Embedded Finite Models, Stability, and the Impact of Order

John T. Baldwin and Michael A. Benedikt

August 98

Embedded Finite Models, Stability Theory, and the Impact of Order.

John T. Baldwin

Department of Mathematics
University of Illinois at Chicago
Chicago IL 60607

Michael A. Benedikt

Bell Laboratories
Lucent Technologies
Naperville, IL 60566

Abstract

We extend bounds on the expressive power of first-order logic over finite structures and over ordered finite structures, by generalizing to the situation where the finite structures are embedded in an infinite structure M, where M satisfies some simple combinatorial properties studied in model-theoretic stability theory. We first consider first-order logic over finite structures embedded in a stable structure, and show that it has the same generic expressive power as first-order logic on unordered finite structures. It follows from this that having the additional structure of, for example, an abelian group or an equivalence relation, does not allow one to define any new generic queries. We also consider first-order logic over finite structures living within any model M that lacks the independence property and show that its expressive power is bounded by first-order logic over finite ordered structures. This latter result gives an enormous class of structures in which the expressive power of first-order logic is sharply limited; it shows that common queries such as parity and connectivity cannot be defined for finite structures living within structures from this huge class. It also gives a pure combinatorial property of an interpreted structure M that is sufficient to extend results on first-order logic on ordered structures to first-order logic on finite structures embedded in M.

1 Introduction

Finite and infinite model theory share an interest in how the expressive power of various logics (first-order, fragments of second-order) is impacted by the presence of certain kinds of combinatorial objects, such as linear orders or trees.

In finite model theory a major focus has been the impact of the existence of an order relation on logics over finite structures. For example, there has been much work comparing the expressivity of various predicate logics on ordered structures with that on unordered structures (see [10]), and these results indicate some dramatic differences between logics on pure first-order structures and logics on ordered structures. There are also results on bounding the expressive power of first-order logic and fragments of second-order logic on ordered structures (e.g. [19]), which show that the presence of an order yields tight links between predicate logics and complexity classes. The gap between ordered and unordered structures has led to an examination of which logics can define a linear order on finite structures [16].

In infinite "classical" model theory there is a research program with a somewhat similar flavor, called *stability theory*. One set of stability-theoretic results focuses on structures in which no definable relation on n-tuples induces a linear order on an infinite subset of its domain. It is shown that these structures (the *stable* structures) have a much better definability theory than structures that do have embedded orders; the former have few types and always admit universal domains. These results give a sense that first-order logic is 'better-behaved' on stable structures than on arbitrary structures. Another dividing line of interest is between structures that do or do not admit an infinite parametrically definable independent family: those that do not are said to be *without the independence property*. This class of structures includes the stable ones, but it also encompasses many ordered structures, such as real closed ordered fields, that have a tame definability theory. Structures without the independence property are of interest not only for their use in model theory, but also for their connection with the V-C dimension that arises in machine learning [23].

This work aims at showing one interesting connection between these two lines of research. We will extend results on the expressiveness of first-order logic on finite models to the situation where finite models are embedded in a fixed infinite structure. Our main theorems will say, roughly, that results on first-order definability on unordered finite structures extend to first-order logic on structures that are embedded in a stable structure (e.g. any equivalence relation, the complex field), and that results on first-order definability on ordered finite structures extend to first-order logic in structures without the independence property (e.g. trees, linear orders, real closed fields).

In particular, we will be extending a line of results that began with [26] and continued with [5], which showed that interesting properties of finite structures, such as the parity of a set or the connectivity of a graph, are not definable within certain ambient structures. The papers in this literature show that under certain conditions on the model (e.g. o-minimal in [5], quasi-o-minimal in [3]), no new queries can be defined other than those already definable on finite ordered structures, provided that we restrict our attention to *permutation invariant* queries (see the definition of generic in Section 2). Here we extend these "collapse" results to structures without the independence property. In addition, we prove a stronger result for stable structures; we show that any permutation-invariant queries that are definable in a stable structure, are pure first-order definable in the sense of finite model theory.

One of the main motivations for the results here comes from *constraint databases* [20, 21], which are used as the basis for geographical and temporal data-models. The framework of constraint databases assumes some underlying model $\mathcal{M} = \langle \mathbb{U}, \Omega \rangle$ where $\mathbb{U}$ is an infinite set and Ω is a signature that consists of a number of interpreted functions and predicates over $\mathbb{U}$. In the original work of [21], (generalized) databases over $\mathcal{M}$ are given by quantifier-free formulae $\phi(x_1, \ldots, x_n)$ in the language of Ω; such a database represents the set

$$\mathcal{M}_\phi \;=\; \{\vec{a} = (a_1, \ldots, a_n) \mid \vec{a} \in \mathbb{U}^n,\ \mathcal{M} \models \phi(\vec{a})\}.$$

Geographical databases can be modeled by considering models such as $\mathcal{M} = \langle \mathbb{R}, +, *, 0, 1, < \rangle$, with the intention that generalized databases represent regions on the real plane. In general, $\mathcal{M}$ is chosen to have quantifier elimination. In this case, one performs the query evaluation process by applying the quantifier elimination procedure; see [21].

The results here give bounds on expressibility in constraint database models based on a wide class of interpreted structures. They extend the results of [5, 25, 26] to show that queries such as parity and graph connectivity cannot be expressed in constraint databases based on any structure lacking the independence property. For structures that are based on the reals, one can then use the techniques of [13] to derive that various topological properties cannot be expressed in constraint structures lacking the independence property. These results are also novel in that they show that many constraint structures, such as trees and abelian groups, which appear incomparable with orders, have generic expressive power bounded by first-order logic on ordered structures.

Organization and quick summary Section 2 introduces the main definitions in the paper, gives a more detailed overview of previous work on embedded finite structures, and states the main theorems of the paper. Section 3 presents the key tool: a quantifier reduction theorem for models with a predicate for an indiscernible set. Section 4 uses these results to prove the main theorems. Section 5 gives conclusions and discusses future work.

2 Preliminaries

The following basic definitions provide a framework for the model-theoretic results dealt with in this paper. Most of the terminology below comes from the database literature; see [1] for a general discussion of the notions of genericity given here and [8] for a discussion of collapse results.

Let $\mathbb{U}$ be an infinite set, and let S be a finite relational language. A *boolean query* on $\mathbb{U}$, Q, is a collection of S-structures with domain a finite subset of $\mathbb{U}$. Since we will not discuss any nonboolean queries, we will drop the word boolean and just talk about queries on $\mathbb{U}$, or $\mathbb{U}$-queries. Technically, we should preface 'query' with an S to indicate which finite set of additional relational symbols are used in the query; then our theorems would be about 'all S-queries' or (see below) 'all order S-queries'. Following common practice we suppress this parameter. A *generic query* on $\mathbb{U}$ is a query that is closed under S-isomorphism. An *abstract query* Q is a collection of isomorphism types of finite S-structures. Given an abstract query Q and set $\mathbb{U}$, Q defines a generic query $Q_{\mathbb{U}}$ by considering all finite S-structures embedded in $\mathbb{U}$ whose isomorphism type is in Q. Clearly a generic query on an infinite $\mathbb{U}$ determines an abstract query. Finite model theory investigates the collection of abstract queries and (equivalently) generic queries that are definable in predicate logics. Well-known results include that the abstract query of all even-cardinality sets (the *parity* query) and the set of connected graphs are not first-order definable.

Now let L be a language disjoint from S, and let M be an L-structure with domain $\mathbb{U}$. For any S-structure A with domain contained in $\mathbb{U}$, $M(A)$ denotes the unique $L \cup S$-structure that expands M and agrees with A on the interpretation of the predicates in S. Such an A is what we mean by an *embedded finite model*. We note that there are interesting frameworks that study "mixed" structures that are quite different from the one dealt with here, particularly the one presented in [12].

Given a first-order sentence ϕ in $L \cup S$, the *query defined by* ϕ is the set of S-structures A with domain a finite subset of $\mathbb{U}$ such that that $M(A) \models \phi$. We say ϕ is generic if the query defined by it is generic. Given a $\mathbb{U}$-query Q we say it is *first-order over* M if there is $\phi \in L \cup S$ that defines Q. An abstract query Q is called first-order over M if there is a first-order $\phi \in L \cup S$ that defines $Q_{\mathbb{U}}$.

For a model M, we let

$FO(M)$ be $\{Q : Q$ is first-order definable over $M\}$ and

$FO_{GEN}(M)$ be $\{Q : Q$ is generic and Q is first-order definable over $M\}$.

These two classes are the basic objects of study in this paper. In particular, we can compare the generic queries definable in different models M, since a generic query "makes sense" in any model.

A *pure first-order query* is an abstract query that is definable over the structure $\langle \mathbb{U}, = \rangle$ by an $\{S, =\}$-formula ψ. The results in [18] imply that any pure first-order abstract query is first-order definable in the sense of finite model theory. That is, if a set of isomorphism types is definable over the trivial infinite structure, then it is the collection of finite models of a first-order sentence. Note that these two notions are a priori incomparable: in the first case the quantifiers range over $\mathbb{U}$ and in the second case only over a finite set. The following easy fact was used in [5] and [24]:

Fact 2.1 *If M' and M are elementarily equivalent L-structures and the abstract query Q is defined by ϕ over M, then Q is also defined by ϕ over M'.*

Using this we see that the collection $FO_{GEN}(M)$ is actually an invariant of the theory of M. In this paper, then, we will focus on $FO_{GEN}(M)$ rather than $FO(M)$. We prove a number of *collapse* results stating that if a generic query is definable with the help of symbols from L, it can be defined without these symbols. The first such result is:

Theorem A *Let $\mathbb{U}$ be infinite and M be a stable L-structure with domain $\mathbb{U}$. If Q is a first-order definable generic query over M, then Q is a pure first-order query.*

We call this a *generic collapse* theorem, and say that an M satisfying the above implication, i.e an M satisfying:

If $Q \in FO_{GEN}(M)$ then Q is pure first-order

exhibits *generic collapse to equality*. Hence our first main result is that stable theories have generic collapse to equality.

A model M is *stable* if it does not admit a parametrically definable linear order. More precisely, there is no $\phi(\vec{x}, \vec{y})$ such that it is *consistent* with the theory of M that there are $\langle \vec{a}_n : n \in \omega \rangle$ such that $\phi(\vec{a}_m, \vec{a}_n) \leftrightarrow m < n$.

Standard model theory references such as [2], [9], [28] or [17] discuss stable theories at length, give several alternative definitions, and give numerous examples of stable structures. Standard examples of stable theories include algebraically closed fields, theories of equivalence relations, and any completion of the theory of abelian groups. The prototypical unstable structures are linearly-ordered structures (e.g. a real closed ordered field) and the random graph.

We now weaken the notion of genericity to deal with properties of linearly-ordered structures. Let S be a finite relational language. An *abstract (boolean) order query* Q is a collection of linearly ordered S-structures which is closed under $\{S, <\}$-isomorphism.

If $(\mathbb{U}, <)$ is a linearly ordered infinite set, then an abstract order query induces a query on this structure which satisfies a weaker kind of genericity. A *local map* is a partial function on $\mathbb{U}$. Given $\langle \mathbb{U}, < \rangle$, an S-structure A and a local map f whose domain contains the domain of A, fA is the image structure induced by f. A query on $\langle \mathbb{U}, < \rangle$, Q is *locally order generic* if it is preserved by local order-preserving maps on $\langle \mathbb{U}, < \rangle$. The query Q_O induced by an abstract order query Q on an infinite linearly ordered structure O is always locally order generic. Conversely, any locally order generic query naturally determines an abstract order query.

Let L be a language disjoint from S and let M be an $L' = L \cup \{<\}$-structure with domain $\mathbb{U}$. For any S-structure A with domain contained in $\mathbb{U}$, $M(A)$ denotes the unique $L \cup \{S, <\}$-structure that expands M and agrees with A on the interpretation of the predicates in S. We say that an abstract order query Q is *first-order over* M if there is a $\phi \in L \cup \{S, <\}$ that defines Q. That is $(A, <) \in Q \leftrightarrow M(A) \models \phi$.

We can now study the class

$FO_{OG}(M) = \{Q : Q$ is an order-generic query first-order definable over $M\}$

This class makes sense for any ordered structure M. Our goal is to show that for certain ordered structures, $FO_{OG}(M)$ is as small as it can be. An abstract order query Q is *order definable* if it is defined over the structure $\langle \mathbb{U}, < \rangle$ by an $\{S, <\}$-formula ψ.

Theorem B *Let $(\mathbb{U}, <)$ be an infinite linear order and M an $L \cup \{<\}$-structure that expands $(\mathbb{U}, <)$, which does not have the independence property. If Q is a first-order definable locally order generic query over M, then Q is an order-definable query.*

Roughly speaking, a structure M does not have the independence property (we also say M *lacks IP* or $M \in NIP$, for short), if there is no infinite parametrically definable independent family of subsets of M. Before we give the formal definition, we discuss the background and significance of the above result. Both Theorem B and the corollaries at the end of this section were proved for the real ordered group in [26]. This result was extended to o-minimal structures in [5] and also to the quasi-o-minimal structures of [4]. In contrast, it is easy to verify that there are models M for which any recursive query can be expressed over M (for example, $\langle N, +, *, < \rangle$ [14]); hence collapse of any sort fails for such structures. In fact, it is shown in [8] that even in the simplest structure not in NIP — the random graph — the parity query is definable and hence collapse fails.

Theorem B generalizes all previous collapse results, since the class NIP is huge. It includes all stable structures, all *o*-minimal and quasi-*o*-minimal structures (e.g. Pressburger arithmetic), all C-minimal structures (e.g. regular trees), all linear orders [27], and ordered abelian groups [15].

Definition 2.2

1. *The formula $\phi(\vec{x}, \vec{y})$ has the m-independence property in T if there exists a set of m, $\lg(x)$-tuples $\vec{b}_1, \ldots \vec{b}_m$ such that for every $X \subseteq \{1, \ldots, m\}$, there is an $\vec{a}_X$ such that $\phi(\vec{b}_i, \vec{a}_X)$ holds if and only if $i \in X$.*

2. *ϕ does not have the independence property in T just if for some integer m, $\phi(\vec{x}, \vec{y})$ does not have the m-independence property in T.*

3. *T does not have the independence property ($T \in NIP$) if no ϕ has the independence property in T. A model does not have the independence property if its theory does not have it.*

NIP is a basic combinatorial class. Like the class of stable structures, it is closed under taking reducts. Further, structures in NIP admit a Vapnik-Chervonenkis dimension, which implies that every parametrically definable family is PAC-learnable [23]. Other studies of NIP include [28], where a fundamental dichotomy between 'strict linear order' and 'the independence property' as sources of instability is proved, and [22].

Our next result compares the generic queries definable over a model M with the generic queries definable over a linear order. It is known that there are generic queries that are expressible by making use of an order on the structure which are not pure first-order. An example, due to Gurevich, is given in [1], and a consequence of this example for embedded finite models is that $FO_{GEN}(\langle \mathbb{U}, = \rangle)$ is a proper subset of $FO_{GEN}(\langle \mathbb{U}, < \rangle)$ for any linear order $<$. However, our next main result shows that adding extra structure beyond order does not yield new generic queries, if the model does not have the independence property. We will establish the following variant of Theorem B:

Theorem C *Let $\mathbb{U}$ be infinite and M be an L-structure with domain $\mathbb{U}$ that does not have the independence property. If Q is a first-order definable generic query over M, then Q is definable over any dense linear order without endpoints.*

A model satisfying the second sentence of Theorem C is said to have *generic collapse to order*. Note that in the collapse result above, and in the corollaries below, there is no assumption that the model M contains a definable (partial or total) order. It is known that an abstract order query definable in an infinite linear order is first-order definable as a class of finite ordered structures. This follows, for example, from the natural/active collapse result for o-minimal structures in [5], but was probably known earlier. This yields:

Corollary *Let $\mathbb{U}$ be infinite and M be an L-structure with domain $\mathbb{U}$ that does not have the independence property. If Q is a first-order definable generic query over M, then the abstract ordered query defined by Q (i.e the class of linearly ordered expansions of isomorphism types in Q) is definable by a formula in $\{S, <\}$.*

As a further corollary we have:

Corollary *Let S contain a single binary relation symbol. If M does not have the independence property, then the following queries are not first-order definable over M: parity, transitive closure of a graph, Hamiltonicity, connectivity of a graph or hypergraph.*

Proof: The examples are known not to be expressible over ordered structures [1]. □

It follows from this that parity and connectivity are not definable in abelian groups, algebraic closed fields, and the many other structures in NIP. It follows from Theorem A that the separating example of Gurevich mentioned above is not expressible in stable structures.

3 Quantifier Reduction in Expansions by Indiscernibles

The main tool for proving collapse results will be a theorem about extending a permutation of a set of indiscernibles to a family of partial isomorphisms. This approach to proving collapse is patterned on [5] and [24], but the extension result we need (Theorem 4.3) requires quite a bit of machinery to prove. In this section we prove a set of results about expansions of structures by a predicate for an indiscernible set or sequence. These results will be used in the next section to prove the key result about extending mappings on indiscernibles.

3.1 The Setting

We fix a first order theory T for a countable language L. Let $<$ be a binary relation symbol. If for some model M of T an infinite subset of M is linearly ordered by an L-formula, we may take $<$ to be this ordering; otherwise it is a new symbol.

For any set of formulas Δ and set A of elements of some model of T, $p = \mathrm{tp}_{\Delta}(\vec{b}/A)$ (the *Δ-type of $\vec{b}$ over A*) denotes the collection of A-instances of formulas from Δ satisfied by $\vec{b}$. If the length of $\vec{b}$ is r, we say p is a (Δ, r)-type. If Δ is just $=$, we say p is an *equality type.* If Δ is just $\{=, <\}$, we say p is an *order type* or $<$-type. If Δ is all of L, we say L-type. For finite sequences $\vec{t}_1, \vec{t}_2$, we write $\vec{t}_1 \equiv_{A,\Delta} \vec{t}_2$, if $\vec{t}_1$ and $\vec{t}_2$ have the same Δ-type over A; we frequently omit Δ if it clear from context. By an *equality formula* we mean a quantifier-free formula in equality and by an *order formula* we mean a quantifier-free formula in order and equality.

Definition 3.1 *A set I is a set of (Δ, L)-*indiscernibles *if any $\vec{a}, \vec{b}$ in I which have the same Δ-type have the same L-type. If $\Delta = \{x = y\}$ then we say I is a* set of pure indiscernibles*; if $\Delta = \{x < y, x = y\}$ then we say I is a* sequence of order indiscernibles.

In the rest of the paper we will prove results about a set of Δ-indiscernibles, where Δ is $\{x = y\}$ or $\{x < y, x = y\}$. In our results about pure indiscernibles we will always assume that T is stable, and in results about order-indiscernibles we assume only $T \in NIP$. Note that in the latter case we are not assuming that $<$ is contained in L. We will call the first case 'the equality case' and the second case 'the order case'; in the order case we will always take I to be an order-indiscernible set relative to some dense order without endpoints.

We note that such indiscernibles always exist: there is a model M of T which contains an infinite linearly ordered subset I of $<$- indiscernibles, and the order can be taken to be a complete dense linear order without endpoints [9]. A basic result of stability theory (e.g. [2] V.1.3) asserts that if T is stable I must be a set of pure indiscernibles.

We will assume, until noted otherwise, that

- If T is stable $\Delta = \{x = y\}$; otherwise $\Delta = \{x = y, x < y\}$
- L^+ is $L \cup \Delta$ plus a unary predicate P.
- M^+ is M expanded to interpret P by I and $<$ by the ordering of the indiscernible sequence.
- T^+ denotes the theory of (M, I).

We will call a structure as above *small* if further

- M is $|I|^+$-saturated.
- if $\Delta = \{=, <\}$, then the order type of $\langle I, <\rangle$ is a complete linear-order.

A structure for L^+ that is elementarily equivalent to a small (M, I) will be called *pseudo-small*, and the theory T^+ of a small model is called a *pseudo-small theory,* since a model of T^+ 'thinks' the set of indiscernibles is small.

The main result stated in this section is a quantifier-reduction result for such theories T^+. Namely, we show that if T is stable and I is a set of indiscernibles in $M \models T$, then every relation on M defined by an L^+-formula can be defined by one in which only *bounded* quantification over the set of indiscernibles is allowed. We show the analogous result for a theory T which does not have the independence property where I is a sequence of order indiscernibles.

Definition 3.2 *1. A formula is* basic *if it is a Boolean combination of L-formulas and Δ-formulas (i.e. equalities, if $\Delta = \{x = y\}$, or inequalities and inequalities if $\Delta = \{x = y, x < y\}$).*

*2. An L^+-formula ϕ is P-*bounded *if it is in the least collection of formulas containing the basic formulas and closed under Boolean operations and the quantifications $\exists x \in P$ and $\forall x \in P$.*

Henceforth $\forall \vec{w} \in P$ abbreviates $\forall w_1 \in P \ldots \forall w_i \in P$.

By the usual method of obtaining prenex normal forms, every P-bounded formula is equivalent to one consisting of a string of P-bounded quantifiers followed by a basic formula. We want to show that every L^+-definable relation is definable using only quantification over the indiscernibles. That is, our goal is to prove

Theorem 3.3 (The Main Reduction Theorem) *Let T be a stable theory (a theory without the independence property) and let T^+ be a pseudo-small extension of T. For each formula $\phi(\vec{x}) \in L^+$, there is an equivalent P-bounded formula $\phi'(\vec{x})$. That is,*

$$T^+ \models \forall \vec{x}\ [\phi(\vec{x}) \leftrightarrow \phi'(\vec{x})].$$

Note that the definition of basic (and hence P-bounded) is more restrictive in the stable case, so that the conclusion of the theorem states something stronger for stable T.

Before talking about the proof of Theorem 3.3, let us mention an important consequence of it, which will be used later:

Corollary 3.4 *For every L^+-formula $\phi(\vec{x}, \vec{y})$ there is a Δ-formula $\psi(\vec{w}, \vec{y})$ such that for every $\vec{m}$ there is a $\vec{c}_{\vec{m}} \in I$ such that*

$$\forall \vec{y} \in P[\psi(\vec{c}_{\vec{m}}, \vec{y}) \leftrightarrow \phi(\vec{m}, \vec{y})].$$

That is, all L^+-types over I are Δ-definable.

Proof. By Theorem 3.3 we can assume that $\phi(\vec{x}, \vec{y})$ is P-bounded. Apply Proposition 3.5 to the quantifier-free part of ϕ to get a P-bounded formula $\delta(\vec{w}, \vec{y})$ so that for each $\vec{m}$ in M there is $\vec{c}_{\vec{m}}$ in I such that:

$$M^+ \models \forall \vec{y} \in P\ [\phi(\vec{m}, \vec{y}) \leftrightarrow \delta(\vec{c}_{\vec{m}}, \vec{y})].$$

Rewrite $\delta(\vec{w}, \vec{y})$ as $(Q_1 z_1)\dots(Q_k z_k)\ \eta(\vec{w}, \vec{y}, \vec{z})$ where the Q_i's are P-bounded quantifiers bounding z_i and $\eta \in L$. Now using the Δ-indiscernibility of I, it easy to show by induction on quantifiers that $(Q_1 z_1)\dots(Q_k z_k)\ \eta(\vec{w}, \vec{y}, \vec{z})$ is equivalent over I to a Δ-formula. $\square_{3.4}$

3.2 Brief sketch of the proof of the 'hard' theorem

Theorem 3.3 is elementary but extremely involved. We include here a brief sketch of how it is proved. The first step it to prove the following proposition, which is used both in proving the reduction result and in applying it. In the proposition, we assume the context of Theorem 3.3: e.g. T is stable for $\Delta = \{x = y\}$ and $T \in NIP$ if $\Delta = \{x = y, x < y\}$.

Proposition 3.5 *For every basic formula $\phi(\vec{x}, \vec{y})$ there is a quantifier free Δ-formula $\psi(\vec{w}, \vec{y})$ such that for every $\vec{m} \in M$ there is a $\vec{c}_{\vec{m}} \in I$ such that*

$$\forall \vec{y} \in P[\psi(\vec{c}_{\vec{m}}, \vec{y}) \equiv \phi(\vec{m}, \vec{y})].$$

For T stable, the above proposition is well-known. In fact, the statement that *all* types are definable in the sense of this proposition is another characterization of stable theories (Theorem II.2.12 of [28]). For $T \in NIP$, the above result asserts the definability of types over sequences of indiscernibles; this result is new here, although the case $|\vec{y}| = 1$ is Theorem II.4.13 of [28]. The proof is a fairly straightforward argument, using completeness of the order.

The proof of Theorem 3.3 is an induction, using Proposition 3.5 in its base case. We now list definitions that will be necessary to understand the induction.

Definition 3.6 *Fix a structure M, finite sets Δ, Δ_1 of formulas, and I a set (Δ, L)-indiscernibles. For A, A_0, A_1 finite subsets of I with $A \subseteq A_1$ and some $\vec{n} \in M$, $p = \mathrm{tp}(\vec{n}/A_1)$* Δ_1-splits in A over A_0 *if there exist sequences $\vec{c}, \vec{d} \in A$ which have the same Δ-type over A_0, but for some $\beta \in \Delta_1$, $\vec{c}$ and $\vec{d}$ disagree on $\vec{n}$ for β. (That is, $\neg[\beta(\vec{n}, \vec{c}) \leftrightarrow \beta(\vec{n}, \vec{d})]$).*

If $A = A_1$, the 'in A' is omitted; frequently A_1 will be I. (In e.g. [28] this is called (Δ, Δ_1)-splitting, but we will always have Δ fixed so we write only one parameter.) When $\Delta_1 = \{\beta\}$ we say there is a β-split for $\vec{n}$ over A_0. Since the second β in the phrase '$p = \mathrm{tp}_\beta(\vec{n}/A_1)$ does not β-split' is redundant, it is often omitted. Note that if $p = \mathrm{tp}_\beta(\vec{n}/A_1)$ does not β-split over a finite subset C of A_1, then there is a quantifier free Δ-formula $\delta(\vec{w}, \vec{y})$ and a $\vec{c} \in C$ such that for any $\vec{a} \in A_1$,

$$\beta(\vec{n}, \vec{a}) \leftrightarrow \delta(\vec{c}, \vec{a}).$$

Here δ can be taken to be the disjunction of the finitely many, necessarily quantifier free (by quantifier elimination in the theories of equality and dense-linear order), complete Δ-types over C of elements that satisfy β. We say $p = \mathrm{tp}_\beta(\vec{n}/A_1)$ is Δ-*defined* over C.

We now show that to prove the main quantifier reduction lemma, it suffices to handle a certain class of $\exists\forall$ formulas.

Proposition 3.7 *Let I be a set of (Δ, L) - indiscernibles in L-structure M. Suppose that (M, I) satisfies the following condition. For every basic formula $\beta(x_1, \dots, x_n, z, \vec{y})$ and for every Δ-formula $\delta(\vec{w}, \vec{y})$ with $y_1, \dots y_R = \vec{y}$, the formula*

$$\gamma(\vec{x}) = \exists z\ \exists \vec{w} \in P\ \forall \vec{y} \in P\ [\beta(\vec{x}, z, \vec{y}) \leftrightarrow \delta(\vec{w}, \vec{y})]$$

is equivalent in M^+ to a P-bounded formula.

Then every L^+-formula is equivalent in $T^+ = \mathrm{Th}(M, I)$ to a P-bounded formula.

Proof: We first observe that the assertion that every L^+-formula is equivalent to a P-bounded formula is preserved by L^+-elementary-equivalence. Thus we need only establish this assertion in (M, I) to get it as a consequence of T^+. For quantifier-free L^+-formulas, the assertion follows by relativizing to P (e.g. replace $\neg P(x)$ by $\forall y \in P(y \neq x)$). To prove the conclusion by induction on quantifier rank, it suffices to remove one existential unbounded quantifier from the front of a P-bounded formula. That is, to convert $\rho(\vec{x}) = \exists z (Q_1 y_1) \ldots (Q_R y_R) \beta(\vec{x}, z, \vec{y})$ to a P-bounded formula. But, by Proposition 3.5, for each L-formula β and for any Δ-schema δ defining β, $\rho(\vec{x})$ is equivalent to:

$\exists z\ \exists \vec{w} \in P\ \forall \vec{y} \in P\ (\beta(\vec{x}, z, \vec{y}) \leftrightarrow \delta(\vec{w}, \vec{y}))$
$\wedge (Q_1 y_1) \ldots (Q_R y_R) \delta(\vec{w}, \vec{y})$.

Applying the hypothesis of this proposition and induction we can convert this statement to a P-bounded formula; hence the proposition is proved. $\square_{3.7}$
Thus, to prove our main quantifier reduction result for a pseudo-small T^+, we need only show the hypothesis of Proposition 3.7 holds in an $|I|^+$-saturated model M of T. We briefly say how this is done for the stable case. The case of $T \in NIP$ is similar, but significantly more tedious. *For the rest of this subsection we assume T is stable.*

The *trace* of a formula $\phi(\vec{m}, \vec{y})$ on I is the set of $\vec{c}$ from I satisfying $\phi(\vec{m}, \vec{c})$. If M is a model of a stable theory, then Proposition 3.5 says that the trace of an L-formula on a set of indiscernibles is uniformly defined with parameters in that set. The next lemma is simply a restatement of 3.5 for a partitioning of the parameter variables into $\vec{x}, z$ (where the partitioning is used in the application).

Lemma 3.8 *For every L-formula $\beta(\vec{x}, z, \vec{y})$ there is a (quantifier-free) equality formula $\psi(\vec{w}, \vec{y})$ satisfying: for every $\vec{m}, b$ there is a $\vec{c}_{\vec{m},b} \in I$ such that*

$$\forall \vec{y} \in P[\beta(\vec{m}, b, \vec{y}) \leftrightarrow \psi(\vec{c}_{\vec{m},b}, \vec{y})].$$

The following definitions establish crucial notation for the proof.

Notation 3.9
(J=J_β) Fix $\beta(x_1, \ldots, x_n, z, y_1, \ldots, y_R)$ in L. Now fix an equality formula $\psi(\vec{w}, \vec{y})$ which provides a uniform definition for the trace of β on I^R. Let $J = J_\beta$ be the length of $\vec{w}$.

Definition 3.10 *Two formulas $\phi(\vec{x}, w_1, \ldots w_i)$ and $\phi'(\vec{x}, w_1, \ldots w_i)$ are called* $\vec{w}, I$-equivalent *if $\forall \vec{x}\ \forall w_1 \in P, \ldots\ \forall w_i \in P$*
$[\phi(\vec{x}, w_1, \ldots, w_i) \leftrightarrow \phi'(\vec{x}, w_1, \ldots, w_i)]$.

Let C and E be subsets of I. We say $\mathrm{tp}_\beta(\vec{m}, d/I)$ *splits in E over every proper subset of C* if for every proper subset C' of C, there are tuples $\vec{t_1}, \vec{t_2} \in E$ such that $\vec{t_1} \equiv_{C'} \vec{t_2}$ and $\neg[\beta(\vec{m}, d, \vec{t_1}) \leftrightarrow \beta(\vec{m}, d, \vec{t_2})]$.

To establish the hypothesis of Proposition 3.7, we show the formula $\exists z\ \exists \vec{w} \in P\ \forall \vec{y} \in P\ [\beta(\vec{x}, z, \vec{y}) \leftrightarrow \delta(\vec{w}, \vec{y})]$ is equivalent on M^+ to a P-bounded formula. We do this by introducing an auxiliary formula $\psi_\delta(\vec{x}, \vec{w})$.

Definition 3.11 *(ψ_δ) For any equality-formula $\delta(w_1, \ldots, w_n, y_1, \ldots, y_R)$, let $\psi_\delta(\vec{x}, w_1, \ldots w_n)$ denote the L^+-formula such that, for arbitrary $\vec{m}$ and $\vec{c} \in I$ $\psi_\delta(\vec{m}, \vec{c})$ holds if and only if the c_i are distinct and for some d,*

1. *$\delta(\vec{c}, \vec{y})$ defines $\mathrm{tp}_\beta(\vec{m}, d/I)$ and*
2. *$\mathrm{tp}_\beta(\vec{m}, d/I)$ splits over every proper subset of C, where C is the range of $\vec{c}$.*

Proposition 3.12 *For every $\vec{m}$ there exists $n \leq J_\beta$ and an equality formula $\delta(w_1, \ldots w_n, \vec{y})$ such that for some $\vec{c}$ from I, $\psi_\delta(\vec{m}, \vec{c})$ holds.*

Proof. First, choose any set $\vec{c}_1$ such that $\psi(\vec{c}_1, \vec{y})$ defines $\mathrm{tp}_\beta(\vec{m}, d/I)$. Then choose a minimal subset of $\vec{c}_1$, enumerated by $\vec{c}$, over which $\mathrm{tp}_\beta(\vec{m}, d/I)$ does not split. Now, $\mathrm{tp}_\beta(\vec{m}, d/I)$ is, in fact, defined over $\vec{c}$ (by the disjunction of the equality types over $\vec{c}$ of sequences $\vec{e}$ satisfying $\beta(\vec{m}, d, \vec{e})$); this formula is the required δ. $\square_{3.12}$

Let D_J be a finite set of equality formulas such that every equality formula $\delta'(w_1, \ldots w_i, \vec{y})$ with $i \leq J = J_\beta$ is equivalent to one in D_J. Applying Proposition 3.12, we see that for each $\delta(\vec{w}, \vec{y})$ there is a subset S of D_J such that for all $\vec{m}$ in M, we have $M^+ \models\ \exists \vec{w} \in P\ \exists z\ \forall \vec{y} \in P\ [\beta(\vec{m}, z, \vec{y}) \leftrightarrow \delta(\vec{w}, \vec{y})]$ if and only if $M^+ \models \bigvee_{\delta' \in S} \exists \vec{w} \in P\ \psi_{\delta'}(\vec{m}, \vec{w})$. It follows that the hypothesis of Proposition 3.7 and thus Theorem 3.3 for the stable case follow from

Lemma 3.13 (Induction Lemma for Stable Case) *For each L-formula $\beta(\vec{x}, z, \vec{y})$ and each equality formula $\delta(w_1, \ldots w_n, \vec{y})$ with $\lg(\vec{w}) \leq J_\beta$, $\psi_\delta(\vec{x}, \vec{w})$ is $\vec{w}$,I equivalent to a P-bounded formula. That is, there is P-bounded $\psi'(\vec{x}, \vec{w})$ such that:*

$$T^+ \models \forall \vec{x}\ \forall \vec{w} \in I\ [\psi_\delta(\vec{x}, \vec{w}) \leftrightarrow \psi'(\vec{x}, \vec{w})].$$

It now only remains to show Lemma 3.13: this is proved by downward induction on the length of $\vec{w}$, starting at J_β. The base case follows easily by using the definition of J_β, while the induction step requires a bit of work. The proof appears in the full paper.

The case of $T \in NIP$ is done via a similar induction. $\square_{3.3}$

4 Applications of the Quantifier Reduction Results to Query Collapse Theorems

We now apply the results on expansions by indiscernibles proved in the previous section to prove the expressivity results on embedded finite structures stated in Section 2. We start in Subsection 4.1 by stating a theorem on permutations of indiscernibles that captures all we need from the previous section. Subsection 4.2 begins the applications to query collapse; we present a lemma allowing us to concentrate a number of finite partial counterexamples into a single infinite counterexample. We then combine that result with the extension result of the Subsection 4.1 to prove the main collapse results stated at the beginning of the paper.

4.1 Extending Maps on Indiscernibles

We need to fix some notation. Recall that a partial L-isomorphism from a structure A to a structure B is an injective map from a subset of A to B which preserves relations from L. We recall the term 'partially isomorphic' from [10] and [9], which is also referred to as 'back-and-forth equivalence' [17].

Definition 4.1 *For any language L, two L-structures A and B are* partially L-isomorphic, *written $A \approx_L B$ if there exists a non-empty family J of partial L-isomorphisms which have the back and forth property: for each $j \in J$ and x in A or B there is $j' \in J$ that extends j with x in the domain or range of the graph of j'.*

Note that if $A \approx_L B$, then A and B are elementarily equivalent in L.

Notation 4.2 *Let I be a subset of N and let L^* be the expansion of L^+ by constants $\{c_\alpha : \alpha < |I|\}$. Fix a permutation f of I and an enumeration $\{a_\alpha : \alpha < |I|\}$ of I. Then N_1^f is the expansion of N where c_α is interpreted as a_α and N_2^f is the expansion of N where c_α is interpreted as $f(a_\alpha)$. We will omit the f if it is fixed in context.*

Note that $N_1^f \approx_{L^*} N_2^f$ implies that the partial isomorphisms in the back-and-forth system can all be taken to extend f. Here is the result that is relevant for collapse theorems:

Theorem 4.3 *Let N be a model that is stable (respectively, does not have the independence property) and I be a set of indiscernibles (resp. sequence of order indiscernibles). Suppose that (N, I) is pseudo-small and that (N, I) is ω_1-saturated.*

For any permutation f on I (order-preserving permutation of I),

$$N_1^f \approx_{L^*} N_2^f.$$

Proof: We use Δ to mark whether we are in the order or equality case. Note that in the order case, the definition of pseudo-small requires that (N, I) is elementary equivalent to (N', I') where I' is an order-indiscernible sequence with complete dense order type. The quantifier-reduction results proved in Section 3 are all properties of the L^+-theory, and hence carry over from small (N', I') to pseudo-small (N, I).

It suffices to show that the collection of pairs of finite sequences $\langle \vec{a}, \vec{b} \rangle$ such that $\langle N_1^f, \vec{a} \rangle$ is L^*-elementary equivalent to $\langle N_2^f, \vec{b} \rangle$ form a back-and-forth system. So suppose $\langle N_1^f, \vec{a} \rangle$ is L^*-elementarily equivalent to $\langle N_2^f, \vec{b} \rangle$, and a' is an element of N_1^f. We shall find b' such that $\langle N_1^f, \vec{a}, a' \rangle$ is L^*-elementarily equivalent to $\langle N_2^f, \vec{b}, b' \rangle$. It suffices to find a b' such that for each $\rho(x, \vec{w}, \vec{v})$ in L^+, the f-image of the $\rho(x, \vec{w}, \vec{v})$-type of $a', \vec{a}$ over I is the same as the $\rho(x, \vec{w}, \vec{v})$-type of $b', \vec{b}$ over I.

By Corollary 3.4, we know that there is a quantifier-free Δ-formula $\gamma_\rho(\vec{v}, \vec{v}')$ such that the ρ-type of $a', \vec{a}$ over I is definable by an instance of γ_ρ. That is, there is $\vec{i}_\rho \in I$ such that for each $\vec{c} \in I$ $\gamma_\rho(\vec{c}, \vec{i}_\rho)$ holds if and only if $\rho(a', \vec{a}, \vec{c})$. Let $\vec{j}_\rho$ be $g(\vec{i}_\rho)$. Since $\langle N_1, \vec{a} \rangle$ is L^*-elementarily equivalent to $\langle N_2, \vec{b} \rangle$, and $\langle N_1, \vec{a} \rangle$ satisfies $\exists x \forall \vec{v} \in P(\gamma_\rho(\vec{v}, \vec{i}_\rho) \leftrightarrow \rho(x, \vec{w}, \vec{v}))$, we have $\langle N_1, \vec{b} \rangle$ satisfies $\exists x \forall \vec{v} \in P(\gamma_\rho(\vec{v}, \vec{j}_\rho) \leftrightarrow \rho(x, \vec{w}, \vec{v}))$. Let I_0 be the union of the ranges of all the sequences $\vec{i}_\rho$ and let $q(x)$ be the type over $I_0 \cup \{\vec{a}\}$ containing the formulas $\forall \vec{v} \in P(\gamma_\rho(\vec{v}, \vec{i}_\rho) \leftrightarrow \rho(x, \vec{a}, \vec{v}))$ for all L^+-formulas ρ. Then I_0 is countable and the type $g_1(q)$ (where $g_1 = g \cup \langle \vec{a}, \vec{b} \rangle$) over the countable set $J_0 \cup \{\vec{b}\}$, where $J_0 = g_1(I_0)$, is also consistent. So by ω_1-saturation of (N, I) we can choose b' to realize $g_1(q)$ and complete this part of the back-and-forth. The other direction of the back-and-forth is done by a symmetric argument, completing the theorem. $\square_{4.3}$

4.2 Connecting Query Collapse to Indiscernibles, and Proof of the Main Collapse Theorems.

We state the connection between indiscernibles and query collapse results. Specifically, we give a general model-theoretic argument (Lemma 4.5) that applies in both the stable and NIP cases, and which connects the query collapse questions mentioned in the introduction to results about indiscernibles. This technique is implicit in [5] and [3]. The formulation and proof of Lemma 4.5 here use elementary extension arguments patterned directly on the proof of Lindstrom's theorem reported in [11] rather than nonstandard analysis as in [5] or the special models used in [3]. Finally, we use the concentration result, Lemma 4.5, and Theorem 4.3 from the previous section to prove the main theorems. As before, we will use a parameter Δ to mark whether we are in the order or unordered case: Δ is either $\{=\}$ or $\{=,<\}$.

Let M be an $L \cup \Delta$-structure with an infinite domain $\mathbb{U}$. Let $\vec{A}_n, \vec{B}_n$ be a sequence of finite relations interpreting the symbols of S in $\mathbb{U}$. We will sometimes write A_n, B_n for the structure whose universe is the union of the domains of the relations $\vec{A}_n, \vec{B}_n$ (the *active domain*).

For any two S-structures A and B we write $A =_{n,S,\Delta} B$ or say the structures are n, S, Δ *equivalent* if the two structures agree on the first n sentences of the language with vocabulary $\{S \cup \Delta\}$. Two frequent uses of this notation will be $(\mathbb{U}, \vec{A}_n) =_{n,S,\Delta} (\mathbb{U}, \vec{B}_n)$ and $A_n =_{n,S,\Delta} B_n$ where $\vec{A}_n, \vec{B}_n$ are finite relations on $\mathbb{U}$. We will often omit Δ when it is equality and just write $=_{n,S}$. Note that the expressions $(M, \vec{A})$ and $M(A)$ mean the same thing.

Remark 4.4 *Since the relativization of any Δ-sentence to* $\operatorname{dom} \vec{A}$ *can be expressed as a sentence about $(\mathbb{U}, \vec{A}_n)$, for any A_n, B_n with $(\mathbb{U}, \vec{A}_n) =_{n,S,\Delta} (\mathbb{U}, \vec{B}_n)$ and $A_n \models \phi$ but $B_n \models \neg\phi$ there are A'_n, B'_n with finite domains contained in $\mathbb{U}$ such that $A'_n =_{n,S,\Delta} B'_n$, $A'_n \models \phi$, and $B'_n \models \neg\phi$. The converse also holds.*

Recall from the first section that $M = \langle \mathbb{U}, \ldots \rangle$ exhibits *collapse to equality* if every Δ-generic query definable over M is definable in the structure $\langle \mathbb{U}, = \rangle$. We will prove collapse by showing that a failure of collapse yields a family of finite structures as in the hypotheses of Lemma 4.5 below. We will apply Lemma 4.5 to produce a single infinite counterexample, and then we'll apply Theorem 4.3 to derive a contradiction from that counterexample.

Lemma 4.5 (Concentration Lemma) *Let I be a set of (Δ, L) indiscernibles in a structure M with universe $\mathbb{U}$. Suppose there are finite $\Delta \cup S$-interpretations $\vec{A}_n$, $\vec{B}_n$ for each n with* $\operatorname{dom} A_n$, $\operatorname{dom} B_n \subseteq I$, *$(I, \Delta, \vec{A}_n) =_{n,S} (I, \Delta, \vec{B}_n)$, $(M, \vec{A}_n) \models \phi$, and $(M, \vec{B}_n) \models \neg\phi$.*

Then there is a model M' of $Th(M)$, a set of Δ-indiscernibles I' in M' with $(M, I) \equiv_{L^+} (M', I')$, two infinite S-interpretations $\vec{T}$ and $\vec{U}$ with domains in I', and a map f on I' such that f is an isomorphism from $(I', \vec{T}, \Delta)$ to $(I', \vec{U}, \Delta)$ but $M'(\vec{T}) \models \phi$ and $M'(\vec{U}) \models \neg\phi$.

Proof. Let $L^+ = L \cup \Delta \cup \{\mathbf{P}\}$ where $\mathbf{P}$ is a new unary predicate symbol, and let $\langle M, I \rangle$ be the L^+-structure where $\mathbf{P}$ is interpreted by I. Form L^{++} by adding further unary predicate symbols $\vec{\mathbf{T}}$, $\vec{\mathbf{U}}$ corresponding to the symbols in S and, for each n, two $n+2$-ary functions symbols $\mathbf{f}_n$ and $\mathbf{g}_n$.

The following sentences describe approximations to a back-and-forth. Let $\chi_i(x_1, \ldots x_n, x)$ for $i < \omega$ enumerate the $n+1$-ary formulas in $\Delta \cup \vec{\mathbf{T}}$ and $\chi'_i(x_1, \ldots x_n, x)$ for $i < \omega$ enumerate the associated $n+1$-ary formulas in $\Delta \cup \vec{\mathbf{U}}$. Then for any $r < \omega$, let $\theta_{n,r}(\mathbf{P}, \vec{\mathbf{T}}, \vec{\mathbf{U}}, \mathbf{f}_n, \mathbf{g}_n)$ contain the assertion that f_n and g_n act within $\mathbf{P}$ and the conjunction of

$\forall \vec{x} \in \mathbf{P}, \forall \vec{y} \in \mathbf{P}, \forall x \in \mathbf{P}$

$$(\exists y \in \mathbf{P}(\bigwedge_{0<i<r} (\chi_i(\vec{x}, x) \leftrightarrow \chi'_i(\vec{y}, y)))$$

$$\rightarrow \bigwedge_{0<i<r} (\chi_i(\vec{x}, x) \leftrightarrow \chi'_i(\vec{y}, \mathbf{f}_n(\vec{x}, \vec{y}, x))))$$

and

$$\forall \vec{x} \in \mathbf{P}, \forall \vec{y} \in \mathbf{P}, \forall x \in \mathbf{P}$$

$$(\exists y \in \mathbf{P}(\bigwedge_{0<i<r} (\chi_i(\vec{x}, x) \leftrightarrow \chi'_i(\vec{y}, y)))$$

$$\rightarrow \bigwedge_{0<i<r} (\chi_i(\vec{x}, \mathbf{g}_n(\vec{x}, \vec{y}, y)) \leftrightarrow \chi'_i(\vec{y}, y))).$$

Let Γ^* be a collection of L^{++}-sentences which assert:

1. $\operatorname{dom} \vec{\mathbf{T}} \subset \mathbf{P}$, $\operatorname{dom} \vec{\mathbf{U}} \subset \mathbf{P}$.
2. the L^+-theory of (M, I, Δ).
3. $\phi(\vec{\mathbf{T}})$, $\neg\phi(\vec{\mathbf{U}})$.
4. for each n, r, $\theta_{n,r}(\mathbf{P}, \vec{\mathbf{T}}, \vec{\mathbf{U}}, \mathbf{f}_n, \mathbf{g}_n)$.

The consistency of finite subsets of item 4 follows from $(I, \Delta, \vec{A}_n) =_{n,S} (I, \Delta, \vec{B}_n)$. Let M' be a countable model of Γ^* and suppose $I, \vec{U}, \vec{T}, f_n, g_n$ interpret $\mathbf{P}, \vec{\mathbf{U}}, \vec{\mathbf{T}}, \mathbf{f}_n, \mathbf{g}_n$. Then, the $\mathbf{f}_n$ and $\mathbf{g}_n$ determine a back-and-forth system which guarantees that $(P, \Delta, \vec{U})$ and

$(P, \Delta, \vec{T})$ are isomorphic $\Delta \cup S$-structures, giving us the required isomorphism f. $\Box_{4.5}$

We now combine Lemma 4.5 with Theorem 4.3 to get the collapse results, beginning with stable T and generic queries:

Theorem 4.6 *Let $M \models T$ where T is stable. Every generic query Q which is $L \cup S$-definable by a formula ϕ in M is S-definable in M.*

Proof. If Q is not defined by any formula θ in S, there are finite relations $\vec{A}_n$, $\vec{B}_n$ with $\operatorname{dom} A_n, \operatorname{dom} B_n \subseteq \mathbb{U}$, $(\mathbb{U}, \vec{A}_n) =_{n,S} (\mathbb{U}, \vec{B}_n)$, and $M(A_n) \models \phi$ and $M(B_n) \models \neg\phi$. By Remark 4.4 we have (possibly rechoosing the A_n, B_n) that $A_n =_{n,S} B_n$. Since T is stable, we can fix a countable *set* of indiscernibles $I = \{i_0, i_1, \ldots\}$ in M. By Fact 2.1, we can replace M by an $|I|^+$ saturated model and thus assume (M, I) is pseudo-small. Let $\vec{A}'_n$ $(\vec{B}'_n)$ be the image of $\vec{A}_n$ $(\vec{B}_n$) under an injective map taking the domain of $\vec{A}_n$ $(\vec{B}_n)$ to finite subsets of I. Then we have $\vec{A}'_n =_{n,S} \vec{B}'_n$, $(M, \vec{A}'_n) \models \phi$ and $(M, \vec{B}'_n) \models \neg\phi$. The first statement follows from the fact that we are taking isomorphic images of all relations in S, and the second from the genericity of ϕ. Reversing the use of Remark 4.4, we may assume $(I, \vec{A}'_n) =_{n,S} (I, \vec{B}'_n)$.

We now have the hypotheses of Lemma 4.5 with $\Delta = \{x = y\}$ (for $(I, \vec{A}'_n), (I, \vec{B}'_n)$); we apply that result to get a model (M', I'). Let the infinite S-structures $\vec{T}, \vec{U}$ and the function f on I' witness the conditions on M' guaranteed in Lemma 4.5. Let N^* be an ω_1-saturated model of $\operatorname{Th}(M', I', \vec{T}, \vec{S}, f)$. So in N^* we have $(I, \vec{U})$ and $(I, \vec{T})$ are S-isomorphic but $\phi(\vec{\mathbf{T}})$ and $\neg\phi(\vec{\mathbf{U}})$ hold; moreover, $N^+ = N^*|L^+$ is pseudo-small. Expand N^+ to models N_1 and N_2 for the language L^* as defined in Notation 4.2.

Now, we apply Theorem 4.3: this tells us exactly that there is a family of partial-isomorphisms between N_1 and N_2 that each extend f. Since f is an S-isomorphism and $\vec{U}, \vec{T}$ are contained in I, the back and forth system shows, in fact, that $(N_1, \vec{T})$ and $(N_2, \vec{U})$ are $L^* \cup S$-elementarily equivalent. Now we have a contradiction, since the reducts of these models to $L \cup S$ disagree on the sentence ϕ. $\Box_{4.6}$

Theorem A follows immediately from Theorem 4.6. We now show the analagous results to Theorem 4.6 for the ordered case. We first have

Proof of Theorem B:

Suppose the Theorem fails and let ϕ witness that M does not have locally-generic collapse to order (i.e. $\phi \in L \cup S$ defines a locally-generic query Q but Q is not definable over the reduct $(\mathbb{U}, <)$). Since Q is not defined over $(\mathbb{U}, <)$, there are finite relations $\vec{A}_n$, $\vec{B}_n$ with $\operatorname{dom} A_n, \operatorname{dom} B_n \subseteq \mathbb{U}$ with $(\mathbb{U}, \vec{A}_n) =_{n,S,<} (\mathbb{U}, \vec{B}_n)$ and $M(\vec{A}_n) \models \phi$ and $M(\vec{B}_n) \models \neg\phi$. As in the proof of Theorem 4.6, we may assume $A_n =_{n,S,<} B_n$. And, again by Fact 2.1, we may assume M is ω_1-saturated and contains a countable set I that is order indiscernible and complete densely ordered by the original order $<$. Thus, (M, I) is pseudo-small. Let $\vec{A}'_n$ $(\vec{B}'_n)$ be the image of $\vec{A}_n$ $(\vec{B}_n$) under an order-preserving injective map taking the domain of $\vec{A}_n$ $(\vec{B}_n)$ to a finite subset of I. Then we have $\vec{A}'_n =_{n,S,<} \vec{B}'_n$, $M(\vec{A}'_n) \models \phi$ and $M(\vec{B}'_n) \models \neg\phi$. The first statement follows from the fact that we are taking images of all relations in S and preserving order, and the second from the local genericity of ϕ.

Invoking Remark 4.4 again, we have the hypotheses of Lemma 4.5 with Δ as $\{x = y, x < y\}$. We get a model with a sequence of order indiscernibles which contains $S \cup \{=, <\}$-isomorphic $(I, \Delta, \vec{T})$ and $(I, \Delta, \vec{U})$, but with $M(T) \models \phi$ and $M(U) \models \neg\phi$. We finish the proof exactly as in the the last paragraph of the proof of Theorem 4.6, using Theorem 4.3 for theories without the independence property. $\Box_B$

Finally, we get the proof of **Theorem C:** Suppose Q is not definable over some dense linear order $\langle D, < \rangle$, but Q is defined by $\phi \in L \cup \{S\}$ over M. Let I be a set of indiscernibles in M that is ordered by a dense linear order without endpoints (not necessarily in L). By the genericity of Q and the completeness of the theory of dense-linear order without endpoints, we know that Q is not definable in the structure $\langle I, < \rangle$.

Once again there are finite relations A_n, B_n with $\operatorname{dom} A_n, \operatorname{dom} B_n \subseteq I$, $(I, <, \vec{A}_n) =_{n,S,<} (I, <, \vec{B}_n)$ and $M(A_n)$ satisfies ϕ while $M(B_n)$ doesn't satisfy ϕ (where the latter follows from the definability of Q by ϕ). Again, we have the hypotheses of Lemma 4.5 and we construct (M', I'). We now have a model (M', I') which is pseudo-small and in which I is linearly ordered by $<$. Moreover, there are $S \cup \{<\}$-isomorphic $(I, \Delta, \vec{T})$ and $(I, \Delta, \vec{U})$, but with $M(T) \models \phi$ and $M(U) \models \neg\phi$. By applying Theorem 4.3, we get that this isomorphism extends to a family of partial-isomorphisms of (M', I'), and this gives a contradiction exactly as in the proof of Theorem 4.6. $\Box_C$

5 Conclusions

The paper connects questions explored in finite model theory and constraint databases with results in stability theory, a connection that we feel is both surprising and fruitful. The results here generalize known results in finite model theory, but the techniques used here could also be used to get new expressive bounds in both finite model theory and in database theory; e.g. to obtain bounds on the expressive power of first and second-order logics on finite models with interpreted functions, along the lines of [6].

We now discuss some possible extensions. We have shown that models without the independence property exhibit generic collapse to order. One can ask if the converse holds: are there models with the independence property that also have generic collapse? The answer is yes; generic collapse requires only that the model be well-behaved *locally*; e.g. within some definable or ∞-definable set. A model that is the disjoint union of a stable structure and the badly-behaved structure $\langle N, +, *, < \rangle$ will thus still exhibit generic collapse. An interesting open problem is to see if generic collapse to order holds for models that fail to have the independence property within any definable set, and if generic collapse to equality holds for models that have empty stable part [22].

Acknowledgements: The goal of extending the collapse-to-order results for o-minimal structures in [5] to a class of models closed under reduct was suggested by Jan Van den Bussche. We also benefited from suggestions from Dave Marker, Anand Pillay, and Leonid Libkin.

References

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

[2] J. T. Baldwin. *Fundamentals of Stability Theory.* Springer-Verlag, 1988.

[3] O.Belagradek, A. Stolboushkin, M. Tsaitlin. Extended order-generic queries. DIMACS technical report 96-01.

[4] O.Belagradek, Ya'acov Peterzil, Frank Wagner. Quasi O-minimal groups. Manuscript. To appear.

[5] M. Benedikt, G. Dong, L. Libkin, L. Wong. Relational expressive power of constraint query languages. To appear in JACM. Extended abstract in PODS '96.

[6] M. Benedikt and L. Libkin. Safe Constraint Queries Manuscript to appear.

[7] M. Benedikt and L. Libkin. On the structure of queries in constraint query languages. In *LICS '96.*

[8] M. Benedikt and L. Libkin. Languages for Relational Databases over Interpreted Structures. In *PODS '97.*

[9] C.C. Chang and H.J. Keisler. *Model Theory.* North-Holland Elsevier, 1990.

[10] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer Verlag, 1995.

[11] J. Flum. Characterizing logics. In *Model-Theoretic Logics, edited by J. Barwise and S. Feferman, pages 78-91, Springer-Verlag, 1985.*

[12] E. Grädel and Y. Gurevich. Metafinite model theory. In *Proc. LCC*, LNCS vol. 960, 1994, pages 313–366.

[13] S. Grumbach and J. Su. Finitely representable databases, In *PODS'94*, pages 289–300.

[14] S. Grumbach and J. Su. First-order definability over constraint databases. In *PCCP'95.*

[15] Y. Gurevich and P. Schmitt The theory of ordered abelian groups does not have the independence property. In *Trans. American Math Society*, 284 (1984) pages 171–182.

[16] L. Hella, P. Kolaitis, and K. Luosto. How to define a linear order on finite models. In *LICS '94.*

[17] W. Hodges. *Model Theory.* Cambridge University Press, 1992.

[18] R. Hull and J. Su. Domain independence and the relational calculus. *Acta Informatica* 31:513–524, 1994.

[19] N. Immerman. Languages that capture complexity classes. *SIAM J. Comput.* 16 (1987), 760–778.

[20] P. Kanellakis. Constraint programming and database languages: A tutorial. In *PODS'95*, pages 46–53.

[21] P. Kanellakis, G. Kuper, and P. Revesz. Constraint query languages. *JCSS* 51 (1995), 26–52.

[22] H.J. Keisler. Measures and Forking. Annals of Pure and Applied Logic 34 (1987), no. 2, 119-169.

[23] C. Laskowski. Vapnik-Chervonenkis classes of definable sets. J. of the London Mathematical Society. Vol. 45 (1992), no. 2, 377-384.

[24] M. Otto and J. Van den Bussche. First-order queries on databases embedded in an infinite structure. *Information Processing Letters*, 60 (1996), no. 1,37-41.

[25] J. Paredaens, J. Van den Bussche, and D. Van Gucht. Towards a theory of spatial database queries. In *PODS'94*, pages 279–288.

[26] J. Paredaens, J. Van den Bussche, and D. Van Gucht. First-order queries on finite structures over the reals. In *LICS'95*, pages 79–87.

[27] Bruno Poizat, *Cours de theories de modéles* Villeurbanne: Nur al-Mantiq wal-Ma'rifah

[28] S. Shelah. *Classification Theory and the Number of Nonisomorphic Models* North-Holland, 1978.