



University of Pennsylvania
ScholarlyCommons

IRCS Technical Reports Series

Institute for Research in Cognitive Science

March 1995

Preservation Theorems in Finite Model Theory

Eric Rosen
University of Pennsylvania

Scott Weinstein
University of Pennsylvania, weinstein@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/ircs_reports

Rosen, Eric and Weinstein, Scott, "Preservation Theorems in Finite Model Theory" (1995). *IRCS Technical Reports Series*. 119.
https://repository.upenn.edu/ircs_reports/119

University of Pennsylvania Institute for Research in Cognitive Science Technical Report No. IRCS-95-04.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ircs_reports/119
For more information, please contact repository@pobox.upenn.edu.

Preservation Theorems in Finite Model Theory

Abstract

We develop various aspects of the finite model theory of L^k (there exists) and $L^k_{\infty\omega}$ (there exists). We establish the optimality of normal forms for $L^k_{\infty\omega}$ (there exists) over the class of finite structures and demonstrate separations over the class of finite structures and demonstrate separations among descriptive complexity classes within $L^k_{\infty\omega}$ (there exists). We establish negative results concerning preservation theorems for L^k (there exists) and $L^k_{\infty\omega}$ (there exists). We introduce a generalized notion of preservation theorem and establish some positive results concerning "generalized preservation theorems" for first-order definable classes of finite structures which are closed under extensions.

Comments

University of Pennsylvania Institute for Research in Cognitive Science Technical Report No. IRCS-95-04.

The Institute For Research In Cognitive Science

Preservation Theorems in Finite Model Theory

(To appear in Daniel Leivant, editor,
Logic and Computational Complexity, Springer-Verlag)


by

Eric Rosen
Scott Weinstein

University of Pennsylvania
3401 Walnut Street, Suite 400C
Philadelphia, PA 19104-6228

March 1995

Site of the NSF Science and Technology Center for
Research in Cognitive Science



Preservation Theorems in Finite Model Theory*

Eric Rosen** and Scott Weinstein***

Department of Philosophy
University of Pennsylvania
Philadelphia PA 19104, USA

Abstract. We develop various aspects of the finite model theory of $L^k(\exists)$ and $L^k_{\infty\omega}(\exists)$. We establish the optimality of normal forms for $L^k_{\infty\omega}(\exists)$ over the class of finite structures and demonstrate separations among descriptive complexity classes within $L^k_{\infty\omega}(\exists)$. We establish negative results concerning preservation theorems for $L^k(\exists)$ and $L^k_{\infty\omega}(\exists)$. We introduce a generalized notion of preservation theorem and establish some positive results concerning “generalized preservation theorems” for first-order definable classes of finite structures which are closed under extensions.

1 Introduction

In this paper we investigate the status of preservation theorems in finite model theory. We focus our attention on classes of finite structures which are closed under extensions and their definability in fragments of the infinitary language $L^{\omega}_{\infty\omega}$. The language $L^{\omega}_{\infty\omega}$ was introduced by Barwise [4] in connection with the investigation of inductive definability over infinite structures. Recently, the study of $L^{\omega}_{\infty\omega}$ has played a central role in analyzing the behavior of fixed-point logics over the class of finite structures (see [5, 13]). Of particular interest from the point of view of our current investigation are the works of Kolaitis and Vardi [12] and Afrati, Cosmadakis, and Yannakakis [1] which exploit existential fragments of $L^{\omega}_{\infty\omega}$ in analyzing the expressive power of Datalog.

The starting point for our investigation is the well-known failure of the preservation theorem of Los and Tarski over finite structures. Recall that the Los-Tarski Theorem states that any first-order definable class of structures which is closed under extensions is definable by a first-order existential sentence. Scott and Suppes conjectured that this theorem generalizes to the finite case, that is, if $\text{Mod}_f(\varphi)$ (the collection of finite models of the first-order sentence φ) is closed under extensions, then $\text{Mod}_f(\varphi) = \text{Mod}_f(\psi)$, for some first-order existential sentence ψ . Tait [18] showed that this conjecture fails; Gurevich and Shelah [9, 10] gave simpler counterexamples employing universal-existential first-order sentences.

* We would like to thank Maria Bonet, Yuri Gurevich, and Steven Lindell for valuable discussions on the subject of this paper.

** Supported in part by NSF-STC SBR-8920230.

*** Supported in part by NSF CCR-9403447.

In light of the failure of the Los-Tarski Theorem over finite structures, it is natural to inquire whether “generalized preservation theorems” might hold in the finite case. In this paper, we investigate the prospects for such a positive approach to preservation properties in the context of finite model theory. In particular, we examine generalized versions of ordinary preservation theorems where an algebraic restriction on a class of structures definable in a given language yields information about the syntactic structure of formulas which define that class in an *extension* of that language. In this spirit, we show that for certain classes of first-order sentences Φ , if $\varphi \in \Phi$ and $\text{Mod}_f(\varphi)$ is closed under extensions, then $\text{Mod}_f(\varphi) = \text{Mod}_f(\psi)$ for some ψ in the existential fragment of $L_{\infty\omega}^\omega$ (or even in $\text{Datalog}(\neq, \neg)$). In contrast, we also establish the failure of the analog of the Los-Tarski Theorem for $L_{\infty\omega}^\omega$ itself, both over finite structures and over arbitrary structures. That is, we show that there is a sentence φ of $L_{\infty\omega}^\omega$ such that both $\text{Mod}_f(\varphi)$ and $\text{Mod}(\varphi)$ are closed under extensions, but neither of these classes is definable by an existential sentence of $L_{\infty\omega}^\omega$.

The paper proceeds as follows. The next section introduces the languages we will study and establishes a simple proposition which characterizes the relative expressive power of their existential fragments. Section 3 develops some finite model theory for the existential fragments of L^k and $L_{\infty\omega}^k$. In particular, we establish the optimality of a normal form for the existential fragment of $L_{\infty\omega}^k$ over finite structures and demonstrate separations among descriptive complexity classes within $L_{\infty\omega}^k(\exists)$. In Section 4, we prove the failure of existential preservation for $L_{\infty\omega}^\omega$. Section 5 is devoted to establishing positive results concerning generalized preservation theorems for fragments of first-order logic over finite structures. In the final section, we discuss a number of open problems and present without proof some related results concerning preservation under homomorphisms. A full treatment of these results will appear in [17].

2 Preliminaries

Let \mathcal{F}_σ be the collection of finite structures of signature σ . We will assume that the universe of any $A \in \mathcal{F}_\sigma$ is an initial segment of $N = \{0, 1, 2, \dots\}$. We will often use A, B, \dots etc. to denote both a structure and its universe when no confusion is likely to result. We assume that the signature σ is finite and contains no function symbols; we suppress mention of σ when no confusion is likely to result. A *boolean query* $\mathcal{C} \subseteq \mathcal{F}$ is a class of finite structures that is closed under isomorphisms. We use \mathcal{C} to range over boolean queries. In what follows, we will focus attention on boolean queries which are closed under extensions.

Definition 1. $\text{EXT} = \{\mathcal{C} \subseteq \mathcal{F} \mid \forall A, B \in \mathcal{C}, \text{ if } A \in \mathcal{C} \text{ and } A \subseteq B, \text{ then } B \in \mathcal{C}\}$.

Let L be a logical language and let φ be a sentence of L . $\text{Mod}(\varphi) = \{A \mid A \models \varphi\}$ is the *L-class* determined by φ and $\text{Mod}_f(\varphi) = \{A \in \mathcal{F} \mid A \models \varphi\}$ is the boolean query expressed by φ . We say that \mathcal{C} is *L-definable*, just in case it is the boolean query expressed by some sentence $\varphi \in L$. We will often use L to denote the set of *L-definable* boolean queries. We let FO denote first-order logic, $L_{\infty\omega}$,

the usual infinitary extension of first-order logic which allows conjunction and disjunction over arbitrary sets of formulas, L^k , the fragment of FO consisting of those formulas all of whose variables both free and bound are among x_1, \dots, x_k , and similarly $L_{\infty\omega}^k$, the k -variable fragment of $L_{\infty\omega}$; $L_{\infty\omega}^\omega = \bigcup_{k \in \omega} L_{\infty\omega}^k$. We let $\text{FO}(\exists)$ denote the set of existential formulas of FO, that is, those formulas obtained by closing the set of atomic formulas and negated atomic formulas under the operations of conjunction, disjunction, and existential quantification. We define $L_{\infty\omega}(\exists)$, the set of existential formulas of $L_{\infty\omega}$, similarly, but require, in addition, closure under infinitary conjunction and disjunction. We let $L^k(\exists)$ consist of the formulas common to $\text{FO}(\exists)$ and L^k and we define $L_{\infty\omega}^k(\exists)$ and $L_{\infty\omega}^\omega(\exists)$ similarly. A Datalog(\neq, \neg) program P is a collection of rules of the form

$$\eta_0 \leftarrow \eta_1, \dots, \eta_k.$$

Such a rule has a *head*, η_0 , and a *body*, η_1, \dots, η_k . Each of the η_i is either an inequality or a literal over the signature $\sigma \cup \tau$ where σ and τ are disjoint; σ consists of the *extensional* relations and constants of P and τ consists of the *intensional* relations of P . The heads of all rules are built from intensional relations and intensional relations occur only positively throughout P . The program contains a distinguished intensional relation R of arity $n \geq 0$ and determines an n -ary query over structures in \mathcal{F}_σ . The value of this query for a given $A \in \mathcal{F}_\sigma$ is the value of R when the program is viewed as determining least-fixed points for each of the intensional relations with respect to a simultaneous induction associated with the program. The reader may consult [1, 12] for further details and discussion. As with logics, we use Datalog(\neq, \neg) to refer to the class of queries computed by Datalog(\neq, \neg) programs as well as to the class of programs themselves. Datalog programs are defined similarly except that all the η_i are restricted to be positive literals, even those built from extensional relations. Observe that Datalog(\neq, \neg) is contained in the least fixed-point extension of first-order logic (FO+LFP).

In our current notation, the failure of the Los-Tarski Theorem over finite structures may be expressed as:

$$\text{FO} \cap \text{EXT} \not\subseteq \text{FO}(\exists).$$

This raises the question of whether $\text{FO} \cap \text{EXT}$ is contained in the existential fragment of some stronger logic. The following proposition completely characterizes the relative expressive power of the existential fragments of the logics in which we are interested.

Proposition 2.

$$\text{FO}(\exists) \subset \text{Datalog}(\neq, \neg) \subset L_{\infty\omega}^\omega(\exists) \subset L_{\infty\omega}(\exists) = \text{EXT}.$$

Proof. It is easy to see that every query in $\text{FO}(\exists)$ can be expressed by a program in Datalog(\neq, \neg) which makes use of no recursion. It is well-known that this inclusion is strict, for example, the query (s, t) -connectivity is expressible in Datalog but not in FO. The inclusion of Datalog(\neq, \neg) in $L_{\infty\omega}^\omega(\exists)$ has been noted by Afrati, Cosmadakis, and Yannakakis [1] (see also [12]); the argument to show

this is a variant of the proof that least fixed-point logic is contained in $L_{\infty\omega}^\omega$ over the class of finite structures (see [14]). Afrati, Cosmadakis, and Yannakakis [1] also exhibit queries which witness the separation of $\text{Datalog}(\neq, \neg)$ and $L_{\infty\omega}^\omega(\exists)$, even over the class of polynomial time computable queries. The identity between $L_{\infty\omega}(\exists)$ and EXT has been noted by Kolaitis and independently by Lo (see [1] and [15]). Finally, it is easy to construct polynomial time computable boolean queries in EXT which are not in $L_{\infty\omega}^\omega$. For example, let \mathcal{C} be the query over the signature $\{E, s, t\}$ of source-target graphs that says that there is an E -path from s to t whose length is less than half the cardinality of the structure. It is clear that $\mathcal{C} \in \text{EXT}$. It is also easy to verify that \mathcal{C} is not in $L_{\infty\omega}^\omega$ (and therefore not in $L_{\infty\omega}^\omega(\exists)$) by a straightforward application of the k -pebble Ehrenfeucht-Fraïssé game which we review below. ■

The above proposition together with the failure of the Los-Tarski Theorem in the finite case suggests the following questions.

1. Is $\text{FO} \cap \text{EXT} \subseteq L_{\infty\omega}^\omega(\exists)$?
2. Is $\text{FO} \cap \text{EXT} \subseteq \text{Datalog}(\neq, \neg)$?
3. Is $L_{\infty\omega}^\omega \cap \text{EXT} \subseteq L_{\infty\omega}^\omega(\exists)$?

Clearly a positive answer to the second or third question would imply a positive answer to the first. In Section 4, we provide a negative answer to the third question. In Section 5, we provide partial positive answers to the first and second questions. Before proceeding to these results, we develop some of the finite model theory of $L^k(\exists)$ and $L_{\infty\omega}^k(\exists)$ in the next section.

3 Basic Finite Model Theory for $L^k(\exists)$ and $L_{\infty\omega}^k(\exists)$

In this section, we present some basic model theory for L^k , $L_{\infty\omega}^k$, $L^k(\exists)$, and $L_{\infty\omega}^k(\exists)$. After a brief discussion of game-theoretic characterizations of equivalence and definability in these languages, we proceed to consider questions of finite axiomatizability and normal forms.

Let L be one of the logical languages we consider. Given a structure A , the L -theory of A is the collection of sentences of L which are satisfied by A . We say that A is L -equivalent to B , if and only if, the L -theory of A is equal to the L -theory of B and we say that A is L -compatible with B , if and only if, the L -theory of A is contained in the L -theory of B . Note that if L is closed under negation, then the relations of L -equivalence and L -compatibility coincide, whereas for languages like $L^k(\exists)$ and $L_{\infty\omega}^k(\exists)$ these relations are distinct. We use the notations \equiv^k , $\equiv_{\infty\omega}^k$, \preceq^k , and $\preceq_{\infty\omega}^k$ for L^k -equivalence, $L_{\infty\omega}^k$ -equivalence, $L^k(\exists)$ -compatibility, and $L_{\infty\omega}^k(\exists)$ -compatibility, respectively. The main tool for studying these relations are refinements of the Ehrenfeucht-Fraïssé game. Barwise [4] characterized $L_{\infty\omega}^k$ -equivalence in terms of partial isomorphisms, while Immerman [11] and Poizat [16] provided related pebble game characterizations of L^k -equivalence. Kolaitis and Vardi [12] characterized compatibility in the negation free fragment of $L_{\infty\omega}^k(\exists)$ both in terms of collections of partial *homomorphisms* as well as in

terms of a one-sided, positive version of the pebble game. Below we use a minor variant of the approach in [12] to characterize $L_{\infty\omega}^k(\exists)$ -compatibility.

A collection I of partial isomorphisms from A to B is said to have the k -[back-and-]forth property if for all $f \in I$ such that the domain of f has cardinality $< k$, and all $a \in A$ [$b \in B$], there is a function $g \in I$ such that $f \subseteq g$ and $a \in \text{dom}(g)$ [$b \in \text{rng}(g)$]. (That is, the k -forth property is the one-sided version, going forth from A , of the k -back-and-forth property.)

Barwise [4] proved the following proposition which gives an algebraic characterization of $L_{\infty\omega}^k$ -equivalence.

Proposition 3 (Barwise [4]). *Let A and B be structures of signature σ and let h be the map with $\text{dom}(h) = \{c^A \mid c \in \sigma\}$ such that $h(c^A) = c^B$ for all $c \in \sigma$. The following conditions are equivalent.*

1. $A \equiv_{\infty\omega}^k B$.
2. *There is a non-empty set I of partial isomorphisms from A to B such that*
 - (a) *I is closed under subfunctions;*
 - (b) *I has the k -back-and-forth property;*
 - (c) *for all $f \in I$, $f \cup h$ is a partial isomorphism from A to B .*

In a similar spirit, Kolaitis and Vardi [12] gave an algebraic characterization of the compatibility relation for the negation free fragment of $L_{\infty\omega}^k(\exists)$ in terms of collections of partial *homomorphisms* with the k -forth property. We adapt their approach to the case of $L_{\infty\omega}^k(\exists)$ in the following theorem.

Proposition 4 (Kolaitis and Vardi [12]). *Let A and B be structures of signature σ and let h be the map with $\text{dom}(h) = \{c^A \mid c \in \sigma\}$ such that $h(c^A) = c^B$ for all $c \in \sigma$. The following conditions are equivalent.*

1. $A \preceq_{\infty\omega}^k B$.
2. *There is a non-empty set I of partial isomorphisms from A to B such that*
 - (a) *I is closed under subfunctions;*
 - (b) *I has the k -forth property;*
 - (c) *for all $f \in I$, $f \cup h$ is a partial isomorphism from A to B .*

Both Propositions 3 and 4 can be expressed more colorfully in terms of pebble games. This approach to L^k -equivalence was introduced by Immerman [11] and Poizat [16] and as an approach to $L_{\infty\omega}^k(\exists)$ -compatibility by Kolaitis and Vardi [12]. In order to state the relevant results in a suitably refined form, we require the notion of the *quantifier rank* of a formula. We state this definition for formulas of $L_{\infty\omega}$ since all the languages we consider are fragments of it.

Definition 5. The quantifier rank of $\varphi \in L_{\infty\omega}$ ($\text{qr}(\varphi)$) is defined by the following induction.

1. $\text{qr}(\varphi) = 0$ if φ is atomic;
2. $\text{qr}(\neg\varphi) = \text{qr}(\varphi)$;
3. $\text{qr}(\bigwedge \Phi) = \text{qr}(\bigvee \Phi) = \sup(\{\text{qr}(\varphi) \mid \varphi \in \Phi\})$;

$$4. \text{qr}(\exists x\varphi) = \text{qr}(\forall x\varphi) = \text{qr}(\varphi) + 1.$$

The n -round, k -pebble Ehrenfeucht-Fraïssé game on A and B is played between two players, Spoiler and Duplicator, with k pairs of pebbles, $(\alpha_1, \beta_1), \dots, (\alpha_k, \beta_k)$. The Spoiler begins each round by choosing a pair of pebbles (α_i, β_i) that may or may not be in play on the boards A and B . He (by convention, the Spoiler is male, the Duplicator female) either places α_i on an element of A , or β_i on an element of B . The Duplicator then plays the remaining pebble on the other model. The Spoiler wins the game if after any round $m \leq n$ the function f from A to B , which sends the element pebbled by α_i to the element pebbled by β_i and preserves the denotations of constants, is not a partial isomorphism; otherwise, the Duplicator wins the game. The n -round \exists^k -game is the one-sided version of the n -round, k -pebble Ehrenfeucht-Fraïssé game in which the Spoiler is restricted to play a pebble α_i into A at every round while the Duplicator responds by playing β_i into B ; the winning condition remains the same. Both the k -pebble Ehrenfeucht-Fraïssé game and its one-sided variant have infinite versions, which we call the *eternal* k -pebble Ehrenfeucht-Fraïssé game and the *eternal* \exists^k -game. In these games, the play continues through a sequence of rounds of order type ω . The Spoiler wins the game, if and only if, he wins at the n^{th} -round for some $n \in \omega$ as above; otherwise, the Duplicator wins. In describing the play of pebble games below, we will often use S to refer to the Spoiler and D to refer to the Duplicator. We will also often use α_i, β_i , etc. to refer to both pebbles and the elements they pebble at a given round of play.

The foregoing n -round games may be used to characterize equivalence and compatibility of structures with respect to L^k sentences and $L^k(\exists)$ sentences of quantifier rank n , and the eternal games may be used to characterize equivalence and compatibility of structures with respect to $L_{\infty\omega}^k$ sentences and $L_{\infty\omega}^k(\exists)$ sentences. Given structures A and B we let $A \equiv^{k,n} B$, if and only if, A and B satisfy the same sentences of L^k of quantifier rank $\leq n$ and we let $A \preceq^{k,n} B$, if and only if, every sentence of $L^k(\exists)$ of quantifier rank $\leq n$, which is true in A , is also true in B . The following two propositions use the n -round pebble games to characterize these relations. The first is due to Immerman [11] and Poizat [16] and the second is essentially due to Kolaitis and Vardi [12].

Proposition 6 (Immerman [11], Poizat [16]). *For all structures A and B , the following conditions are equivalent.*

1. $A \equiv^{k,n} B$.
2. *The Duplicator has a winning strategy for the n -round, k -pebble Ehrenfeucht-Fraïssé game on A and B .*

Proposition 7 (Kolaitis and Vardi [12]). *For all structures A and B , the following conditions are equivalent.*

1. $A \preceq^{k,n} B$.
2. *The Duplicator has a winning strategy for the n -round \exists^k -game on A and B , with the Duplicator playing on B .*

The next proposition gives a characterization of the infinitary equivalence and compatibility relations in terms of the eternal games. It is essentially due to Kolaitis and Vardi [14, 12].

- Proposition 8 (Kolaitis and Vardi [14, 12]).** 1. For all structures A and B , the following conditions are equivalent.
- (a) $A \equiv_{\infty\omega}^k B$.
 - (b) The Duplicator has a winning strategy for the eternal k -pebble Ehrenfeucht-Fraïssé game on A and B .
2. For all structures A and B , the following conditions are equivalent.
- (a) $A \preceq_{\infty\omega}^k B$.
 - (b) The Duplicator has a winning strategy for the eternal \exists^k -game on A and B , with the Duplicator playing on B .

Kolaitis and Vardi [14, 12] observed that over finite structures infinitary equivalence and compatibility coincide with their finitary analogs.

- Proposition 9 (Kolaitis and Vardi [14, 12]).** 1. Let A or B be a finite structure. Then, the following conditions are equivalent.
- (a) $A \equiv_{\infty\omega}^k B$.
 - (b) $A \equiv^k B$.
2. Let B be a finite structure. Then, the following conditions are equivalent.
- (a) $A \preceq_{\infty\omega}^k B$.
 - (b) $A \preceq^k B$.

The foregoing propositions yield the following corollaries concerning definability.

Proposition 10 (Kolaitis and Vardi [12]). For all $\mathcal{C} \subseteq \mathcal{F}$, the following conditions are equivalent.

1. \mathcal{C} is $L_{\infty\omega}^k(\exists)$ -definable.
2. For all $A \in \mathcal{C}$ and $B \notin \mathcal{C}$, $A \not\equiv_{\infty\omega}^k B$.
3. For all $A \in \mathcal{C}$ and $B \notin \mathcal{C}$, $A \not\preceq_{\infty\omega}^k B$.
4. For all $A \in \mathcal{C}$ and $B \notin \mathcal{C}$, there is an $n \in \omega$ such that the Spoiler has a winning strategy for the n -round \exists^k -game on A and B with the Spoiler playing on A .

Let L and L' be logical languages and let T be a collection of sentences of L . We say that T is *finitely axiomatizable in L'* , if and only if, there is a sentence $\varphi \in L'$ such that $\text{Mod}_f(T) = \text{Mod}_f(\varphi)$. Dawar, Lindell and Weinstein [5] prove that the $L_{\infty\omega}^k$ -theory of any finite model is finitely axiomatizable in L^k . As a corollary, they obtain a simple normal form for $L_{\infty\omega}^k$ over \mathcal{F} , in particular, they show that every sentence of $L_{\infty\omega}^k$ is equivalent to a countable disjunction of sentences of L^k and is also equivalent to a countable conjunction of sentences of L^k . In contrast, we show below that there are finite models whose $L^k(\exists)$ -theories are not finitely axiomatizable in $L^k(\exists)$. Building on this result, we prove that the normal form for $L_{\infty\omega}^k$ over \mathcal{F} (every sentence of $L_{\infty\omega}^k$ is equivalent over \mathcal{F} to a

countable disjunction of countable conjunctions of sentences of L^k) exhibited by Kolaitis and Vardi [14] is optimal when considered as a normal form for $L^k_{\infty\omega}(\exists)$ sentences over $L^k(\exists)$.

We begin by proving that there are models whose $L^k(\exists)$ -theories are not finitely axiomatizable in $L^k(\exists)$. Our argument exploits the *k-extension axioms*, which we now describe briefly. Let σ be a purely relational, finite signature. A *basic k-type* π over the signature σ is a maximal consistent set of literals over σ in the variables x_1, \dots, x_k . A *k-extension axiom* of signature σ is a sentence of the form $\forall x_1 \dots x_{k-1} \exists x_k (\bigwedge \pi \rightarrow \bigwedge \pi')$, where π is a basic $(k-1)$ -type of signature σ , π' is a basic k -type of signature σ , and $\pi \subseteq \pi'$. Over a fixed signature σ , the *k-Gaifman theory*, Γ_k , is the set of all k -extension axioms of signature σ . It is easy to see that, for each k , there are only finitely many k -extension axioms. Gaifman [7] showed that the theory $T = \bigcup_k \Gamma_k$ axiomatizes an ω -categorical model called the *random structure*. Fagin [6] proved the 0-1 law for first-order logic by showing that every extension axiom is *almost surely true* over \mathcal{F} . Fagin's result implies that almost every $A \in \mathcal{F}$ satisfies the k -Gaifman theory. Immerman [11] showed that any two models of the k -Gaifman theory are L^k -equivalent and Kolaitis and Vardi [14] made use of the k -Gaifman theory in their proof of the 0-1 law for $L^{\omega}_{\infty\omega}$. We make the following easy observation.

Proposition 11. *Let $A \models \Gamma_k$, and let B be any (finite or infinite) model. Then $B \preceq^k_{\infty\omega} A$. Equivalently, for all $\varphi \in L^k_{\infty\omega}(\exists)$, if φ is satisfiable, then $A \models \varphi$.*

Proof. The proof follows easily from Proposition 8 by considering the eternal \exists^k -game on B and A with the Duplicator playing on A . The k -Gaifman axioms essentially say that D can extend a partial isomorphism with domain of size $< k$ in every possible way. Therefore, she has a winning strategy for the game. ■

We observe that this result yields a compactness theorem over finite structures and a finitary analog of the Löwenheim-Skolem Theorem for $L^k_{\infty\omega}(\exists)$.

Corollary 12. *For every $k \in \omega$, there is an $n_k \in \omega$ such that for every set Φ of sentences of $L^k_{\infty\omega}(\exists)$, Φ is satisfiable, if and only if, every finite subset of Φ is satisfiable, if and only if, Φ is satisfied in a model of size n_k .*

The next proposition establishes that there are finite structures whose $L^k(\exists)$ -theory is not finitely axiomatizable in $L^k(\exists)$.

Proposition 13. *For all $k \geq 2$, there is a model $A_k \in \mathcal{F}$ such that the $L^k(\exists)$ -theory of A_k is not finitely axiomatizable in $L^k(\exists)$.*

Proof. Let A_k be any finite model of the k -Gaifman theory over the language of graphs. We show that for any $n \in \omega$, there is a B_k^n such that $A_k \preceq^{k,n} B_k^n$ and $A_k \not\preceq^{k,n+1} B_k^n$. This implies that the theory of A_k cannot be axiomatized by $L^k(\exists)$ sentences of quantifier rank $\leq n$ and, therefore, that it is not finitely axiomatizable in $L^k(\exists)$.

For the purpose of defining the models B_k^n , we require the following notion and notation. A basic k -type π satisfies the *distinctness condition* if for every

$l < k$, the formula $x_l \neq x_k \in \pi$. Let $\{\pi_1, \dots, \pi_s\}$ be a set of basic $(k-1)$ -types such that

1. every basic $(k-1)$ -type is equivalent to some π_i and
2. if $i \neq j$, then π_i is not equivalent to π_j .

Similarly, for each $1 \leq i \leq s$, let $\{\pi_{i,1}, \dots, \pi_{i,n(i)}\}$ be a set of basic k -types each of which extends π_i and satisfies the distinctness condition such that

1. every basic k -type which extends π_i and satisfies the distinctness condition is equivalent to some $\pi_{i,j}$ and
2. if $j \neq j'$, then $\pi_{i,j}$ is not equivalent to $\pi_{i,j'}$.

We proceed to define the models B_k^n . Let B_k^1 be the graph on two vertices with exactly one loop and no other edges. Thus B_k^1 realizes both basic 1-types. Given that B_k^n has been defined, we now define B_k^{n+1} as an extension of B_k^n . For each $(k-1)$ -tuple \bar{b} of elements of B_k^n , let $\tau(\bar{b})$ be the unique i such that $B_k^n \models \pi_i(\bar{b})$, and let $X_{\bar{b}} = \{b_{\bar{b},j}^{n+1} \mid 1 \leq j \leq n(\tau(\bar{b}))\}$ be a set of distinct objects disjoint from B_k^n . We suppose that for any distinct pair of $(k-1)$ -tuples \bar{a} and \bar{b} of elements of B_k^n , $X_{\bar{a}} \cap X_{\bar{b}} = \emptyset$. Let X be the union of all the sets $X_{\bar{b}}$. We let the universe of $B_k^{n+1} = B_k^n \cup X$. The edge relation of B_k^{n+1} is obtained from that of B_k^n by adding the minimal number of edges so that each k -tuple $\bar{b} * b_{\bar{b},j}^{n+1}$ satisfies $\pi_{\tau(\bar{b}),j}$. It is easy to see that each B_k^{n+1} is well-defined. We say that the *height* of an element b introduced in this construction is the least n such that $b \in B_k^n$.

We first show that $A_k \preceq^{k,n} B_k^n$. By Proposition 7, it suffices to describe a winning strategy for D in the n -round \exists^k -game with D playing on B_k^n and S playing on A_k . The strategy we describe for D will allow her to play her m^{th} move on some $b \in B_k^m$, for each $m \leq n$. In round 1, D answers the first move of S by playing her pebble on the appropriate element of $B_k^1 \subseteq B_k^n$ to create a partial isomorphism. Suppose that D has played only onto elements of B_k^m through round m , where $m < n$. Let S choose pebble pair (α_l, β_l) to play in round $(m+1)$. We consider two cases. If S plays α_l on the same element as some $\alpha_{l'}$, for $l \neq l'$, then D must play β_l onto the element pebbled by $\beta_{l'}$. Doing so, she obviously maintains a partial isomorphism and succeeds in playing within B_k^{m+1} . On the other hand, suppose that S plays α_l on a distinct element such that the elements pebbled by $\bar{\alpha} * \alpha_l$ on A after the round satisfy $\pi_{i,j}$ (we may need to pad the tuple pebbled by $\bar{\alpha}$ to a tuple of length $(k-1)$ by repeating its last element, if all the pebbles are not in play at this round). Before D plays her $(m+1)^{\text{st}}$ move, the pebbles $\bar{\beta}$ are on a tuple \bar{b} (similarly padded, if necessary) that satisfies π_i . She then plays β_l on the element $b_{\bar{b},j}^{m+1} \in B_k^{m+1}$, thereby maintaining a partial isomorphism. This strategy enables her to win the n -round game.

Next, we show that $A_k \not\preceq^{k,n+1} B_k^n$. By Proposition 7, it suffices to show that S can win the $(n+1)$ -round game with D playing on B_k^n and S playing on A_k . We describe a strategy for play by S which forces D to pebble an element of height at least m by the end of round m to avoid losing at that round. It follows

that S wins the $(n + 1)$ -round game since all elements of B_k^n have height $\leq n$. S plays as follows. He first places his k -pebbles on a set of k distinct elements which form a k -clique, that is, for every pair of distinct pebbled elements a and a' , $A_k \models E(a, a')$. S may play in this way since $A_k \models \Gamma_k$. By our construction above, if $b, b' \in B_k^n$ are distinct elements of the same height, $B_k^n \not\models E(b, b')$. It follows immediately that any r -clique in B_k^n contains an element of height at least r . Therefore, if S has not won by round k , D has pebbled an element of height at least k by the end of that round. Note that in case $(n + 1) \leq k$, we are done, since at round $(n + 1)$, D will be unable to play onto an element of height at least $(n + 1)$ to form an $(n + 1)$ -clique.

We proceed to describe the strategy for S's continuing play under the assumption that $k < (n + 1)$. Suppose that through round $m, k \leq m < (n + 1)$, D has played a pebble onto an element of height at least m , and that the k pebbles S has played lie on distinct elements of A_k which form a k -clique. We show how S can play to ensure that D must play onto an element of height at least $(m + 1)$ at round $(m + 1)$, if she is to prevent S from winning at this round, and leave the round with a k -clique pebbled. Suppose that β_i is pebbling an element b of height greater than the height of any other element pebbled in B_k^n at round m . By our hypothesis, the height of b is at least m . Pick $j \neq i$ (recall that $2 \leq k$) and let $a \in A_k$ be the element pebbled by α_j . S picks up α_j and places it on an $a' \in A_k$ such that

1. $A_k \models E(a, a) \leftrightarrow \neg E(a', a')$ and
2. for every $a'' \in A_k$ on which one of the remaining $(k - 1)$ pebbles lies, $a' \neq a''$ and $A_k \models E(a', a'') \wedge E(a'', a')$.

The existence of such an a' follows from the fact that $A_k \models \Gamma_k$. We claim that to avoid losing at this round, D must play her pebble β_j onto an element b' of height greater than the height of b , and hence of height at least $(m + 1)$. Let b'' be the element pebbled by β_j at round m . By our construction, each element of B_k^n is connected to at most $(k - 1)$ elements of lesser height. Therefore, from the hypotheses that S had pebbled a k -clique at round m , and that b is an element of maximal height pebbled by D at that round, we may conclude that the only element of height \leq the height of b adjacent to b onto which D could play β_j is b'' itself. But this play would fail to maintain a partial isomorphism with the elements S has now pebbled at round $(m + 1)$ by the first condition we have imposed on the choice of a' above. Therefore, to avoid losing at round $(m + 1)$, D must pebble an element of height at least $(m + 1)$. ■

The next result follows immediately.

Corollary 14. *There are infinitely many formulas of $L^k(\exists)$ which are pairwise inequivalent over \mathcal{F} .*

We now consider $L_{\infty\omega}^k(\exists)$ -theories and normal forms for $L_{\infty\omega}^k(\exists)$ sentences over \mathcal{F} . We let $\text{Th}_{\exists}^k(A)$ denote the $L_{\infty\omega}^k(\exists)$ -theory of A . Before proceeding, we define the following fragments of $L_{\infty\omega}^k(\exists)$.

1. Let $\bigwedge L^k(\exists) = \{\theta \mid \theta = \bigwedge \Phi, \text{ for some } \Phi \subseteq L^k(\exists)\}$.
2. Let $\bigvee L^k(\exists) = \{\theta \mid \theta = \bigvee \Phi, \text{ for some } \Phi \subseteq L^k(\exists)\}$.
3. Let $\bigwedge(\bigvee L^k(\exists)) = \{\theta \mid \theta = \bigwedge \Phi, \text{ for some countable } \Phi \subseteq \bigvee L^k(\exists)\}$.
4. Let $\bigvee(\bigwedge L^k(\exists)) = \{\theta \mid \theta = \bigvee \Phi, \text{ for some countable } \Phi \subseteq \bigwedge L^k(\exists)\}$.

Proposition 15. *For all finite structures A , there is a $\theta \in \bigwedge L^k(\exists)$ such that $\text{Mod}_f(\theta) = \text{Mod}_f(\text{Th}_{\exists}^k(A))$.*

Proof. Observe that $\text{Mod}_f(\text{Th}_{\exists}^k(A)) = \{B \in \mathcal{F} \mid A \preceq_{\infty\omega}^k B\}$. Let $\mathcal{C}_A = \mathcal{F} - \text{Mod}_f(\text{Th}_{\exists}^k(A))$. By Proposition 9, for each $B \in \mathcal{C}_A$, there is a sentence $\varphi_B \in L^k(\exists)$ such that $A \models \varphi_B$ and $B \not\models \varphi_B$. Let $\theta = \bigwedge_{B \in \mathcal{C}_A} \varphi_B$. It is easy to verify that $\text{Mod}_f(\theta) = \text{Mod}_f(\text{Th}_{\exists}^k(A))$. ■

Kolaitis and Vardi [12] obtained a normal form for the negation free fragment of $L_{\infty\omega}^k(\exists)$ over \mathcal{F} . It is easy to extend their result to $L_{\infty\omega}^k(\exists)$ and to provide a dual normal form as well. We codify these normal forms in the next proposition.

Proposition 16 (Kolaitis and Vardi [12]). *For each $\varphi \in L_{\infty\omega}^k(\exists)$, there is a $\theta \in \bigvee(\bigwedge L^k(\exists))$ and a $\zeta \in \bigwedge(\bigvee L^k(\exists))$ such that $\text{Mod}_f(\varphi) = \text{Mod}_f(\theta) = \text{Mod}_f(\zeta)$.*

Proof. Let $\mathcal{C} = \text{Mod}_f(\varphi)$. By Proposition 10, for each $A \in \mathcal{C}, B \in \mathcal{F} - \mathcal{C}$, there is a sentence $\theta_{A,B} \in L^k(\exists)$ such that $A \models \theta_{A,B}$ and $B \not\models \theta_{A,B}$. Let $\theta = \bigvee_{A \in \mathcal{C}} (\bigwedge_{B \notin \mathcal{C}} \theta_{A,B})$ and let $\zeta = \bigwedge_{B \notin \mathcal{C}} (\bigvee_{A \in \mathcal{C}} \theta_{A,B})$. It is easy to verify that the proposition holds for this choice of θ and ζ . ■

Next we show that the fragments $\bigwedge L^k(\exists)$ and $\bigvee L^k(\exists)$ are closed under finite conjunction, finite disjunction, and existential quantification over \mathcal{F} . This means that if an $L_{\infty\omega}^k(\exists)$ -definable query cannot be expressed in either $\bigwedge L^k(\exists)$ or $\bigvee L^k(\exists)$, then it is only definable using both an infinitary conjunction and an infinitary disjunction.

Proposition 17. *The languages $\bigwedge L^k(\exists)$ and $\bigvee L^k(\exists)$ are both closed under finite conjunction, finite disjunction, and existential quantification over \mathcal{F} .*

Proof. Let $\Phi = \{\varphi_i(x, \bar{y}) \mid i \in \omega\}$ be a set of formulas of $L^k(\exists)$. We show that if $\theta(\bar{y}) = \exists x \bigwedge \Phi$, then $\theta(\bar{y})$ is equivalent over \mathcal{F} to some formula $\theta'(\bar{y}) \in \bigwedge L^k(\exists)$. (The other closure conditions may be easily verified.) Let $\psi_m = \bigwedge_{0 \leq l \leq m} \varphi_l(x, \bar{y})$ and let $\theta'(\bar{y}) = \bigwedge_{m \in \omega} \exists x \psi_m$. We show θ' is equivalent to θ . It is obvious that θ implies θ' . Let $A \in \mathcal{F}$ and $\bar{a} \in A$ be such that $A \models \theta'[\bar{a}]$. Because A is finite, there is some $a' \in A$ such that for arbitrarily large m , $A \models \psi_m[a', \bar{a}]$. Therefore $A \models \bigwedge_{m \in \omega} \psi_m[a', \bar{a}]$, and θ' implies θ . ■

Below we show that the query classes $\bigwedge L^k(\exists)$ and $\bigvee L^k(\exists)$ are proper subsets of $\bigwedge(\bigvee L^k(\exists))$ and that neither of $\bigwedge L^k(\exists)$ and $\bigvee L^k(\exists)$ is a subset of the other. We first give necessary and sufficient conditions for classes to be definable in $\bigwedge L^k(\exists)$ and $\bigvee L^k(\exists)$.

Proposition 18. 1. A class \mathcal{C} is definable in $\bigwedge L^k(\exists)$ iff for all $B \notin \mathcal{C}$, there is a $\varphi_B \in L^k(\exists)$ such that $B \not\models \varphi_B$ and for all $A \in \mathcal{C}$, $A \models \varphi_B$.

2. A class \mathcal{C} is definable in $\bigvee L^k(\exists)$ iff for all $A \in \mathcal{C}$, there is a $\varphi_A \in L^k(\exists)$ such that $A \models \varphi_A$ and for all $B \notin \mathcal{C}$, $B \not\models \varphi_A$.

Proof. To prove 1., suppose that \mathcal{C} is defined by the sentence $\bigwedge_{n \in \omega} \psi_n$, and that $B \notin \mathcal{C}$. Then there is some ψ_m such that $B \not\models \psi_m$. Let φ_B be this ψ_m . In the other direction, observe that the sentence $\varphi = \bigwedge_{B \notin \mathcal{C}} \varphi_B$ defines \mathcal{C} . The proof of 2. is similar. ■

Proposition 19. For each $k \geq 2$, there is a polynomial time computable boolean query $\mathcal{C} \in \bigwedge L^k(\exists) - \bigvee L^k(\exists)$.

Proof. Let $k \geq 2$ be given and let the graph A_k be a model of the k -Gaifman theory. Let T be the $L^k(\exists)$ -theory of A_k and let $\theta = \bigwedge T$. Clearly, $\theta \in \bigwedge L^k(\exists)$. Let $\mathcal{C} = \text{Mod}_f(\theta)$. It is easy to see that $\mathcal{C} = \{B \in \mathcal{F} \mid A_k \preceq^k B\}$. It then follows immediately from the fact that the relation \preceq^k is polynomial time computable (see Kolaitis and Vardi [12]), that \mathcal{C} is polynomial time computable. In the proof of Proposition 13, we showed that for every satisfiable $\varphi \in L^k(\exists)$, $\text{Mod}_f(\varphi) \not\subseteq \mathcal{C}$. It follows immediately that $\mathcal{C} \neq \text{Mod}_f(\psi)$ for every sentence $\psi \in \bigvee L^k(\exists)$. ■

Proposition 20. There is a polynomial time computable boolean query $\mathcal{C} \in \bigvee L^2(\exists)$ such that for all $k \in \omega$, $\mathcal{C} \notin \bigwedge L^k(\exists)$. In consequence, for each $k \geq 2$, there is a class $\mathcal{C} \in \bigvee L^k(\exists) - \bigwedge L^k(\exists)$.

Proof. Over the signature $\sigma = \{E, s, t\}$, let $\mathcal{C} = \{A \mid \text{there is a path from } s \text{ to } t\}$, the class of (s, t) -connected graphs. This class is clearly in $\bigvee L^2(\exists)$. As noted earlier, it is in Datalog, and, hence, polynomial time computable. From Proposition 18, to show that $\mathcal{C} \notin \bigwedge L^k(\exists)$, it suffices to show that there is a $B \notin \mathcal{C}$ such that for all $n \in \omega$, there is an $A_n \in \mathcal{C}$ such that $A_n \preceq^{k,n} B$. This latter condition is equivalent to D's possessing a winning strategy for the n -round \exists^k -game on A_n and B . We construct B to give her the greatest possible freedom in choosing her moves. Let M be any graph such that $M \models \Gamma_{k+1}$, and let M_s (resp. M_t) be obtained from M by requiring that s (resp. t) denote a loop-free element. We define B to be the disjoint union of M_s and M_t , thus insuring that $B \notin \mathcal{C}$.

For each n , let A_n be the simple chain from s to t of length 2^{n+2} . The basic idea is that by choosing the chain to be long enough, S will not be able to witness the existence of a path from s to t in only n moves. Let $d(x, y)$ be the natural distance function on A_n .

We now describe D's strategy. In each round m , D chooses to play on an element of M_s iff S just played a pebble on $a \in A_n$ such that either (i) $d(s, a) \leq 2^{(n+2)-m}$; or (ii) there is a j such that β_j is on an element of M_s and $d(\alpha_j, a) \leq 2^{(n+2)-m}$. She then plays her pebble on an element of the appropriate component of B so that she maintains a partial isomorphism among the pebbles on that component. It is easy to see that this is possible because M_s and M_t are models of Γ_{k+1} .

In order to establish that this is a winning strategy, it suffices to verify the following two claims.

1. In each round $l \leq n$, if D plays a pebble β_i on M_s , then α_i is not adjacent to t on A_n . Similarly for M_t and s .
2. After each round l , for all pairs of pebbles $\{\alpha_i, \alpha_j\}$, if $A_n \models E(\alpha_i, \alpha_j)$, then β_i and β_j are on the same component of B .

We argue, by induction, that if D plays β_i on M_s in round m , then $d(s, \alpha_i) \leq (2^{(n+2)^{-1}} + 2^{(n+2)^{-2}} + \dots + 2^{(n+2)^{-m}}) < 2^{n+2} - 1$. Since $d(s, t) = 2^{n+2}$, this establishes that $A_n \not\models E(\alpha_i, t)$. In round 1, D plays β_i on M_s iff $d(s, \alpha_i) \leq 2^{(n+2)^{-1}}$. Suppose that in round $m+1$ D plays β_i on M_s . Then either $d(s, \alpha_i) \leq 2^{(n+2)^{-m}}$ or there is an α_j such that β_j is on M_s , $d(\alpha_i, \alpha_j) \leq 2^{(n+2)^{-(m+1)}}$, and, by induction hypothesis, $d(s, \alpha_j) \leq (2^{(n+2)^{-1}} + 2^{(n+2)^{-2}} + \dots + 2^{(n+2)^{-m}})$. In both cases, the induction condition is maintained. The second part of Claim 1 follows from the fact that in round m , if D plays β_i on M_t , then S must have played α_i such that $d(s, \alpha_i) > 2^{(n+2)^{-m}} > 1$. To prove Claim 2, observe that at each round m , if $\beta_i \in M_s$, and $\beta_j \in M_t$, then $d(\alpha_i, \alpha_j) \geq 2^{(n+2)^{-m}} > 1$. The details are similar to the previous argument. \blacksquare

The next result shows that the normal form for $L_{\infty\omega}^k(\exists)$ over \mathcal{F} given in Proposition 16 is optimal.

Proposition 21. *For all $k \geq 2$, there is a class $\mathcal{C} \subseteq \mathcal{F}$ such that $\mathcal{C} \in \bigvee (\bigwedge L^k(\exists)) - (\bigwedge L^k(\exists) \cup \bigvee L^k(\exists))$.*

Proof. The proof of this proposition is a synthesis of the proofs of the preceding two results. We define a set of models $\{A_1, A_2, \dots\}$ which are pairwise $L^k(\exists)$ -incompatible such that for each i , the $L^k(\exists)$ -theory of A_i is not finitely axiomatizable in $L^k(\exists)$. We then let $\mathcal{C} = \{B \mid \exists i(A_i \preceq^k B)\}$. The arguments to show that this class is neither in $\bigvee L^k(\exists)$ nor in $\bigwedge L^k(\exists)$ are minor variants of the proofs of Propositions 19 and 20.

We define each model A_i as an expansion of a homeomorphic image of a graph which is a model of the $(k+1)$ -Gaifman theory. Let R be a finite graph that satisfies Γ_{k+1} ; observe that R also verifies Γ_k . Each A_i is obtained from R by replacing all edges which are not loops by pairwise disjoint paths of length i . Where there is a two-way, undirected edge, a single undirected path is inserted, rather than two directed paths. To clarify the exposition, we also add a unary predicate V to the signature to label the original ‘vertices’ of R .

To verify that \mathcal{C} is not in $\bigvee L^k(\exists)$, it suffices to show that there is a model $A \in \mathcal{C}$ and a sequence B^1, B^2, \dots , disjoint from \mathcal{C} , such that for each n , $A \preceq^{k,n} B^n$. Let A be A_1 , and let each B^n be obtained from the model B_k^n from the proof of Proposition 13 by putting every element into the extension of the predicate V . From that proof it is immediate that, for all n , $A_1 \preceq^{k,n} B^n$ but $A_1 \not\preceq^k B^n$. For all $2 \leq i$, $A_i \models \exists x \neg Vx$ and, consequently, $A_i \not\preceq^k B^n$. This establishes that each B^n is not in \mathcal{C} .

In order to show that $\mathcal{C} \notin \bigwedge L^k(\exists)$, we now define a single $B' \notin \mathcal{C}$ such that for all n , there is an $A_{f(n)}$ such that $A_{f(n)} \preceq^{k,n} B'$. By Proposition 18, this will establish that $\mathcal{C} \notin \bigwedge L^k(\exists)$. Let R^+ be an expansion of R obtained by labeling exactly one looped element with the predicate V ; and let R^- be

obtained similarly by labeling a loop-free element. Here the predicate V plays the same role as the constants s and t in the proof of Proposition 20. We define B' to be the disjoint union of k copies of both R^+ and R^- , and let $f(x) = 2^{x+2}$. It is easy to see that $B' \notin \mathcal{C}$. As in the proof of Proposition 20, the Duplicator wins the n -move \exists^k -game on $A_{2^{n+2}}$ and B' because the labeled vertices of $A_{2^{n+2}}$ are too far apart for S to distinguish the models by witnessing that they are actually connected. ■

Finally, we prove the following separation.

Proposition 22. *Over \mathcal{F} , for $k \geq 3$, $L^k(\exists) \subset (\bigwedge L^k(\exists)) \cap (\bigvee L^k(\exists))$.*

Proof. Let $\text{Path}(x, y)$ express the binary query ‘there is an E -path from x to y .’ For signature $\sigma = \{E, s\}$, we define $\mathcal{C} = \{A \mid \exists x (\text{Path}(s, x) \text{ and } \text{Path}(x, x))\}$. Let $\theta_n(x, y)$ be an $L^3(\exists)$ formula that defines the binary query ‘there is a path of length n from x to y .’ It is easy to see that \mathcal{C} is in $\bigvee L^k(\exists)$. Also observe that $\varphi = \bigwedge_{n \in \omega} \exists x \exists y (s = x \wedge \theta_n(x, y))$ defines \mathcal{C} . Finally, there are arbitrarily large minimal models in \mathcal{C} , that is, models $A \in \mathcal{C}$ such that for all proper submodels $B \subset A$, $B \notin \mathcal{C}$. This immediately implies that $\mathcal{C} \notin \text{FO}(\exists)$ and, *a fortiori*, not in $L^k(\exists)$. ■

4 The Failure of Existential Preservation for $L_{\infty\omega}^\omega$

In this section we prove that $L_{\infty\omega}^\omega \cap \text{EXT} \not\subseteq L_{\infty\omega}^\omega(\exists)$. Indeed, we establish that there is a sentence $\theta \in L_{\infty\omega}^\omega$ such that $\text{Mod}(\theta)$ is closed under extensions, but there is no $\psi \in L_{\infty\omega}^\omega(\exists)$ such that $\text{Mod}_f(\theta) = \text{Mod}_f(\psi)$. Thus, θ witnesses the failure of existential preservation for $L_{\infty\omega}^\omega$ simultaneously over the class of finite structures and over the class of all structures. The central lemma on which this result relies is of interest in itself. It says that for all $k \geq 3$, the finitary language L^k fails in a strong way to satisfy an existential preservation property. Andreka, van Benthem, and Nemeti [3] showed that for every $k \geq 3$, there is a sentence $\varphi_k \in L^k$ which is preserved under extensions, but which is not equivalent to any sentence of $L^k(\exists)$. For $k \geq 3$, the sentence φ_k they construct uses a relation symbol of arity $k - 1$ and has the property that it is equivalent to a sentence of $L^{k+1}(\exists)$. They state the following open problems.

- For any $k \geq 3$ and $n \in \omega$, find sentences $\varphi_n \in L^k$ which are preserved under extensions, but which are not equivalent to any sentence of $L^{k+n}(\exists)$.
- For $k > 3$, is there a formula of L^k containing only (one) binary relation symbols which is preserved under extensions, but is not equivalent to any sentence of $L^k(\exists)$?

The next proposition settles both these open problems. The main result of the section follows easily from the proof of this proposition.

Proposition 23. *For each $k < \omega$, there is a sentence $\theta_k \in L^3$, containing a single binary relation, such that*

1. $\text{Mod}(\theta_k)$ is closed under extensions, but
2. $\text{Mod}_f(\theta_k) \neq \text{Mod}_f(\varphi)$ for all $\varphi \in L^k(\exists)$.

Proof. Before presenting the full proof, we sketch the basic outline. Let the k -pyramid of B , $\mathcal{P}^k(B)$, be the smallest class of (finite and infinite) models containing B that is closed under substructures and L^k -equivalence. For each $k \geq 3$, we define finite structures A_k and B_k with the following properties:

1. $A_k \preceq_{\infty\omega}^k B_k$;
2. $\mathcal{P}^3(B_k)$ is L^3 -definable;
3. $A_k \notin \mathcal{P}^3(B_k)$.

Let $\varphi_k \in L^3$ be such that $\text{Mod}(\varphi_k) = \mathcal{P}^3(B_k)$, and let $\theta_k = \neg\varphi_k$. It is obvious that $\text{Mod}(\theta_k)$ is closed under extensions, that $A_k \models \theta_k$, and that $B_k \not\models \theta_k$. Suppose $\varphi \in L^k(\exists)$ is such that $A_k \models \varphi$. Since $A_k \preceq_{\infty\omega}^k B_k$, this implies that $B_k \models \varphi$, and therefore that φ is not equivalent to θ_k .

We define structures A_k and B_k in terms of simpler submodels. For $f \leq t$, let the $[t, f]$ -flag, $F[t, f]$, be the directed chain of length t with one additional vertex attached to the f^{th} link. That is, the vertex set of $F[t, f]$ is $\{0, 1, \dots, t, t+1\}$, and the edge relation is $\{(i, i+1) \mid i < t\} \cup \{(f, t+1)\}$. A_k is the disjoint union of the $k+1$ flags— $F[2k+2, k+1], F[2k+2, k+2], \dots, F[2k+2, 2k+1]$. Let the $[k, j]$ -tree, $T[k, j]$, be the tree obtained from A_k by fusing the i^{th} nodes of each flag, for all $i \leq j$. This tree has height $2k+2$ and the node at height j has outdegree $k+1$. Then B_k is the disjoint union of the k trees— $T[k, 0], T[k, 1], \dots, T[k, k-1]$.

First we show that $A_k \preceq_{\infty\omega}^k B_k$ by describing a winning strategy for D in the eternal \exists^k -game on A_k and B_k . A *component* of a model is a maximal connected submodel. Observe that every component of A_k is embeddable in every component of B_k . Call a component of either A_k or B_k *vacant* at round n if there is no pebble located on any element of that component before the players make their n^{th} moves. We consider two cases of moves for S. First, suppose that in some round n , S plays pebble α_i on a vacant component A^n of A_k . Since there are only k pairs of pebbles, and since pebble β_i is not on the board, there is a vacant component B^n of B_k , and an isomorphic injection $h_n : A^n \hookrightarrow B^n$. D will play pebble β_i on $h_n(\alpha_i)$. In the other case, S plays on a non-vacant component A^n . There is some $m < n$ such that A^n has been occupied continuously since round m and either $m = 1$ or A^n was vacant at round $m-1$. Thus $A^n = A^m$, and there are previously defined B^m and h_m . D now plays β_i on $h_m(\alpha_i)$. By this condition, every pair of pebbles (α_l, β_l) on components A^m and B^m satisfies the condition that $h_m(\alpha_l) = \beta_l$. In both cases, it is clear that D has maintained a partial isomorphism. By Proposition 8, it now follows immediately that $A_k \preceq_{\infty\omega}^k B_k$.

Next, we show that $\mathcal{P}^3(B_k)$ is definable in L^3 . Consider the following properties:

1. A contains no chains of length $\geq 2k+2$.
2. A contains no cycles of length $\leq 2k+2$.

3. No element $a \in A$ has indegree ≥ 2 , that is, $A \models \neg \exists x \exists y \exists z (x \neq y \wedge Exz \wedge Eyz)$.

It is easy to show that each property is expressible in L^3 , is closed under substructures, and holds of B_k . From this it follows immediately that each $B' \in \mathcal{P}^3(B_k)$ possesses all three properties. Consequently, every member of $\mathcal{P}^3(B_k)$ is a forest consisting of directed trees of height $\leq 2k + 2$.

Next we note the following facts:

Lemma 24. *Let A and B be the disjoint unions of components $\{A_1, \dots, A_m\}$ and $\{B_1, \dots, B_n\}$, respectively. For $k \geq 3$, $A \equiv_{\infty\omega}^k B$ if and only if for each component $A_i [B_i]$, either the number of components of A that are L^k -equivalent to it is equal to the number of components of B that are L^k -equivalent to it or both numbers are $\geq k$.*

This result can be proved by a simple pebble game argument.

Lemma 25. *For each h , and each $k \geq 3$, up to equivalence in L^k there are only finitely many trees of height $\leq h$.*

The proof proceeds by induction on h . The case where $h = 1$ is obvious. Given a tree T , call a proper subtree that contains a node t of height 1 and all of its descendants a 1-tree of T . For $h > 1$, we claim that two trees T_1 and T_2 of height at most h are L^k -equivalent if and only if for each 1-tree $T' \subset T_i$, the number of 1-trees of T_1 that are L^k -equivalent to T' equals the number of 1-trees of T_2 that are L^k -equivalent to T' , or both numbers are $\geq k$. The argument is just like the proof of the preceding lemma. From the claim, the lemma follows immediately.

Corollary 26. *For each h , and each $k \geq 3$, up to equivalence in L^k there are only finitely many forests of height $\leq h$.*

This is an immediate consequence of the preceding lemmas.

These observations establish that there are only finitely many complete L^k -theories that are satisfiable in $\mathcal{P}^3(B_k)$. Moreover, each such theory has a finite model. By [5], every such theory is axiomatized by a single L^k sentence. Hence, if we let φ_k be the disjunction of these sentences, we have $\text{Mod}(\varphi_k) = \mathcal{P}^3(B_k)$ as desired.

Finally, we argue that $A_k \notin \mathcal{P}^3(B_k)$. By the definition of $\mathcal{P}^3(B_k)$, for every $B' \in \mathcal{P}^3(B_k)$, there is an $m \in \omega$ and a sequence $(E_0, D_1, E_1, \dots, D_m, E_m)$ of structures, with $B_k = E_0$ and $B' = E_m$, such that:

1. For all $1 \leq i \leq m$, $D_i \subseteq E_{i-1}$.
2. For all $1 \leq i \leq m$, $D_i \equiv^3 E_i$.

It suffices to show that for any such sequence, A_k cannot be embedded in any E_i . Let $g : \mathcal{P}^3(B_k) \mapsto \{0, 1, \dots, k + 1\}$ be the function such that $g(D)$ is the maximum number of components of A_k that can be embedded in D pairwise disjointly. We show that for each $i \leq m$, $g(E_i) < k + 1$. In fact, we show that g is

monotonically decreasing on the aforementioned sequence. Because each D_i is a submodel of E_{i-1} , it is clear that $g(D_i) \leq g(E_{i-1})$. It remains to establish that $g(B_k) < k + 1$ and that $g(E_i) \leq g(D_i)$.

Observe that any embedding of a flag $F[2k + 2, f]$ into a component C of any $B' \in \mathcal{P}^3(B_k)$ must map the root of the flag to the root of C . This implies that no two flags of A_k can be disjointly embedded into any such component and, since B_k has only k components, that $g(B_k) < k + 1$.

From Lemma 24, it follows that every E_i can be obtained from D_i by repeated application of the following three operations. First, replace some component with a component that is L^3 -equivalent to it. Second, add a disjoint copy of a tree that is L^3 -equivalent to at least 3 components. Third, remove a component that is L^3 -equivalent to at least 3 other components. Thus, it suffices to argue that no such operation performed on some $B' \in \mathcal{P}^3(B_k)$ can yield a B'' such that $g(B'') > g(B')$. It is obvious that removing a component cannot increase the value of g .

We claim that it suffices to consider the effect of the other two operations on components of height $= 2k + 2$. If trees T and T' are L^3 -equivalent, then they have the same height. Also, no component $F[2k + 2, f]$ of A_k can be embedded in any tree of height $< 2k + 2$. This establishes that the presence of shorter components in a model B does not affect the value of $g(B)$.

Observe that for all trees T and T' such that $T \equiv^3 T'$, $F[t, f]$ can be embedded in T iff it can be embedded in T' . This is because the following property can be expressed in L^3 : there is an element x such that (i) there is a y such that there is a path of length f from y to x ; (ii) x has outdegree 2; (iii) there is a y such that there is a path of length $t - f$ from x to y . Over trees, this property says that the model embeds $F[t, f]$. Consequently the operation of replacement cannot increase the value of g .

It remains to establish that adding an additional component to a model $B' \in \mathcal{P}^3(B_k)$ does not change the value of g . We observe that B_k has the following properties:

1. For each $(2k + 2)$ -chain contained in B_k there is at most one $j, 0 \leq j \leq k - 1$, such that the j^{th} link of the chain has outdegree > 1 .
2. For each $(2k + 2)$ -chain contained in B_k there is at most one $j, k + 1 \leq j \leq 2k + 1$, such that the j^{th} link of the chain has outdegree > 1 .

These properties are closed under substructures and L^3 -equivalence; consequently, they hold of every model $B' \in \mathcal{P}^3(B_k)$. Let C_1, C_2 , and C_3 be L^3 -equivalent components of B' of height $2k + 2$. The above argument establishes that each C_i is either some $F[2k + 2, f]$, or the simple $(2k + 2)$ -chain. Let B'' be the extension of B' obtained by adding a component C_4 . Observe that, in fact, all four components must be isomorphic, and embed at most one isomorphism type of flag. Therefore, the image of any embedding $h : A_k \mapsto B''$ can contain vertices from at most one of these four components. This demonstrates that $g(B') = g(B'')$, and completes the proof. \blacksquare

The following result establishes the failure of existential preservation for $L_{\infty\omega}^\omega$.

Theorem 27. *There is a sentence $\theta \in L_{\infty\omega}^\omega$ such that both*

1. $\text{Mod}(\theta)$ is closed under extensions.
2. For all $\varphi \in L_{\infty\omega}^\omega(\exists)$, $\text{Mod}_f(\theta) \neq \text{Mod}_f(\varphi)$.

Proof. We claim that it suffices to show that for each $k \in \omega$ there is a sentence $\theta_k \in L^3$ and a pair of finite models A_k and B_k such that

1. $\text{Mod}(\theta_k)$ is closed under extensions.
2. $A_k \models \theta_k$ and $B_k \not\models \theta_k$.
3. $A_k \preceq_{\infty\omega}^k B_k$.
4. For all j , $A_j \models \theta_k$.

Let $\theta = \bigwedge_k \theta_k$. It is clear that θ is closed under extensions and that it has finite models, since it is true in each A_k . Suppose that φ is a sentence in $L_{\infty\omega}^k(\exists)$ such that θ implies φ . Then $A_k \models \varphi$, and therefore $B_k \models \varphi$. But for all l , $B_l \not\models \varphi$. Therefore, $\text{Mod}_f(\theta) \neq \text{Mod}_f(\varphi)$.

The sentences θ_k and the models A_k and B_k from the proof of Proposition 23 fail to meet condition 4 because for $j < k$, $A_j \not\models \theta_k$. To see this, observe that A_j will always be a submodel of B_k . To fix this defect, it suffices to construct A'_k, B'_k , and θ'_k as in the proof of Proposition 23 that also satisfy the additional condition that, for all j and k , $A'_j \notin \mathcal{P}^3(B'_k)$. In order to accomplish this, we add simple ‘gadgets’ to the models. Let the k -cycle, C_k , be the graph on k vertices whose edge relation forms a simple, directed cycle of length k . Then let A'_k and B'_k be obtained from A_k and B_k , respectively, by adding a disjoint copy of C_k . By slightly modifying the proof of Proposition 23, we can show that $A'_k \preceq_{\infty\omega}^k B'_k$, and that there is a $\theta'_k \in L^3$ satisfied by exactly the models in the complement of $\mathcal{P}^3(B'_k)$ such that $A'_k \models \theta'_k$. Finally, it is easy to verify that for $j \neq k$, the j -cycle cannot be embedded in any $B \in \mathcal{P}^3(B'_k)$ and, therefore, $A'_j \models \theta'_k$. ■

5 Generalized Preservation Theorems in the Finite Case

In this section, we prove some generalized preservation theorems for fragments of FO. Our results are of the form

$$L \cap \text{EXT} \subseteq L'$$

for certain quantifier prefix classes $L \subset \text{FO}$ and $L' = L_{\infty\omega}^\omega(\exists)$ or $\text{Datalog}(\neq, \neg)$. Recall that Tait [18] showed

$$\text{FO} \cap \text{EXT} \not\subseteq \text{FO}(\exists),$$

and that Gurevich and Shelah [9, 10] gave examples showing that

$$\text{FO}[\forall^* \exists^*] \cap \text{EXT} \not\subseteq \text{FO}(\exists).$$

Compton observed that

$$\text{FO}[\exists^* \forall^*] \cap \text{EXT} \subseteq \text{FO}(\exists),$$

which shows that these examples are best possible in terms of quantifier alternation prefix (see [9]). Kolaitis and Vardi (see [2]) observed that the example of Gurevich and Shelah [9] can be defined in $\text{Datalog}(\neq, \neg)$. Theorem 29 below establishes that

$$\text{FO}[\exists^*\forall\exists] \cap \text{EXT} \subseteq \text{Datalog}(\neq, \neg).$$

It follows that all the examples in the literature witnessing the failure of the Los-Tarski Theorem in the finite case are definable in $\text{Datalog}(\neq, \neg)$, since all these examples are in the prefix class $\text{FO}[\exists^*\forall\exists]$ (a sequence of existential quantifiers followed by one universal quantifier followed by one existential quantifier). The next theorem establishes a slightly more general result with $L_{\infty\omega}^\omega(\exists)$ in place of $\text{Datalog}(\neq, \neg)$.

Theorem 28. $\text{FO}[\exists^*\forall\exists^*] \cap \text{EXT} \subseteq L_{\infty\omega}^\omega(\exists)$.

Proof. Let $\varphi \in \text{FO}[\exists^*\forall\exists^*] \cap \text{EXT}$. That is, $\varphi \in \text{FO}[\exists^*\forall\exists^*]$ and $\text{Mod}_f(\varphi) \in \text{EXT}$. Let $\mathcal{C} = \text{Mod}_f(\varphi)$. We proceed to show that $\mathcal{C} \in L_{\infty\omega}^\omega(\exists)$. By Proposition 10, it suffices to show that there is a k such that, for each $A \in \mathcal{C}$ and $B \notin \mathcal{C}$, there is a $\theta_{A,B} \in L_{\infty\omega}^k(\exists)$ such that $A \models \theta_{A,B}$ and $B \not\models \theta_{A,B}$.

Let $\varphi = \exists x_1 \dots x_i \forall y \exists z_1 \dots z_j \psi(\bar{x}, y, \bar{z})$, where ψ is quantifier free, and let $k = i + j + 1$ (we suppose, without loss of generality, that $i > 0$). We now describe a winning strategy for S in the eternal \exists^k -game on A and B , for $A \in \mathcal{C}$ and $B \notin \mathcal{C}$, which establishes, by Proposition 8, the existence of $\theta_{A,B} \in L_{\infty\omega}^k(\exists)$ with the desired properties. There are two stages. Let $\bar{a} = (a_1, \dots, a_i)$ be a sequence of elements of A such that $A \models \forall y \exists \bar{z} \psi(\bar{a}, y, \bar{z})$. If D has not lost after h rounds, for $h < i$, S plays pebble α_{h+1} on element a_{h+1} . If S has not won after i moves, and D has played her pebbles on $\bar{b} = (b_1, \dots, b_i)$, then $B \models \exists y \forall \bar{z} \neg \psi(\bar{b}, y, \bar{z})$ (since $B \not\models \varphi$).

The goal of the second part of S's strategy is to force D to play a pebble on some element b' such that $B \models \forall \bar{z} \neg \psi(\bar{b}, b', \bar{z})$, without removing any of the pebbles $\alpha_1, \dots, \alpha_i$ which 'fix the interpretation' of the variables x_1, \dots, x_i on both A and B . Regardless of the element a' on which S will have played his corresponding pebble, $A \models \exists \bar{z} \psi(\bar{a}, a', \bar{z})$, so that he can then win easily. In order to describe S's strategy, we first define a sequence of subsets of the universe of B . Let $\Gamma_0 = \{b' \mid b' \in B \text{ and } B \models \forall \bar{z} \neg \psi(\bar{b}, b', \bar{z})\}$. Observe that $B \models \exists y \forall \bar{z} \neg \psi(\bar{b}, y, \bar{z})$, and therefore Γ_0 is non-empty. Given $\Gamma_0, \dots, \Gamma_m$, if $(\bigcup_{i \leq m} \Gamma_i) \cap \bar{b} = \emptyset$, then let B_{m+1} be the submodel of B whose universe is $(B - \bigcup_{i \leq m} \Gamma_i)$. Let $\Gamma_{m+1} = \{b' \mid b' \in B_{m+1} \text{ and } B_{m+1} \models \forall \bar{y} \neg \psi(\bar{b}, b', \bar{y})\}$. For each B_m , since $B_m \subseteq B$, we have that $B_m \models \forall \bar{x} \exists y \forall \bar{z} \neg \psi(\bar{x}, y, \bar{z})$. In particular, $B_m \models \exists y \forall \bar{z} \neg \psi(\bar{b}, y, \bar{z})$ and thus, as above, Γ_{m+1} is non-empty. Since B is finite, there is some n such that $\Gamma_n \cap \bar{b} \neq \emptyset$, and some element $b_f \in \Gamma_n \cap \bar{b}$ pebbled by β_f . Then B is partitioned into the sets $\Gamma_0, \dots, \Gamma_{n-1}, B_n$. We also have that $A \models \exists \bar{z} \psi(\bar{a}, a_f, \bar{z})$, and $B_n \models \forall \bar{z} \neg \psi(\bar{b}, b_f, \bar{z})$.

The Spoiler can win by executing a substrategy that compels D to play in sets Γ_m of successively smaller index. Let \bar{c} be a sequence of elements of length j such that, $A \models \psi(\bar{a}, a_f, \bar{c})$. S plays his next j moves on this sequence, until D makes a losing move or plays a pebble β_g onto an element in Γ_m , for $m \leq n - 1$.

We claim that one of these two possibilities must occur. For suppose that D plays on a sequence $\bar{d} \subseteq B_n$. Then $B_n \models \neg\psi(\bar{b}, b_j, \bar{d})$, and $\psi(\bar{x}, y, \bar{z})$ witnesses that the function that takes $\bar{a} * a_j * \bar{c}$ to $\bar{b} * b_j * \bar{d}$ and preserves the denotations of constants is not a partial isomorphism.

Suppose that D has played some pebble β_g into some set Γ_m . By the same argument as above, reusing pebbles $\{\alpha_{i+1}, \dots, \alpha_k\} - \{\alpha_g\}$, S can either win or force D to play into some $\Gamma_{m'}$, for some $m' < m$. Iterating this procedure, S can force D to play into Γ_0 , and then win by using the same procedure one more time. \blacksquare

We remark the following two refinements of the foregoing theorem.

1. For each $B \notin \mathcal{C}$, there is a number m_B such that for all $A \in \mathcal{C}$, S wins the m_B -round \exists^k -game on A and B . (Here, m_B is determined by the maximum number of sets Γ that get defined on B , for any choice of D's first i moves.) It follows easily from Proposition 7 that this condition is equivalent to there being a $\theta_B \in L^k(\exists)$, with quantifier rank $\leq m_B$, such that for all $A \in \mathcal{C}$, $A \models \theta_B$, and $B \not\models \theta_B$. Then $\theta' = \bigwedge_{B \notin \mathcal{C}} \theta_B$ is equivalent to φ and is a single infinite conjunction of $L^k(\exists)$ sentences. We know by Proposition 20 that not all sentences of $L^k_{\infty\omega}(\exists)$ can be expressed in this form. Indeed, it follows from Theorem 29 below that if $\varphi \in \text{FO}[\exists^* \forall \exists] \cap \text{EXT}$, then φ is equivalent to a formula in $\bigwedge L^k(\exists) \cap \bigvee L^k(\exists)$ for some k .
2. Suppose that φ is an L^k sentence with quantifier type $\forall \exists^*$ (this notion of quantifier type may be defined straightforwardly, and is distinct from the notion of prefix class). In this case, we can show, by a modification of the proof of Theorem 28, that φ is equivalent to an $L^k_{\infty\omega}(\exists)$ sentence. This contrasts with Proposition 23 above which established that for all k , there is a sentence $\varphi_k \in L^3$ such that $\text{Mod}_f(\varphi_k) \in \text{EXT}$, but φ_k is not equivalent over \mathcal{F} to any sentence in $L^k_{\infty\omega}(\exists)$.

Theorem 29. $\text{FO}[\exists^* \forall \exists] \cap \text{EXT} \subseteq \text{Datalog}(\neq, \neg)$.

Proof. Let $\varphi = \exists x_1 \dots x_j \forall y \exists z \beta(\bar{x}, y, z)$, with $\beta(\bar{x}, y, z)$ quantifier free. Let $\bar{c} = (c_1, \dots, c_p)$ be the sequence of constants in the signature of φ and let $\mathcal{C} = \text{Mod}_f(\varphi)$. For $a \in A$, we say that a closes with parameters \bar{a} iff there is a sequence $a_0 (= a), a_1, \dots, a_n$ such that for all $l < n$, $A \models \beta(\bar{a}, a_l, a_{l+1})$ and there is an $m \leq n$ such that $A \models \beta(\bar{a}, a_n, a_m)$. Note that this is equivalent to there being an a' such that there is a $\beta(\bar{a}, y, z)$ -path from a to a' , and a $\beta(\bar{a}, y, z)$ -cycle including a' .

We claim that $A \models \varphi$ iff there is a j -tuple \bar{a} such that every element of $\bar{a} \cup \bar{c}$ closes with parameters \bar{a} . Suppose that A does not satisfy these conditions. We prove that $A \models \forall \bar{x} \exists y \forall z \neg \beta(\bar{x}, y, z)$ where the latter sentence is equivalent to $\neg \varphi$. Let $\bar{a} \subseteq A$ be a sequence of length j . By hypothesis, there is an $a' \in \bar{a} \cup \bar{c}$ such that a' does not close with parameters \bar{a} . Since A is finite, this implies that there is an $m \geq 0$ and a sequence $a' = a'_0, \dots, a'_m$ such that for all $l < m$, $A \models \beta(\bar{a}, a'_l, a'_{l+1})$ and $A \models \forall z \neg \beta(\bar{a}, a'_m, z)$, as desired.

In the other direction, let \bar{a} be such that every member of $\bar{a} \cup \bar{c}$ closes with parameters \bar{a} . Let $\bar{s}_h = \langle a_{h0}(= a_h), \dots, a_{hm_h} \rangle$ and $\bar{t}_h = \langle e_{h0}(= c_h), \dots, e_{hn_h} \rangle$ be sequences witnessing that each element of $\bar{a} \cup \bar{c}$ closes with parameters \bar{a} . Let B be the submodel of A with universe $\bigcup_i \bar{s}_i \cup \bigcup_j \bar{t}_j$. Then it is easy to verify that $B \models \varphi$ and, since $\text{Mod}_f(\varphi) \in \text{EXT}$, it follows that $A \models \varphi$.

The following program, with $\bar{x} = (x_1, \dots, x_j)$, computes φ :

$$\begin{aligned} P(\bar{x}, y, z) &\leftarrow \beta(\bar{x}, y, z) \\ P(\bar{x}, y, z) &\leftarrow P(\bar{x}, y, w), P(\bar{x}, w, z) \\ Q &\leftarrow P(\bar{x}, x_1, y_1), P(\bar{x}, y_1, y_1), \dots, P(\bar{x}, x_j, y_j), P(\bar{x}, y_j, y_j), \\ &\quad P(\bar{x}, c_1, w_1), P(\bar{x}, w_1, w_1), \dots, P(\bar{x}, c_p, w_p), P(\bar{x}, w_p, w_p) \end{aligned}$$

This can be easily converted into a Datalog(\neq, \neg) program. Let $\beta(\bar{x}, y, z) = \bigvee_i \delta_i$, where each δ_i is a conjunction of literals. Replace the clause $P(\bar{x}, y, z) \leftarrow \beta(\bar{x}, y, z)$ with the clauses $P(\bar{x}, y, z) \leftarrow \delta_i$, for all i . \blacksquare

6 Conclusion

In this section we discuss some open problems that are naturally suggested by our investigations and we present some further results bearing on the problem of preservation under homomorphisms in the finite case.

6.1 Open Problems

The first and most obvious question is the extent to which our results can be generalized from fragments of FO to the entire language. In this connection, we restate two of the problems mentioned earlier which remain open in light of our study.

Problem 1. Is $\text{FO} \cap \text{EXT} \subseteq \text{Datalog}(\neq, \neg)$?

Problem 2. Is $\text{FO} \cap \text{EXT} \subseteq L_{\infty\omega}^\omega(\exists)$?

Obviously, a positive answer to the first of these questions implies a positive answer to the second. Should the answer to these questions be negative, it would be of interest to characterize the classes $\text{FO} \cap \text{Datalog}(\neq, \neg)$ and $\text{FO} \cap L_{\infty\omega}^\omega(\exists)$ in some informative way. An example of a characterization of this kind is the following theorem of Ajtai and Gurevich [2]. $\text{FO}^+(\exists)$ denotes the positive existential fragment of FO.

Proposition 30 (Ajtai and Gurevich [2]). $\text{FO} \cap \text{Datalog} = \text{FO}^+(\exists)$.

As remarked above, the Gurevich-Shelah counterexample to the Los-Tarski Theorem in the finite case witnesses that $\text{FO} \cap \text{Datalog}(\neq, \neg) \neq \text{FO}(\exists)$. Might $\text{FO} \cap \text{Datalog}(\neq, \neg)$ be contained in some level of the first-order quantifier alternation hierarchy, be it not the lowest level? Should, on the other hand, the answer to Problem 1 be positive, we might try to establish even stronger results such as a positive answer to

Problem 3. Is $(\text{FO} + \text{LFP}) \cap \text{EXT} \subseteq \text{Datalog}(\neq, \neg)$?

6.2 Preservation under Homomorphisms

In this subsection we briefly turn our attention to a different preservation property. A *homomorphism* from A to B is a map $h : A \mapsto B$ such that for all n -ary relation symbols $R(\bar{x})$, and for all n -tuples $\bar{a} \subseteq A$, if $A \models R(\bar{a})$, then $B \models R(h(\bar{a}))$. A class of models \mathcal{C} is closed under homomorphisms iff for all A and B such that there is a homomorphism from A to B , if $A \in \mathcal{C}$, then $B \in \mathcal{C}$. Let HOM denote the set of classes in \mathcal{F} that are closed under homomorphisms. A sentence φ in FO , $L_{\infty\omega}^w$, etc. is *positive*, if and only if, it does not contain any negations. The following well-known classical result is a direct consequence of the Los-Tarski Theorem: for all $\varphi \in \text{FO}$, $\text{Mod}(\varphi)$ is closed under homomorphisms, if and only if, φ is equivalent to a positive existential sentence. This theorem is one of a few classical results whose validity over \mathcal{F} remains unknown. In our current notation, we can formulate the question as the following open problem, the interest of which has been emphasized by Gurevich [10] and Kolaitis (see [8]).

Problem 4. Is $\text{FO} \cap \text{HOM} \subseteq \text{FO}^+(\exists)$?

(To avoid confusion, it should be remarked that although [10] announces a solution to Problem 4, this claim has been withdrawn.)

The following proposition yields some information about the homomorphism preservation question. We direct the reader to [17] for its proof.

Proposition 31. $\text{Datalog}(\neq, \neg) \cap \text{HOM} \subseteq \text{Datalog}$.

Propositions 29, 30, and 31 yield as an immediate corollary the following special case of the homomorphism preservation theorem.

Corollary 32. $\text{FO}[\exists^*\forall\exists] \cap \text{HOM} = \text{FO}^+(\exists)$.

References

1. F. Afrati, S. Cosmadakis, and M. Yannakakis. On datalog vs. polynomial time. In *Proceedings of the 10th ACM Symposium on Principles of Database Systems*, 1991.
2. M. Ajtai and Y. Gurevich. Datalog vs. first-order logic. In *Proceedings of the 30th IEEE Symposium on Foundations of Computer Science*, pages 142–146, 1989.
3. H. Andreka, J. van Benthem, and I. Nemeti. Submodel preservation theorems in finite variable fragments. In A. Ponse, M. de Rijke, and Y. Venema, editors, *Modal Logic and Process Algebra*. Cambridge University Press, 1994.
4. J. Barwise. On Moschovakis closure ordinals. *Journal of Symbolic Logic*, 42:292–296, 1977.
5. A. Dawar, S. Lindell, and S. Weinstein. Infinitary logic and inductive definability over finite structures. *Information and Computation*, to appear.
6. R. Fagin. Probabilities on finite models. *Journal of Symbolic Logic*, 41(1):50–58, March 1976.
7. H. Gaifman. Concerning measures in first-order calculi. *Israel Journal of Mathematics*, 2:1–18, 1964.
8. E. Grädel and J. Tyszkiewicz, editors. *Problems in Finite Model Theory*. Available via anonymous ftp from ftp.informatik.rwth-aachen.de, version of May 25, 1994.

9. Y. Gurevich. Toward logic tailored for computational complexity. In M. Richter et al., editors, *Computation and Proof Theory*, pages 175–216. Springer Lecture Notes in Mathematics, 1984.
10. Y. Gurevich. On finite model theory (extended abstract). In S. R. Buss and P. J. Scott, editors, *Feasible Mathematics*, pages 211–219. Birkhauser, 1990.
11. N. Immerman. Upper and lower bounds for first-order expressibility. *Journal of Computer and System Sciences*, 25:76–98, 1982.
12. Ph. G. Kolaitis and M. Y. Vardi. On the expressive power of datalog: tools and a case study. In *Proceedings of the 9th ACM Symposium on Principles of Database Systems*, pages 61–71, 1990.
13. Ph. G. Kolaitis and M. Y. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In *Proceedings of the 7th IEEE Symposium on Logic in Computer Science*, pages 46–57, 1992.
14. Ph. G. Kolaitis and M. Y. Vardi. Infinitary logics and 0-1 laws. *Information and Computation*, 98(2):258–294, 1992.
15. Libo Lo. Preservation theorems of finite models (abstract). *Journal of Symbolic Logic*, 58:376, 1993.
16. B. Poizat. Deux ou trois choses que je sais de L_n . *Journal of Symbolic Logic*, 47(3):641–658, 1982.
17. E. Rosen. Finite model theory and finite variable logics, 1995. In preparation.
18. W. Tait. A counterexample to a conjecture of Scott and Suppes. *Journal of Symbolic Logic*, 24(1):15–16, 1959.