $\sqcup$. For sets P and Q, the disjoint union $P \sqcup Q$ is formally the set $P \cup \{\{P, \{P, q\}\} : q \in Q\}$ (according to one common definition). While $\sqcup$ is neither truly commutative nor associative and Q is not generally a subset of $P \sqcup Q$, with a modicum of care one may harmlessly pretend that $\sqcup$ is both commutative and associate and Q is an actual subset of $P \sqcup Q$ (identified with $\{\{P, \{P, q\}\} : q \in Q\}$).

(2) $(\mathbf{A} \oplus_W \mathbf{B}) \to_W (\mathbf{A} \oplus_X \mathbf{B})$
(3) $\mathbf{A} \to_X \mathbf{C}$ and $\mathbf{B} \to_X \mathbf{C} \iff (\mathbf{A} \oplus_X \mathbf{B}) \to_X \mathbf{C}$
(4) $\mathbf{C} \to_X \mathbf{A}$ and $\mathbf{C} \to_X \mathbf{B} \iff \mathbf{C} \to_X (\mathbf{A} \otimes_X \mathbf{B})$ □

PROOF. Statement (1) amounts to the (obvious) claim that any homomorphism which fixes X pointwise also fixes W pointwise. For statement (2), we define a function $h : W \sqcup (A \setminus W) \sqcup (B \setminus W) \longrightarrow X \sqcup (A \setminus X) \sqcup (B \setminus X)$ by patching the map $W \hookrightarrow X$ (i.e., the identity map on W) together with natural embeddings $A \setminus W \xrightarrow{\cong} (X \setminus W) \sqcup (A \setminus X) \hookrightarrow X \sqcup (A \setminus X)$ and $B \setminus W \xrightarrow{\cong} (X \setminus W) \sqcup (B \setminus X) \hookrightarrow X \sqcup (B \setminus X)$. One can check that h is a homomorphism from $\mathbf{A} \oplus_W \mathbf{B}$ to $\mathbf{A} \oplus_X \mathbf{B}$ over W. Statements (3) and (4) are standard facts about the coproduct and product. □

There are many similar observations that one might make such as, for instance, the fact that if $\mathbf{A} \to_X \mathbf{B}$ and $\mathbf{B} \to_Y \mathbf{C}$ then $\mathbf{A} \to_{X \cap Y} \mathbf{C}$. We have included in Lemma 2.4 only those facts which we will explicitly need later on.

Notation 2.5. We generalize the homomorphism-over-a-set notation $\to_X$ in order to assert the existence of homomorphisms satisfying certain other constraints. For structures $\mathbf{A}$ and $\mathbf{B}$ and tuples $\vec{a} \in A^k$ and $\vec{b} \in B^k$, notation $(\mathbf{A}, \vec{a}) \to_X (\mathbf{B}, \vec{b})$ asserts that there exists a homomorphism $h : \mathbf{A} \to_X \mathbf{B}$ such that $h(a_i) = b_i$ for all $i \in \{1, \dots, k\}$. If π is a one-to-one partial function from $\mathbf{A}$ to $\mathbf{B}$, then notation $\mathbf{A} \to_\pi \mathbf{B}$ asserts the existence of a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ which extends π. We write $\mathbf{A} \rightleftarrows_\pi \mathbf{B}$ if both $\mathbf{A} \to_\pi \mathbf{B}$ and $\mathbf{B} \to_{\pi^{-1}} \mathbf{A}$.

2.3 Gaifman Graphs and Tree-Depth

Unless otherwise stated, *graphs* are simple graphs (i.e., undirected and without self-loops). A *finite rooted forest* is a disjoint union of finitely many finite rooted trees. For vertices v and w in a finite rooted forest $\mathcal{F}$, we write $v \dashrightarrow w$ if v is the parent of w. The *height* of $\mathcal{F}$ is number of vertices in the longest path from a root to a leaf. The *closure* $\overline{\mathcal{F}}$ of $\mathcal{F}$ is a graph having the same vertex set as $\mathcal{F}$, in which vertices v and w are adjacent if, and only if, one of the pairs (v, w) and (w, v) belongs to the transitive closure of $\dashrightarrow$ (i.e., $v \neq w$ and v and w are ancestors in $\mathcal{F}$).

As defined in Nešetřil and de Mendez [2006], the *tree-depth* $\mathrm{td}(\mathcal{G})$ of a finite graph $\mathcal{G}$ is the minimal height of a finite rooted forest whose closure contains $\mathcal{G}$ as a subgraph. Tree-depth has several combinatorial equivalents, including *minimum elimination tree height* and *vertex ranking number* [Deogun et al. 1994; Nešetřil and de Mendez 2007].) An inductive form of this definition (Lemma 2.2 of [Nešetřil and de Mendez 2006]) is given by:

$$\mathrm{td}(\mathcal{G}) = \begin{cases} 1 & \text{if } \mathcal{G} \text{ has a single vertex,} \\ 1 + \min\limits_{\text{vertex } v \text{ of } \mathcal{G}} \mathrm{td}(\mathcal{G} \setminus v) & \text{if } \mathcal{G} \text{ is connected and has multiple vertices,} \\ \max\limits_{\text{component } \mathcal{G}' \text{ of } \mathcal{G}} \mathrm{td}(\mathcal{G}') & \text{if } \mathcal{G} \text{ is disconnected.} \end{cases}$$

Here $\mathcal{G} \setminus v$ denotes the graph obtained from $\mathcal{G}$ by removing vertex v and all edges incident to v.

Remark 2.6. We do not define tree-depth for infinite graphs, since present purposes do not require it. We mention, however, that there is a natural definition of tree-depth for certain infinite graphs. This definition uses transfinite induction with inf instead of min and sup instead of max above (cp. [Nešetřil and Shelah 2003]).

Examples 2.7. We give some examples of tree-depth of finite graphs.

(a) The "kite" graph pictured below has tree-depth 3.

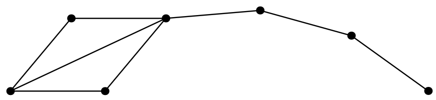

(b) The complete graph on n vertices has tree-depth n.

(c) The path on n vertices has tree-depth $\lfloor \log_2 n \rfloor + 1$.

(d) Tree-depth is related to tree-width via the following inequality (from [Nešetřil and de Mendez 2007]):

$$\mathrm{tw}(\mathcal{G}) + 1 \le \mathrm{td}(\mathcal{G}) \le \mathrm{tw}(\mathcal{G}) \log_2 |\mathcal{G}|.$$

The *Gaifman graph* $\mathcal{G}(\mathbf{A})$ of a structure $\mathbf{A}$ is the graph with vertex set A in which two elements are adjacent if, and only if, they appear together in some tuple in a relation of $\mathbf{A}$. The Gaifman graph captures precisely the "metric" information in a structures (concerning distance and connectivity among elements and subsets). For a proper subset $X \subset A$, let $\mathcal{G}(\mathbf{A}) \setminus X$ denote the induced subgraph of $\mathcal{G}(\mathbf{A})$ with vertex set $A \setminus X$. (Do not confuse $\mathcal{G}(\mathbf{A}) \setminus X$ with $\mathcal{G}(\mathbf{A} \setminus X)$, that is, the Gaifman graph of the induced substructure of $\mathbf{A}$ with universe $A \setminus X$. For vocabularies σ containing relation symbols of arity ≥ 3, these two graphs can be different.)

The *tree-depth* $\mathrm{td}_X(\mathbf{A})$ of a finite structure $\mathbf{A}$ over a subset $X \subseteq A$ is defined

$$\mathrm{td}_X(\mathbf{A}) = \begin{cases} 0 & \text{if } X = A, \\ \mathrm{td}(\mathcal{G}(\mathbf{A}) \setminus X) & \text{if } X \subset A. \end{cases}$$

(Just as we do not define tree-depth of infinite graphs, neither do we define $\mathrm{td}_X(\mathbf{A})$ when $\mathbf{A}$ is an infinite structure.) The next lemma, which follows directly from definitions, lists some basic properties of tree-depth of finite structures.

LEMMA 2.8 (PROPERTIES OF TREE-DEPTH). *Let $\mathbf{A}$ and $\mathbf{B}$ be finite structures over X.*

(*1*) $\mathrm{td}_X(\mathbf{A} \oplus_X \mathbf{B}) = \max\{\mathrm{td}_X(\mathbf{A}), \mathrm{td}_X(\mathbf{B})\}$

(*2*) $\mathbf{B} \subseteq \mathbf{A} \Longrightarrow \mathrm{td}_X(\mathbf{B}) \le \mathrm{td}_X(\mathbf{A})$

(*3*) $\mathrm{td}_X(\mathbf{A}) \le \mathrm{td}_{X \cup Y}(\mathbf{A}) + |Y|$ *for all* $Y \subseteq A$.

(*4*) *If $X \ne A$ and $\mathcal{G}(\mathbf{A}) \setminus X$ is connected, then* $\mathrm{td}_X(\mathbf{A}) = 1 + \min_{y \in A \setminus X} \mathrm{td}_{X \cup \{y\}}(\mathbf{A})$. □

We remark that Lemma 2.8(1,4), together with the assertion that $\mathrm{td}_X(\mathbf{A}) = 0$ if $A = X$, completely axiomatizes tree-depth.

2.4 Retractions and Cores

Let $\mathbf{B}$ be a substructure of $\mathbf{A}$. A homomorphism from $\mathbf{A}$ to $\mathbf{B}$ is a *retraction* if it restricts to the identity map on B. That is, a retraction is a homomorphism $\mathbf{A} \to_B \mathbf{B}$. As an alternative to the notation $\mathbf{A} \to_B \mathbf{B}$, we write $\mathbf{A} \xrightarrow{\text{retr}} \mathbf{B}$ to assert the existence of a retraction from $\mathbf{A}$ to $\mathbf{B}$. In the event that $\mathbf{A} \xrightarrow{\text{retr}} \mathbf{B}$, we call $\mathbf{B}$ a *retract* of $\mathbf{A}$, we call $\mathbf{A}$ a *co-retract* of $\mathbf{B}$, and we call the inclusion map $\mathbf{B} \hookrightarrow \mathbf{A}$ (i.e., the identity map on B) a *co-retraction* from $\mathbf{B}$ to $\mathbf{A}$. Note that retracts are always induced substructures, that is, $\mathbf{A} \xrightarrow{\text{retr}} \mathbf{B}$ implies $\mathbf{A}|_B = \mathbf{B}$. The next lemma lists some basic facts about retractions.

LEMMA 2.9 (PROPERTIES OF RETRACTIONS).

(*1*) $\mathbf{A} \xrightarrow{\text{retr}} \mathbf{B} \xrightarrow{\text{retr}} \mathbf{C} \Longrightarrow \mathbf{A} \xrightarrow{\text{retr}} \mathbf{C}$

(*2*) $(\mathbf{A} \oplus_X \mathbf{B}) \xrightarrow{\text{retr}} \mathbf{A} \Longleftrightarrow \mathbf{B} \to_X \mathbf{A}$

(*3*) $\bigoplus_{i \in I}{}_A\, \mathbf{B}_i \xrightarrow{\text{retr}} \mathbf{A} \Longleftrightarrow \bigwedge_{i \in I} \Big(\mathbf{B}_i \xrightarrow{\text{retr}} \mathbf{A}\Big)$

(*4*) $\bigwedge_{n \in \mathbb{N}} \Big(\mathbf{A}_{n+1} \xrightarrow{\text{retr}} \mathbf{A}_n\Big) \Longrightarrow \bigcup_{n \in \mathbb{N}} \mathbf{A}_n \xrightarrow{\text{retr}} \mathbf{A}_0$ $\square$

In statement (4), $\bigcup_{n\in\mathbb{N}} \mathbf{A}_n$ is the union of the chain of co-retracts $\mathbf{A}_0 \xleftarrow{\text{retr}} \mathbf{A}_1 \xleftarrow{\text{retr}} \mathbf{A}_2 \xleftarrow{\text{retr}} \cdots$, that is, the structure with universe $\bigcup_{n\in\mathbb{N}} A_n$ in which relation symbol R has the interpretation $\bigcup_{n\in\mathbb{N}} R^{\mathbf{A}_n}$.

Another key notion in this article is that of a core over a set X. A structure $\mathbf{A}$ over X is a *core over* X if every homomorphism $\mathbf{A} \to_X \mathbf{A}$ is an automorphism (i.e., an isomorphism from $\mathbf{A}$ onto itself). The following lemma gives a useful characterization of finite cores and associates a unique core with every finite structure.

LEMMA 2.10. *Let* $\mathbf{A}$ *be a finite structure and let* $X \subseteq A$.

(*1*) $\mathbf{A}$ *is a core over* X *if, and only if, it has no proper retract over* X *(i.e., for every retract* $\mathbf{B}$ *of* $\mathbf{A}$*, either* $A = B$ *or* $X \nsubseteq B$*).*

(*2*) $\mathbf{A}$ *has a retract which is a core over* X*. Moreover, if* $\mathbf{A} \xrightarrow{\text{retr}} \mathbf{B}_1$ *and* $\mathbf{A} \xrightarrow{\text{retr}} \mathbf{B}_2$ *where* $\mathbf{B}_1$ *and* $\mathbf{B}_2$ *are both cores over* X*, then* $\mathbf{B}_1 \cong_X \mathbf{B}_2$*.* $\square$

The notion of cores has a long history in the context of graph theory, where Lemma 2.10 originated (see [Hell and Nešetřil 1992]).

2.5 Canonical Cores and the Homomorphism Lattice

For every finite set X, we fix a class $\mathscr{C}_X$ of finite cores over X containing exactly one representative from each $\cong_X$-equivalence class of finite structures.[5] Since $\mathscr{C}_X$ contains only finite structures, each unique up to $\cong_X$, it follows that $\mathscr{C}_X$ is a countably infinite set. We call members of $\mathscr{C}_X$ *canonical cores over* X.

LEMMA 2.11. *For every finite structure* $\mathbf{A}$ *and* $X \subseteq A$*, there exists a unique* $\mathbf{C} \in \mathscr{C}_X$ *such that* $\mathbf{A} \rightleftarrows_X \mathbf{C}$*. Moreover,* $\mathrm{td}_X(\mathbf{C}) \le \mathrm{td}_X(\mathbf{A})$ *and every homomorphism* $h : \mathbf{C} \to_X \mathbf{A}$ *is injective and satisfies* $\mathbf{A} \xrightarrow{\text{retr}} h(\mathbf{C})$*.* $\square$

[5] One can avoid using the axiom of choice here by appealing to the lexicographic order on graphs whose vertex set is an initial segment of $\mathbb{N}$.

The structure $\mathbf{C}$ in Lemma 2.11 is called the *(canonical) core of* $\mathbf{A}$ *over* X and denoted by $\mathbf{Core}_X(\mathbf{A})$. In the special case where $X = \emptyset$, we write $\mathscr{C}$ instead of $\mathscr{C}_\emptyset$ and $\mathbf{Core}(\mathbf{A})$ instead of $\mathbf{Core}_\emptyset(\mathbf{A})$.

The relation $\rightarrow_X$ of homomorphism over X partially orders the set $\mathscr{C}_X$ of canonical cores over X. (This is tantamount to the assertion that $\rightarrow_X$ quasi-orders the $\rightleftarrows_X$-equivalence classes of finite structures over X.) Moreover, the poset $(\mathscr{C}_X, \rightarrow_X)$ is easily seen to be a lattice: Lemma 2.4(3,4), together with the observation that both $\mathbf{A} \oplus_X \mathbf{B}$ and $\mathbf{A} \otimes_X \mathbf{B}$ are finite whenever $\mathbf{A}$ and $\mathbf{B}$ are finite, implies that $\mathbf{Core}_X(\mathbf{C}_1 \oplus_X \mathbf{C}_2)$ and $\mathbf{Core}_X(\mathbf{C}_1 \otimes_X \mathbf{C}_2)$ are respectively the least upper bound and greatest lower bound of canonical cores $\mathbf{C}_1, \mathbf{C}_2 \in \mathscr{C}_X$.

The poset $(\mathscr{C}_X, \rightarrow_X)$ is known as the *homomorphism lattice*. The homomorphism lattice is a well-studied object in combinatorics with a large literature on its maximal antichains, gaps, density properties, and so forth (see [Hell and Nešetřil 2004] for an overview).

2.6 Primitive-Positive and Existential-Positive Formulas

It is intuitively obvious that every existential-positive formula can be expressed as a finite disjunction of primitive-positive formulas. Indeed, there is a purely syntactic procedure for extracting primitive-positive disjuncts out of a given existential-positive formula ψ: simply apply rules $\phi_1 \wedge (\phi_2 \vee \phi_3) \rightsquigarrow (\phi_1 \wedge \phi_2) \vee (\phi_1 \wedge \phi_3)$ and $\exists x\ (\phi_1 \vee \phi_2) \rightsquigarrow (\exists x\ \phi_1) \vee (\exists x\ \phi_2)$, etc., to all subformulas of ψ. In every such rule, the quantifier-count (resp. quantifier-rank) of the lefthand formula is at least the quantifier-count (resp. quantifier-rank) of each disjunct in the righthand formula. This observation leads to the following lemma.

Lemma 2.12. *Every existential-positive formula ψ is logically equivalent to a finite disjunction $\theta_1 \vee \cdots \vee \theta_m$ of primitive-positive formulas θ_i such that* $\mathrm{qcount}(\psi) \geq \max_i \mathrm{qcount}(\theta_i)$ *and* $\mathrm{qrank}(\psi) \geq \max_i \mathrm{qrank}(\theta_i)$. □

We now turn our attention to a fundamental one-to-one correspondence between primitive-positive sentences (up to logical equivalence) and finite structures (up to homomorphic equivalence), attributed to Chandra and Merlin [1977] in the context of relational databases and conjunctive queries (see also [Kolaitis and Vardi 2000]). This correspondence extends to a bijection between existential-positive sentences (up to logical equivalence) and finite antichains in the homomorphism lattice. An important feature of this correspondence, which seems to have been previously overlooked or underemphasized, is the tight relationship between quantifier-rank and tree-depth. The next few lemmas formally define this correspondence, thereby translating from the logical framework of implication and logical equivalence to the combinatorial framework of homomorphisms and homomorphic equivalence.

Lemma 2.13. *For every primitive-positive sentence θ, there is a finite structure $\mathbf{A}_\theta$ of size* $\mathrm{qcount}(\theta)$ *and tree-depth* $\mathrm{qrank}(\theta)$ *such that $\mathbf{A}_\theta \rightarrow \mathbf{B} \Longleftrightarrow \mathbf{B} \models \theta$ for all structures $\mathbf{B}$.*

Proof. Without loss of generality, we may assume that no variable in θ is quantified more than once (by appropriately renaming variables) and that θ contains no equality subformula $x = y$ (these subformulas can be eliminated by purely syntactic means in a manner that preserves both quantifier-count and quantifier-rank). Suppose $x_1, \ldots, x_m$ are the distinct variable symbols occurring in θ. We define the structure $\mathbf{A}_\theta$ as follows. The universe A_θ is the set $\{x_1, \ldots, x_m\}$. For each relation symbol R of arity k, the tuple $(x_{i_1}, \ldots, x_{i_k})$ belongs to $R^{\mathbf{A}_\theta}$ if and only if $Rx_{i_1} \ldots x_{i_k}$ occurs as a subformula of θ. One can check that $\mathbf{B} \models \theta \iff \mathbf{A}_\theta \to \mathbf{B}$ for all structures $\mathbf{B}$. Since obviously $|\mathbf{A}_\theta| = \text{qcount}(\theta)$, it remains to show that $\text{td}(\mathbf{A}_\theta) = \text{qrank}(\theta)$. To that end, consider the "directed arc" relation on A_θ defined by $x_i \dashrightarrow\!\!\!\!\!\dashrightarrow x_j$ if, and only if, the quantification $\exists x_j$ lies within the scope of quantification $\exists x_i$ in θ. For example, if θ is the sentence $(\exists x_1 (\exists x_2 R x_1 x_2) \wedge (\exists x_3 R x_3 x_1)) \wedge (\exists x_4 R x_4 x_4)$, we have arcs $x_1 \dashrightarrow\!\!\!\!\!\dashrightarrow x_2$ and $x_1 \dashrightarrow\!\!\!\!\!\dashrightarrow x_3$. The relation $\dashrightarrow\!\!\!\!\!\dashrightarrow$ describes a rooted forest $\mathcal{F}$ on the set A_θ (namely, $\dashrightarrow\!\!\!\!\!\dashrightarrow$ is the "ancestor of" relation on $\mathcal{F}$, that is, the transitive closure of "parent of" relation $\dashrightarrow$). Clearly, $\mathcal{G}(\mathbf{A}_\theta)$ is a subgraph of the closure $\overline{\mathcal{F}}$. Therefore, $\text{td}(\mathcal{A}_\theta) \leq \text{height}(\mathcal{F}) = \text{qrank}(\theta)$. □

Lemma 2.14. *For every finite structure* $\mathbf{A}$*, there is a primitive-positive sentence* $\theta_{\mathbf{A}}$ *with quantifier-count* $|\mathbf{A}|$ *and quantifier-rank* $\text{td}(\mathbf{A})$ *such that* $\mathbf{A} \to \mathbf{B} \iff \mathbf{B} \models \theta_{\mathbf{A}}$ *for all structures* $\mathbf{B}$*.*

Proof. Suppose $a_1, \ldots, a_m$ enumerate the elements of A and let $x_1, \ldots, x_m$ be a corresponding sequence of variable symbols. By definition of tree-depth, there exists a rooted forest $\mathcal{F}$ of height $\text{td}(\mathbf{A})$ such that $\mathcal{G}(\mathbf{A})$ is a subgraph of $\overline{\mathcal{F}}$. We will inductively define primitive-positive formulas θ_i for all $i \in \{1, \ldots, m\}$. For every i such that a_i is a leaf in $\mathcal{F}$, let θ_i be the conjunction of all atomic formulas $Rx_{j_1} \ldots x_{j_k}$ where R is a k-ary relation symbols and $(a_{j_1}, \ldots, a_{j_k}) \in R^{\mathbf{A}}$ such that elements $a_{j_1}, \ldots, a_{j_k}$ all lie in the unique branch of $\mathcal{F}$ containing the leaf a_i. For every i such that a_i is not a leaf, let θ_i be the formula $\exists x_i\ \psi_i$ where ψ_i is the conjunction of formulas θ_j for all j such that $a_i \dashrightarrow a_j$ (i.e., a_j is a child of a_i) in $\mathcal{F}$. Note that θ_i is a sentence for all i such that a_i is a root in $\mathcal{F}$. Let $\theta_{\mathbf{A}}$ be the conjunction of sentences θ_i for all roots a_i. One can check that $\theta_{\mathbf{A}}$ has the desired properties, that is, $\text{qcount}(\theta_{\mathbf{A}}) = |\mathbf{A}|$ and $\text{qrank}(\theta_{\mathbf{A}}) = \text{td}(\mathbf{A})$ and $\mathbf{A} \to \mathbf{B} \iff \mathbf{B} \models \theta_{\mathbf{A}}$ for all structures $\mathbf{B}$. □

The operators $\theta \longmapsto \mathbf{A}_\theta$ and $\mathbf{A} \longmapsto \theta_{\mathbf{A}}$ described by Lemmas 2.13 and 2.14 are easily seen to preserve quasi-orders $\vdash$ (logical implication) and $\to$ (homomorphism). That is, for all primitive-positive sentences θ_1, θ_2 and finite structures $\mathbf{A}_1, \mathbf{A}_2$, we have $\theta_1 \vdash \theta_2 \iff \mathbf{A}_{\theta_1} \to \mathbf{A}_{\theta_2}$ and $\mathbf{A}_1 \to \mathbf{A}_2 \iff \theta_{\mathbf{A}_1} \vdash \theta_{\mathbf{A}_2}$. Taking cores, we obtain an order-preserving bijection between the homomorphism lattice $(\mathscr{C}, \to)$ and the set of logical-equivalence classes of primitive-positive sentences partially ordered by $\vdash$.

The following definition and proposition characterize existential-positive sentences in terms of certain finite models.

Definition 2.15. A structure $\mathbf{M}$ is an *hom-minimal model* of an existential-positive sentence ψ if

- $\mathbf{M} \models \psi$ and
- $(\mathbf{A} \models \psi$ and $\mathbf{A} \to \mathbf{M}) \Longrightarrow \mathbf{M} \to \mathbf{A}$ for all structures $\mathbf{A}$.

Proposition 2.16. *Suppose* $\mathbf{M}$ *is a finite hom-minimal model of an existential-positive sentence* ψ*. Then* $\text{qcount}(\psi) \geq |\mathbf{Core}(\mathbf{M})|$ *and* $\text{qrank}(\psi) \geq \text{td}(\mathbf{Core}(\mathbf{M}))$*.*

PROOF. By Lemma 2.12, ψ is logically equivalent to a finite disjunction $\theta_1 \vee \cdots \vee \theta_m$ of primitive-positive sentences θ_i such that $\mathrm{qcount}(\psi) \geq \max_i \mathrm{qcount}(\theta_i)$ and $\mathrm{qrank}(\psi) \geq \max_i \mathrm{qrank}(\theta_i)$. By Lemma 2.13, there exist finite structures $\mathbf{A}_1, \ldots, \mathbf{A}_m$ such that $|\mathbf{A}_i| \leq \mathrm{qcount}(\theta_i)$ and $\mathrm{td}(\mathbf{A}_i) \leq \mathrm{qrank}(\theta_i)$ and $\mathbf{A}_i \to \mathbf{B} \iff \mathbf{B} \models \theta_i$ for all structures $\mathbf{B}$. Note that $\mathbf{A}_i \models \psi$ since $\mathbf{A}_i \models \theta_i$ (as $\mathbf{A}_i \to \mathbf{A}_i$) for all i. Since $\mathbf{M} \models \psi$, there exists $j \in \{1, \ldots, m\}$ such that $\mathbf{M} \models \theta_j$ and hence $\mathbf{A}_j \to \mathbf{M}$. As $\mathbf{M}$ is hom-minimal with respect to ψ, it follows that $\mathbf{M} \to \mathbf{A}_j$, that is, $\mathbf{A}_j \rightleftarrows \mathbf{M}$. Since $\mathbf{A}_j$ and $\mathbf{M}$ are homomorphically equivalent finite structures, they have isomorphic cores $\mathbf{Core}(\mathbf{A}_j) \cong \mathbf{Core}(\mathbf{M})$. Recall that $|\mathbf{Core}(\mathbf{A}_j)| \leq |\mathbf{A}_j|$ and $\mathrm{td}(\mathbf{Core}(\mathbf{A}_j)) \leq \mathrm{td}(\mathbf{A}_j)$ by Lemma 2.11. Therefore, we have

$$|\mathbf{Core}(\mathbf{M})| = |\mathbf{Core}(\mathbf{A}_j)| \leq |\mathbf{A}_j| \leq \mathrm{qcount}(\theta_i) \leq \mathrm{qcount}(\psi),$$
$$\mathrm{td}(\mathbf{Core}(\mathbf{M})) = \mathrm{td}(\mathbf{Core}(\mathbf{A}_j)) \leq \mathrm{td}(\mathbf{A}_j) \leq \mathrm{qrank}(\theta_i) \leq \mathrm{qrank}(\psi). \quad \square$$

Notice that any two hom-minimal models $\mathbf{M}$ and $\mathbf{N}$ of an existential-positive sentence ψ are either homomorphically equivalent ($\mathbf{M} \rightleftarrows \mathbf{N}$) or homomorphically incomparable ($\mathbf{M} \not\to \mathbf{N}$ and $\mathbf{N} \not\to \mathbf{M}$). Consequently, the set $\{\mathbf{Core}(\mathbf{M}) : \mathbf{M}$ is a finite hom-minimal model of $\psi\}$ is an antichain in the homomorphism lattice $(\mathscr{C}, \to)$ (in fact, a finite antichain). This antichain completely characterizes ψ up to logical equivalence (in fact, there is a bijection between finite antichains in $(\mathscr{C}, \to)$ and existential-positive sentences up to logical equivalence).

2.7 Back-and-Forth Equivalence

For structures $\mathbf{A}, \mathbf{B}$ and $n \in \mathbb{N}$, we write $\mathbf{A} \equiv^n \mathbf{B}$ and say that $\mathbf{A}$ and $\mathbf{B}$ are *n-back-and-forth equivalent* if they satisfy exactly the same first-order sentences of quantifier-rank n. The sequence of equivalence relations $\equiv^0, \equiv^1, \equiv^2, \ldots$, each one a refinement of the previous, measures the extent to which $\mathbf{A}$ and $\mathbf{B}$ look alike from the perspective of first-order logic.

Remark 2.17. Another name for $\equiv^n$ is *elementary equivalence up to quantifier-rank n*. (Our nonstandard choice of terminology "n-back-and-forth equivalence" refers to the combinatorial characterization of $\equiv^n$ in terms of n-back-and-forth systems of partial isomorphisms, soon to be defined.) Structures $\mathbf{A}$ and $\mathbf{B}$ are said to be *elementarily equivalent* if $\mathbf{A} \equiv^n \mathbf{B}$ for all $n \in \mathbb{N}$, that is, if no first-order sentence (of any quantifier-rank) distinguishes between $\mathbf{A}$ and $\mathbf{B}$. Elementary equivalence is a rather subtle equivalence relation on infinite structures: classical model theory attempts to understand (count, parameterize, classify, etc.) the often very rich collection of non-isomorphic structures within a single elementary equivalence class. On finite structures, the story is much simpler: elementary equivalence entails isomorphism. In fact, any two finite structures $\mathbf{A}$ and $\mathbf{B}$ are isomorphic if, and only if, $\mathbf{A} \equiv^{\min(|A|,|B|)+1} \mathbf{B}$.

Remark 2.18. It is well-known folklore that there are only finitely many $\equiv^n$-equivalence classes of structures for each n (although the exact number depends on the finite relational vocabulary σ). This is closely related to the fact that there are only finitely many first-order sentences of quantifier-rank n up to logical equivalence.

There are a number of useful combinatorial characterizations of $\equiv^n$. One popular characterization is via Ehrenfeucht-Fraïssé games. We prefer working with the (essentially identical) concept of an n-back-and-forth system of partial isomorphisms, due to Fraïssé. Recall that a *partial isomorphism* from $\mathbf{A}$ to $\mathbf{B}$ is a partial function π from A to B which restricts to an isomorphism from $\mathbf{A}|_{\mathrm{Dom}(\pi)}$ to $\mathbf{B}|_{\mathrm{Range}(\pi)}$.

Definition 2.19. An *n-back-and-forth system* on structures $\mathbf{A}$ and $\mathbf{B}$ is a sequence $\emptyset \neq \Pi_0 \subseteq \Pi_1 \subseteq \cdots \subseteq \Pi_n$ of sets Π_i of partial isomorphisms from $\mathbf{A}$ to $\mathbf{B}$ such that for all $i \in \{1, \dots, n\}$ and $\pi \in \Pi_{i-1}$,

$$\begin{array}{lll} \text{(forth)} & \forall a \in A\ \exists \pi_a \in \Pi_i & \text{such that } \pi_a \text{ extends } \pi \text{ and } a \in \mathrm{Dom}(\pi_a), \\ \text{(back)} & \forall b \in B\ \exists \pi_b \in \Pi_i & \text{such that } \pi_b \text{ extends } \pi \text{ and } b \in \mathrm{Range}(\pi_b). \end{array}$$

We now state the key lemma characterizing $\equiv^n$ in terms of this definition. Proof of this fundamental lemma can be found in any (finite) model theory text.

LEMMA 2.20. *Structures $\mathbf{A}$ and $\mathbf{B}$ are n-back-and-forth equivalent if, and only if, there exists an n-back-and-forth system on $\mathbf{A}$ and $\mathbf{B}$.* $\square$

Remark 2.21. Extending $\equiv^n$ to a relation between structures with distinguished tuples (of the same arity), a simple inductive characterization emerges:

- $(\mathbf{A}, \vec{a}) \equiv^0 (\mathbf{B}, \vec{b})$ if and only if $\vec{a} \longmapsto \vec{b}$ is a legitimate partial isomorphism from $\mathbf{A}$ to $\mathbf{B}$;
- for $n \geq 1$, $(\mathbf{A}, \vec{a}) \equiv^n (\mathbf{B}, \vec{b})$ if and only if $\forall \alpha \in A\ \exists \beta \in B\ (\mathbf{A}, \vec{a}\alpha) \equiv^{n-1} (\mathbf{B}, \vec{b}\beta)$ and $\forall \beta \in B\ \exists \alpha \in A\ (\mathbf{A}, \vec{a}\alpha) \equiv^{n-1} (\mathbf{B}, \vec{b}\beta)$.

3. BOUNDED TREE-DEPTH

In this section, we take a fresh look at the concepts of homomorphisms and cores from the perspective of structures with bounded tree-depth. For every natural number n, we define a relation of *n-homomorphism* between structures which approximates the usual relation of homomorphism "up to tree-depth n". This leads to the notion of the *n-core* of a structure over a set. We then introduce the property of *n-freeness* of sets Y and Z over a set X in a structure $\mathbf{A}$. This property is the bounded tree-depth analogue of Y and Z being separated by X in the Gaifman graph $\mathcal{G}(\mathbf{A})$.

Proviso 3.1. In this section and throughout the rest of this article, $X, Y, Z, \dots$ will always be finite sets (in particular, finite subsets of structures $\mathbf{A}, \mathbf{B}, \mathbf{C}, \dots$ even when these structures are infinite). Whenever we mention $\mathbf{A}, \mathbf{B}, \mathbf{C}, \dots, X, Y, Z, \dots$ in the same breath, it is assumed that $X, Y, Z, \dots$ lie inside the intersection of universes $A, B, C, \dots$. For instance, the statement $(\mathbf{A} \oplus_X \mathbf{B}) \to (\mathbf{C} \oplus_Y \mathbf{D})$ includes the assumption that $(X \cup Y) \subseteq (A \cap B \cap C \cap D)$.

3.1 n-Homomorphism

Definition 3.2. Let $\mathbf{A}$ and $\mathbf{B}$ be structures over a set X and let $n \in \mathbb{N}$. We write $\mathbf{A} \to^n_X \mathbf{B}$ and say that $\mathbf{A}$ is *n-homomorphic to* $\mathbf{B}$ *over* X if $\mathbf{C} \to_X \mathbf{A} \Longrightarrow \mathbf{C} \to_X \mathbf{B}$ for all finite structures $\mathbf{C}$ of tree-depth at most n over X. We write $\mathbf{A} \rightleftarrows^n_X \mathbf{B}$ and say that $\mathbf{A}$ and $\mathbf{B}$ are *n-homomorphically equivalent over* X if both $\mathbf{A} \to^n_X \mathbf{B}$ and $\mathbf{B} \to^n_X \mathbf{A}$. As usual, we write $\mathbf{A} \to^n \mathbf{B}$ (resp. $\mathbf{A} \rightleftarrows^n \mathbf{B}$) if $\mathbf{A} \to^n_\emptyset \mathbf{B}$ (resp. $\mathbf{A} \rightleftarrows^n_\emptyset \mathbf{B}$).

For use later on (in §4 and §5), we provide a special notation for the generalized concepts of n-homomorphism relative to a partial isomorphism π, as well as n-homomorphism over X between structures with distinguished k-tuples (cp. Notation 2.5).

Notation 3.3. Let π be a partial isomorphism from $\mathbf{A}$ to $\mathbf{B}$ with domain $X \subseteq A$. We write $\mathbf{A} \to_\pi^n \mathbf{B}$ if $\mathbf{C} \to_X \mathbf{A} \Longrightarrow \mathbf{C} \to_\pi \mathbf{B}$ for all finite structures $\mathbf{C}$ of tree-depth at most n over X. We write $\mathbf{A} \rightleftarrows_\pi^n \mathbf{B}$ if both $\mathbf{A} \to_\pi^n \mathbf{B}$ and $\mathbf{B} \to_{\pi^{-1}}^n \mathbf{A}$. For tuples $\vec{a} \in A^k$ and $\vec{b} \in B^k$, we write $(\mathbf{A}, \vec{a}) \to_X^n (\mathbf{B}, \vec{b})$ if $(\mathbf{C}, \vec{c}) \to_X (\mathbf{A}, \vec{a}) \Longrightarrow (\mathbf{C}, \vec{c}) \to_X (\mathbf{B}, \vec{b})$ for all finite structures $\mathbf{C}$ and tuples $\vec{c} \in C^k$ such that $\mathbf{C}$ has tree-depth at most n over $X \cup \{c_1, \ldots, c_k\}$. For example, if π is the identity function on X and $\vec{x}$ is a tuple whose coordinates enumerate the elements of X, then the three expressions $\mathbf{A} \to_X^n \mathbf{B}$ and $\mathbf{A} \to_\pi^n \mathbf{B}$ and $(\mathbf{A}, \vec{x}) \to^n (\mathbf{B}, \vec{x})$ have the same meaning. Each of these notations will be convenient at different times.

All lemmas in this section which involve the notion of n-homomorphism over a set ($\to_X^n$) generalize in straightforward manner to statements in terms of n-homomorphism relative to a partial isomorphism ($\to_\pi^n$). As a matter of taste, we prefer dealing with homomorphisms over sets rather than relative to partial isomorphisms. However, in a few cases where the partial-isomorphism version of a result is explicitly needed, both versions are presented (with the partial-isomorphism version stated as a corollary).

Before turning to the main lemma of this subsection, we point out a few basic properties of the n-homomorphism relation $\to_X^n$. First, $\to_X^n$ is transitive. In fact, just like $\to_X$, the relation $\to_X^n$ is a quasi-order on structures. Clearly, $\mathbf{A} \to_X \mathbf{B}$ implies $\mathbf{A} \to_X^n \mathbf{B}$ for all n. Just as clearly, $\mathbf{A} \to_X^n \mathbf{B}$ implies $\mathbf{A} \to_{X'}^{n'} \mathbf{B}$ for all $n' \leq n$ and $X' \subseteq X$. Only slightly less obvious is the fact that $\mathbf{A} \to_X^n \mathbf{B}$ if, and only if, $\mathbf{C} \to_X \mathbf{A} \Longrightarrow \mathbf{C} \to_X \mathbf{B}$ for all finite structures $\mathbf{C}$ of tree-depth at most n over X such that $\mathcal{G}(\mathbf{C}) \setminus X$ is connected.

The following lemma, which we will invoke repeatedly throughout this article, says that $\to_X^n$ behaves essentially just like $\to_X$ with respect to X-sums. (It is an easy exercise to check that Lemma 3.4 remains valid if $\to_X^n$ is replaced by $\to_X$.)

Lemma 3.4. *Let $\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_1, \mathbf{B}_2$ be structures over X.*

(1) *If $\mathbf{A}_1 \to_X^n \mathbf{B}_1$ and $\mathbf{A}_2 \to_X^n \mathbf{B}_2$, then $(\mathbf{A}_1 \oplus_X \mathbf{A}_2) \to_X^n (\mathbf{B}_1 \oplus_X \mathbf{B}_2)$.*

(2) *If $\mathbf{A}_1 \to_{X_1}^n \mathbf{B}_1$ and $\mathbf{A}_2 \to_{X_2}^n \mathbf{B}_2$, then*

$$(\mathbf{A}_1 \oplus_{X_1 \cap X_2} \mathbf{A}_2) \to_{X_1 \cup X_2}^n (\mathbf{B}_1 \oplus_{X_1 \cap X_2} \mathbf{B}_2).$$

(3) *If $\mathbf{A} \to_X^n \mathbf{B}$, then $(\mathbf{A} \oplus_X \mathbf{B}) \rightleftarrows_B^n \mathbf{B}$.*

Proof. For a set W, we denote by $\varnothing_W$ the "empty" structure with universe W (in which all relation symbols are interpreted as the empty set). In our proof of statement (1), below, we will use the following general fact: for every homomorphism $f : \mathbf{A} \to_X \mathbf{B}$, there exist a structure $\mathbf{C}$ and homomorphisms $g : \mathbf{A} \to_X \mathbf{C}$ and $h : \mathbf{C} \to_X \mathbf{B}$ such that $\mathrm{td}_X(\mathbf{C}) \leq \mathrm{td}_X(\mathbf{A})$ and $f = h \circ g$ and $h^{-1}(X) = X$. To see this, define function $g : A \to A$ by letting $g(a) = f(a)$ if $f(a) \in X$ and letting $g(a) = a$ otherwise. Now let $\mathbf{C} = g(\mathbf{A})$ be the image of $\mathbf{A}$ under g (so clearly g is a homomorphism $\mathbf{A} \to_X \mathbf{C}$) and let $h = f \restriction C$. While $\mathbf{C}$ is not necessarily a substructure of $\mathbf{A}$, it does holds that $X \subseteq C$ and $\mathbf{C}|_{C \setminus X} \subseteq \mathbf{A}|_{A \setminus X}$, from which it follows that $\mathrm{td}_X(\mathbf{C}) \leq \mathrm{td}_X(\mathbf{A})$. Finally, it is easy to check that h is a homomorphism $\mathbf{C} \to_X \mathbf{B}$ and $h^{-1}(X) = X$.

(1) Assume that $\mathbf{A}_1 \to_X^n \mathbf{B}_1$ and $\mathbf{A}_2 \to_X^n \mathbf{B}_2$. Let $\mathbf{C}$ be a finite structure such that $\mathrm{td}_X(\mathbf{C}) \leq n$ and $\mathbf{C} \to_X (\mathbf{A}_1 \oplus_X \mathbf{A}_2)$. We must show that $\mathbf{C} \to_X (\mathbf{B}_1 \oplus_X \mathbf{B}_2)$. By the aforementioned general fact, there exist a structure $\mathbf{D}$ and homomorphisms $\mathbf{C} \to_X \mathbf{D}$ and $h : \mathbf{D} \to_X (\mathbf{A}_1 \oplus_X \mathbf{A}_2)$ such that $\mathrm{td}_X(\mathbf{D}) \leq n$ and $h^{-1}(X) = X$. We will show that $\mathbf{D} \to_X (\mathbf{B}_1 \oplus_X \mathbf{B}_2)$, thus proving $\mathbf{C} \to_X (\mathbf{B}_1 \oplus_X \mathbf{B}_2)$. We first decompose $\mathbf{D}$ as a finite X-sum $\mathbf{D} = \mathbf{D}_1 \oplus_X \cdots \oplus_X \mathbf{D}_m$ such that graphs $\mathcal{G}(\mathbf{D}_i) \setminus X$ are connected for all $i \in \{1, \dots, m\}$ (such a decomposition exists for all finite structures over X). The fact that $\mathcal{G}(\mathbf{D}_i) \setminus X$ is connected and $h^{-1}(X) = X$ implies that $h(D_i) \subseteq A_1$ or $h(D_i) \subseteq A_2$ (hence $\mathbf{D}_i \to_X \mathbf{A}_1$ or $\mathbf{D}_i \to_X \mathbf{A}_2$). Note that $\mathrm{td}_X(\mathbf{D}_i) \leq n$ since $\mathbf{D}_i \subseteq \mathbf{D}$. As $\mathbf{A}_1 \to_X^n \mathbf{B}_1$ and $\mathbf{A}_2 \to_X^n \mathbf{B}_2$, it follows that $\mathbf{D}_i \to_X \mathbf{B}_1$ or $\mathbf{D}_i \to_X \mathbf{B}_2$. Either way, we have $\mathbf{D}_i \to_X (\mathbf{B}_1 \oplus_X \mathbf{B}_2)$ for $i = 1, \dots, m$. Therefore, $\mathbf{D} \to_X (\mathbf{B}_1 \oplus_X \mathbf{B}_2)$, as required.

(2) Assume $\mathbf{A}_1 \to_{X_1}^n \mathbf{B}_1$ and $\mathbf{A}_2 \to_{X_2}^n \mathbf{B}_2$. Let $Y_1 = X_2 \setminus X_1$ and $Y_2 = X_1 \setminus X_2$, and let $\mathbf{A}_i \oplus \boldsymbol{\varnothing}_{Y_i}$ (resp. $\mathbf{B}_i \oplus \boldsymbol{\varnothing}_{Y_i}$) denote the disjoint union of $\mathbf{A}_i$ (resp. $\mathbf{B}_i$) and the empty structure $\boldsymbol{\varnothing}_{Y_i}$ with universe Y_i. Clearly, we have $(\mathbf{A}_i \oplus \boldsymbol{\varnothing}_{Y_i}) \to_{X_1 \cup X_2}^n (\mathbf{B}_i \oplus \boldsymbol{\varnothing}_{Y_i})$ for $i = 1, 2$. Therefore, by part (1),

$$((\mathbf{A}_1 \oplus \boldsymbol{\varnothing}_{Y_1}) \oplus_{X_1 \cup X_2} (\mathbf{A}_2 \oplus \boldsymbol{\varnothing}_{Y_2})) \to_{X_1 \cup X_2}^n ((\mathbf{B}_1 \oplus \boldsymbol{\varnothing}_{Y_1}) \oplus_{X_1 \cup X_2} (\mathbf{B}_2 \oplus \boldsymbol{\varnothing}_{Y_2})).$$

The result now follows, as

$$((\mathbf{A}_1 \oplus \boldsymbol{\varnothing}_{Y_1}) \oplus_{X_1 \cup X_2} (\mathbf{A}_2 \oplus \boldsymbol{\varnothing}_{Y_2})) \cong_{X_1 \cup X_2} (\mathbf{A}_1 \oplus_{X_1 \cap X_2} \mathbf{A}_2),$$
$$((\mathbf{B}_1 \oplus \boldsymbol{\varnothing}_{Y_1}) \oplus_{X_1 \cup X_2} (\mathbf{B}_2 \oplus \boldsymbol{\varnothing}_{Y_2})) \cong_{X_1 \cup X_2} (\mathbf{B}_1 \oplus_{X_1 \cap X_2} \mathbf{B}_2).$$

(3) Assume $\mathbf{A} \to_X^n \mathbf{B}$. Let $\mathbf{B}' \cong_X \mathbf{B}$, that is, let $\mathbf{B}'$ be a structure over X which is isomorphic over X to $\mathbf{B}$. Let $\mathbf{A}_1 = \mathbf{A}$ and $\mathbf{B}_1 = \mathbf{B}'$ and $\mathbf{A}_2 = \mathbf{B}_2 = \mathbf{B}$ and $X_1 = X$ and $X_2 = B$. Note that $\mathbf{A}_1 \to_{X_1} \mathbf{B}_1$ and $\mathbf{A}_2 \to_{X_2} \mathbf{B}_2$ and $X_1 \cap X_2 = X$ and $X_1 \cup X_2 = B$ (since $X \subseteq B$). By part (2), we have $\mathbf{A} \oplus_X \mathbf{B} \to_B^n \mathbf{B}' \oplus_X \mathbf{B}$. Combining this observation with the obvious homomorphisms $\mathbf{B}' \oplus_X \mathbf{B} \to_B \mathbf{B}$ and $\mathbf{B} \to_B \mathbf{A} \oplus_X \mathbf{B}$, we conclude that $\mathbf{A} \oplus_X \mathbf{B} \rightleftarrows_B^n \mathbf{B}$. □

For future reference (specifically in the proof of Lemma 5.12), we state without proof the analogue of Lemma 3.4(2) for n-homomorphism relative to a partial isomorphism (see Notation 3.3).

Corollary 3.5. *Suppose π is a partial isomorphism from a structure $\mathbf{A} = \mathbf{A}_1 \oplus_X \mathbf{A}_2$ to structure $\mathbf{B} = \mathbf{B}_1 \oplus_Y \mathbf{B}_2$ such that $X \subseteq \mathrm{Dom}(\pi)$ and $\pi(X) = Y$. For $i = 1, 2$, let $\pi \restriction A_i$ denote the restriction of π to $\mathrm{Dom}(\pi) \cap A_i$. If $\mathbf{A}_1 \to_{\pi \restriction A_1}^n \mathbf{B}_1$ and $\mathbf{A}_2 \to_{\pi \restriction A_2}^n \mathbf{B}_2$, then $\mathbf{A} \to_\pi^n \mathbf{B}$. Moreover, if $\mathbf{A}_1 \rightleftarrows_{\pi \restriction A_1}^n \mathbf{B}_1$ and $\mathbf{A}_2 \rightleftarrows_{\pi \restriction A_2}^n \mathbf{B}_2$, then $\mathbf{A} \rightleftarrows_\pi^n \mathbf{B}$.* □

In light of a connection (to be discussed in §3.4) between $\rightleftarrows^n$-equivalence classes of structures and existential-positive theories up to quantifier-rank n, Lemma 3.4 and Corollary 3.5 can be viewed as "composition theorems" in the spirit of the Feferman-Vaught theorem from classical model theory [Feferman and Vaught 1959].

3.2 n-Cores

Definition 3.6. For every finite set X and natural number n, we denote by $\mathscr{C}_X^n$ the set of finite canonical cores over X with tree-depth at most n over X. That is, $\mathscr{C}_X^n = \{\mathbf{C} \in \mathscr{C}_X : \mathrm{td}_X(\mathbf{C}) \leq n\}$. Members of $\mathscr{C}_X^n$ are called *n-cores over X*.

A few observations are in order. First and most obvious, $\mathscr{C}_X$ is the union of sets $\mathscr{C}_X^n$ over all $n \in \mathbb{N}$. Second, we have $\mathbf{A} \rightarrow_X^n \mathbf{B}$ if, and only if, $\mathbf{C} \rightarrow_X \mathbf{A} \Longrightarrow \mathbf{C} \rightarrow_X \mathbf{B}$ for every $\mathbf{C} \in \mathscr{C}_X^n$. Third, $\rightarrow_X$ coincides with $\rightarrow_X^n$ on the class $\mathscr{C}_X^n$. That is, $\mathbf{C}_1 \rightarrow_X \mathbf{C}_2 \Longleftrightarrow \mathbf{C}_1 \rightarrow_X^n \mathbf{C}_2$ for all $\mathbf{C}_1, \mathbf{C}_2 \in \mathscr{C}_X^n$. Fourth, $\rightarrow_X$ is clearly a partial order on $\mathscr{C}_X^n$, since $\mathscr{C}_X^n$ is a subset of the homomorphism lattice $(\mathscr{C}_X, \rightarrow_X)$. So what about the poset $(\mathscr{C}_X^n, \rightarrow_X)$? Is it also a lattice?

LEMMA 3.7. *$(\mathscr{C}_X^n, \rightarrow_X)$ is an upper semilattice. That is, every two structures in $\mathscr{C}_X^n$ have a least upper bound (l.u.b.) with respect to $\rightarrow_X$. Moreover, the l.u.b. of two structures in $(\mathscr{C}_X^n, \rightarrow_X)$ coincides with their l.u.b. in the lattice $(\mathscr{C}_X, \rightarrow_X)$.*

PROOF. Consider any $\mathbf{C}_1, \mathbf{C}_2 \in \mathscr{C}_X^n$. Recall from §2.4 that $\mathbf{Core}_X(\mathbf{C}_1 \oplus_X \mathbf{C}_2)$ is the l.u.b. of $\mathbf{Core}_X(\mathbf{C}_1)$ $(= \mathbf{C}_1)$ and $\mathbf{Core}_X(\mathbf{C}_2)$ $(= \mathbf{C}_2)$ in the homomorphism lattice $(\mathscr{C}_X, \rightarrow)$. Therefore, it suffices to show that $\mathbf{Core}_X(\mathbf{C}_1 \oplus_X \mathbf{C}_2)$ has tree-depth $\leq n$ over X, that is, $\mathbf{Core}_X(\mathbf{C}_1 \oplus_X \mathbf{C}_2) \in \mathscr{C}_X^n$. This follows from Lemma 2.8(1,2) as

$$\mathrm{td}_X(\mathbf{Core}_X(\mathbf{C}_1 \oplus_X \mathbf{C}_1)) \leq \mathrm{td}_X(\mathbf{C}_1 \oplus_X \mathbf{C}_1) = \max\{\mathrm{td}_X(\mathbf{C}_1), \mathrm{td}_X(\mathbf{C}_1)\} \leq n,$$

where the first inequality is by Lemma 2.11. □

Remark 3.8. $(\mathscr{C}_X^n, \rightarrow_X)$ is in fact a lattice (not just an upper semilattice). However, it fails to be a sublattice of $(\mathscr{C}_X, \rightarrow_X)$ since the greatest lower bound (g.l.b.) of two structures in $\mathscr{C}_X^n$ does not always coincide with their g.l.b. in $\mathscr{C}_X$. We give an example of this phenomenon when $n = 2$ and $X = \emptyset$. In the example, the vocabulary σ consists of binary relations R_1, R_2, R_3, R_4. Structures $\mathbf{C}_1$ and $\mathbf{C}_2$ and their product $\mathbf{C}_1 \otimes \mathbf{C}_2$ are depicted below on the left. Ordered pairs of elements belonging to the relations R_1, R_2, R_3, R_3 are indicated by arrows; each line style corresponds to one of the four relations.

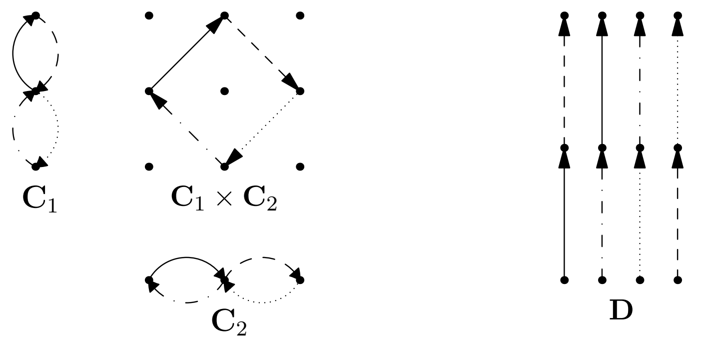

$\mathbf{C}_1$ and $\mathbf{C}_2$ are easily seen to be cores of tree-depth 2 (over $\emptyset$). Their product $\mathbf{C}_1 \otimes \mathbf{C}_2$ is not a core; however, $\mathbf{Core}(\mathbf{C}_1 \otimes \mathbf{C}_2)$ is isomorphic to the motley 4-cycle inside $\mathbf{C}_1 \otimes \mathbf{C}_2$ (i.e., what remains after removing the five isolated elements). We have seen that $\mathbf{Core}(\mathbf{C}_1 \otimes \mathbf{C}_2)$ is the g.l.b. of $\mathbf{Core}(\mathbf{C}_1)$ $(= \mathbf{C}_1)$ and $\mathbf{Core}(\mathbf{C}_2)$ $(= \mathbf{C}_2)$ in the homomorphism lattice $(\mathscr{C}_X, \to_X)$ (see §2.4). However, $\mathbf{Core}(\mathbf{C}_1 \otimes \mathbf{C}_2)$ has tree-depth 3 and thus cannot be the g.l.b. of $\mathbf{C}_1$ and $\mathbf{C}_2$ in $(\mathscr{C}^2_X, \to_X)$. That structure is $\mathbf{D}$, pictured above on the right. ($\mathbf{D}$ is what we call the 2-*core* of $\mathbf{C}_1 \otimes \mathbf{C}_2$, to be defined shortly.)

The next proposition gives a key finiteness property of bounded tree-depth structures. (Here is one place where the assumption that σ is a finite relational vocabulary is crucial. Proposition 3.9 is false for vocabularies which contain infinitely many relation symbols.)

PROPOSITION 3.9. *Up to $\rightleftarrows_X$, there are only finitely many finite structures over X with tree-depth $\leq n$ over X. Equivalently, there are only finitely many canonical cores over X of tree-depth $\leq n$ over X (i.e., $\mathscr{C}^n_X$ is a finite set).*

In a graph-theoretic context, Proposition 3.9 is proved as Corollary 3.4 of [Nešetřil and de Mendez 2006]. The proof for relational structures is essentially the same. By a straightforward induction, one can show that $|\mathscr{C}^n_X| \leq 2^{2^{\cdots 2^{\kappa}}}$ (a tower of exponentials of height $n+1$) where $\log_2 \kappa = \sum_{R\in\sigma}(|X|+n)^{\mathrm{arity}(R)}$; here κ is the number of σ-structures with universe $X \sqcup \{1,\dots,n\}$. (Proposition 3.9 may also be deduced from the folklore theorem that there are only finite many first-order sentences of quantifier-rank n up to logical equivalence; see Remark 2.18.)

The fact that $(\mathscr{C}^n_X, \to_X)$ is finite means it is complete as an upper semilattice, that is, every subset of $\mathscr{C}^n_X$ has a least upper bound. This leads to the notion of the *n-core* of a (finite or infinite) structure over a finite subset of its universe.

Definition 3.10 *(n-Core)*. For a structure $\mathbf{A}$ and a finite set $X \subseteq A$, the *n-core* $\mathbf{Core}^n_X(\mathbf{A})$ *of* $\mathbf{A}$ *over* X is the least upper bound of $\{\mathbf{C} \in \mathscr{C}^n_X : \mathbf{C} \to_X \mathbf{A}\}$ in $(\mathscr{C}^n_X, \to_X)$.

We emphasize that $\mathbf{Core}^n_X(\mathbf{A})$ is defined for all structures $\mathbf{A}$ (so long as $X \subseteq A$ is a finite set). This is in contrast to $\mathbf{Core}_X(\mathbf{A})$, which is defined only for finite structures $\mathbf{A}$.

We point out two useful characterizations of $\mathbf{Core}^n_X(\mathbf{A})$, both easy exercises. On one hand, we have the explicit expression

$$\mathbf{Core}^n_X(\mathbf{A}) = \mathbf{Core}_X(\textstyle\bigoplus_X \{\mathbf{C} \in \mathscr{C}^n_X : \mathbf{C} \to_X \mathbf{A}\}).$$

Alternatively, one can check that $\mathbf{Core}^n_X(\mathbf{A})$ is the unique $\mathbf{C} \in \mathscr{C}^n_X$ such that $\mathbf{A} \rightleftarrows^n_X \mathbf{C}$.

The following lemma underscores the utility of the operator $\mathbf{Core}^n_X(\cdot)$: it allows us to answer questions about $\to^n_X$ by answering questions about $\to_X$ (a relationship between structures for which there is a simple explicit existential witness, namely a homomorphism fixing X pointwise).

LEMMA 3.11. $\mathbf{A} \to^n_X \mathbf{B}$ *if, and only if,* $\mathbf{Core}^n_X(\mathbf{A}) \to_X \mathbf{B}$.

PROOF. Let $\mathscr{D}$ be the set $\{\mathbf{C} \in \mathscr{C}^n_X : \mathbf{C} \to_X \mathbf{A}\}$, so in particular $\mathbf{Core}^n_X(\mathbf{A}) = \mathbf{Core}_X(\bigoplus_X \mathscr{D})$. We have $\mathbf{A} \to^n_X \mathbf{B} \Longleftrightarrow \bigwedge_{\mathbf{C}\in\mathscr{D}}(\mathbf{C} \to_X \mathbf{B})$ by virtue of our earlier observation that $\mathbf{A} \to^n_X \mathbf{B}$ if, and only if, $\mathbf{C} \to_X \mathbf{A} \Longrightarrow \mathbf{C} \to_X \mathbf{B}$ for all $\mathbf{C} \in \mathscr{C}^n_X$. Now notice that $\bigwedge_{\mathbf{C}\in\mathscr{D}}(\mathbf{C} \to_X \mathbf{B}) \Longleftrightarrow \bigoplus_X \mathscr{D} \to_X \mathbf{B} \Longleftrightarrow \mathbf{Core}_X(\bigoplus_X \mathscr{D}) \to_X \mathbf{B} \Longleftrightarrow \mathbf{Core}^n_X(\mathbf{A}) \to_X \mathbf{B}$. Putting these statements together, we get $\mathbf{A} \to^n_X \mathbf{B} \Longleftrightarrow \mathbf{Core}^n_X(\mathbf{A}) \to_X \mathbf{B}$ as required. □

As a special case of Lemma 3.11, we have $\mathbf{Core}^n_X(\mathbf{A}) \to_X \mathbf{A}$. The next lemma describes an important relationship between the parameters n and X in the definition of $\mathbf{Core}^n_X(\mathbf{A})$. This lemma will come in handy at a few critical junctures later on.

LEMMA 3.12. *For every structure* $\mathbf{A}$ *and finite sets* $X, Y \subseteq A$, *there is a homomorphism* $\mathbf{Core}^n_{X\cup Y}(\mathbf{A}) \to_X \mathbf{Core}^{n+|Y|}_X(\mathbf{A})$.

PROOF. By definition of $\mathbf{Core}^m_X(\mathbf{A})$, we have $\mathbf{B} \to_X \mathbf{Core}^m_X(\mathbf{A})$ for all $m \in \mathbb{N}$ and all structures $\mathbf{B}$ such that $\mathbf{B} \to_X \mathbf{A}$ and $\mathrm{td}_X(\mathbf{B}) \le m$. To prove the lemma, we simply apply this observation with $\mathbf{B} = \mathbf{Core}^n_{X\cup Y}(\mathbf{A})$ and $m = n + |Y|$, using the fact that

$$\mathrm{td}_X(\mathbf{Core}^n_{X\cup Y}(\mathbf{A})) \le \mathrm{td}_{X\cup Y}(\mathbf{Core}^n_{X\cup Y}(\mathbf{A})) + |Y| \le n + |Y|$$

where the first inequality is by Lemma 2.8(3). □

3.3 Core-Size Bound

The core-size bound, defined below, will play an important role later on. It comes up in the freeness lemma (Lemma 3.22), which features prominently in the proof of the finite homomorphism preservation theorem (Theorem 5.16).

Definition 3.13 *(Core-Size Bound).* For all $m, n \in \mathbb{N}$, let

$$\beta^n_m = \max\{|C| : \mathbf{C} \in \mathscr{C}^n_X \text{ such that } |X| = m\}.$$

That is, β^n_m is the size of the largest n-core over a set of size m.

Since $\mathscr{C}^n_X$ is a finite set (by Proposition 3.9) and each $\mathbf{C} \in \mathscr{C}^n_X$ is a finite structure, we see that β^n_m is a finite number. (We remark that the core-size bound, like practically everything else in this article, depends on the finite relational vocabulary σ.) Note that the core-size bound satisfies $\beta^{n'}_{m'} \le \beta^n_m$ for all $n' \le n$ and $m' \le m$.

Even though we view β^n_m as a function from $\mathbb{N} \times \mathbb{N}$ to $\mathbb{N}$, we have chosen notation β^n_m (as opposed say to $\beta(n, m)$) for the sake of consistency with related notations like $\mathscr{C}^n_X$ and $\to^n_X$. In these related notations, the superscript n corresponds to tree-depth over the subscript X. By analogy, the subscript m in β^n_m stands for the size $|X|$ of the set X.

3.4 Logical Characterization of n-Homomorphism

We now link up the notions of n-homomorphism and existential-positive definability via the correspondence described in §2.6 between primitive-positive sentences (up to logical equivalence) and finite structures (up to homomorphic equivalence).

LEMMA 3.14. $\mathbf{A} \to^n \mathbf{B}$ *if, and only if,* $\mathbf{A} \models \theta \Longrightarrow \mathbf{B} \models \theta$ *for every primitive-positive sentence* θ *of quantifier-rank* n.

PROOF. ($\Longrightarrow$) Assume $\mathbf{A} \to^n \mathbf{B}$ and suppose θ is a primitive-positive sentence of quantifier-rank n such that $\mathbf{A} \models \theta$. By Lemma 2.13, there exists a finite structure $\mathbf{C}$ (called in $\mathbf{A}_\theta$ in Lemma 2.13) such that $\mathrm{td}(\mathbf{C}) \le n$ and $\mathbf{C} \to \mathbf{D} \Longleftrightarrow \mathbf{D} \models \theta$ for all structures $\mathbf{D}$. Since $\mathbf{A} \models \theta$, we have $\mathbf{C} \to \mathbf{A}$. By definition of $\to^n$, it follows that $\mathbf{C} \to \mathbf{B}$. Therefore, $\mathbf{B} \models \theta$.

($\Longleftarrow$) Assume that $\mathbf{A} \models \theta \Longrightarrow \mathbf{B} \models \theta$ for every primitive-positive sentence θ of quantifier-rank n. Suppose $\mathbf{C}$ is a finite structure of tree-depth $\leq n$ such that $\mathbf{C} \to \mathbf{A}$. By Lemma 2.14, there exists a primitive-positive sentence $\theta_{\mathbf{C}}$ such that $\mathrm{qrank}(\theta_{\mathbf{C}}) \leq n$ and $\mathbf{D} \models \theta_{\mathbf{C}} \Longleftrightarrow \mathbf{C} \to \mathbf{D}$ for all $\mathbf{D}$. Since $\mathbf{C} \to \mathbf{A}$, we have $\mathbf{A} \models \theta_{\mathbf{C}}$. As $\mathrm{qrank}(\theta_{\mathbf{C}}) \leq n$, we have $\mathbf{B} \models \theta_{\mathbf{C}}$ and hence $\mathbf{C} \to \mathbf{B}$. Since $\mathbf{C} \to \mathbf{A} \Longrightarrow \mathbf{C} \to \mathbf{B}$ for all finite $\mathbf{C}$ of tree-depth $\leq n$, we conclude that $\mathbf{A} \to^n \mathbf{B}$. □

The next two lemmas rely on the fact that there are only finitely many $\rightleftarrows^n$-equivalence classes by Proposition 3.9.

Lemma 3.15. *For any class $\mathcal{P}$ of [finite] structures, statements (i)–(iii) are equivalent:*

(i) *$\mathcal{P}$ is definable by an existential-positive sentence of quantifier-rank n [on finite structures],*
(ii) *$\mathcal{P}$ is closed under $\to^n$ [on finite structures],*
(iii) *$\mathcal{P}$ is closed under $\to$ as well as closed under $\rightleftarrows^n$ [on finite structures].* □

Since we will not use this lemma, we leave its proof as an exercise. Instead, we prove a closely related lemma, which we will explicitly use later on.

Lemma 3.16. *Let $\mathcal{P}$ and $\mathcal{Q}$ be classes of [finite] structures. Suppose that for all [finite] structures $\mathbf{A}$ and $\mathbf{B}$, if $\mathbf{A} \in \mathcal{P}$ and $\mathbf{A} \to^n \mathbf{B}$ then $\mathbf{B} \in \mathcal{Q}$. Then there exists an existential-positive sentence Ψ of quantifier-rank at most n such that $\mathcal{P} \subseteq \mathrm{Mod}_{[\mathrm{fin}]}(\Psi) \subseteq \mathcal{Q}$.*

In logical terms, we say that Ψ is an "interpolant" between $\mathcal{P}$ and $\mathcal{Q}$.

Proof. Let $\mathscr{P}$ be the set of n-cores of structures in $\mathcal{P}$, that is, $\mathscr{P} = \{\mathbf{Core}^n(\mathbf{A}) : \mathbf{A} \in \mathcal{P}\}$. Since $\mathscr{P}$ is a subset of $\mathscr{C}^n$, which is finite by Proposition 3.9, $\mathscr{P}$ is clearly also finite. Let $\mathbf{C}_1, \dots, \mathbf{C}_m$ enumerate the members of $\mathscr{P}$. By Lemma 2.14, there exist primitive-positive sentences $\theta_1, \dots, \theta_m$ such that $\mathrm{qrank}(\theta_i) \leq \mathrm{td}(\mathbf{C}_i)$ and $\mathbf{C}_i \to \mathbf{D} \Longleftrightarrow \mathbf{D} \models \theta_i$ for all structures $\mathbf{D}$. Let Ψ be the finite disjunction $\theta_1 \vee \dots \vee \theta_m$. Note that Ψ is an existential-positive sentence of quantifier-rank $\max_{i=1}^m \mathrm{qrank}(\theta_i) = \max_{i=1}^m \mathrm{td}(\mathbf{C}_i) \leq n$.

Consider any $\mathbf{A} \in \mathcal{P}$. We have $\mathbf{Core}^n(\mathbf{A}) \in \mathscr{P}$ and hence $\mathbf{Core}^n(\mathbf{A}) = \mathbf{C}_i$ for some $i \in \{1, \dots, m\}$. Since $\mathbf{Core}^n(\mathbf{A}) \to \mathbf{A}$ (by Lemma 3.11), it follows that $\mathbf{C}_i \to \mathbf{A}$ and thus $\mathbf{A} \models \theta_i$ and consequently $\mathbf{A} \models \Psi$. Therefore, $\mathcal{P} \subseteq \mathrm{Mod}_{[\mathrm{fin}]}(\Psi)$.

Now consider any $\mathbf{B} \in \mathrm{Mod}_{[\mathrm{fin}]}(\Psi)$. There exists $i \in \{1, \dots, m\}$ such that $\mathbf{B} \models \theta_i$ and thus $\mathbf{C}_i \to \mathbf{B}$. It follows that $\mathbf{Core}^n(\mathbf{A}) \to \mathbf{B}$ for some $\mathbf{A} \in \mathcal{P}$. By Lemma 3.11, we have $\mathbf{A} \to^n \mathbf{B}$. Therefore, $\mathbf{B} \in \mathcal{Q}$ and so $\mathrm{Mod}_{[\mathrm{fin}]}(\Psi) \subseteq \mathcal{Q}$. □

3.5 Freeness

We define a relationship of *freeness* between two sets Y and Z over a third set X in a structure $\mathbf{A}$. Intuitively, freeness captures the notion of separation (in the graph-theoretic sense) from the standpoint of homomorphisms. (Recall that two sets U and V are *separated* in a graph $\mathcal{G}$ if there is no path in $\mathcal{G}$ between U and V.) We also define a relationship of *n-freeness*, which is the natural bounded tree-depth analogue of freeness. A heads up to the reader: the results of this section are not used in the proof of the equirank h.p.t. in §4, but will play an essential role when we prove the finite h.p.t. in §5.

Definition 3.17. Let $\mathbf{A}$ be a structure and let $X, Y, Z \subseteq A$ such that $Y \cap Z \subseteq X$. We say that Y and Z are:

separated over X if sets $Y \setminus X$ and $Z \setminus X$ are separated in the graph $\mathcal{G}(\mathbf{A}) \setminus X$;
free over X if $\mathbf{A} \to_{Y \cup Z} \mathbf{B}$ for all $\mathbf{B}$ such that $\mathbf{A} \to_{X \cup Y} \mathbf{B}$ and $\mathbf{A} \to_{X \cup Z} \mathbf{B}$;
n-free over X if $\mathbf{A} \to^n_{Y \cup Z} \mathbf{B}$ for all $\mathbf{B}$ such that $\mathbf{A} \to^n_{X \cup Y} \mathbf{B}$ and $\mathbf{A} \to^n_{X \cup Z} \mathbf{B}$.

The statement that Y and Z are free (resp. n-free) over X is expressed via notation $Y \perp\!\!\!\perp_X Z$ (resp. $Y \perp\!\!\!\perp^n_X Z$).

A little later on, we will give some examples to help the reader internalize this definition. But first, let us prove a few preliminary lemmas.

Notation 3.18. Let $(\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$ denote the structure $\mathbf{A} \oplus_X \mathbf{A}$ (viewed as a structure over $X \cup Y \cup Z$) in which the set $Y \setminus X$ is identified with its lefthand copy and $Z \setminus X$ is identified with its righthand copy, as illustrated below (in the simple case where X, Y, Z are disjoint).

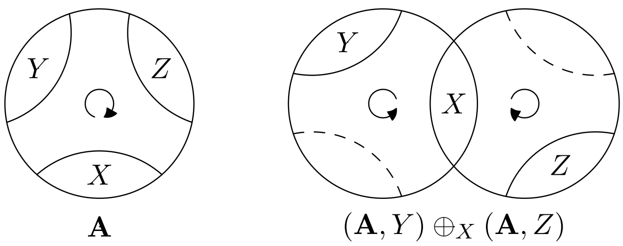

Notice that sets Y and Z are separated over X in the structure $(\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$. Also note that there is an obvious homomorphism $(\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z) \to_{X \cup Y \cup Z} \mathbf{A}$. (In the picture above, this homomorphism can be visualized as folding the lefthand side of $\mathbf{A} \oplus_X \mathbf{A}$ onto the righthand side.)

The following lemma characterizes freeness and n-freeness via the existence of certain homomorphisms. In the statements of Lemmas 3.19 and 3.20 below, we assume that X, Y, Z are subsets of a structure $\mathbf{A}$ such that $Y \cap Z \subseteq X$.

LEMMA 3.19. *Sets Y and Z are free over X in $\mathbf{A}$ if, and only if, there exists a homomorphism $\mathbf{A} \to_{Y \cup Z} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$. They are n-free over X in $\mathbf{A}$ if, and only if, $\mathbf{A} \to^n_{Y \cup Z} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$.*

PROOF. We will prove the statement for n-freeness; the argument for freeness is similar (and simpler). Assume $Y \perp\!\!\!\perp^n_X Z$ in $\mathbf{A}$. The natural lefthand and righthand inclusion maps are homomorphisms $\mathbf{A} \to_{X \cup Y} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$ and $\mathbf{A} \to_{X \cup Z} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$. Since $\to_{X \cup Y}$ implies $\to^n_{X \cup Y}$ and $\to_{X \cup Z}$ implies $\to^n_{X \cup Z}$, we have $\mathbf{A} \to^n_{X \cup Y} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$ and $\mathbf{A} \to^n_{X \cup Z} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$. Therefore, we have $\mathbf{A} \to^n_{Y \cup Z} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$ by n-freeness of Y and Z over X in $\mathbf{A}$.

In the other direction, assume $\mathbf{A} \to^n_{Y \cup Z} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$. Suppose $\mathbf{B}$ is a structure such that $\mathbf{A} \to^n_{X \cup Y} \mathbf{B}$ and $\mathbf{A} \to^n_{X \cup Z} \mathbf{B}$. Applying Lemma 3.4(2) with $\mathbf{A}_1 = \mathbf{A}_2 = \mathbf{A}$ and $\mathbf{B}_1 = \mathbf{B}_2 = \mathbf{B}$ and $X_1 = X \cup Y$ and $X_2 = X \cup Z$, we get

$$(\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z) \to^n_{X \cup Y \cup Z} (\mathbf{B}, Y) \oplus_X (\mathbf{B}, Z).$$

Since $(\mathbf{B},Y) \oplus_X (\mathbf{B},Z) \to_{X\cup Y\cup Z} \mathbf{B}$ (as noted earlier), we have $\mathbf{A} \to^n_{Y\cup Z} \mathbf{B}$ (since the composition of relations $\mathbf{A} \to^n_{Y\cup Z} \cdot \to^n_{X\cup Y\cup Z} \cdot \to_{X\cup Y\cup Z} \mathbf{B}$ implies the relation $\to^n_{Y\cup Z}$). We have thus shown $Y \perp\!\!\!\perp^n_X Z$ in $\mathbf{A}$, as required. □

The next lemma describes the basic relationships among separation, freeness and n-freeness.

Lemma 3.20 (Properties of Separation, Freeness and n-Freeness).

(1) *If Y and Z are separated over X in $\mathbf{A}$, then they are free over X. If Y and Z are free over X, then they are n-free over X for every $n \in \mathbb{N}$.*

(2) *If Y and Z are separated (resp. free, n-free) over X in $\mathbf{A}$, then they are separated (resp. free, n-free) over every set X' such that $X \subseteq X' \subseteq A$.*

(3) *If Y and Z are free (resp. n-free) over X in $\mathbf{A}$, then they are free (resp. n-free) over X in every structure $\mathbf{A}'$ such that $\mathbf{A} \rightleftarrows_{X\cup Y\cup Z} \mathbf{A}'$ (resp. $\mathbf{A} \rightleftarrows^n_{X\cup Y\cup Z} \mathbf{A}'$).*

(4) *If Y and Z are free (resp. n-free) over X in $\mathbf{A}$, then they are free (resp. n-free) over X in every co-retract of $\mathbf{A}$.*

Proof. The various parts of this lemma follow almost directly from definitions (with help from Lemma 3.19).

(1) If Y and Z are separated over X in $\mathbf{A}$, we can find a homomorphism $\mathbf{A} \to_{X\cup Y\cup Z} (\mathbf{A},Y) \oplus_X (\mathbf{A},Z)$, which by Lemma 3.19 proves that $Y \perp\!\!\!\perp_X Z$ in $\mathbf{A}$. The statement that free implies n-free follows from the fact that $\mathbf{A} \to_{Y\cup Z} (\mathbf{A},Y) \oplus_X (\mathbf{A},Z)$ implies $\mathbf{A} \to^n_{Y\cup Z} (\mathbf{A},Y) \oplus_X (\mathbf{A},Z)$.

(2) In the case of separation, the statement is obvious. In the case of n-freeness, suppose $Y \perp\!\!\!\perp^n_X Z$ in $\mathbf{A}$ and $X \subseteq X' \subseteq A$. Let $\mathbf{B}$ be any structure such that $\mathbf{A} \to^n_{X'\cup Y} \mathbf{B}$ and $\mathbf{A} \to^n_{X'\cup Z} \mathbf{B}$. It clearly follows that $\mathbf{A} \to^n_{X\cup Y} \mathbf{B}$ and $\mathbf{A} \to^n_{X\cup Z} \mathbf{B}$. So by n-freeness of Y and Z over X, we have $\mathbf{A} \to^n_{Y\cup Z} \mathbf{B}$. We conclude that $Y \perp\!\!\!\perp^n_{X'} Z$ in $\mathbf{A}$, thus proving the statement for n-freeness. The same basic argument works in the case of freeness.

(3) Suppose $Y \perp\!\!\!\perp^n_X Z$ in $\mathbf{A}$ and $\mathbf{A} \rightleftarrows^n_{X\cup Y\cup Z} \mathbf{A}'$ (so in particular $\mathbf{A}$ and $\mathbf{A}'$ are equivalent under $\rightleftarrows^n_{X\cup Y}$ and $\rightleftarrows^n_{X\cup Z}$ and $\rightleftarrows^n_{Y\cup Z}$). Let $\mathbf{B}$ be any structure such that $\mathbf{A}' \to^n_{X\cup Y} \mathbf{B}$ and $\mathbf{A}' \to^n_{X\cup Z} \mathbf{B}$. We have $\mathbf{A} \to^n_{X\cup Y} \mathbf{B}$ and $\mathbf{A} \to^n_{X\cup Z} \mathbf{B}$ by transitivity of $\to^n_{X\cup Y}$ and $\to^n_{X\cup Z}$. By n-freeness of Y and Z over X, it follows that $\mathbf{A} \to^n_{Y\cup Z} \mathbf{B}$. We now get $\mathbf{A}' \to^n_{Y\cup Z} \mathbf{B}$ by transitivity of $\to^n_{Y\cup Z}$. Therefore, $Y \perp\!\!\!\perp^n_X Z$ in $\mathbf{A}'$. The same basic argument works in the case of freeness.

(4) This follows from part (3), as $\mathbf{A}' \xrightarrow{\text{retr}} \mathbf{A}$ implies $\mathbf{A} \rightleftarrows_{X\cup Y\cup Z} \mathbf{A}'$. □

Examples 3.21. Having described the basic properties of freeness and n-freeness, we present several examples to help the reader internalize these notions (and also, indirectly, the notion of tree-depth).

(a) Sets Y and Z are 0-free over X in $\mathbf{A}$ if, and only if, there is no edge between $Y \setminus X$ and $Z \setminus X$ in the Gaifman graph $\mathcal{G}(\mathbf{A}) \setminus X$. (Easy exercise.)

(b) Sets Y and Z are n-free over $\emptyset$ in $\mathbf{A}$ if, and only if, the distance between Y and Z in the Gaifman graph $\mathcal{G}(\mathbf{A})$ is at least $2^n + 1$. To see this, suppose that Y and Z have distance $\geq 2^n + 1$ in $\mathcal{G}(\mathbf{A})$. Let $\mathbf{C}$ be any structure with tree-depth $\leq n$ over $Y \cup Z$ such that there exists a homomorphism $h : \mathbf{C} \to_{Y\cup Z} \mathbf{A}$.

Clearly, Y and Z have distance $\geq 2^n + 1$ in $\mathcal{G}(\mathbf{C})$. From this and the fact that $\mathrm{td}_{Y \cup Z}(\mathbf{C}) \leq n$, it follows that $\mathbf{C}$ is a disjoint union of substructures $\mathbf{C}_1 \oplus \mathbf{C}_2$ such that $Y \subseteq C_1$ and $Z \subseteq C_2$. We now get a homomorphism $\mathbf{C} = \mathbf{C}_1 \oplus \mathbf{C}_2 \to_{Y \cup Z} (\mathbf{A}, Y) \oplus (\mathbf{A}, Z)$ by mapping $\mathbf{C}_1$ (resp. $\mathbf{C}_2$) to the lefthand (resp. righthand) copy of $\mathbf{A}$ via the homomorphism h. Thus, we have shown that $\mathbf{A} \to^n_{Y \cup Z} (\mathbf{A}, Y) \oplus (\mathbf{A}, Z)$. Lemma 3.19 now implies that $Y \perp\!\!\!\perp^n_\emptyset Z$ in $\mathbf{A}$. Conversely, suppose that Y and Z have distance $\leq 2^n$ in $\mathcal{G}(\mathbf{A})$. There exists a path P in $\mathcal{G}(\mathbf{A})$ of length $\leq 2^n$ with endpoints in Y and Z. The substructure $\mathbf{A}|_{Y \cup Z \cup P}$ has tree-depth $\leq n$ over $Y \cup Z$ and satisfies $\mathbf{A}|_{Y \cup Z \cup P} \to_{Y \cup Z} \mathbf{A}$ and $\mathbf{A}|_{Y \cup Z \cup P} \not\to_{Y \cup Z} (\mathbf{A}, Y) \oplus (\mathbf{A}, Z)$. Therefore, $\mathbf{A} \not\to^n_{Y \cup Z} (\mathbf{A}, Y) \oplus (\mathbf{A}, Z)$. By Lemma 3.19, it follows that $Y \not\perp\!\!\!\perp^n_\emptyset Z$ in $\mathbf{A}$.

(c) Generalizing one direction of examples (a,b): if the distance between $Y \setminus X$ and $Z \setminus X$ in $\mathcal{G}(\mathbf{A}) \setminus X$ is at least $2^n + 1$, then $Y \perp\!\!\!\perp^n_X Z$ in $\mathbf{A}$. The converse, however, is false whenever $X \neq \emptyset$ or $n \neq 0$ (see example (g)).

(d) Sets Y and Z are free over $\emptyset$ in $\mathbf{A}$ if, and only if, they are separated over $\emptyset$ in $\mathbf{A}$. To see this, suppose $Y \perp\!\!\!\perp_\emptyset Z$ in $\mathbf{A}$ and let $\mathbf{B}$ be the structure $(\mathbf{A}, Y) \oplus (\mathbf{A}, Z)$, that is, $\mathbf{A} \oplus \mathbf{A}$ where the set Y (resp. Z) is identified with its lefthand (resp. righthand) copy. We have $\mathbf{A} \to_Y \mathbf{B}$ and $\mathbf{A} \to_Z \mathbf{B}$ via the left and right embeddings of set A into $A \sqcup A$. By freeness, we have $\mathbf{A} \to_{Y \cup Z} \mathbf{B}$. It follows that Y and Z are separated in $\mathbf{A}$, since these sets are separated in $\mathbf{B}$. The reverse implication (separated over $\emptyset \Longrightarrow$ free over $\emptyset$) is a special case of Lemma 3.20(1).

(e) Consider a directed graph $\mathbf{D}$ (viewed as a structure in a vocabulary containing a single binary relation) with vertices y, x_1, x_2, z and arcs (y, x_1), (y, x_2), (x_1, z), (x_2, z). Sets $\{y\}$ and $\{z\}$ are separated over $\{x_1, x_2\}$, but not separated over $\{x_1\}$. However, $\{y\}$ and $\{z\}$ are free over $\{x_1\}$. (This can be seen by noting that there is a retraction $\mathbf{D} \xrightarrow{\text{retr}} \mathbf{D}|_{\{x_1, y, z\}}$.)

(f) Consider a different directed graph $\mathbf{D}'$ consisting of vertices y and z together with a directed path from y to z of length ℓ (i.e., with $\ell - 1$ intermediate vertices) for every integer $\ell \geq 2$. (To be clear, these paths are mutually disjoint except at endpoints y and z.) In this case, $\{y\}$ and $\{z\}$ are n-free over a set X in $\mathbf{D}'$ if, and only if, X intersects the path of length ℓ for every $\ell \in \{2, 3, \dots, 2^n\}$.

(g) Suppose we modify the directed graph $\mathbf{D}'$ from example (f) so that there are exactly two directed paths from y to z for each length $\ell \geq 2$. Now $\{y\}$ and $\{z\}$ are n-free over a set X if, and only if, X intersects at least one of each pair of paths of length ℓ for every $\ell \in \{2, 3, \dots, 2^n\}$. (In particular, if X intersects exactly one of each pair of paths of length $2, 3, \dots, 2^n$, then y and z have distance 2 in $\mathcal{G}(\mathbf{D}') \setminus X$ and yet $\{y\} \perp\!\!\!\perp^n_X \{z\}$ in $\mathbf{D}'$; cp. example (c).)

Let $\mathbf{A}$ be any structure and consider any subsets $Y, Z \subseteq A$. Clearly Y and Z are separated and hence n-free in $\mathbf{A}$ over the entire universe A. This raises an interesting question: what is the size of the (not necessarily unique) smallest set W over which Y and Z are n-free in $\mathbf{A}$? The next lemma answers this question by showing that $|W|$ is bounded by a function of $|Y \cup Z|$ and n (independent of the structure $\mathbf{A}$). The fact that there exists a bound on $|W|$ independent of $\mathbf{A}$ is not at all obvious from the definition of n-freeness. Here is where the core-size bound β^n_m makes its entrance. Recall that β^n_m is the size of the largest n-core over a set of size m.

Lemma 3.22 (Freeness Lemma). *Suppose that sets Y and Z are n-free over X in a structure $\mathbf{A}$. Then there exists a subset $W \subseteq X$ of size $|W| \leq \beta^n_{|Y \cup Z|}$ such that Y and Z are n-free over W in $\mathbf{A}$.*

Proof. Suppose $Y \perp\!\!\!\perp^n_X Z$ in $\mathbf{A}$. By Lemma 3.19, we have $\mathbf{A} \to^n_{Y \cup Z} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$. So by Lemma 3.11, there exists a homomorphism $h : \mathbf{C} \to_{Y \cup Z} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$ where $\mathbf{C} = \mathbf{Core}^n_{Y \cup Z}(\mathbf{A})$.

Let $h(\mathbf{C})$ denote the homomorphic image of $\mathbf{C}$ and let $W = X \cap h(C)$. Sets Y and Z are clearly separated and hence free over W in $h(\mathbf{C})$. There is thus a chain of homomorphisms

$$\mathbf{C} \to_{Y \cup Z} h(\mathbf{C}) \to_{Y \cup Z} (h(\mathbf{C}), Y) \oplus_W (h(\mathbf{C}), Z) \to_{Y \cup Z} (\mathbf{A}, Y) \oplus_W (\mathbf{A}, Z)$$

where the middle homomorphism exists by Lemma 3.19 and the righthand homomorphism is the natural embedding of $h(\mathbf{C}) \oplus_W h(\mathbf{C})$ into $\mathbf{A} \oplus_W \mathbf{A}$. By composition, we get a homomorphism $\mathbf{Core}^n_{Y \cup Z}(\mathbf{A}) \to_{Y \cup Z} (\mathbf{A}, Y) \oplus_W (\mathbf{A}, Z)$. Lemma 3.11 now implies $\mathbf{A} \to^n_{Y \cup Z} (\mathbf{A}, Y) \oplus_W (\mathbf{A}, Z)$. Therefore, by Lemma 3.19 we have Y and Z are n-free over W in $\mathbf{A}$.

To complete the proof, note that $|W| \leq |h(\mathbf{C})| \leq |\mathbf{C}| \leq \beta^n_{|Y \cup Z|}$ since $\mathbf{C}$ is an n-core over $Y \cup Z$. □

Remark 3.23. We might say that sets Y and Z are *strongly n-free over X* in a structure $\mathbf{A}$ if $\mathbf{A} \to^n_{X \cup Y \cup Z} \mathbf{B}$ for all structures $\mathbf{B}$ such that $\mathbf{A} \to^n_{X \cup Y} \mathbf{B}$ and $\mathbf{A} \to^n_{X \cup Z} \mathbf{B}$. (Equivalently, Y and Z are strongly n-free over X if there exists a homomorphism $\mathbf{A} \to^n_{X \cup Y \cup Z} (\mathbf{A}, Y) \oplus_X (\mathbf{A}, Z)$.) As the name would suggest, strong n-freeness over X clearly implies n-freeness over X. While the converse is false, strong n-freeness over X is implied by $(n + |X|)$-freeness over X. Indeed, the notions of n-freeness and strong n-freeness are superficially quite similar. The definition of strong n-freeness might even appear more natural and therefore more worth studying than n-freeness. The critical difference (and the reason we study n-freeness instead of strong n-freeness) is the freeness lemma (Lemma 3.22), which becomes false when restated with "strongly n-free" replacing "n-free".

3.6 Extension Cores

Let $c_1, c_2, c_3, \dots$ be a fixed sequence of arbitrary but distinct dummy elements. For concreteness, we can take c_i to be the ith ordinal. We are now, for the time being, no longer interested in cores over a finite set X, but rather in cores over the disjoint union $X \sqcup \{c_1, \dots, c_m\}$ where $m \in \mathbb{N}$.

Definition 3.24. For every structure $\mathbf{A}$, finite subset $X \subseteq A$ and $m, n \in \mathbb{N}$, let

$$\mathscr{E}^n_{X,m}(\mathbf{A}) = \left\{ \mathbf{C} \in \mathscr{C}^n_{X \sqcup \{c_1, \dots, c_m\}} : \mathbf{C} \to_X \mathbf{A} \right\}$$

where $X \sqcup \{c_1, \dots, c_m\}$ is the disjoint union of sets X and $\{c_1, \dots, c_m\}$. Members of $\mathscr{E}^n_{X,m}(\mathbf{A})$ are called *m-extension n-cores of* $\mathbf{A}$ *over* X. The union $\mathscr{E}^n_{X,0}(\mathbf{A}) \cup \dots \cup \mathscr{E}^n_{X,m}(\mathbf{A})$ is denoted by $\mathscr{E}^n_{X,\leq m}(\mathbf{A})$.

The key properties of extension cores are enumerated below.

Lemma 3.25 (Properties of Extension Cores).

(*1*) $(\mathbf{A} \oplus_X \mathbf{C}) \xrightarrow{\text{retr}} \mathbf{A}$ *for every* $\mathbf{C} \in \mathscr{E}^n_{X,m}(\mathbf{A})$.

(2) $\mathscr{E}^n_{X,\le m}(\mathbf{A})$ *is a finite set and every* $\mathbf{C} \in \mathscr{E}^n_{X,\le m}(\mathbf{A})$ *is a finite structure.*

(3) If $\mathbf{A} \to^{m+n}_X \mathbf{B}$*, then* $\mathscr{E}^n_{X,m}(\mathbf{A}) \subseteq \mathscr{E}^n_{X,m}(\mathbf{B})$.

(4) If $\mathbf{A} \rightleftarrows^{m+n}_X \mathbf{B}$ *(in particular, if* $\mathbf{A} \xrightarrow{\text{retr}} \mathbf{B}$ *and* $X \subseteq B$*), then* $\mathscr{E}^n_{X,m}(\mathbf{A}) = \mathscr{E}^n_{X,m}(\mathbf{B})$.

Proof. Statement (1) follows from Lemma 2.9(2) since $\mathbf{C} \to_X \mathbf{A}$ for every $\mathbf{C} \in \mathscr{E}^n_{X,m}(\mathbf{A})$. Statement (2) is a direct consequence of Proposition 3.9. As for (3), suppose $\mathbf{A} \to^{m+n}_X \mathbf{B}$ and let $\mathbf{C} \in \mathscr{E}^n_{X,m}(\mathbf{A})$. By Lemma 2.8(3), we have

$$\mathrm{td}_X(\mathbf{C}) \le \mathrm{td}_{X \sqcup \{c_1,\dots,c_m\}}(\mathbf{C}) + |\{c_1,\dots,c_m\}| \le n + m.$$

Therefore, $\mathbf{A} \to^{m+n}_X \mathbf{B}$ yields $\mathbf{C} \to_X \mathbf{B}$. It follows that $\mathbf{C} \in \mathscr{E}^n_{X,m}(\mathbf{B})$ and hence $\mathscr{E}^n_{X,m}(\mathbf{A}) \subseteq \mathscr{E}^n_{X,m}(\mathbf{B})$. This proves statement (3), which implies (4) straightaway. □

For future reference (specifically in the proof of Lemma 5.12), we state a generalization of Lemma 3.25(4) in terms of $n{+}m$-homomorphism relative to a partial isomorphism π (recall Notation 3.3).

Corollary 3.26. *Suppose* $\mathbf{A} \rightleftarrows^{n+m}_\pi \mathbf{B}$ *where* π *is a partial isomorphism from* $\mathbf{A}$ *to* $\mathbf{B}$ *with domain* $S \subseteq A$ *and range* $T \subseteq B$*. Let* $\mathbf{C} \in \mathscr{E}^n_{S,m}(\mathbf{A})$ *and let* $\mathbf{D}$ *be the isomorphic image of* $\mathbf{C}$ *under the function* $\pi \cup \mathrm{id}_{C \setminus S} : C \longrightarrow (C \setminus S) \cup T$ *defined by* $c \longmapsto \pi(c)$ *if* $c \in S$ *and* $c \longmapsto c$ *if* $c \in C \setminus S$*. Then* $\mathbf{D} \in \mathscr{E}^n_{T,m}(\mathbf{B})$. □

Corollary 3.26 specializes to Lemma 3.25(4) in the case where $S = T = X$ and π is the identity function on X.

4. EQUIRANK HOMOMORPHISM PRESERVATION THEOREM

In this section, we prove our first main result: the equirank h.p.t. This section is in principle independent from the following section (§5), in which we prove the finite h.p.t. (In particular, the n-extension property plays no role in that proof.) However, a close reading of this section will serve as a good warm-up for the more complicated construction in §5.

4.1 The n-Extension Property

For structures $\mathbf{A}$ and $\mathbf{B}$ and tuples $\vec{a} \in A^k$ and $\vec{b} \in B^k$, recall that notation $(\mathbf{A}, \vec{a}) \to^n_X (\mathbf{B}, \vec{b})$ expresses $(\mathbf{C}, \vec{c}) \to_X (\mathbf{A}, \vec{a}) \Longrightarrow (\mathbf{C}, \vec{c}) \to_X (\mathbf{B}, \vec{b})$ for every finite structure $\mathbf{C}$ and tuple $\vec{c} \in C^k$ such that $\mathbf{C}$ has tree-depth at most n over $X \cup \{c_1,\dots,c_k\}$ (Notation 3.3).

Definition 4.1. A structure $\mathbf{A}$ is *n-extendable* if, for every set $X \subseteq A$ of size $< n$ and every structure $\mathbf{B}$ such that $\mathbf{A} \rightleftarrows^{n-|X|}_X \mathbf{B}$, it holds that $\forall b \in B\ \exists a \in A$ s.t. $(\mathbf{A}, a) \rightleftarrows^{n-|X|-1}_X (\mathbf{B}, b)$.

Cast into different notation (see Notation 3.3), $\mathbf{A}$ is n-extendable if, and only if, for every structure $\mathbf{B}$ and partial isomorphism π from $\mathbf{A}$ to $\mathbf{B}$ such that $|\pi| < n$ and $\mathbf{A} \rightleftarrows^{n-|\pi|}_\pi \mathbf{B}$, it holds that $\forall b \in B\ \exists a \in A$ s.t. $\mathbf{A} \rightleftarrows^{n-|\pi|-1}_{\pi \cup \{(a,b)\}} \mathbf{B}$.

Remark 4.2. The n-extension property was first defined in a preliminary version of this article [Rossman 2005], where it was called *n-existential-positive saturation.*

The terminology was changed to avoid confusion with the common model-theoretic notion of saturation. The new terminology bears an intentional similarity to the "n-extension axioms" that show up in the theory of the infinite random graph [Blass and Rossman 2005; Grädel et al. 2007].

LEMMA 4.3. *Suppose structures* $\mathbf{A}$ *and* $\mathbf{B}$ *are* n*-extendable and* $\mathbf{A} \rightleftarrows^n \mathbf{B}$. *Then* $\mathbf{A} \equiv^n \mathbf{B}$.

PROOF. For $k \in \{0, \dots, n\}$, let Π_k be the set of partial isomorphisms π from $\mathbf{A}$ to $\mathbf{B}$ such that $|\pi| \leq k$ and $\mathbf{A} \rightleftarrows^{n-k}_{\pi} \mathbf{B}$. The n-extendability of $\mathbf{A}$ and $\mathbf{B}$ implies that the sequence $\Pi_0, \dots, \Pi_n$ is an n-back-and-forth system on $\mathbf{A}$ and $\mathbf{B}$. Therefore, $\mathbf{A} \equiv^n \mathbf{B}$ by Lemma 2.20. □

The next lemma gives a nice finitary criterion for n-extendability of a structure $\mathbf{A}$ by reducing the class of structures $\mathbf{B}$ and elements $b \in B$ in Definition 4.1 to finitely many possibilities.

LEMMA 4.4. *A structure* $\mathbf{A}$ *is* n*-extendable if, and only if, for every set* $X \subseteq A$ *of size* $< n$ *and* 1*-extension core* $\mathbf{C} \in \mathscr{E}^{n-|X|-1}_{X,1}(\mathbf{A})$, *there exists* $a \in A$ *such that* $(\mathbf{A}, a) \rightleftarrows^{n-|X|-1}_X (\mathbf{A} \oplus_X \mathbf{C}, c_1)$.

PROOF. The ($\Longrightarrow$) direction is trivial, so we prove only the ($\Longleftarrow$) direction. Assume that for every every set $X \subseteq A$ of size $< n$ and 1-extension core $\mathbf{C} \in \mathscr{E}^{n-|X|-1}_{X,1}(\mathbf{A})$, there exists $a \in A$ such that $(\mathbf{A}, a) \rightleftarrows^{n-|X|-1}_X (\mathbf{A} \oplus_X \mathbf{C}, c_1)$. Let $X \subseteq A$ be such that $|X| < n$, let $\mathbf{B}$ be any structure such that $\mathbf{A} \rightleftarrows^{n-|X|}_X \mathbf{B}$, and let b be any element of B. We must find $a \in A$ such that $(\mathbf{A}, a) \rightleftarrows^{n-|X|-1}_X (\mathbf{B}, b)$. We may assume that $b \notin X$, since otherwise $(\mathbf{A}, b) \rightleftarrows^{n-|X|-1}_X (\mathbf{B}, b)$ and there is nothing further to show. We have

$$\begin{array}{cl} \mathbf{Core}^{n-|X|-1}_{X\cup\{b\}}(\mathbf{B}) & \\ \downarrow_X & \text{by Lemma 3.12,} \\ \mathbf{Core}^{n-|X|}_X(\mathbf{B}) & \\ \| & \text{since } \mathbf{A} \rightleftarrows^{n-|X|}_X \mathbf{B}, \\ \mathbf{Core}^{n-|X|}_X(\mathbf{A}) & \\ \downarrow_X & \\ \mathbf{A}. & \end{array}$$

Let $\mathbf{C}$ be the structure obtained from $\mathbf{Core}^{n-|X|-1}_{X\cup\{b\}}(\mathbf{B})$ by substituting b with the dummy element c_1 (assuming $c_1 \notin \mathbf{Core}^{n-|X|-1}_{X\cup\{b\}}(\mathbf{B})$ without loss of generality). Thus, we have $(\mathbf{C}, c_1) \cong_X (\mathbf{Core}^{n-|X|-1}_{X\cup\{b\}}(\mathbf{B}), b)$ and $(\mathbf{C}, c_1) \rightleftarrows^{n-|X|-1}_X (\mathbf{B}, b)$. Note that $\mathbf{C} \in \mathscr{C}^{n-|X|-1}_{X\sqcup\{c_1\}}$ and $\mathbf{C} \rightarrow_X \mathbf{A}$ (by the above). Thus, $\mathbf{C}$ is a 1-extension core in the set $\mathscr{E}^{n-|X|-1}_{X,1}(\mathbf{A})$. By our initial assumption, it follows that there exists an element $a \in A$ such that $(\mathbf{A}, a) \rightleftarrows^{n-|X|-1}_X (\mathbf{A} \oplus_X \mathbf{C}, c_1)$. Therefore, it suffices to show that $(\mathbf{A} \oplus_X \mathbf{C}, c_1) \rightleftarrows_X^{n-|X|-1} (\mathbf{B}, b)$. By Lemma 3.4(1), this follows from the previously established facts that $\mathbf{A} \rightleftarrows_X^{n-|X|-1} \mathbf{B}$ and $(\mathbf{C}, c_1) \rightleftarrows_X^{n-|X|-1} (\mathbf{B}, b)$. $\square$

4.2 Extendable Co-retracts

Definition 4.5. For every structure $\mathbf{A}$ and $n \in \mathbb{N}$, we define a new structure $\Xi_n(\mathbf{A})$ by

$$\Xi_n(\mathbf{A}) = \bigoplus_{\substack{X \subseteq A \\ |X|<n \\ \mathbf{C} \in \mathscr{E}_{X,1}^{n-|X|-1}(\mathbf{A})}}{}_{A} \Big(\mathbf{A} \oplus_X \mathbf{C}\Big).$$

For $\ell \in \mathbb{N}$, let

$$\Xi_n^\ell(\mathbf{A}) = \underbrace{\Xi_n(\Xi_n(\ldots \Xi_n}_{\ell \text{ times}}(\mathbf{A}) \ldots)).$$

Also, let

$$\Xi_n^\star(\mathbf{A}) = \bigcup_{\ell \in \mathbb{N}} \Xi_n^\ell(\mathbf{A}).$$

That is, $\Xi_n^\star(\mathbf{A})$ is the union of the chain of structures $\mathbf{A} \subseteq \Xi_n(\mathbf{A}) \subseteq \Xi_n(\Xi_n(\mathbf{A})) \subseteq \cdots$.

Example 4.6. Suppose that the vocabulary σ contains just a single binary relation (so that σ-structures are directed graphs). Let $\mathbf{A}$ be a σ-structure with universe $\{a, b, c\}$ and binary relation $\{(b, a), (b, c), (c, b)\}$. The structure $\Xi_2(\mathbf{A})$ is depicted below.

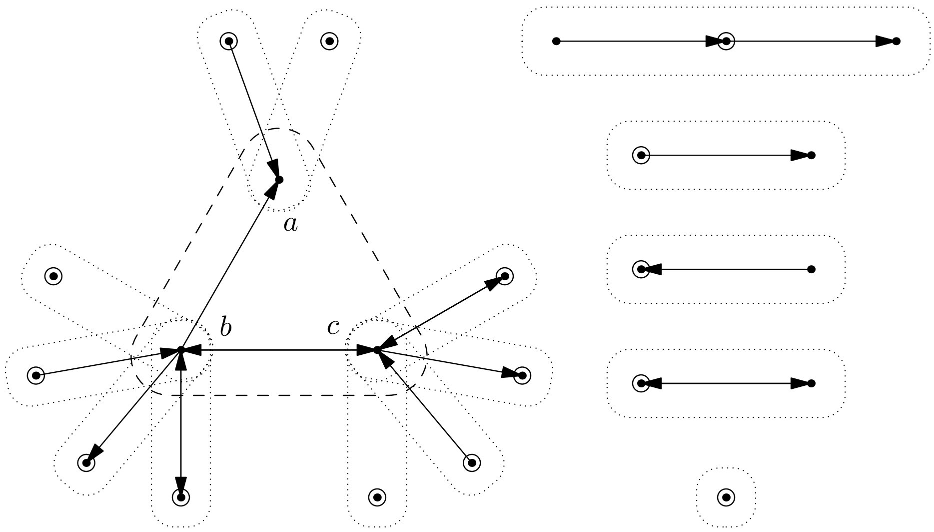

In this picture, $\mathbf{A}$ is clearly delineated as a substructure of $\Xi_2(\mathbf{A})$ (it lies inside the dashed triangle) with elements a, b, c explicitly labeled. The rest of $\Xi_2(\mathbf{A})$ consists of various 1-extension cores, each attached to a subset of $\{a, b, c\}$ of size

< 2. Extension cores attached to a single element $x \in \{a,b,c\}$ are canonical cores in $\mathscr{C}^0_{\{x,c_1\}}$, while those attached to the empty set $\emptyset$ are canonical cores in $\mathscr{C}^1_{\{c_1\}}$. Each extension core is surrounded by a dotted rectangle with the distinguished dummy element c_1 encircled.

LEMMA 4.7. *If* $\mathbf{A}$ *is a finite structure, then so is* $\Xi_n(\mathbf{A})$.

PROOF. Suppose $\mathbf{A}$ is finite. There are finitely many subsets $X \subseteq A$ (of size $< n$). Moreover, for every X, the set $\mathscr{E}^{n-|X|-1}_{X,1}(\mathbf{A})$ is finite and every each $\mathbf{C} \in \mathscr{E}^{n-|X|-1}_{X,1}(\mathbf{A})$ is a finite structure by Lemma 3.25(2). Thus, $\Xi_n(\mathbf{A})$ is clearly finite. □

On the other hand, note that $\Xi^\star_n(\mathbf{A})$ is always an infinite structure (for $n > 0$), even if $\mathbf{A}$ is finite. This is a main obstacle to proving that the equirank h.p.t. (Theorem 4.12) holds on finite structures.

LEMMA 4.8. *For every structure* $\mathbf{A}$ *and* $n \in \mathbb{N}$,

(1) $\Xi_n(\mathbf{A}) \xrightarrow{\text{retr}} \mathbf{A}$,

(2) $\Xi^\star_n(\mathbf{A}) \xrightarrow{\text{retr}} \Xi^\ell_n(\mathbf{A})$ *for all* $\ell \in \mathbb{N}$, *and*

(3) $\Xi_n(\mathbf{A}) \xrightarrow{\text{retr}} \mathbf{A} \oplus_X \mathbf{C}$ *for all* $X \subseteq A$ *of size* $< n$ *and all* $\mathbf{C} \in \mathscr{C}^{n-|X|-1}_X(\mathbf{A})$.

PROOF.

(1) For every index pair $(X, \mathbf{C})$ in the indexed A-sum defining $\Xi_n(\mathbf{A})$, we have $\mathbf{A} \oplus_X \mathbf{C} \xrightarrow{\text{retr}} \mathbf{A}$ by Lemma 3.25(1) (since $\mathbf{C}$ is an extension core of $\mathbf{A}$ over X). Therefore, $\Xi_n(\mathbf{A}) \xrightarrow{\text{retr}} \mathbf{A}$ by Lemma 2.9(3).

(2) Statement (1) implies that $\Xi^\star_n(\mathbf{A})$ is the union of the chain of co-retracts $\mathbf{A} \xleftarrow{\text{retr}} \Xi_n(\mathbf{A}) \xleftarrow{\text{retr}} \Xi_n(\Xi_n(\mathbf{A})) \xleftarrow{\text{retr}} \cdots$. It follows by Lemma 2.9(4) that $\Xi^\star_n(\mathbf{A})$ is a co-retract of $\Xi^\ell_n(\mathbf{A})$ for all $\ell \in \mathbb{N}$.

(3) Given X and $\mathbf{C}$, it holds that
$$\Xi_n(\mathbf{A}) = \bigoplus_{\substack{Y \subseteq A \\ |Y|<n \\ \mathbf{D} \in \mathscr{E}^{n-|Y|-1}_{Y,1}(\mathbf{A}) \\ (Y,\mathbf{D}) \neq (X,\mathbf{C})}}^{A \oplus_X C} \Big((\mathbf{A} \oplus_X \mathbf{C}) \oplus_Y \mathbf{D}\Big).$$

For every index pair $(Y, \mathbf{D})$, we have $\mathbf{D} \to_Y \mathbf{A}$ (since $\mathbf{D}$ is an extension core of $\mathbf{A}$ over Y) and thus $\mathbf{D} \to_Y \mathbf{A} \oplus_X \mathbf{C}$. It follows that $(\mathbf{A} \oplus_X \mathbf{C}) \oplus_Y \mathbf{D} \xrightarrow{\text{retr}} \mathbf{A} \oplus_X \mathbf{C}$ by Lemma 2.9(2). Therefore, $\Xi_n(\mathbf{A}) \xrightarrow{\text{retr}} \mathbf{A} \oplus_X \mathbf{C}$ by Lemma 2.9(3). □

LEMMA 4.9. $\Xi^\star_n(\mathbf{A})$ *is* n*-extendable.*

PROOF. By virtue of Lemma 4.4, it suffices to show that for every $X \subseteq \Xi^\star_n(\mathbf{A})$ of size $< n$ and every $\mathbf{C} \in \mathscr{E}^{n-|X|-1}_{X,1}(\Xi^\star_n(\mathbf{A}))$, there exists $a \in \Xi^\star_n(\mathbf{A})$ such that $(\Xi^\star_n(\mathbf{A}), a) \rightleftarrows^{n-|X|-1}_X (\Xi^\star_n(\mathbf{A}) \oplus_X \mathbf{C}, c_1)$. In fact, we will show—even stronger—that there exists $a \in \Xi^\star_n(\mathbf{A})$ such that $(\Xi^\star_n(\mathbf{A}), a) \rightleftarrows_X (\Xi^\star_n(\mathbf{A}) \oplus_X \mathbf{C}, c_1)$. To this end, we will define homomorphisms

$$f : (\Xi^\star_n(\mathbf{A}) \oplus_X \mathbf{C}) \to_X \Xi^\star_n(\mathbf{A}), \qquad g : \Xi^\star_n(\mathbf{A}) \to_X (\Xi^\star_n(\mathbf{A}) \oplus_X \mathbf{C})$$

such that $g(f(c_1)) = c_1$. The result then follows by setting $a = f(c_1)$.

Because X is finite, there exists $\ell \in \mathbb{N}$ such that $X \subseteq \Xi^{\ell}_n(\mathbf{A})$. Lemma 4.8(2) implies $\Xi^{\star}_n(\mathbf{A}) \xrightarrow{\text{retr}} \Xi^{\ell}_n(\mathbf{A})$. It follows that $\mathscr{E}^{n-|X|-1}_{X,1}(\Xi^{\ell}_n(\mathbf{A})) = \mathscr{E}^{n-|X|-1}_{X,1}(\Xi^{\star}_n(\mathbf{A}))$ by Lemma 3.25(4). Therefore, $\mathbf{C} \in \mathscr{E}^{n-|X|-1}_{X,1}(\Xi^{\ell}_n(\mathbf{A}))$.

We now wish to consider the structure $\Xi^{\ell}_n(\mathbf{A}) \oplus_X \mathbf{C}$, viewed as a substructure of $\Xi^{\ell+1}_n(\mathbf{A})$. However, we will instead consider the structure $\Xi^{\ell}_n(\mathbf{A}) \oplus_X \mathbf{C}'$ where $\mathbf{C}'$ is a structure which is isomorphic to $\mathbf{C}$ over X via an isomorphism $i : \mathbf{C} \xrightarrow{\cong}_X \mathbf{C}'$. This avoids potential confusion, as c_1 cannot be seen as an element of $\Xi^{\ell}_n(\mathbf{A}) \oplus_X \mathbf{C}'$.

By Lemma 4.8(3), we have $\Xi^{\ell+1}_n(\mathbf{A}) = \Xi_n(\Xi^{\ell}_n(\mathbf{A})) \xrightarrow{\text{retr}} \Xi^{\ell}_n(\mathbf{A}) \oplus_X \mathbf{C}'$. In particular, we have a chain of retractions

$$\Xi^{\star}_n(\mathbf{A}) \xrightarrow{\text{retr}} \Xi^{\ell+1}_n(\mathbf{A}) \xrightarrow{\text{retr}} \Xi^{\ell}_n(\mathbf{A}) \oplus_X \mathbf{C}' \xrightarrow{\text{retr}} \Xi^{\ell}_n(\mathbf{A}).$$

We now fix retractions

$$r : \Xi^{\star}_n(\mathbf{A}) \xrightarrow{\text{retr}} \Xi^{\ell}_n(\mathbf{A}), \qquad s : \Xi^{\star}_n(\mathbf{A}) \xrightarrow{\text{retr}} \Xi^{\ell}_n(\mathbf{A}) \oplus_X \mathbf{C}'.$$

Let $\overline{r}$ and $\overline{s}$ be the corresponding co-retractions

$$\overline{r} : \Xi^{\ell}_n(\mathbf{A}) \hookrightarrow \Xi^{\star}_n(\mathbf{A}), \qquad \overline{s} : \Xi^{\ell}_n(\mathbf{A}) \oplus_X \mathbf{C}' \hookrightarrow \Xi^{\star}_n(\mathbf{A}).$$

We form the $\oplus_X$-sum of homomorphisms

$$r \oplus_X i : \Xi^{\star}_n(\mathbf{A}) \oplus_X \mathbf{C} \longrightarrow_X \Xi^{\ell}_n(\mathbf{A}) \oplus_X \mathbf{C}',$$
$$\overline{r} \oplus_X i^{-1} : \Xi^{\ell}_n(\mathbf{A}) \oplus_X \mathbf{C}' \longrightarrow_X \Xi^{\star}_n(\mathbf{A}) \oplus_X \mathbf{C}.$$

(These maps are defined in the obvious way, viz. the value of $r \oplus_X i$ on an element $\alpha \in \Xi^{\star}_n(A) \oplus_X C$ equals $r(\alpha)$ if $\alpha \in \Xi^{\star}_n(A)$ and $i(\alpha)$ if $\alpha \in C$.) Homomorphisms $f : (\Xi^{\star}_n(\mathbf{A}) \oplus_X \mathbf{C}) \rightarrow_X \Xi^{\star}_n(\mathbf{A})$ and $g : \Xi^{\star}_n(\mathbf{A}) \rightarrow_X (\Xi^{\star}_n(\mathbf{A}) \oplus_X \mathbf{C})$ are now defined by

$$f = \overline{s} \circ (r \oplus_X i), \qquad g = (\overline{r} \oplus_X i^{-1}) \circ s.$$

It is easy to see that f and g fix X pointwise, since each of the constituent homomorphisms $r, s, i, \overline{r}, \overline{s}, i^{-1}$ fixes X pointwise. Finally, we have $g(f(c_0)) = c_0$ as:

$$\begin{aligned}
g(f(c_0)) &= g(\overline{s}((r \oplus_X i)(c_0))) \\
&= g(\overline{s}(i(c_0))) && \text{since } c_0 \in C, \\
&= g(i(c_0)) && \text{since } \overline{s} \text{ is a co-retraction and } i(c_0) \in C' \subseteq \mathrm{Dom}(\overline{s}), \\
&= (\overline{r} \oplus_X i^{-1})(i(c_0)) \\
&= i^{-1}(i(c_0)) && \text{since } i(c_0) \in C', \\
&= c_0.
\end{aligned}$$

□

Remark 4.10. Every structure $\mathbf{A}$ has a co-retract which is simultaneously n-extendable for every $n \in \mathbb{N}$, for example, $\bigcup_{\ell \in \mathbb{N}} \Xi_\ell(\Xi_{\ell-1}(\dots \Xi_2(\Xi_1(\mathbf{A}))\dots))$.

We are now ready for the main theorem of this section. The result is slightly more general than the equirank h.p.t., which follows straightaway as a corollary.

THEOREM 4.11. *Suppose $\mathcal{P}$ and $\mathcal{Q}$ are classes of structures and Φ is a first-order sentence such that for all structures $\mathbf{A}$ and $\mathbf{B}$,*

- *if $\mathbf{A} \in \mathcal{P}$ and $\mathbf{A} \rightarrow \mathbf{B}$ then $\mathbf{B} \models \Phi$, and*

- *if* $\mathbf{A} \models \Phi$ *and* $\mathbf{A} \to \mathbf{B}$ *then* $\mathbf{B} \in \mathcal{Q}$.

Then there exists an existential-positive sentence Ψ *such that* $\text{qrank}(\Psi) \leq \text{qrank}(\Phi)$ *and* $\mathcal{P} \subseteq \text{Mod}(\Psi) \subseteq \mathcal{Q}$.

Proof. Let $n = \text{qrank}(\Phi)$ and suppose $\mathbf{A}$ and $\mathbf{B}$ are structures and such that $\mathbf{A} \in \mathcal{P}$ and $\mathbf{A} \to^n \mathbf{B}$. We will prove that $\mathbf{B} \in \mathcal{Q}$. The conclusion that $\mathcal{P} \subseteq \text{Mod}(\Psi) \subseteq \mathcal{Q}$ for some existential-positive sentence Ψ of quantifier-rank at most n then follows by Lemma 3.16.

Let $\mathbf{C} = \mathbf{Core}^n(\mathbf{A})$ and note that $\mathbf{A} \rightleftarrows^n \mathbf{C}$ and $\mathbf{C} \to \mathbf{B}$. Structures $\Xi^\star_n(\mathbf{A})$ and $\Xi^\star_n(\mathbf{C})$ are both n-extendable by Lemma 4.9. Since $\Xi^\star_n(\mathbf{A}) \xrightarrow{\text{retr}} \mathbf{A}$ and $\Xi^\star_n(\mathbf{C}) \xrightarrow{\text{retr}} \mathbf{C}$ (by Lemma 4.8), we have $\mathbf{A} \rightleftarrows^n \Xi^\star_n(\mathbf{A})$ and $\mathbf{C} \rightleftarrows^n \Xi^\star_n(\mathbf{C})$. Therefore, $\Xi^\star_n(\mathbf{A}) \rightleftarrows^n \Xi^\star_n(\mathbf{C})$ by transitivity of $\rightleftarrows^n$. Lemma 4.3 now yields $\Xi^\star_n(\mathbf{A}) \equiv^n \Xi^\star_n(\mathbf{C})$. The full picture is

$$\begin{array}{ccccc} \Xi^\star_n(\mathbf{A}) & \equiv^n & \Xi^\star_n(\mathbf{C}) & & \\ \big\downarrow{\scriptstyle \text{retr}} & & \big\downarrow{\scriptstyle \text{retr}} & & \\ \mathbf{A} & \rightleftarrows^n & \mathbf{C} & \to & \mathbf{B}. \end{array}$$

In particular, we have $\mathbf{A} \to \Xi^\star_n(\mathbf{A}) \equiv^n \Xi^\star(\mathbf{C}) \to \mathbf{B}$.

We now argue that $\mathbf{B} \in \mathcal{Q}$. Since $\mathbf{A} \in \mathcal{P}$ and $\mathbf{A} \to \Xi^\star_n(\mathbf{A})$, it follows that $\Xi^\star_n(\mathbf{A}) \models \Phi$ because of the hypothesis concerning $\mathcal{P}$ and Φ. Since $\Xi^\star_n(\mathbf{A}) \equiv^n \Xi^\star_n(\mathbf{C})$, we have $\Xi^\star_n(\mathbf{C}) \models \Phi$. Since $\Xi^\star_n(\mathbf{C}) \to \mathbf{B}$, we conclude that $\mathcal{B} \in \mathcal{Q}$ by the hypothesis concerning $\mathcal{Q}$ and Φ. $\square$

Theorem 4.12 (Equirank Homomorphism Preservation Theorem). *A first-order sentence is preserved under homomorphisms on all structures if, and only if, it is equivalent to an existential-positive sentence of equal quantifier-rank.*

Proof. The easy ($\Longleftarrow$) direction was given earlier as Lemma 2.3. The ($\Longrightarrow$) direction is nothing but the special case of Theorem 4.11 where $\mathcal{P} = \mathcal{Q} = \text{Mod}(\Phi)$ for a first-order sentence Φ. $\square$

We once again emphasize that $\Xi^\star_n(\mathbf{A})$ is an infinite structure, even if $\mathbf{A}$ is finite. The proof of Theorem 4.12 thus uses, in an essential way, the fact that Φ is preserved under homomorphisms on all structures, not just on finite structures. For this reason, Theorem 4.12 does automatically yield a valid result when restricted to finite structures. Our interpolation and preservation theorems on finite structures (Theorems 5.15 and 5.16) do not include this strong "equirank" condition. However, we believe that Theorem 4.11 and 4.12 are likely true on finite structures. We advance a conjecture to this effect in §7.

5. FINITE HOMOMORPHISM PRESERVATION THEOREM

Our proof of the finite h.p.t. (Theorem 5.16) follows the same basic plan as our proof of the equirank h.p.t. in the previous section.

Definition 5.1. For every structure $\mathbf{A}$ and parameters $r, s, t, u \in \mathbb{N}$, we define a new structure $\Delta^{r,s,t,u}(\mathbf{A})$ by

$$\Delta^{r,s,t,u}(\mathbf{A}) = \bigoplus_{\substack{S\subseteq A \\ |S|\leq s \\ \mathbf{C}\in\mathscr{E}^r_{S,\leq t}(\mathbf{A})}} {}_A \left(\mathbf{A} \oplus_S \underbrace{\mathbf{C} \oplus_S \cdots \oplus_S \mathbf{C}}_{u \text{ times}}\right).$$

In this expression, S ranges over subsets of A of size $\leq s$, and $\mathbf{C}$ ranges over $(\leq t)$-extension r-cores of $\mathbf{A}$ over S. For each index pair $(S, \mathbf{C})$, u disjoint copies of $\mathbf{C}$ are glued onto $\mathbf{A}$ via the operation $\oplus_S$.

Let $\Delta^{r,s,t,u}(A)$ denote the universe of $\Delta^{r,s,t,u}(\mathbf{A})$ (by mild abuse of notation). For each element $\alpha \in \Delta^{r,s,t,u}(A) \setminus A$, there is a unique index pair $(S_\alpha, \mathbf{C}_\alpha)$ and $i_\alpha \in \{1, \ldots, u\}$ such that α belongs to the i_αth copy of $\mathbf{C}_\alpha$.

—S_α is called the *support* of α and denoted by $Supp(\alpha)$.

—The i_αth copy of $\mathbf{C}_\alpha$, viewed as a substructure of $\Delta^{r,s,t,u}(\mathbf{A})$, is called the *clan* of α and denoted by $\mathbf{Clan}(\alpha)$. The universe of $\mathbf{Clan}(\alpha)$ is denoted by $Clan(\alpha)$.

Note that $Supp(\alpha) = Clan(\alpha) \cap A$. We remark that $Supp(a)$ and $Clan(a)$ are undefined for elements $a \in A$.

The structural operator $\Delta^{r,s,t,u}(\cdot)$ has similar properties to $\Xi_n(\cdot)$. Lemma 5.2, below, corresponds to Lemmas 4.7 and 4.8 about $\Xi_n(\cdot)$. We omit the proof of Lemma 5.2, as the arguments are virtually identical.

LEMMA 5.2.

(1) *If* $\mathbf{A}$ *is finite, then so is* $\Delta^{r,s,t,u}(\mathbf{A})$.

(2) $\Delta^{r,s,t,u}(\mathbf{A}) \xrightarrow{\text{retr}} \mathbf{A} \oplus_{Supp(\alpha)} \mathbf{Clan}(\alpha) \xrightarrow{\text{retr}} \mathbf{A}$ *for all* $\alpha \in \Delta^{r,s,t,u}(A) \setminus A$. □

From now on, let n be a fixed positive integer. We define integers $r(\ell), s(\ell), t(\ell), u(\ell)$ for $\ell = 0, \ldots, n-1$, as well as $r(n)$ and $t(n)$, by the following inductive scheme:

$$r(0) = 0 \qquad t(0) = 1$$
$$r(\ell+1) = r(\ell) + t(\ell) \qquad t(\ell+1) = t(\ell) \cdot s(\ell)$$
$$s(\ell) = \beta^{r(\ell)}_{u(\ell)} \qquad u(\ell) = (n-\ell) \cdot t(\ell).$$

This induction easily seen to be well-founded.[6]

Remark 5.3. This definition of sequences r, s, t, u is precisely what our main technical lemma (Lemma 5.12) requires. However, any other choice of sequences

[6]The first few values of r, s, t, u are given by:

$r(0) = 0$	$r(1) = 1$	$r(2) = n+1$
$s(0) = \beta^0_n = n$	$s(1) = \beta^1_{n(n-1)}$	$s(2) = \beta^{n+1}_{n(n-2)\beta^1_{n(n-1)}}$
$t(0) = 1$	$t(1) = n$	$t(2) = n\beta^1_{n(n-1)}$
$u(0) = n$	$u(1) = n(n-1)$	$u(2) = n(n-2)\beta^1_{n(n-1)}$

r, s, t, u is acceptable, so long as all of the above equations hold after each equality $=$ is replaced by an inequality $\geq$. The reader may therefore disregard the equations defining r, s, t, u (at least for the time being) and simply regard

$$r(0), t(0), u(0), s(0), r(1), t(1), u(1), s(1), \dots \\ \dots, r(n{-}1), t(n{-}1), u(n{-}1), s(n{-}1), r(n), t(n)$$

as a sufficiently fast increasing sequence, since any fast enough increasing sequence satisfies the system of inequalities with $\geq$ replacing $=$, above.

Definition 5.4. For every structure $\mathbf{A}$ and $\ell \in \{0, \dots, n{-}1\}$, let

$$\Delta_\ell(\mathbf{A}) = \Delta^{r(\ell),s(\ell),t(\ell),u(\ell)}(\mathbf{A}) = \bigoplus_{\substack{S \subseteq A \\ |S| \leq s(\ell) \\ \mathbf{C} \in \mathscr{E}^{r(\ell)}_{S, \leq t(\ell)}(\mathbf{A})}} {}_A \left(\mathbf{A} \oplus_S \underbrace{\mathbf{C} \oplus_S \cdots \oplus_S \mathbf{C}}_{u(\ell) \text{ times}} \right),$$

$$\Gamma_\ell(\mathbf{A}) = \Delta_\ell(\Delta_{\ell+1}(\dots \Delta_{n-2}(\Delta_{n-1}(\mathbf{A})) \dots)).$$

In addition, let $\Gamma_n(\mathbf{A}) = \mathbf{A}$ and $\widetilde{\mathbf{A}} = \Gamma_0(\mathbf{A})$, and let $\widetilde{A}$ (resp. $\Gamma_\ell(A)$) denote the universe of $\widetilde{\mathbf{A}}$ (resp. $\Gamma_\ell(\mathbf{A})$). Note that $\Gamma_\ell(\mathbf{A}) = \Delta_\ell(\Gamma_{\ell+1}(\mathbf{A}))$ for all $\ell \in \{0, \dots, n{-}1\}$.

Here, roughly, is the picture of $\widetilde{\mathbf{A}}$ for the reader to keep in mind. $\widetilde{\mathbf{A}}$ is built up around $\mathbf{A}$ in what we will call "levels", similar to the layers of an onion. Levels are labeled $n, n{-}1, \dots, 0$ in decreasing order from the inside out. The nth and innermost level of $\widetilde{\mathbf{A}}$ is the original structure $\mathbf{A}$. For $\ell < n$, the ℓth level of $\widetilde{\mathbf{A}}$ consists of a collection of extension cores ("clans"), each glued onto a set of elements in higher levels (a "support"). $\widetilde{\mathbf{A}}$ is thus constructed from the inside out. Moving outward from higher to lower levels, the arity of clans and the size of supports rapidly decreases. An element of level 0 has a far smaller clan and support than a typical element of level $n{-}1$. The central idea behind this construction is to ensure that for every $\ell \in \{0, \dots, n\}$ and every $s(\ell)$-tuple $\vec{\alpha}$ of elements of level $\geq \ell$, the $\equiv^\ell$-class of $(\widetilde{\mathbf{A}}, \vec{\alpha})$ (i.e., the first-order type of $\vec{\alpha}$ in $\widetilde{\mathbf{A}}$ up to quantifier-rank ℓ) is an invariant of the $\rightleftarrows^{r(\ell)}$-class of $(\widetilde{\mathbf{A}}, \vec{\alpha})$ (i.e., the "existential-positive type" of $\vec{\alpha}$ in $\widetilde{\mathbf{A}}$ up to quantifier-rank $r(\ell)$).

We will now precisely define the terms *level*, *support* and *clan* in the structure $\widetilde{\mathbf{A}}$. (Earlier we defined *support* and *clan* in structures $\Delta^{r,s,t,u}(\mathbf{A})$.) These notions give us a convenient way of indexing elements of $\widetilde{\mathbf{A}}$.

Definition 5.5 *(Level, Support and Clan).* The *level* $\lambda(\alpha)$ of an element $\alpha \in \widetilde{A}$ is defined by

$$\lambda(\alpha) = \max\{\ell : \alpha \in \Gamma_\ell(A)\}.$$

Thus, $\Gamma_\ell(A) = \{\alpha \in \widetilde{A} : \lambda(\alpha) \geq \ell\}$. For every $\alpha \in \widetilde{A} \setminus A$, the level $\lambda(\alpha)$ is the unique $\ell \in \{0, \dots, n{-}1\}$ such that $\alpha \in \Gamma_\ell(A) \setminus \Gamma_{\ell+1}(A)$.

For an element $\alpha \in \widetilde{A} \setminus A$ of level ℓ, let *Supp*(α) and $\mathbf{Clan}(\alpha)$ denote the support and clan of α as defined in Definition 5.1 for α considered as an element of

$\Delta_\ell(\Gamma_{\ell+1}(A)) \setminus \Gamma_{\ell+1}(A)$ in the structure $\Delta_\ell(\Gamma_{\ell+1}(\mathbf{A}))$. We remark that $\mathit{Supp}(a)$ and $\mathbf{Clan}(a)$ are undefined for elements $a \in A$.

Lemma 5.2 has an immediate corollary.

COROLLARY 5.6.

(1) *If* $\mathbf{A}$ *is finite, then so is* $\widetilde{\mathbf{A}}$.

(2) $\widetilde{\mathbf{A}} \xrightarrow{\text{retr}} \Gamma_{\lambda(\alpha)}(\mathbf{A}) \xrightarrow{\text{retr}} \Gamma_{\lambda(\alpha)+1}(\mathbf{A}) \oplus_{\mathit{Supp}(\alpha)} \mathbf{Clan}(\alpha) \xrightarrow{\text{retr}} \Gamma_{\lambda(\alpha)+1}(\mathbf{A})$ *for all* $\alpha \in \widetilde{A} \setminus A$. □

It is helpful to think of $\widetilde{A}$ as an acyclic directed graph, in which there is an arc from α to α' (denoted by $\alpha \dashrightarrow \alpha'$) for all $\alpha \in \widetilde{A} \setminus A$ and $\alpha' \in \mathit{Supp}(\alpha)$. This relation is clearly acyclic, since $\lambda(\alpha) < \lambda(\alpha')$ for every arc $\alpha \dashrightarrow \alpha'$. Note that the out-degree of an element $\alpha \in \widetilde{A} \setminus A$ is bounded by $s(\lambda(\alpha))$. Since $s(\ell)$ is a (rapidly) increasing function, this means that elements of higher (inner) levels potentially have much greater out-degree than elements of lower (outer) levels.

Definition 5.7 *(ℓ-Closure and ℓ-Frontier).* Let $X \subseteq \widetilde{A}$ and $\ell \in \{0, \ldots, n\}$.

—We say that X is *ℓ-closed* if $\mathit{Clan}(\alpha) \subseteq X$ for all $\alpha \in X$ such that $\lambda(\alpha) < \ell$.

—The *ℓ-closure* of X, denoted by $\mathrm{cl}_\ell(X)$, is the unique minimal ℓ-closed set containing X.

—The *ℓ-frontier* of X, denoted by ${}^\ell X$, is the set $\mathrm{cl}_\ell(X) \cap \Gamma_\ell(A)$.

Note that the intersection of ℓ-closed sets is ℓ-closed (so the ℓ-closure of a set X is well-defined).

Remark 5.8.

(1) It is easy to see that $\mathrm{cl}_\ell(X) = \bigcup_{\alpha \in X} \mathrm{cl}_\ell(\{\alpha\})$ for every $X \subseteq \widetilde{A}$ and $\ell \in \{0, \ldots, n\}$. It follows that ${}^\ell X = \bigcup_{\alpha \in X} {}^\ell\{\alpha\}$.

(2) If X is ℓ-closed, then it contains the endpoint α' of every arc $\alpha \dashrightarrow \alpha'$ such that $\alpha \in X$ and $\lambda(\alpha) < \ell$. (Indeed, we have $X \supseteq \mathit{Clan}(\alpha) \supseteq \mathit{Supp}(\alpha) \ni \alpha'$.)

(3) The ℓ-frontier of a set X consists of precisely the elements of $\widetilde{A}$ that one can reach by starting at an element in X and following arcs of $\dashrightarrow$ until crossing into $\Gamma_\ell(A)$ (i.e., until reaching an element in $\widetilde{A}$ of level $\geq \ell$). Recall that the level of elements increases along arcs of $\dashrightarrow$, so that eventually any path reaches the set $\Gamma_\ell(A)$. (Having explained the ℓ-frontier, we will no longer speak about the acyclic arc relation $\dashrightarrow$.)

The next lemma bounds the size of the ℓ-frontier of a set X in terms of $|X|$.

LEMMA 5.9. $|{}^\ell X| \leq |X| \cdot t(\ell)$ *for all* $X \subseteq \widetilde{A}$ *and* $\ell \in \{0, \ldots, n\}$.

PROOF. We argue by induction on ℓ. In the base case $\ell = 0$, the claim is obvious as ${}^0 X = X$ and $t(0) = 1$. For the induction step, let $\ell \geq 1$ and assume $|{}^{\ell-1}X| \leq |X| \cdot t(\ell - 1)$. Notice that

$$
{}^\ell X = \bigcup_{\alpha \in {}^{\ell-1}X} \begin{cases} \mathit{Supp}(\alpha) & \text{if } \lambda(\alpha) = \ell - 1, \\ \{\alpha\} & \text{if } \lambda(\alpha) \geq \ell. \end{cases}
$$

We see that each $\alpha \in {}^{\ell-1}X$ "contributes" to ${}^{\ell}X$ either one element (itself) or else $|Supp(\alpha)|$ many elements. Since $|Supp(\alpha)| \le s(\ell-1)$ for all α of level $\ell - 1$,

$$|{}^{\ell}X| \le |{}^{\ell-1}X| \cdot (\text{maximal "contribution" of any } \alpha \in {}^{\ell-1}X) \le |X| \cdot t(\ell-1) \cdot s(\ell-1) = |X| \cdot t(\ell).$$

Therefore, the lemma holds by induction. □

So far we have considered only a single structure $\mathbf{A}$ and its corresponding $\widetilde{\mathbf{A}}$. We will now focus on the relationship between $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$ given two structures $\mathbf{A}$ and $\mathbf{B}$.

Definition 5.10 *(ℓ-Kosher).* For $\ell \in \{0, \dots, n\}$, a partial isomorphism f from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$ is *ℓ-kosher* if it satisfies conditions (K1)–(K4), below.

- ***K1.*** $\widetilde{\mathbf{A}} \rightleftarrows_f^{r(\ell)} \widetilde{\mathbf{B}}$. (See Notation 3.3 regarding the meaning of $\rightleftarrows_f^{r(\ell)}$.)
- ***K2.*** $\mathrm{Dom}(f)$ and $\mathrm{Range}(f)$ are ℓ-closed sets in $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$, respectively.
- ***K3.*** $\min\{\ell, \lambda(\alpha)\} = \min\{\ell, \lambda(f(\alpha))\}$ for all $\alpha \in \mathrm{Dom}(f)$.
- ***K4.*** $f(Clan(\alpha)) = Clan(f(\alpha))$ for all $\alpha \in \mathrm{Dom}(f)$ such that $\lambda(\alpha) < \ell$.

To be clear: in (K3), $\lambda(\alpha)$ is the level of α in $\widetilde{\mathbf{A}}$ and $\lambda(f(\alpha))$ is the level of $f(\alpha)$ in $\widetilde{\mathbf{B}}$. Similarly, in (K4), $Clan(\alpha)$ is the clan of α in $\widetilde{\mathbf{A}}$ and $Clan(f(\alpha))$ is the clan of $f(\alpha)$ in $\widetilde{\mathbf{B}}$.

Note that ℓ-kosher is a stronger condition as ℓ increases (so n-kosher implies $n{-}1$-kosher, etc.). This is crucial for defining an n-back-and-forth system between structures $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$ (in Lemma 5.12 and Corollaries 5.13 and 5.14).[7] The next lemma states some obvious properties of ℓ-kosher partial isomorphisms.

LEMMA 5.11.

(*1*) *f is an ℓ-kosher partial isomorphism from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$ if, and only if, its inverse f^{-1} is an ℓ-kosher partial isomorphism from $\widetilde{\mathbf{B}}$ to $\widetilde{\mathbf{A}}$.*

(*2*) *The empty map (with domain $\emptyset$) is an n-kosher partial isomorphism from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$ if, and only if, $\mathbf{A} \rightleftarrows^{r(n)} \mathbf{B}$.*

(*3*) *If f is ℓ-kosher, then $f(\mathrm{cl}_\ell(X)) = \mathrm{cl}_\ell(f(X))$ for all $X \subseteq \mathrm{Dom}(f)$.*

PROOF. For statement (1), conditions (K1)–(K3) are evidently symmetric in f and f^{-1}; condition (K4) is symmetric in f and f^{-1} once one assumes (K3). For statement (2), simply notice that conditions (K2)–(K4) are trivial when f is the empty map. Statement (3) is easily deduced from definitions. □

Now comes our main technical lemma.

[7] From the standpoint of the n-round Ehrenfeucht-Fraïssé game, we begin with a pair of structures $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$ for which the empty partial isomorphism is n-kosher. Duplicator (the player attempting to show that $\widetilde{\mathbf{A}} \equiv^n \widetilde{\mathbf{B}}$) is able to maintain the condition that, after k rounds of the game, the k pairs of elements from $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$ (representing the plays in rounds $1, \dots, k$) extend to an $(n-k)$-kosher partial isomorphism.

LEMMA 5.12 (MAIN LEMMA). *Let $\ell \in \{0, \dots, n-1\}$ and suppose f is an ℓ+1-kosher partial isomorphism from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$ with domain $\mathrm{cl}_{\ell+1}(X)$ for some $X \subseteq \widetilde{A}$ of size $< n - \ell$. Then for every $\alpha \in \widetilde{A}$, there exists an ℓ-kosher partial isomorphism extending $f \restriction \mathrm{cl}_\ell(X)$ to domain $\mathrm{cl}_\ell(X \cup \{\alpha\})$.*

At a high level, the proof works as follows. Given an element $\alpha \in \widetilde{A}$, we consider the ℓ-frontier ${}^\ell(X \cup \{\alpha\})$, which we partition into disjoint subsets $U = {}^\ell X$ and $V = {}^\ell\{\alpha\} \setminus {}^\ell X$. (Note that f is defined on U but not on V.) We first construct an ℓ-kosher partial isomorphism h which extends $f \restriction U$ to the ℓ-frontier ${}^\ell(X \cup \{\alpha\})$. We then extend h an ℓ-kosher partial isomorphism on the entire ℓ-closure $\mathrm{cl}_\ell(X \cup \{\alpha\})$.

The proof is divided into eight steps, where Steps 1–7 cover the construction of h and Step 8 gives the extension of h to $\mathrm{cl}_\ell(X \cup \{\alpha\})$. We begin with the observation that sets U and V are free over ${}^{\ell+1}X$ in $\widetilde{A}$. Using the freeness lemma (Lemma 3.22), we show that there exists a subset $W \subseteq {}^{\ell+1}X$ of size $\leq s(\ell)$ such that U and V are $r(\ell)$-free over W in $\widetilde{\mathbf{A}}$. We then consider $\mathbf{Core}^{r(\ell)}_{W \cup V}(\widetilde{\mathbf{A}})$, the $r(\ell)$-core of $\widetilde{\mathbf{A}}$ over $W \cup V$. As $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell+1)}_{f \restriction W} \widetilde{\mathbf{B}} \xrightarrow{\text{retr}} \Gamma_{\ell+1}(\mathbf{B})$ (by ℓ+1-kosherness of f) and $r(\ell+1) = t(\ell) + s(\ell) \geq |V| + |W|$ (since $|V| \leq t(\ell)$ by Lemma 5.9), it follows that $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$ have isomorphic $|V|$-extension $r(\ell)$-cores over W and $f(W)$, respectively (by Lemma 3.25(4)). In particular, there exists an extension core $(\mathbf{C}, c_1, \dots, c_{|V|}) \in \mathscr{E}^{r(\ell)}_{f(W),|V|}(\Gamma_{\ell+1}(\mathbf{B}))$ such that $\mathbf{Core}^{r(\ell)}_{W \cup V}(\widetilde{\mathbf{A}}) \cong \mathbf{C}$ via an isomorphism which extends $f \restriction W$ and maps V onto the set $\{c_1, \dots, c_{|V|}\}$ (by Corollary 3.26). As $\widetilde{\mathbf{B}} \xrightarrow{\text{retr}} \Gamma_\ell(\mathbf{B}) \xrightarrow{\text{retr}} \Gamma_{\ell+1}(\mathbf{B}) \oplus \mathbf{C} \xrightarrow{\text{retr}} \Gamma_{\ell+1}(\mathbf{B})$ (by Lemma 5.6(2)), we obtain a function from V to $\widetilde{B}$ via the isomorphism $\mathbf{Core}^{r(\ell)}_{W \cup V}(\widetilde{\mathbf{A}}) \xrightarrow{\cong} \mathbf{C}$ where we view $\mathbf{C}$ as a substructure of $\Gamma_{\ell+1}(\mathbf{B}) \oplus \mathbf{C}$ (which in turn is a retract of $\widetilde{\mathbf{B}}$). This gives the map h on V. Taking $h \restriction U = f \restriction U$, we get $h : {}^\ell(X \cup \{\alpha\}) \longrightarrow \widetilde{B}$. Steps 1–7 define h more formally and show that it is an ℓ-kosher partial isomorphism from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$. Step 8 extends h to the entire ℓ-closure $\mathrm{cl}_\ell(X \cup \{\alpha\})$ in a manner essentially governed by ℓ-kosher conditions (K2)–(K4).

PROOF. Fix $\alpha \in \widetilde{A}$ (we can assume $\alpha \notin \mathrm{cl}_\ell(X)$, since otherwise the lemma is trivial) and let $U = {}^\ell X$ and $V = {}^\ell\{\alpha\} \setminus {}^\ell X$. Note that $U \cup V = {}^\ell(X \cup \{\alpha\})$ (see Remark 5.8(1)). We proceed in eight steps.

Step 1. We claim that U and V are free over ${}^{\ell+1}X$ in $\widetilde{\mathbf{A}}$. It follows from definitions that U and V are separated over ${}^{\ell+1}X$ in $\Gamma_\ell(\mathbf{A})$. Therefore, U and V are free over ${}^{\ell+1}X$ in $\Gamma_\ell(\mathbf{A})$ by Lemma 3.20(1). Since $\widetilde{\mathbf{A}} \xrightarrow{\text{retr}} \Gamma_\ell(\mathbf{A})$, Lemma 3.20(4) implies that U and V are free over ${}^{\ell+1}X$ in $\widetilde{\mathbf{A}}$.

Step 2. Lemma 5.9 yields $|U \cup V| = |{}^\ell(X \cup \{\alpha\})| \leq (n-\ell) \cdot t(\ell)$. By the freeness lemma (Lemma 3.22), there exists $W \subseteq {}^{\ell+1}X$ such that $|W| \leq \beta^{r(\ell)}_{(n-\ell)\cdot t(\ell)} = s(\ell)$ and $U \perp\!\!\!\perp^{r(\ell)}_W V$ in $\widetilde{\mathbf{A}}$. Fix any such set W.

We claim that $W \subseteq \mathrm{Dom}(f)$ and $f(W) \subseteq \Gamma_{\ell+1}(B)$. For the first claim: we have $W \subseteq {}^{\ell+1}X \subseteq \mathrm{cl}_{\ell+1}(X) = \mathrm{Dom}(f)$. For the second claim: note that $W \subseteq {}^{\ell+1}X \subseteq \Gamma_{\ell+1}(A)$ and hence $f(W) \subseteq \Gamma_{\ell+1}(B)$ by ℓ+1-kosher condition (K3).

Step 3. We claim that $\mathbf{Core}^{r(\ell)}_{W\cup V}(\widetilde{\mathbf{A}}) \to_{f\restriction W} \Gamma_{\ell+1}(\mathbf{B})$, that is, there exists a homomorphism from $\mathbf{Core}^{r(\ell)}_{W\cup V}(\widetilde{\mathbf{A}})$ to $\Gamma_{\ell+1}(\mathbf{B})$ which extends $f \restriction W$ (note that $W \subseteq {}^{\ell+1}X \subseteq \mathrm{Dom}(f)$). Such a homomorphism is formed by composition as follows:

$$\begin{array}{cl}
\mathbf{Core}^{r(\ell)}_{W\cup V}(\widetilde{\mathbf{A}}) & \\
\downarrow_W & \text{by Lemma 3.12,} \\
\mathbf{Core}^{r(\ell)+|V|}_{W}(\widetilde{\mathbf{A}}) & \\
\downarrow_W & \text{since } r(\ell)+|V| = r(\ell)+|^{\ell}\{\alpha\}| \le r(\ell)+t(\ell) = r(\ell+1), \\
\mathbf{Core}^{r(\ell+1)}_{W}(\widetilde{\mathbf{A}}) & \\
\cong_{f\restriction W} & \text{since } \widetilde{\mathbf{A}} \rightleftarrows^{r(\ell+1)}_{f} \widetilde{\mathbf{B}} \text{ by } \ell{+}1\text{-kosher condition (K1),} \\
\mathbf{Core}^{r(\ell+1)}_{f(W)}(\widetilde{\mathbf{B}}) & \\
\cong_{f(W)} & \text{since } \widetilde{\mathbf{B}} \xrightarrow{\text{retr}} \Gamma_{\ell+1}(\mathbf{B}) \text{ and hence } \widetilde{\mathbf{B}} \rightleftarrows^{r(\ell+1)}_{f(W)} \Gamma_{\ell+1}(\mathbf{B}), \\
\mathbf{Core}^{r(\ell+1)}_{f(W)}(\Gamma_{\ell+1}(\mathbf{B})) & \\
\downarrow_{f(W)} & \\
\Gamma_{\ell+1}(\mathbf{B}). &
\end{array}$$

Step 4. Let $m = |V \setminus W|$ and fix an arbitrary enumeration $v_1, \dots, v_m$ of $V \setminus W$. It is obvious that there exists a (unique) canonical core $\mathbf{C}' \in \mathscr{C}^{r(\ell)}_{f(W)\sqcup\{c_1,\dots,c_m\}}$ such that

$$(\mathbf{Core}^{r(\ell)}_{W\cup V}(\widetilde{\mathbf{A}}), v_1, \dots, v_m) \cong_{f\restriction W} (\mathbf{C}', c_1, \dots, c_m).$$

That is, $\mathbf{Core}^{r(\ell)}_{W\cup V}(\widetilde{\mathbf{A}})$ is isomorphic to $\mathbf{C}'$ via a bijection which extends partial isomorphism $f \restriction W$ and maps v_i to c_i for all $i \in \{1, \dots, m\}$.

We claim that $\mathbf{C}'$ is an extension core in the set $\mathscr{E}^{r(\ell)}_{f(W),m}(\Gamma_{\ell+1}(\mathbf{B}))$. By Step 3 and our choice of $\mathbf{C}'$, we have

$$\mathbf{C}' \cong_{f^{-1}\restriction f(W)} \mathbf{Core}^{r(\ell)}_{W\cup V}(\widetilde{\mathbf{A}}) \to_{f\restriction W} \Gamma_{\ell+1}(\mathbf{B}).$$

By composition of the lefthand isomorphism and the righthand homomorphism, we have $\mathbf{C}' \to_{f(W)} \Gamma_{\ell+1}(\mathbf{B})$. Therefore, $\mathbf{C}' \in \mathscr{E}^{r(\ell)}_{f(W),m}(\Gamma_{\ell+1}(\mathbf{B}))$ as claimed.

Note that $|f(W)| = |W| \le s(\ell)$ and $m \le |V| = |^{\ell}\{\alpha\}| \le t(\ell)$. Thus, we see that $(f(W), \mathbf{C}')$ shows up among index pairs $(S, \mathbf{C})$ in the following expression defining $\Gamma_\ell(\mathbf{B})$:

$$\Gamma_\ell(\mathbf{B}) = \Delta_\ell(\Gamma_{\ell+1}(\mathbf{B})) = \bigoplus_{\substack{S\subseteq\Gamma_{\ell+1}(B),\\ |S|\le s(\ell)\\ \mathbf{C}\in\mathscr{E}^{r(\ell)}_{S,\le t(\ell)}(\Gamma_{\ell+1}(\mathbf{B}))}} {}_{\Gamma_{\ell+1}(B)} \left(\Gamma_{\ell+1}(\mathbf{B}) \oplus_S \underbrace{\mathbf{C} \oplus_S \cdots \oplus_S \mathbf{C}}_{u(\ell)\text{ times}}\right).$$

Let $\mathbf{C}_1, \dots, \mathbf{C}_{u(\ell)}$ be the distinct clans corresponding to the index pair $(f(W), \mathbf{C}')$.

It holds that $\mathbf{C}' \cong_{f(W)} \mathbf{C}_1 \cong_{f(W)} \cdots \cong_{f(W)} \mathbf{C}_{u(\ell)}$ and

$$\Gamma_\ell(\mathbf{B}) \xrightarrow{\text{retr}} \Gamma_{\ell+1}(\mathbf{B}) \oplus_{f(W)} \mathbf{C}_1 \oplus_{f(W)} \cdots \oplus_{f(W)} \mathbf{C}_{u(\ell)} \xrightarrow{\text{retr}} \Gamma_{\ell+1}(\mathbf{B}).$$

Step 5. The set ${}^\ell X \setminus {}^{\ell+1}X$ (consisting of elements of $\mathrm{cl}_\ell(X)$ of level exactly ℓ) has at most $|{}^\ell X| \leq |X| \cdot t(\ell) < (n-\ell)\cdot t(\ell) = u(\ell)$ elements by Lemma 5.9. The set

$$\{\mathbf{C}_1, \ldots, \mathbf{C}_{u(\ell)}\} \setminus \{\mathbf{Clan}(f(x)) : x \in {}^\ell X \setminus {}^{\ell+1}X\}$$

is therefore nonempty. We now fix a choice of $\mathbf{C}$ in this set, as well as an isomorphism

$$g : (\mathbf{Core}^{r(\ell)}_{W\cup V}(\widetilde{\mathbf{A}}), v_1, \ldots, v_m) \xrightarrow{\cong}_{f\restriction W} (\mathbf{C}, c_1, \ldots, c_m).$$

As for the isomorphism g, we will be interested in the image $g(V)$ of the set V.

Step 6. Recall that $U \perp\!\!\!\perp^{r(\ell)}_W V$ in $\widetilde{\mathbf{A}}$ by Step 2. We now prove that $f(U) \perp\!\!\!\perp^{r(\ell)}_{f(W)} g(V)$ in $\widetilde{\mathbf{B}}$. We claim that $f(U)$ and $g(V)$ are free (not just $r(\ell)$-free) over $f(W)$ in $\widetilde{\mathbf{B}}$. To see this, first notice that sets $f(U)$ and C are separated over $f(W)$ in $\Gamma_\ell(\mathbf{B})$. Since $g(V) \subseteq C$, it follows that $f(U)$ and $g(V)$ are separated (and therefore free) over $f(W)$ in $\Gamma_\ell(\mathbf{B})$. Finally, since $\widetilde{\mathbf{B}} \xrightarrow{\text{retr}} \Gamma_\ell(\mathbf{B})$, Lemma 3.20(4) yields $f(U) \perp\!\!\!\perp_{f(W)} g(V)$ in $\widetilde{\mathbf{B}}$.

Step 7. Let function $h : U \cup V \cup W \longrightarrow \widetilde{B}$ be defined by

$$h(y) = \begin{cases} f(y) & \text{if } y \in U \cup W, \\ g(y) & \text{if } y \in V. \end{cases}$$

We advance four claims:

- ***i.*** U and V are $r(\ell)$-free over W in $\widetilde{\mathbf{A}}$,
- ***ii.*** $h(U)$ and $h(V)$ are $r(\ell)$-free over $h(W)$ in $\widetilde{\mathbf{B}}$,
- ***iii.*** $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{h\restriction(W\cup U)} \widetilde{\mathbf{B}}$,
- ***iv.*** $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{h\restriction(W\cup V)} \widetilde{\mathbf{B}}$.

Once we prove claims (i)–(iv), it follows that $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{h\restriction(U\cup V)} \widetilde{\mathbf{B}}$. Indeed, by the definition of freeness, we have

$$\left.\begin{aligned} U \perp\!\!\!\perp^{r(\ell)}_W V \text{ in } \widetilde{\mathbf{A}}, \\ \widetilde{\mathbf{A}} \to^{r(\ell)}_{h\restriction(W\cup U)} \widetilde{\mathbf{B}}, \\ \widetilde{\mathbf{A}} \to^{r(\ell)}_{h\restriction(W\cup V)} \widetilde{\mathbf{B}} \end{aligned}\right\} \Longrightarrow \widetilde{\mathbf{A}} \to^{r(\ell)}_{h\restriction(U\cup V)} \widetilde{\mathbf{B}},$$

$$\left.\begin{aligned} h(U) \perp\!\!\!\perp^{r(\ell)}_{h(W)} h(V) \text{ in } \widetilde{\mathbf{B}}, \\ \widetilde{\mathbf{B}} \to^{r(\ell)}_{h^{-1}\restriction h(W\cup U)} \widetilde{\mathbf{A}}, \\ \widetilde{\mathbf{B}} \to^{r(\ell)}_{h^{-1}\restriction h(W\cup V)} \widetilde{\mathbf{A}} \end{aligned}\right\} \Longrightarrow \widetilde{\mathbf{B}} \to^{r(\ell)}_{h^{-1}\restriction h(U\cup V)} \widetilde{\mathbf{A}}.$$

Claims (i) and (ii) were proved in Steps 2 and 6, respectively. As for claim (iii), ℓ+1-kosherness of f yields $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell+1)}_f \widetilde{\mathbf{B}}$. Since $r(\ell) \leq r(\ell+1)$ and $(W \cup U) \subseteq \mathrm{Dom}(f)$, it follows that $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{f\restriction(W\cup U)} \widetilde{\mathbf{B}}$. This proves claim (iii), since $h \restriction (W \cup U) = f \restriction (W \cup U)$.

Finally, we prove claim (iv). First, we have $\mathbf{C} \rightleftarrows^{r(\ell)}_{f(W)} \Gamma_{\ell+1}(\mathbf{B})$, derived as follows:

$$\mathbf{C} \rightleftarrows^{r(\ell)}_{f^{-1}\restriction W} \widetilde{\mathbf{A}} \rightleftarrows^{r(\ell+1)}_{f} \widetilde{\mathbf{B}} \xrightarrow{\text{retr}} \Gamma_{\ell+1}(\mathbf{B}).$$

Therefore, $\mathbf{C} \rightleftarrows^{r(\ell)}_{C} \Gamma_{\ell+1}(\mathbf{B}) \oplus_{f(W)} \mathbf{C}$ by Lemma 3.4(3). In particular,

$$\mathbf{C} \rightleftarrows^{r(\ell)}_{f(W)\cup g(V)} \big(\Gamma_{\ell+1}(\mathbf{B}) \oplus_{f(W)} \mathbf{C}\big)$$

as $f(W) \cup g(V) \subseteq C$. Putting pieces together, we have

$$\begin{aligned}
\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{W\cup V} \mathbf{Core}^{r(\ell)}_{W\cup V}(\widetilde{\mathbf{A}}) \cong_{(f\restriction W)\cup(g\restriction V)} \mathbf{C} & \qquad \text{by Step 5,} \\
\upharpoonleft\downharpoonright^{r(\ell)}_{f(W)\cup g(V)} & \qquad \text{by previous equation,} \\
\widetilde{\mathbf{B}} \xrightarrow{\text{retr}} \big(\Gamma_{\ell+1}(\mathbf{B}) \oplus_{f(W)} \mathbf{C}\big) & \qquad \text{by Corollary 5.6(2).}
\end{aligned}$$

It follows that $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{(f\restriction W)\cup(g\restriction V)} \widetilde{\mathbf{B}}$, which proves claim (iv) as $h \restriction (W \cup V) = (f \restriction W) \cup (g \restriction V)$.

Step 8. This is the last step in the proof. While the notation gets rather intensive, the basic arguments have been developed already in previous steps. The reader who has followed the action so far should find things fairly straightforward. The plan is roughly as follows. Having defined the function $h : U \cup V \cup W \longrightarrow \widetilde{B}$, we now construct an ℓ-kosher partial isomorphism $h_0 : \mathrm{cl}_\ell(X \cup \{\alpha\}) \longrightarrow \widetilde{B}$ which extends both $h \restriction U \cup V$ and $f \restriction \mathrm{cl}_\ell(X)$ (thus completing the proof). The function h_0 is built up in stages. We begin with the map $h \restriction U \cup V$, which we rename h_ℓ. Note that $\mathrm{Dom}(h_\ell) = U \cup V = \mathrm{cl}_\ell(X \cup \{\alpha\}) \cap \Gamma_\ell(A)$ and $\mathrm{Range}(h_\ell) \subseteq \Gamma_\ell(B)$. Using the results of previous steps, we easily show that h_ℓ is ℓ-kosher and agrees with f on their common subdomain U. We then extend maps h_ℓ and $f \restriction \mathrm{cl}_\ell(X) \cap \Gamma_{\ell-1}(A)$ to an ℓ-kosher partial isomorphism $h_{\ell-1}$ with $\mathrm{Dom}(h_{\ell-1}) = \mathrm{cl}_\ell(X \cup \{\alpha\}) \cap \Gamma_{\ell-1}(A)$ and $\mathrm{Range}(h_{\ell-1}) \subseteq \Gamma_{\ell-1}(B)$. There is a straightforward method for extending ℓ-kosher maps in this way; the ℓ-kosher condition almost dictates how $h_{\ell-1}$ must be defined. Continuing in the same manner, we extend maps $h_{\ell-1}$ and $f \restriction \mathrm{cl}_\ell(X) \cap \Gamma_{\ell-2}(X)$ to an ℓ-kosher partial isomorphism $h_{\ell-2}$ with $\mathrm{Dom}(h_{\ell-2}) = \mathrm{cl}_\ell(X \cup \{\alpha\}) \cap \Gamma_{\ell-2}(A)$ and $\mathrm{Range}(h_{\ell-2}) \subseteq \Gamma_{\ell-2}(B)$. This process continues until we at last obtain h_0.

We now give the formal argument. For $k \in \{0, \dots, \ell\}$, let

$$\begin{aligned}
Y_k &= \mathrm{cl}_\ell(X) \cap \Gamma_k(A) &&= \{y \in \mathrm{cl}_\ell(X) : \lambda(y) \geq k\}, \\
Z_k &= \big(\mathrm{cl}_\ell(\{\alpha\}) \setminus \mathrm{cl}_\ell(X)\big) \cap \Gamma_k(A) &&= \{z \in \mathrm{cl}_\ell(\{\alpha\}) \setminus \mathrm{cl}_\ell(X) : \lambda(z) \geq k\}.
\end{aligned}$$

Note that $Y_\ell = U$ and $Z_\ell = V$ and $Y_0 = \mathrm{cl}_\ell(X)$ and $Y_0 \cup Z_0 = \mathrm{cl}_\ell(X \cup \{\alpha\})$.

Proceeding inductively, we construct a sequence of partial isomorphisms h_ℓ, $h_{\ell-1}$, …, h_0 from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$ satisfying conditions (H1)–(H3) for all $k \in \{0, \dots, \ell\}$.

H1. h_k is ℓ-kosher.

H2. $\mathrm{Dom}(h_k) = Y_k \cup Z_k$ $(= \mathrm{cl}_\ell(X \cup \{\alpha\}) \cap \Gamma_k(A))$.

H3. If $k < \ell$, then h_k extends both h_{k+1} and $f \restriction Y_k$.

Once we construct h_0, we will have completed the proof of the lemma.

For the base case $k = \ell$, we let $h_\ell = h \restriction (U \cup V)$. We have $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{h_\ell} \widetilde{\mathbf{B}}$ by Step 7. Thus, h_ℓ satisfies ℓ-kosher condition (K1). It trivially satisfies conditions (K2)–(K4) since $\mathrm{Dom}(h_\ell) \subseteq \Gamma_\ell(A)$ and $\mathrm{Range}(h_\ell) \subseteq \Gamma_\ell(B)$. Therefore, h_ℓ is ℓ-kosher and thus satisfies (H1). It satisfies (H2) since its domain $U \cup V$ equals $Y_\ell \cup Z_\ell$, and it trivially satisfies (H3). This completes our argument in the base case $k = \ell$.

We now handle the induction step. Let $k \in \{0, \dots, \ell - 1\}$ and suppose we have constructed functions $h_\ell, h_{\ell-1}, \dots, h_{k+1}$ satisfying (H1)–(H3). We give a construction of h_k. Let $\mathbf{D}_1, \dots, \mathbf{D}_q$ enumerate (in any order) the distinct clans of elements of $\mathrm{cl}_\ell(\{\alpha\}) \setminus \mathrm{cl}_\ell(X)$ with level exactly k, and let $S_1, \dots, S_q$ be the corresponding supports. That is, let $\{\mathbf{D}_1, \dots, \mathbf{D}_q\} = \{\mathbf{Clan}(z) : z \in Z_k \setminus Z_{k+1}\}$ and let $S_i = D_i \cap \Gamma_{k+1}(A)$ for all $i \in \{1, \dots, q\}$.

The arguments over the next four paragraphs are carried out for each $i \in \{1, \dots, q\}$. We begin by claiming that $S_i \subseteq \mathrm{Dom}(h_{k+1})$. Indeed, S_i equals $\mathit{Supp}(z_i)$ for some $z_i \in Z_k \setminus Z_{k+1}$. Since $\lambda(z_i) = k$, every element of $\mathit{Supp}(z_i)$ has level $\geq k+1$ and thus belongs to the set $\{z' \in \mathrm{cl}_\ell(X \cup \{\alpha\}) : \lambda(z') \geq k+1\} = \mathrm{cl}_\ell(X \cup \{\alpha\}) \cap \Gamma_{k+1}(A)$. But this set is precisely $Y_{k+1} \cup Z_{k+1}$, which is $\mathrm{Dom}(h_{k+1})$ by (H2). Therefore, $S_i \subseteq \mathrm{Dom}(h_{k+1})$ as claimed.

Let $T_i = h_{k+1}(S_i)$. Partial isomorphism h_{k+1} is ℓ-kosher by (H1). Since all elements of S_i have level $\geq k+1$ in $\widetilde{\mathbf{A}}$, ℓ-kosher condition (K3) of h_{k+1} implies that all elements of T_i have level $\geq k+1$ in $\widetilde{\mathbf{B}}$, that is, $T_i \subseteq \Gamma_{k+1}(B)$. Also, note that $|T_i| \leq s(k)$ since $|T_i| = |S_i|$ and $|S_i| \leq s(k)$ as S_i is the support of an element of level k in $\widetilde{\mathbf{A}}$.

Next, we show that $\Gamma_{k+1}(\mathbf{A}) \rightleftarrows^{r(k)+t(k)}_{h_{k+1}\restriction S_i} \Gamma_{k+1}(\mathbf{B})$. By ℓ-kosher condition (K1) of h_{k+1}, we have $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{h_{k+1}} \widetilde{\mathbf{B}}$. It now follows:

$$
\begin{array}{ll}
\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{h_{k+1}\restriction S_i} \widetilde{\mathbf{B}} & \text{as we merely restrict partial isomorphism } h_{k+1}, \\
\Gamma_{k+1}(\mathbf{A}) \rightleftarrows^{r(\ell)}_{h_{k+1}\restriction S_i} \widetilde{\mathbf{B}} & \text{since } \widetilde{\mathbf{A}} \xrightarrow{\text{retr}} \Gamma_{k+1}(\mathbf{A}) \text{ and so } \widetilde{\mathbf{A}} \rightleftarrows_{S_i} \Gamma_{k+1}(\mathbf{A}), \\
\Gamma_{k+1}(\mathbf{A}) \rightleftarrows^{r(\ell)}_{h_{k+1}\restriction S_i} \Gamma_{k+1}(\mathbf{B}) & \text{since } \widetilde{\mathbf{B}} \xrightarrow{\text{retr}} \Gamma_{k+1}(\mathbf{B}) \text{ and so } \widetilde{\mathbf{B}} \rightleftarrows_{T_i} \Gamma_{k+1}(\mathbf{B}), \\
\Gamma_{k+1}(\mathbf{A}) \rightleftarrows^{r(k)+t(k)}_{h_{k+1}\restriction S_i} \Gamma_{k+1}(\mathbf{B}) & \text{since } r(k) + t(k) = r(k+1) \leq r(\ell).
\end{array}
$$

Note that by the definition of $\Gamma_k(\mathbf{A}) = \Delta_k(\Gamma_{k+1}(\mathbf{A}))$, the clan $\mathbf{D}_i$ is $\cong_{S_i}$-isomorphic to some extension core $\mathbf{D}'_i$ in the set $\mathscr{E}^{r(k)}_{S_i, \leq t(k)}(\Gamma_{k+1}(\mathbf{A}))$. Let $\mathbf{E}_i$ be the isomorphic image of $\mathbf{D}'_i$ under the function $D'_i \longrightarrow (D'_i \setminus S_i) \cup T_i$ defined by $d \longmapsto h_{k+1}(d)$ if $d \in S_i$ and $d \longmapsto d$ if $d \in D'_i \setminus S_i$. By Corollary 3.26, $\mathbf{E}_i$ is an extension core in the set $\mathscr{E}^{r(k)}_{T_i, \leq t(k)}(\Gamma_{k+1}(\mathbf{B}))$.

The upshot of the last three paragraphs is that the pair $(T_i, \mathbf{E}_i)$ shows up among index pairs $(T, \mathbf{E})$ in the following expression defining $\Gamma_k(\mathbf{B})$:

$$
\Gamma_k(\mathbf{B}) = \Delta_k(\Gamma_{k+1}(\mathbf{B})) = \bigoplus_{\substack{T \subseteq \Gamma_{k+1}(B), \\ |T| \leq s(k) \\ \mathbf{E} \in \mathscr{E}^{r(k)}_{T, \leq t(k)}(\Gamma_{k+1}(\mathbf{B}))}} {}_{\Gamma_{k+1}(B)} \left(\Gamma_{k+1}(\mathbf{B}) \oplus_T \underbrace{\mathbf{E} \oplus_T \cdots \oplus_T \mathbf{E}}_{u(k) \text{ times}} \right)
$$

For all $i \in \{1, \dots, q\}$, let $\mathbf{E}_{i,1}, \dots, \mathbf{E}_{i,u(k)}$ be the distinct copies of extension core $\mathbf{E}_i$ in this expression (modulo the common subset T_i). Note that q, the number of different clans of elements of Z_k with level k, is at most $|^k\{\alpha\}| \le t(k)$ by Lemma 5.9. Let $\mathbf{F}_1, \dots, \mathbf{F}_p$ enumerate (in any order) the set $\{\mathbf{Clan}(f(y)) : y \in Y_k \setminus Y_{k+1}\}$ of clans (in $\widetilde{\mathbf{B}}$) of elements $f(y)$ for $y \in \mathrm{cl}_\ell(X)$ with level k (in $\widetilde{\mathbf{A}}$). Note that p, the number of such clans, is at most $|Y_k \setminus Y_{k+1}| \le |^kX|$, which by Lemma 5.9 is at most $|X| \cdot t(k) \le (n-\ell-1) \cdot t(k) \le (n-k-1) \cdot t(k) = u(k) - t(k)$. By a simple pigeonhole argument, there exists a function $j : \{1, \dots, q\} \longrightarrow \{1, \dots, u(k)\}$ such that for all $i \in \{1, \dots, q\}$,

$$\mathbf{E}_{i,j(i)} \notin \{\mathbf{F}_1, \dots, \mathbf{F}_p, \mathbf{E}_{1,j(1)}, \dots, \mathbf{E}_{i-1,j(i-1)}\}.$$

For each i, we fix any isomorphism $g_i : \mathbf{D}_i \cong_{h_{k+1} \restriction S_i} \mathbf{E}_{i,j(i)}$. We now define function $h_k : Y_k \cup Z_k \longrightarrow \widetilde{B}$ by

$$h_k(x) \stackrel{\text{def}}{=} \begin{cases} h_{k+1}(x) & \text{if } x \in Y_{k+1} \cup Z_{k+1} \\ f(x) & \text{if } x \in Y_k \setminus Y_{k+1} \\ g_i(x) & \text{if } x \in D_i \setminus S_i \end{cases} = \begin{cases} h_{k+1}(x) & \text{if } x \in Y_{k+1} \cup Z_{k+1} \\ f(x) & \text{if } x \in Y_k \\ g_i(x) & \text{if } x \in D_i. \end{cases}$$

In the lefthand expression defining h_k, note that sets $Y_{k+1} \cup Z_{k+1}$ and $Y_k \setminus Y_{k+1}$ and $D_1 \setminus S_1, \dots, D_q \setminus S_q$ partition the set $Y_k \cup Z_k$. In the equivalent righthand expression, note that h_k is consistent on overlapping cases $(Y_{k+1} \cup Z_{k+1}) \cap Y_k$ and $(Y_{k+1} \cup Z_{k+1}) \cap D_i$ and $Y_k \cap D_i$; in particular, $f(x) = h_{k+1}(x)$ for all $x \in Y_{k+1}$ by hypothesis (H3) and $g_i(x) = h_{k+1}(x)$ for all $x \in S_i$ since by our choice of g_i (i.e., the fact that g_i is an isomorphism extending $h_{k+1} \restriction S_i$). Note that h_k indeed has domain $Y_k \cup Z_k$, as $\bigcup_{i=1}^q D_i \setminus S_i = Z_k \setminus Z_{k+1}$; thus, h_k satisfies hypothesis (H2). Moreover, h_k clearly extends both h_{k+1} and $f \restriction Y_k$; it thus satisfies hypothesis (H3).

It remains only to show that h_k satisfies hypothesis (H1), i.e., that it is ℓ-kosher. ℓ-kosher conditions (K2)–(K4) are easy to check. We will show that h_k satisfies ℓ-kosher condition (K1), i.e., $\widetilde{\mathbf{A}} \rightleftarrows_{h_k}^{r(\ell)} \widetilde{\mathbf{B}}$. Define substructures $\mathbf{G} \subseteq \Gamma_k(\mathbf{A})$ and $\mathbf{F} \subseteq \Gamma_k(\mathbf{B})$ by

$$\mathbf{G} = \bigcup_{y \in Y_k \setminus Y_{k+1}} \mathbf{Clan}(y), \qquad \mathbf{F} = \mathbf{F}_1 \cup \dots \cup \mathbf{F}_p = \bigcup_{y \in Y_k \setminus Y_{k+1}} \mathbf{Clan}(f(y)).$$

Note that $G \subseteq \mathrm{Dom}(f)$. By ℓ+1-kosherness of f and Lemma 5.11(3), we see that $F = f(G)$. Therefore, $f \restriction G$ is an isomorphism of structures $\mathbf{G} \cong \mathbf{F}$.

Let $S' = G \cap \Gamma_{k+1}(A)$ and $T' = F \cap \Gamma_{k+1}(B)$. We now define substructures $\mathbf{A}' \subseteq \widetilde{\mathbf{A}}$ and $\mathbf{B}' \subseteq \widetilde{\mathbf{B}}$ by

$$\mathbf{A}' = \big(\cdots\big(\big(\big(\Gamma_{k+1}(\mathbf{A}) \oplus_{S'} \mathbf{G}\big) \oplus_{S_1} \mathbf{D}_1\big) \oplus_{S_2} \mathbf{D}_2\big) \cdots \oplus_{S_{q-1}} \mathbf{D}_{q-1}\big) \oplus_{S_q} \mathbf{D}_q,$$

$$\mathbf{B}' = \big(\cdots\big(\big(\big(\Gamma_{k+1}(\mathbf{B}) \oplus_{T'} \mathbf{F}\big) \oplus_{T_1} \mathbf{E}_{1,j(1)}\big) \oplus_{T_2} \mathbf{E}_{2,j(2)}\big) \cdots \cdots \oplus_{T_{q-1}} \mathbf{E}_{q-1,j(q-1)}\big) \oplus_{T_q} \mathbf{E}_{q,j(q)}.$$

Note that $\mathrm{Dom}(h_k) \subseteq A'$ and $\mathrm{Range}(h_k) \subseteq B'$. In addition, note that $\widetilde{\mathbf{A}} \xrightarrow{\text{retr}} \Gamma_k(\mathbf{A}) \xrightarrow{\text{retr}} \mathbf{A}'$ and $\widetilde{\mathbf{B}} \xrightarrow{\text{retr}} \Gamma_k(\mathbf{B}) \xrightarrow{\text{retr}} \mathbf{B}'$. Therefore, once we prove $\mathbf{A}' \rightleftarrows_{h_k}^{r(\ell)} \mathbf{B}'$, it follows that $\widetilde{\mathbf{A}} \rightleftarrows_{h_k}^{r(\ell)} \widetilde{\mathbf{B}}$.

The final claim is that $\mathbf{A}' \rightleftarrows_{h_k}^{r(\ell)} \mathbf{B}'$. This is accomplished by repeatedly applying Corollary 3.5.

—First, note that $\Gamma_{k+1}(\mathbf{A}) \rightleftarrows^{r(\ell)}_{h_k \restriction \Gamma_{k+1}(A)} \Gamma_{k+1}(\mathbf{B})$. Indeed, $h_k \restriction \Gamma_{k+1}(A) = h_{k+1}$ and $\widetilde{\mathbf{A}} \rightleftarrows^{r(\ell)}_{h_{k+1}} \widetilde{\mathbf{B}}$ by ℓ-kosher condition (K1) of h_k; now use $\widetilde{\mathbf{A}} \xrightarrow{\text{retr}} \Gamma_{k+1}(\mathbf{A})$ and $\widetilde{\mathbf{B}} \xrightarrow{\text{retr}} \Gamma_{k+1}(\mathbf{B})$.

—Second, note that $\mathbf{G} \cong_{h_k \restriction G} \mathbf{F}$ and $S' = \Gamma_{k+1}(A) \cap G$ and $T' = h_k(S')$. By Corollary 3.5, this implies

$$\Gamma_{k+1}(\mathbf{A}) \oplus_{S'} \mathbf{G} \rightleftarrows^{r(\ell)}_{h_k \restriction (\Gamma_{k+1}(A) \oplus_{S'} G)} \Gamma_{k+1}(\mathbf{B}) \oplus_{T'} \mathbf{F}.$$

—Next, we have $\mathbf{D}_1 \cong_{h_k \restriction D_1} \mathbf{E}_{1,j(1)}$. Also, $S_1 = (\Gamma_{k+1}(A) \oplus_{S'} G) \cap D_1$ and $T_1 = h_k(S_1)$. Corollary 3.5 implies

$$\big(\Gamma_{k+1}(\mathbf{A}) \oplus_{S'} \mathbf{G}\big) \oplus_{S_1} \mathbf{D}_1 \rightleftarrows^{r(\ell)}_{h_k \restriction (\Gamma_{k+1}(A) \oplus_{S'} G) \oplus_{S_1} D_1} \big(\Gamma_{k+1}(\mathbf{B}) \oplus_{T'} \mathbf{F}\big) \oplus_{T_1} \mathbf{E}_{1,j(1)}.$$

—We proceed in this fashion for $i = 2, \ldots, q$, each time proving

$$\begin{aligned}&\Big(\big(\Gamma_{k+1}(\mathbf{A}) \oplus_{S'} \mathbf{G}\big) \oplus_{S_1} \mathbf{D}_1\Big) \cdots \oplus_{S_i} \mathbf{D}_i \\ &\qquad \rightleftarrows^{r(\ell)}_{h_k \restriction ((\Gamma_{k+1}(A) \oplus_{S'} G) \oplus_{S_1} D_1) \cdots \oplus_{S_i} D_i} \\ &\qquad\qquad \Big(\big(\Gamma_{k+1}(\mathbf{B}) \oplus_{T'} \mathbf{F}\big) \oplus_{T_1} \mathbf{E}_{1,j(1)}\Big) \cdots \oplus_{T_i} \mathbf{E}_{i,j(i)}.\end{aligned}$$

Once we reach $i = q$, we have shown $\mathbf{A}' \rightleftarrows^{r(\ell)}_{h_k} \mathbf{B}'$.

With this, the proof of Lemma 5.12 is concluded. □

The following corollary shows that Lemma 5.12 is really symmetric in $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$.

Corollary 5.13. *Let $\ell \in \{0, \ldots, n-1\}$ and suppose f is an $\ell+1$-kosher partial isomorphism from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$ with domain $\mathrm{cl}_{\ell+1}(X)$ for some $X \subseteq \widetilde{A}$ of size $< n - \ell$. Then for every $\beta \in \widetilde{B}$, there exists an ℓ-kosher partial isomorphism f' extending $f \restriction \mathrm{cl}_\ell(X)$ to range $\mathrm{cl}_\ell(f(X) \cup \{\beta\})$. Moreover, $\mathrm{Dom}(f') = \mathrm{cl}(X \cup \{f'^{-1}(\beta)\})$ for any such f'.*

Proof. We have $f(\mathrm{cl}_{\ell+1}(X)) = \mathrm{cl}_{\ell+1}(f(X))$ by Lemma 5.11(4). Thus, f^{-1} is an $\ell+1$-kosher partial isomorphism from $\widetilde{\mathbf{B}}$ to $\widetilde{\mathbf{A}}$ with domain $\mathrm{cl}_{\ell+1}(f(X))$ where $|f(X)| < n - \ell$. By Lemma 5.12, for every $\beta \in \widetilde{B}$ there is an ℓ-kosher partial isomorphism g from $\widetilde{B}$ to $\widetilde{A}$ extending $f^{-1} \restriction \mathrm{cl}_\ell(f(X))$ to domain $\mathrm{cl}_\ell(f(X) \cup \{\beta\})$. By Lemma 5.11(1), the inverse g^{-1} is an ℓ-kosher partial isomorphism from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$ extending $f \restriction \mathrm{cl}_\ell(X)$ to range $\mathrm{cl}_\ell(f(X) \cup \{\beta\})$. Finally, if f' is any ℓ-kosher partial isomorphism extending $f \restriction \mathrm{cl}_\ell(X)$ to range $\mathrm{cl}_\ell(f(X) \cup \{\beta\})$, then

$$\begin{aligned}\mathrm{Dom}(f') &= f'^{-1}(\mathrm{cl}_\ell(f(X) \cup \{\beta\})) \\ &= \mathrm{cl}_\ell(f'^{-1}(f(X)) \cup \{f'^{-1}(\beta)\}) = \mathrm{cl}_\ell(X \cup \{f'^{-1}(\beta)\})\end{aligned}$$

where the middle equality is by Lemma 5.11(4). □

Lemmas 2.20, 5.11(2), 5.12 and Corollary 5.13 yield:

Corollary 5.14. *Suppose $\mathbf{A} \rightleftarrows^{r(n)} \mathbf{B}$. For $k \in \{0, \ldots, n\}$, let Π_k be the set of partial isomorphisms f from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$ such that $|\mathrm{Dom}(f)| \leq k$ and f extends to an $(n-k)$-kosher partial isomorphism from $\widetilde{\mathbf{A}}$ to $\widetilde{\mathbf{B}}$. Then the sequence $\Pi_0, \ldots, \Pi_n$ is an n-back-and-forth system on $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$. Consequently, $\widetilde{\mathbf{A}} \equiv^n \widetilde{\mathbf{B}}$.* □

Until now, n has been a fixed positive integer. The value of n has determined the sequences $r(\ell), s(\ell), t(\ell), u(\ell)$ for $\ell = 0, \dots, n-1$, as well as the number $r(n)$. We now make this dependence explicit by writing $r_n(\ell), s_n(\ell), t_n(\ell), u_n(\ell)$. Let $\rho : \mathbb{N} \longrightarrow \mathbb{N}$ be the function defined by $\rho(n) = r_n(n)$.

THEOREM 5.15. *Suppose $\mathcal{P}$ and $\mathcal{Q}$ are classes of finite structures and Φ is a first-order sentence such that for all finite structures $\mathbf{A}$ and $\mathbf{B}$,*

- *if $\mathbf{A} \in \mathcal{P}$ and $\mathbf{A} \to \mathbf{B}$ then $\mathbf{B} \models \Phi$, and*
- *if $\mathbf{A} \models \Phi$ and $\mathbf{A} \to \mathbf{B}$ then $\mathbf{B} \in \mathcal{Q}$.*

Then there exists an existential-positive sentence Ψ such that $\mathrm{qrank}(\Psi) \leq \rho(\mathrm{qrank}(\Phi))$ and $\mathcal{P} \subseteq \mathrm{Mod_{fin}}(\Psi) \subseteq \mathcal{Q}$.

The proof follows the same scheme as the proof of Theorem 4.11.

PROOF. Let $n = \mathrm{qrank}(\Phi)$ and suppose $\mathbf{A}$ and $\mathbf{B}$ are finite structures such that $\mathbf{A} \in \mathcal{P}$ and $\mathbf{A} \to^{\rho(n)} \mathbf{B}$. We will show that $\mathbf{B} \in \mathcal{Q}$. The conclusion that $\mathcal{P} \subseteq \mathrm{Mod_{fin}}(\Psi) \subseteq \mathcal{Q}$ for some existential-positive sentence Ψ of quantifier-rank at most $\rho(n)$ then follows by Lemma 3.16.

Let $\mathbf{C} = \mathbf{Core}^{\rho(n)}(\mathbf{A})$ and note that $\mathbf{A} \rightleftarrows^{\rho(n)} \mathbf{C}$ and $\mathbf{C} \to \mathbf{B}$. Corollary 5.14 yields $\widetilde{\mathbf{A}} \equiv^n \widetilde{\mathbf{C}}$, and so we have the diagram

$$\begin{array}{ccccc} \widetilde{\mathbf{A}} & \equiv^n & \widetilde{\mathbf{C}} & & \\ \downarrow{\scriptstyle \mathrm{retr}} & & \downarrow{\scriptstyle \mathrm{retr}} & & \\ \mathbf{A} & \rightleftarrows^{\rho(n)} & \mathbf{C} & \to & \mathbf{B} \end{array}$$

in which all structures are finite. Since $\mathbf{A} \in \mathcal{P}$ and $\mathbf{A} \to \widetilde{\mathbf{A}}$, we have $\widetilde{\mathbf{A}} \models \Phi$ (by the hypothesis involving $\mathcal{P}$). As $\widetilde{\mathbf{A}} \equiv^n \widetilde{\mathbf{C}}$, it follows that $\widetilde{\mathbf{C}} \models \Phi$. Since $\widetilde{\mathbf{C}} \models \Phi$ and $\widetilde{\mathbf{C}} \to \mathbf{B}$, we have $\mathbf{B} \in \mathcal{Q}$ (by the hypothesis involving $\mathcal{Q}$). Having shown that $\mathbf{B} \in \mathcal{Q}$, we are done. □

The main result of this article, the homomorphism preservation theorem on finite structures, now follows as a corollary.

THEOREM 5.16 (FINITE HOMOMORPHISM PRESERVATION THEOREM). *A first-order sentence of quantifier-rank n is preserved under homomorphisms on finite structures if, and only if, it is equivalent in the finite to an existential-positive sentence of quantifier-rank $\rho(n)$.*

PROOF. The easy ($\Longleftarrow$) direction was given earlier as Lemma 2.3. The hard ($\Longrightarrow$) direction is nothing but the special case of Theorem 5.15 where $\mathcal{P} = \mathcal{Q} = \mathrm{Mod_{fin}}(\Phi)$ for a first-order sentence Φ. □

6. NON-ELEMENTARY LOWER BOUNDS

In this section we present a previously unpublished result of Yuri Gurevich and Saharon Shelah (announced but not proved in [Gurevich 1990]). The proof we give here was adapted from unpublished notes. Recall that a function from $\mathbb{N}$ to $\mathbb{N}$ is *non-elementary* if it is not bounded by any tower-exponential function $2^{2^{\cdot^{\cdot^{\cdot^{2^n}}}}}$ of fixed height.

THEOREM 6.1. *There is a sequence $\Phi_1, \Phi_2, \Phi_3, \ldots$ of homomorphism-preserved first-order sentences (in a suitable finite relational vocabulary) such that the minimal quantifier-count of any existential-positive sentence equivalent to Φ_n is a non-elementary function of the length of Φ_n.*

In a slightly different formulation, Theorem 6.1 states that

$$\max_{\substack{\text{hom-preserved first-order } \Phi \\ \text{length}(\Phi)=n}} \ \min_{\substack{\text{ex-pos } \Psi \\ \Psi \equiv \Phi}} \text{qcount}(\Psi) = \text{non-elementary}(n).$$

Since quantifier-count is a lower bound on length, this statement remains true if we replace length(Φ) by qcount(Φ). This theorem neatly complements our equirank h.p.t., which tells us that there is no blow-up in quantifier-rank, that is,

$$\max_{\substack{\text{hom-preserved first-order } \Phi \\ \text{qrank}(\Phi)=n}} \ \min_{\substack{\text{ex-pos } \Psi \\ \Psi \equiv \Phi}} \text{qrank}(\Psi) \leq n.$$

The "suitable" finite relational vocabulary in Theorem 6.1 consists of three unary relation symbols and one binary relation symbol; we believe, however, that Theorem 6.1 should be true for any finite relational vocabulary containing at least one relation symbol of arity at least 2. In the course of proving Theorem 6.1, we show (in Corollary 6.3) that the core-size bound β_0^n is a non-elementary function of n. It follows easily that β_m^n is a non-elementary function of n for every constant m.

6.1 The Existential-Positive Sentence $\Theta_n \vee \Psi_n$

Let σ be the vocabulary $\{Zero, Gen^{♂}, Gen^{♀}, E\}$ where $Zero, Gen^{♂}, Gen^{♀}$ are unary relation symbols and E is a binary relation symbol. Let $(S_n)_{n\in\mathbb{N}}$ be the sequence of sets defined by $S_0 = \{\emptyset\}$ and $S_n = \{♂, ♀\}^{S_{n-1}}$ (i.e., the set of functions from S_{n-1} to $\{♂, ♀\}$) for all $n \geq 1$. Let tower$(n) = |S_n|$, so that tower$(0) = 1$ and tower$(n) = 2^{\text{tower}(n-1)}$ for all $n \geq 1$.

For all $f \in \bigcup_{n\in\mathbb{N}} S_n$, we define a primitive-positive formula $\theta^f(x)$ by

$$\theta^f(x) \triangleq \begin{cases} Zero(x) & \text{if } f \in S_0 \text{ (i.e., } f = \emptyset), \\ \displaystyle\bigwedge_{f' \in S_{n-1}} \Big(\exists y\ Exy \wedge \theta^{f'}(y) \wedge Gen^{f(f')}(y)\Big) & \text{if } f \in S_n \text{ for } n \geq 1. \end{cases}$$

(We use the notation $\triangleq$, above, to distinguish from the internal equality symbol $=$ of first-order logic.) For all $n \geq 1$, we define a primitive-positive sentence Θ_n and an existential-positive sentence Ψ_n by

$$\Theta_n \triangleq \bigwedge_{f \in S_{n-1}} \Big(\exists x\ \theta^f(x) \wedge Gen^{♂}(x)\Big) \wedge \Big(\exists x\ \theta^f(x) \wedge Gen^{♀}(x)\Big),$$

$$\Psi_n \triangleq \Big(\exists x\ Gen^{♂}(x) \wedge Gen^{♀}(x)\Big) \vee \bigvee_{\substack{f,f' \in S_0 \cup \cdots \cup S_{n-1} \\ f \neq f'}} \Big(\exists x\ \theta^f(x) \wedge \theta^{f'}(x)\Big).$$

Notice that qrank$(\Theta_n) =$ qrank$(\Psi_n) = n$ and qcount(Θ_n), qcount$(\Psi_n) \geq |S_{n-1}| =$ tower$(n-1)$.

We introduce some terminology for speaking about structures (with vocabulary σ). Elements of S_n are *n-colors*; by extension, elements of $S_0 \cup \cdots \cup S_{n-1}$ are

<n-colors. Given a structure $\mathbf{A}$, an element $a \in A$ and an n-color $f \in S_n$, we say that *a has n-color f in* $\mathbf{A}$ if $\mathbf{A} \models \theta^f(a)$. Entities ♂ and ♀ are *genders*. We say that a is *male* if $\mathbf{A} \models \mathit{Gen}^{♂}(a)$ and *female* if $\mathbf{A} \models \mathit{Gen}^{♀}(a)$. We say that *a has a gender* if it is male or female (or both). In this terminology, the sentence Θ_n says that for all $f \in S_{n-1}$, there exists a male element with $n{-}1$-color f as well as a female element with $n{-}1$-color f (possibly the same element). The negation $\neg\Psi_n$ of Ψ_n says that all elements have at most one gender and at most one $<n$-color. Thus, if $\mathbf{A} \models \neg\Psi_n$ and an element $a \in A$ has a gender (resp. $<n$-color), then we may unambiguously speak of its (unique) gender (resp. $<n$-color). Bear in mind that an element $a \in A$ has a 0-color in $\mathbf{A}$ if, and only if, $\mathbf{A} \models \mathit{Zero}(a)$.

Notice that genders and colors are preserved under homomorphisms. By this, we mean that if h is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ and $a \in A$ has gender $g \in \{♂, ♀\}$ in $\mathbf{A}$, then $h(a)$ has gender g in $\mathbf{B}$. Similarly, if a has n-color $f \in S_n$ in $\mathbf{A}$, then $h(a)$ has n-color f in $\mathbf{B}$.

6.2 The Hom-Minimal Core $\mathbf{C}_n$

We now define a particular finite structure $\mathbf{C}_n$, which we will show is both a core and a hom-minimal model of $\Theta_n \vee \Psi_n$ (see Definition 2.15). The universe C_n consists of all sequences $\langle f_i, g_i, f_{i+1}, g_{i+1}, \dots, f_{n-1}, g_{n-1}\rangle$ where $i \in \{0, \dots, n-1\}$ and $f_j \in S_j$ and $g_j \in \{♂, ♀\}$ subject to the constraint that $f_j(f_{j-1}) = g_{j-1}$ for all $i < j < n$. (Recall that j-colors are functions from $j-1$-colors to genders.) We regard C_n as a rooted forest in which an element $\langle f_i, g_i, f_{i+1}, g_{i+1}, \dots, f_{n-1}, g_{n-1}\rangle$ where $i < n-1$ has as its parent the element $\langle f_{i+1}, g_{i+1}, \dots, f_{n-1}, g_{n-1}\rangle$. Relation symbols $\mathit{Zero}, \mathit{Gen}^{♂}, \mathit{Gen}^{♀}, E$ are interpreted in $\mathbf{C}_n$ as follows.

- E is the child-of relation in C_n, that is, the set of pairs
$$(\langle f_i, g_i, f_{i+1}, g_{i+1}, \dots, f_{n-1}, g_{n-1}\rangle, \langle f_{i+1}, g_{i+1}, f_{i+2}, g_{i+2}, \dots, f_{n-1}, g_{n-1}\rangle)$$
where $i < n-1$.
- Zero is the set of leaves of C_n, that is, the set of $\langle f_0, g_0, f_1, g_1, \dots, f_{n-1}, g_{n-1}\rangle$ having length $2n$.
- $\mathit{Gen}^{♂}$ (resp. $\mathit{Gen}^{♀}$) is the set of elements $\langle f_i, g_i, f_{i+1}, g_{i+1}, \dots, f_{n-1}, g_{n-1}\rangle$ such that $g_i = ♂$ (resp. $g_i = ♀$).

As an aside, we point out that the structure $\mathbf{C}_n$ is (up to isomorphism) precisely the structure $\mathbf{A}_{\Theta_n}$ as defined in the proof of Lemma 2.13.

LEMMA 6.2. *$\mathbf{C}_n$ is a core.*

PROOF. Even stronger, we will prove that $\mathbf{C}_n$ is *rigid*, meaning that the only homomorphism $\mathbf{C}_n \to \mathbf{C}_n$ is the identity map. We first observe that every homomorphism $h : \mathbf{C}_n \to \mathbf{C}_n$ fixes each root $\langle f_{n-1}, g_{n-1}\rangle$; this is because h preserves both genders and colors and each $\langle f_{n-1}, g_{n-1}\rangle$ is the unique element of $\mathbf{C}_n$ with color f_{n-1} and gender g_{n-1}. We next claim that h fixes all children $\langle f_{n-2}, g_{n-2}, f_{n-1}, g_{n-1}\rangle$ of roots; indeed, $\langle f_{n-2}, g_{n-2}, f_{n-1}, g_{n-1}\rangle$ is the unique element of $\mathbf{C}_n$ with color f_{n-2} and gender g_{n-2} whose parent has color f_{n-1} and gender g_{n-1} (this "second generation" property is similarly preserved under homomorphisms). We proceed inductively to show that h fixes elements at distance k

from a root for all k from 2 to $n-2$. This argument establishes that $\mathbf{C}_n$ is rigid and is therefore a core. $\square$

It is evident from our description of $\mathbf{C}_n$ as rooted forest that it has tree-depth at most n. Since $\mathbf{C}_n$ is a core, it follows that core-size bound β_0^n is at least the size of $\mathbf{C}_n$. As $|\mathbf{C}_n| \geq |S_{n-1}| = \text{tower}(n-1)$, we have obtained a non-elementary lower bound on the core-size bound.

Corollary 6.3. *The core-size bound β_0^n is a non-elementary function of n.* $\square$

We claim that $\mathbf{C}_n \models \Theta_n \wedge \neg\Psi_n$. Each element $\langle f_i, g_i, f_{i+1}, g_{i+1}, \dots, f_{n-1}, g_{n-1}\rangle$ has gender g_i and i-color f_i in $\mathbf{C}_n$ (this is seen by induction on i). In particular, we have $\mathbf{C}_n \models \Theta_n$ since $\mathbf{C}_n$ contains both a male element and a female element of each $n{-}1$-color. On the other, we have $\mathbf{C}_n \models \neg\Psi_n$ since no element of $\mathbf{C}_n$ is both male and female, nor does any element have more than one $<n$-color.

Lemma 6.4. *$\mathbf{C}_n$ is a hom-minimal model of $\Theta_n \vee \Psi_n$.*

Proof. We just observed that $\mathbf{C}_n \models \Theta_n$, so therefore $\mathbf{C}_n \models \Theta_n \vee \Psi_n$. Suppose $\mathbf{B}$ is a model of $\Theta_n \vee \Psi_n$ such that $\mathbf{B} \to \mathbf{C}_n$. To prove that $\mathbf{C}_n$ is a hom-minimal model (Definition 2.15), we must show that $\mathbf{C}_n \to \mathbf{B}$. We begin by observing that $\mathbf{B} \not\models \Psi_n$. Indeed, were it the case that $\mathbf{B} \models \Psi_n$, then $\mathbf{C}_n \models \Psi_n$ since Ψ_n is existential-positive and hence preserved under homomorphisms; but we know that $\mathbf{C}_n \not\models \Psi_n$. Therefore, $\mathbf{B} \not\models \Psi_n$ and so it must be the case that $\mathbf{B} \models \Theta_n$. We now construct a homomorphism h from $\mathbf{C}_n$ to $\mathbf{B}$ starting with the roots of $\mathbf{C}_n$ and working down to the leaves. For each root $\langle f_{n-1}, g_{n-1}\rangle$, there exists an element $b \in B$ with $n{-}1$-color f_{n-1} and gender g_{n-1} (because $\mathbf{B} \models \Theta_n$); let h map $\langle f_{n-1}, g_{n-1}\rangle$ to any such b. For each child $\langle f_{n-2}, g_{n-2}, f_{n-1}, g_{n-1}\rangle$ of $\langle f_{n-1}, g_{n-1}\rangle$, the fact that $h(\langle f_{n-1}, g_{n-1}\rangle)$ has $n{-}1$-color f_{n-1} means that it has a child b' with gender g_{n-2} and $n{-}2$-color f_{n-2}; let h map $\langle f_{n-2}, g_{n-2}, f_{n-1}, g_{n-1}\rangle$ to any such b'. Continuing in this manner, we eventually extend h to the entire universe of $\mathbf{C}_n$. $\square$

Lemma 6.5. *$\Theta_n \vee \Psi_n$ is not logically equivalent to any existential-positive sentence with fewer than* $\text{tower}(n-1)$ *quantifiers.*

Proof. Let Φ be an existential-positive sentence equivalent to $\Theta_n \vee \Psi_n$. Then Φ and $\Theta_n \vee \Psi_n$ have the same hom-minimal models; in particular, $\mathbf{C}_n$ is a finite hom-minimal model of Φ. By Proposition 2.16, we have

$$\text{qcount}(\Phi) \geq |\mathbf{Core}(\mathbf{C}_n)| = |\mathbf{C}_n| \geq \text{tower}(n-1)$$

where the middle equality uses fact that $\mathbf{C}_n$ is a core (Lemma 6.2) and so $\mathbf{C}_n \cong \mathbf{Core}(\mathbf{C}_n)$. $\square$

6.3 An Exponentially Concise First-Order Sentence Equivalent to $\Theta_n \vee \Psi_n$

We now complete the proof of Theorem 6.1 by defining a sequence of first-order sentences Φ_n equivalent to $\Theta_n \vee \Psi_n$ but only exponentially long in n. First, we define various subformulas. These formulas are listed below, along with the intended meaning in models of $\neg\Psi_{n+1}$:

$HasG(x)$ $\Leftrightarrow$ x has a (unique) gender,
$SameG(x, x')$ $\Leftrightarrow$ x and x' have the same gender,

$$\begin{array}{lcl}
\mathit{OppG}(x,x') & \Leftrightarrow & x \text{ and } x' \text{ have opposite genders,}\\
\mathit{HasC}_n(x) & \Leftrightarrow & x \text{ has a (unique) } n\text{-color,}\\
\mathit{SameC}_n(x,x') & \Leftrightarrow & x \text{ and } x' \text{ have the same } n\text{-color,}\\
\mathit{Edge}_n(x,y) & \Leftrightarrow & y \text{ has both an } n{-}1\text{-color and a gender and there is an edge from } x \text{ to } y,\\
\mathit{Swap}_n(x,x',y) & \Leftrightarrow & x,x' \text{ have } n\text{-colors (say } f,f' \in S_n),\ y \text{ has an } n{-}1\text{-color (say } f_0 \in S_{n-1}), \text{ and } f,f' \text{ agree (as functions } S_{n-1} \longrightarrow \{♂,♀\})\\
 & & \text{on all } n{-}1\text{-colors except } f_0, \text{ i.e., } f(f_0) = f'(f_1) \Longleftrightarrow f_0 \neq f_1 \text{ for all } f_1 \in S_{n-1}.
\end{array}$$

In other words, $\mathit{Swap}_n(x,x',y)$ says that the n-colors of x and x' swap values on the $n{-}1$-color of y, but are identical on all other $n{-}1$-colors.

Formulas $\mathit{HasG}(x)$, $\mathit{SameG}(x,y)$ and $\mathit{OppG}(x,y)$ are defined easily enough. (In the following definitions, $\phi \to \psi$ and $\phi \leftrightarrow \psi$ and $\perp$ respectively abbreviate $\neg\phi \vee \psi$ and $(\phi \to \psi) \wedge (\psi \to \phi)$ and $\exists x\ \neg(x = x)$.)

$$\begin{array}{lcl}
\mathit{HasG}(x) & \triangleq & \mathit{Gen}^{♂}(x) \vee \mathit{Gen}^{♀}(x)\\
\mathit{SameG}(x,x') & \triangleq & \big(\mathit{Gen}^{♂}(x) \wedge \mathit{Gen}^{♂}(x')\big) \vee \big(\mathit{Gen}^{♀}(x) \wedge \mathit{Gen}^{♀}(x')\big)\\
\mathit{OppG}(x,x') & \triangleq & \big(\mathit{Gen}^{♂}(x) \wedge \mathit{Gen}^{♀}(x')\big) \vee \big(\mathit{Gen}^{♀}(x) \wedge \mathit{Gen}^{♂}(x')\big)
\end{array}$$

The other formulas are defined recursively. In the base case $n = 0$:

$$\begin{array}{lcl@{\qquad}lcl}
\mathit{HasC}_0(x) & \triangleq & \mathit{Zero}(x) & \mathit{Swap}_0(x,x',y) & \triangleq & \perp\\
\mathit{SameC}_0(x,x') & \triangleq & \mathit{Zero}(x) \wedge \mathit{Zero}(x') & \mathit{Edge}_0(x,y) & \triangleq & \perp
\end{array}$$

For all $n \geq 1$, we define:

$$\begin{array}{lcl}
\mathit{Edge}_n(x,y) & \triangleq & Exy \wedge \mathit{HasC}_{n-1}(y) \wedge \mathit{HasG}(y)\\
\mathit{HasC}_n(x) & \triangleq & \exists y_0\ \mathit{Edge}_n(x,y_0) \wedge\\
 & & \forall y \forall z \begin{pmatrix} \mathit{Edge}_n(x,y) \wedge \\ \mathit{Edge}_{n-1}(y,z) \end{pmatrix} \to \exists y' \begin{pmatrix} \mathit{Edge}_n(x,y') \wedge \\ \mathit{Swap}_{n-1}(y,y',z) \end{pmatrix}\\
\mathit{SameC}_n(x,x') & \triangleq & \mathit{HasC}_n(x) \wedge \mathit{HasC}_n(x') \wedge\\
 & & \forall y \forall y' \begin{pmatrix} \mathit{Edge}_n(x,y) \wedge \\ \mathit{Edge}_n(x',y') \wedge \\ \mathit{SameC}_{n-1}(y,y') \end{pmatrix} \to \mathit{SameG}(y,y')\\
\mathit{Swap}_n(x,x',y_0) & \triangleq & \mathit{HasC}_n(x) \wedge \mathit{HasC}_n(x') \wedge \mathit{HasC}_{n-1}(y_0) \wedge\\
 & & \forall y \forall y' \begin{pmatrix} \mathit{Edge}_n(x,y) \wedge \\ \mathit{Edge}_n(x',y') \wedge \\ \mathit{SameC}_{n-1}(y,y') \end{pmatrix} \to \begin{pmatrix} \mathit{OppG}(y,y') \leftrightarrow \\ \mathit{SameC}_{n-1}(y_0,y) \end{pmatrix}
\end{array}$$

One can check that this recursion is well-founded and that these formulas have the intended meaning on models of $\neg\Psi_{n+1}$.

For all $n \geq 1$, we define first-order sentences Θ_n^* and Ψ_n^* as follows. Let $\Theta_1^* \triangleq \Theta_1$ and $\Psi_1^* \triangleq \Psi_1$. For all $n \geq 2$, let

$$\begin{aligned}\Theta_n^* \triangleq\ & \exists x_0 \Big(\mathit{HasC}_{n-1}(x_0) \wedge \mathit{HasG}(x_0)\Big) \wedge \\ & \forall x \Big(\mathit{HasC}_{n-1}(x) \wedge \mathit{HasG}(x)\Big) \rightarrow \exists x' \Big(\mathit{SameC}_{n-1}(x, x') \wedge \mathit{OppG}(x, x')\Big) \wedge \\ & \forall x \forall y \Big(\mathit{HasC}_{n-1}(x) \wedge \mathit{HasC}_{n-2}(y)\Big) \rightarrow \exists x' \Big(\mathit{Swap}_{n-1}(x, x', y) \wedge \mathit{HasG}(x')\Big),\end{aligned}$$

$$\begin{aligned}\Psi_n^* \triangleq\ & \Psi_{n-1}^* \vee \exists x' \Big(\mathit{HasC}_{n-1}(x') \wedge \textstyle\bigvee_{i=0}^{n-2} \mathit{HasC}_i(x')\Big) \vee \\ & \exists x \exists y \exists y' \begin{pmatrix} \mathit{HasC}_{n-1}(x) \wedge \mathit{Edge}_{n-1}(x, y) \wedge \mathit{Edge}_{n-1}(x, y') \wedge \\ \mathit{SameC}_{n-2}(y, y') \wedge \mathit{OppG}(y, y') \end{pmatrix}.\end{aligned}$$

LEMMA 6.6. *For all $n \geq 1$, let $\Phi_n \triangleq \Theta_n^* \vee \Psi_n^*$.*

(1) *The length of Φ_n is exponential in n.*

(2) *Φ_n is logically equivalent to $\Theta_n \vee \Psi_n$.*

PROOF. Straightforward induction. □

We conclude this section with the proof of Theorem 6.1.

PROOF OF THEOREM 6.1. The sentence Φ_n of Lemma 6.6 has length exponential in n and yet (since Φ_n is equivalent to $\Theta_n \vee \Psi_n$) any existential-positive sentence equivalent to Φ_n has at least tower$(n-1)$ quantifiers by Lemma 6.5. □

7. EXTENSIONS AND OPEN QUESTIONS

We conclude by stating some corollaries of our main results and mentioning a few questions left open by our work.

7.1 Extensions of our Results

7.1.1 *Beyond finite relation vocabularies.* Both of our homomorphism preservation theorems (Theorems 4.12 and 5.16) were stated in terms of first-order sentences in a finite relational vocabulary σ. In fact, these results extend to first-order formulas in vocabularies consisting of arbitrarily many relation symbols as well as constant symbols. We first claim that the addition of finitely many constant symbols $c_1, \dots, c_k$ to σ is no problem, once definitions are suitably modified. *Structures* $\mathbf{A}$ by definition now include interpretations $c_1^{\mathbf{A}}, \dots, c_k^{\mathbf{A}} \in A$ for all constant symbols c_i. *Homomorphisms* $h : \mathbf{A} \longrightarrow \mathbf{B}$ now additionally satisfy $h(c_i^{\mathbf{A}}) = c_i^{\mathbf{B}}$ for all c_i. Elements $c_i^{\mathbf{A}}$ and $c_i^{\mathbf{B}}$ are now identified in the X-sum $\mathbf{A} \oplus_X \mathbf{B}$. Less obvious,$\mathrm{td}_X(\mathbf{A})$ is now defined as tree-depth of the graph $\mathcal{G}(\mathbf{A}) \setminus (X \cup \{c_1^{\mathbf{A}}, \dots, c_k^{\mathbf{A}}\})$. Other definitions, such as those of $\rightarrow_X^n$ and $\mathbf{Core}_X^n(\mathbf{A})$, remain essentially unchanged. It can be checked that all lemmas and proofs now go through exactly as before (when σ is a finite relational vocabulary), including the key proposition that $\mathscr{C}_X^n$ is finite for every n and finite set X (Proposition 3.9).

Having established that our results extend to finite vocabularies containing constant as well as relation symbols, there is a simple argument which extends our main results from first-order sentences to first-order formulas: a formula $\phi(x_1, \dots, x_k)$

is preserved under homomorphisms if, and only if, the corresponding sentence $\phi(c_1, \dots, c_k)$ in the vocabulary extended by constant symbols $c_1, \dots, c_k$ is preserved under homomorphisms. Finally, our main theorems extend to vocabularies with infinitely many relation and constant symbols for the simple reason that any first-order formula involves only finitely many symbols. Hence, it suffices to consider only the relevant finite fragment of the infinite vocabulary. (Of course, the upper bound $\rho(n)$ in the finite h.p.t. will depend on the particular relevant finite fragment.)

Extending our two homomorphism preservation theorems to vocabularies which contain function symbols (in addition to relation and constant symbols) appears to be a messier exercise. We believe that some version of both preservation theorems should hold once the notion of (quantifier-)rank is adjusted to account for nesting of function symbols as well as quantifiers.

7.1.2 *Relativized homomorphism preservation theorems.* The finite h.p.t. (Theorem 5.16) relativizes on any class $\mathcal{K}$ of finite structures which is co-homomorphism closed (meaning that $\mathbf{A} \in \mathcal{K}$ whenever $\mathbf{A} \to \mathbf{B}$ and $\mathbf{B} \in \mathcal{K}$ for all finite $\mathbf{A}$ and $\mathbf{B}$). One example is the class of finite k-partite structures (i.e., finite structures with k-partite Gaifman graphs).

Corollary 7.1. *Suppose $\mathcal{K}$ is a co-homomorphism closed class of finite structures. Then every first-order sentence of quantifier-rank n which is preserved under homomorphisms on $\mathcal{K}$ is logically equivalent on $\mathcal{K}$ to an existential-positive sentence of quantifier-rank $\rho(n)$.*

Proof. This claim boils down to a simple observation. Recall the following diagram from the proof of Theorem 5.15 (of which the finite h.p.t. is a corollary):

$$\begin{array}{ccccc} \widetilde{\mathbf{A}} & \equiv^n & \widetilde{\mathbf{C}} & & \\ \big\downarrow{\scriptstyle \text{retr}} & & \big\downarrow{\scriptstyle \text{retr}} & & \\ \mathbf{A} & \rightleftarrows^{\rho(n)} & \mathbf{C} & \to & \mathbf{B} \end{array}$$

The observation is that if finite structures $\mathbf{A}$, $\mathbf{B}$ are both in class $\mathcal{K}$, then so are finite structures $\widetilde{\mathbf{A}}$, $\mathbf{C}$, $\widetilde{\mathbf{C}}$ since $\mathcal{K}$ is co-homomorphism closed. □

The equirank h.p.t. similarly relativizes on co-homomorphism closed classes.

7.1.3 *Preservation theorems for primitive-positive sentences.* The equirank and finite homomorphism preservation theorems yield corresponding equirank and finite preservation theorems for primitive-positive sentences.

Corollary 7.2. *A first-order sentence of quantifier-rank n is preserved under homomorphisms and products* [*on finite structures*] *if, and only if, it is logically equivalent* [*in the finite*] *to a primitive-positive sentence of quantifier-rank n* [$\rho(n)$].

Proof. The ($\Longleftarrow$) direction, that every primitive-positive formula is preserved under homomorphisms and products, is an easy exercise. In the ($\Longrightarrow$) direction, let Φ be a first-order sentence of quantifier-rank n that is preserved under homomorphisms and products [on finite structures]. By the equirank [finite] h.p.t.,

Φ is equivalent to an existential-positive sentence of quantifier-rank n $[\rho(n)]$. Finally, every existential-positive sentence that is preserved under products has an equivalent primitive-positive sentence with the same quantifier-rank (another easy exercise). □

7.2 Failure of the Classical Homomorphism Interpolation Theorem on Finite Structures

The classical h.p.t. can be seen as a special case of a more general interpolation theorem.

THEOREM 7.3 (CLASSICAL HOMOMORPHISM INTERPOLATION THEOREM). *Let Φ_1 and Φ_2 be first-order sentences and suppose that for all structures $\mathbf{A}$ and $\mathbf{B}$, if $\mathbf{A} \models \Phi_1$ and $\mathbf{A} \to \mathbf{B}$, then $\mathbf{B} \models \Phi_2$. Then there exists an existential-positive "interpolant" between Φ_1 and Φ_2, that is, an existential-positive sentence Ψ such that* $\mathrm{Mod}(\Phi_1) \subseteq \mathrm{Mod}(\Psi) \subseteq \mathrm{Mod}(\Phi_2)$.

The classical h.p.t. is precisely the special case of Theorem 7.3 where $\Phi_1 = \Phi_2$. In light of the finite h.p.t., it is natural to ask whether Theorem 7.3 also survives when restricted to finite structure. Eric Rosen and the author discovered that, unlike the classical h.p.t., but like other classical interpolation theorems, Theorem 7.3 indeed fails on finite structures.

THEOREM 7.4 (FAILURE ON FINITE STRUCTURES). *There exist first-order sentences Φ_1 and Φ_2 such that*

- *for all finite structures $\mathbf{A}$ and $\mathbf{B}$, if $\mathbf{A} \models \Phi_1$ and $\mathbf{A} \to \mathbf{B}$ then $\mathbf{B} \models \Phi_2$, and*
- *there is no existential-positive sentence Ψ such that* $\mathrm{Mod_{fin}}(\Phi_1) \subseteq \mathrm{Mod_{fin}}(\Psi) \subseteq \mathrm{Mod_{fin}}(\Phi_2)$.

PROOF. Consider the vocabulary consisting of a binary relation E and a unary relation P. We regard structures in this vocabulary as directed graphs with a distinguished subset of vertices (defined by P). Let Φ_1 be a sentence expressing "E is anti-reflexive, symmetric and 2-regular (i.e., the edge relation of a 2-regular simple graph) and $|\{x, y\} \cap P| = 1$ for all but a unique undirected edge $\{x, y\}$". Let Φ_2 express "either there exist at least three vertices, or some vertex has a self-loop". Both Φ_1 and Φ_2 are clearly first-order statements.

Notice that every finite model of Φ_1 contains a unique odd cycle (however, Φ_1 has infinite models without odd cycles). Let $\mathbf{A}$ and $\mathbf{B}$ be finite structures and suppose that $\mathbf{A} \models \Phi_1$ and $\mathbf{A} \to \mathbf{B}$. We claim that $\mathbf{B} \models \Phi_2$. We may assume that $\mathbf{B}$ has at most two vertices, since otherwise $\mathbf{B}$ clearly satisfies Φ_2. But then two consecutive vertices in the odd cycle of $\mathbf{A}$ must map to the same vertex of $\mathbf{B}$ under any homomorphism from $\mathbf{A}$ to $\mathbf{B}$. It follows that this vertex of $\mathbf{B}$ contains a self-loop and hence $\mathbf{B} \models \Phi_2$.

It remains to show that Φ_1 and Φ_2 have no existential-positive interpolant on finite structures. Toward a contradiction, assume that Ψ is an existential-positive sentence such that $\mathrm{Mod_{fin}}(\Phi_1) \subseteq \mathrm{Mod_{fin}}(\Psi) \subseteq \mathrm{Mod_{fin}}(\Phi_2)$. Let n be the quantifier-rank of Ψ. Let $\mathbf{A}$ be the structure with universe $A = \{0, \dots, 2^n\}$ and relations $E^{\mathbf{A}} = \{(i, j) : i - j \equiv 1 \text{ or } -1 \text{ modulo } 2^n + 1\}$ and $P^{\mathbf{A}} = \{0, 2, 4, \dots, 2^n - 2, 2^n\}$. The underlying simple graph of $\mathbf{A}$ (throwing away unary relation $P^{\mathbf{A}}$) is a $(2^n + 1)$-cycle and hence 2-regular. With the only exception of $\{0, 2^n\}$, every other undirected edge $\{i, i{+}1\}$ satisfies $|\{i, i{+}1\} \cap P^{\mathbf{A}}| = 1$. So we see that $\mathbf{A} \models \Phi_1$ and consequently $\mathbf{A} \models \Psi$ since $\mathrm{Mod}_{\mathrm{fin}}(\Phi_1) \subseteq \mathrm{Mod}_{\mathrm{fin}}(\Psi)$. Let $\mathbf{B}$ be the two-element structure with $B = P^{\mathbf{B}} = \{0,1\}$ and $E^{\mathbf{B}} = \{(0,1),(1,0)\}$. One can show that $\mathbf{A} \rightarrow^n \mathbf{B}$ (in fact, $\mathbf{B} \cong \mathbf{Core}^n(\mathbf{A})$) by arguing that if $\mathrm{td}(\mathbf{C}) \leq n$ and $\mathbf{C} \rightarrow \mathbf{A}$, then $\mathcal{G}(\mathbf{C})$ is bipartite, the two parts describing a homomorphism $\mathbf{C} \rightarrow \mathbf{B}$ (details are left to the reader). It follows that $\mathbf{B} \models \Psi$ and hence $\mathbf{B} \models \Phi_2$ as $\mathrm{Mod}_{\mathrm{fin}}(\Psi) \subseteq \mathrm{Mod}_{\mathrm{fin}}(\Phi_2)$. But $\mathbf{B}$ has fewer than three vertices and yet contains no self-loop, which implies $\mathbf{B} \models \neg\Phi_2$, yielding a contradiction. Therefore, no such Ψ exists. □

7.3 Open Questions

Our work raises an obvious question about the status of the equirank h.p.t. on finite structures.

Question 7.5. What is the minimal function $\rho : \mathbb{N} \longrightarrow \mathbb{N}$ for which the finite h.p.t. (Theorem 5.16) is valid? In particular, does the finite h.p.t. hold with $\rho(n) = n$? Equivalently, is the equirank h.p.t. (Theorem 4.12) valid on finite structures?

Recall that every structure has an infinite n-extendable co-retract for all $n \in \mathbb{N}$ (Lemma 4.9).

Question 7.6. Is it true that for all $n \in \mathbb{N}$, every finite structure has a finite n-extendable co-retract?

It follows from the proof of Theorem 4.11 that if the answer to Question 7.6 is "yes", then the equirank h.p.t. holds on finite structures.

Finally, we would like to know whether the equirank h.p.t. has analogues among other classical preservation theorems.

Question 7.7. Which classical preservation theorems besides the h.p.t. have valid "equirank" versions?

ACKNOWLEDGMENTS

I am grateful to Anuj Dawar, Ron Fagin, Alan Nash, Jarik Nešetřil and Madhu Sudan for helpful comments on various drafts of this article, to Eric Rosen for discussions resulting in Theorem 7.4 (failure of the classical homomorphism interpolation theorem) and to Yuri Gurevich and Saharon Shelah for granting me access to unpublished notes which provided the basis for the results of §6. Not least, I thank the anonymous referees for their many helpful suggestions.

REFERENCES

Abiteboul, S., Hull, R., and Vianu, V. 1995. *Foundations of Databases*. Addison-Wesley.

Ajtai, M. and Gurevich, Y. 1987. Monotone versus positive. *J. ACM 34*, 1004–1015.

Ajtai, M. and Gurevich, Y. 1994. Datalog vs. first-order logic. *J. Comput. Syst. Sci. 49*, 562–588.

Alechina, N. and Gurevich, Y. 1997. Syntax vs. semantics on finite structures. In *Structures in Logic and Computer Science*, J. Mychielski, G. Rozenberg, and A. Salomaa, Eds. Springer-Verlag, 14–33.

Atserias, A. 2005. On digraph coloring problems and treewidth duality. In *Proceedings of the 20th IEEE Symposium on Logic in Computer Science*. 106–115.

Atserias, A., Dawar, A., and Kolaitis, P. 2006. On preservation under homomorphisms and unions of conjunctive queries. *J. ACM 53,* 2, 208–237.

BALDWIN, J. 2000. Finite and infinite model theory – a historical perspective. *Log. J. IGPL 8,* 5, 605–628.

BLASS, A. AND ROSSMAN, B. 2005. Explicit graphs with extension properties. *Bull. Eur. Assoc. Theor. Comput. Sci.* 86, 166–175.

CHANDRA, A. AND MERLIN, P. 1977. Optimal implementation of conjunctive queries in relational databases. In *Proceedings of the 9th ACM Symposium on Theory of Computing.* 77–90.

CODD, E. 1970. A relational model of data for large shared data banks. *Commun. ACM 13,* 6, 377–387.

DEOGUN, J., KLOKS, T., KRATSCH, D., AND MULLER, H. 1994. On vertex ranking for permutation and other graphs. In *Proceedings of the 11th Symposium on Theoretical Aspects of Computer Science.* Lecture Notes in Computer Science, vol. 775. Springer, 747–758.

EBBINGHAUS, H. AND FLUM, J. 1996. *Finite Model Theory.* Springer-Verlag.

FEDER, T. AND VARDI, M. 2003. Homomorphism closed vs. existential positive. In *Proceedings of the 18th IEEE Symposium on Logic in Computer Science.* 310–320.

FEFERMAN, S. AND VAUGHT, R. L. 1959. The first order properties of products of algebraic systems. *Fund. Math. 47*, 57–103.

GRÄDEL, E., KOLAITIS, P., LIBKIN, L., MARX, M., SPENCER, J., VARDI, M., VENEMA, Y., AND WEINSTEIN, S. 2007. *Finite Model Theory and its Applications.* Springer.

GRÄDEL, E. AND ROSEN, E. 1999. On preservation theorems for two-variable logic. *Math. Logic Quart. 45*, 315–325.

GUREVICH, Y. 1984. Toward logic tailored for computational complexity. In *Computation and Proof Theory*, M. M. R. et al., Ed. Springer Lecture Notes in Mathematics, 175–216.

GUREVICH, Y. 1988. Logic and the challenge of computer science. In *Current Trends in Theoretical Computer Science*, E. Boerger, Ed. 1–57.

GUREVICH, Y. 1990. On finite model theory (extended abstract). In *Feasible Mathematics*, S. R. Buss and P. J. Scott, Eds. Birkhäuser, 211–219.

HELL, P. AND NEŠETŘIL, J. 1992. The core of a graph. *Discrete Math. 109*, 117–126.

HELL, P. AND NEŠETŘIL, J. 2004. *Graphs and Homomorphisms.* Oxford University Press.

HODGES, W. 1993. *Model Theory.* Cambridge University Press.

KOLAITIS, P. AND VARDI, M. 2000. Conjunctive query containment and constraint satisfaction. *J. Comput. Syst. Sci. 61*, 302–332.

KOLAITIS, P. G. 1993. A tutorial on finite model theory (abstract). In *Proceedings of the 8th IEEE Symposium on Logic in Computer Science.* 122.

LIBKIN, L. 2004. *Elements of Finite Model Theory.* Springer-Verlag.

LYNDON, R. 1959. Properties preserved under homomorphism. *Pacific J. Math. 9*, 129–142.

NEŠETŘIL, J. AND DE MENDEZ, P. O. 2006. Tree depth, subgraph coloring and homomorphism bounds. *European J. Combin. 27,* 6, 1022–1041.

NEŠETŘIL, J. AND DE MENDEZ, P. O. 2007. Grad and classes with bounded expansion I. Decompositions. *European J. Combin.*. (to appear).

NEŠETŘIL, J. AND SHELAH, S. 2003. On the order of countable graphs. *European J. Combin. 24*, 649–663.

ROSEN, E. 1995. Finite model theory and finite variable logic. Ph.D. thesis, University of Pennsylvania.

ROSEN, E. 2002. Some aspects of model theory and finite structures. *Bull. Symbolic Logic 8*, 380–403.

ROSEN, E. AND WEINSTEIN, S. 1995. Preservation theorems in finite model theory. In *Logic and Computational Complexity*, D. Leivant, Ed. Spring-Verlag, 480–502.

ROSSMAN, B. 2005. Existential positive formulas and preservation under homomorphisms. In *Proceedings of the 20th IEEE Symposium on Logic in Computer Science.* 467–476.

STOLBOUSHKIN, A. 1995. Finite monotone properties. In *Proceedings of the 10th IEEE Symposium on Logic in Computer Science.* 324–330.

TAIT, W. 1959. A counterexample to a conjecture of Scott and Suppes. *J. Symbolic Logic 24*, 15–16.

Received Month Year; revised Month Year; accepted Month Year