Stability theory, Permutations of Indiscernibles, and Embedded Finite Models*

John Baldwin[†]
Department of Mathematics, Statistics, and Computer Science
University of Illinois at Chicago
E-mail: jbaldwin@math.uic.edu

Michael Benedikt
Bell Laboratories
1000 E. Warrenville Rd.
Naperville, IL 60566
E-mail: benedikt@research.bell-labs.com

March 10, 1998

Abstract

We show that the expressive power of first-order logic over finite models embedded in a model M is determined by stability-theoretic properties of M. In particular, we show that if M is stable, then every class of finite structures that can be defined by embedding the structures in M, can be defined in pure first-order logic. We also show that if M does not have the independence property, then any class of finite structures that can be defined by embedding the structures in M, can be defined in first-order logic over a dense linear order. This extends known results on the definability of classes of finite structures and ordered finite structures in the setting of embedded finite models. These results depend on several results in infinite model theory. Let I be a set of indiscernibles in a model M and suppose (M, I) is elementarily equivalent to (M_1, I_1) where M_1 is $|I_1|^+$-saturated. If M is stable and (M, I) is saturated then every permutation of I extends to an automorphism of M and the theory of (M, I) is stable. Let I be a sequence of $<$-indiscernibles in a model M, which does not have the independence property, and suppose (M, I) is elementarily equivalent to (M_1, I_1) where $(I_1, <)$ is a complete dense linear order and M_1 is $|I_1|^+$-saturated. Then (M, I)-types over I are order-definable and if (M, I) is $\aleph_1$-saturated every order preserving permutation of I can be extended to a back-and-forth system.

1 Introduction

A major research current in both finite and infinite model theory has been the classification of structures via the combinatorial objects that can be defined within them. In infinite 'classical' model theory one prototypical classification program is *stability theory.* One set of stability-theoretic results focuses on structures that admit no definable order. It is shown that these structures (the *stable* structures) are much better behaved than structures that do admit orders: e.g. the former have few types, and always admit universal domains. 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 nicely behaved definability theory.

*primary classification: 03C45, secondary classification: 68P15, 03C40

[†]partially supported by NSF-93-8768

In finite model theory there has been a line of work with a somewhat similar flavor: the investigation of 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 [12], with bounding the expressive power of first-order logic and fragments of second-order logic on ordered structures [12], and with proving the presence of an almost sure linear order according to various probability measures on finite structures [17].

This work aims at showing one connection between these two lines of research: we 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. Logics that deal with this 'hybrid setting' have been studied recently in connection with database theory. In particular, it has been shown that many interesting properties of finite structures that are not first-order definable (such as the parity of a set or the connectivity of a graph) remain undefinable for finite structures that are embedded in structures satisfying appropriate conditions. Results in the literature, starting with [29] and continuing through [6] and [4], show that under certain conditions on the model (e.g. *o*-minimal in [6], quasi o-minimal in [4]) 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 broader 'collapse' results as well to structures without the independence property, a very broad model-theoretic class that includes the examples of [6] and [4] and also includes all stable structures. The connection between generic query expressibility and this class is interesting not only due to the scope of the class, but also because of the connection of this class with many probabilistic and algorithmic issues, such as machine learning [25].

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 definable not just over order but over the trivial structure, and hence are pure first-order definable in the sense of finite model theory.

The collapse theorems above follow from results concerning 'model theory over a predicate'. amd the paper can also be seen as a contribution to this topic. Fix languages $L^- \subset L$ with the predicate P in L but not L^-. Much of this work (e.g. [31], [19], [24]) focuses on the extent to which the properties of an L^--structure N determine the properties of an L-structure M with $P(M) = N$. Our work varies in two ways. To some extent we switch the emphasis by starting with M and asking about properties of the expansion of M obtained by adding a predicate for a subset. More important, instead of considering an arbitrary subset as in the work just described or a submodel as in ([33],[30]), we insist that the set be a collection of (order) indiscernibles. We prove that in a theory without the independence property, types over an indiscernible sequence (of a complete linear order) are order-definable with parameters from the sequence. We show that permutations of sets of indiscernibles in 'pseudo-small' stable L^+-saturated models extend to automorphisms of the models, and we show that a pseudo-small expansion of a stable theory is stable.

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. In Sections 4 and 6 we prove this quantifier reduction theorem for stable theories and theories without the independence property, respectively. As a prelude to the second proof, a key definability result for types over ordered sequences of indiscernibles in theories without the independence property is proved in Section 5. Section 7 uses these results to prove the results for embedded finite model theory. Section 8 contains the results on infinite model theory over a predicate. Section 9 gives conclusions and discusses future work. The reader primarily interested in results on query languages may want to start by reading just Section 2 and Section 7.

2 Embedded finite model theory

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 below, 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*.

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 $M\}$

$FO_{GEN}(M)$ 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 [20] 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 [6] and [28]:

Remark 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'.*

For, if Q were not defined by ϕ over M', there would be a finite S-structure A' witnessing the fact. Let $n = |A|$, and assume $A' \in Q$ but $M'(A') \models \neg\phi$. The statement that there is an S-structure A' of size n with the isomorphism type of A' such that $M'(A') \models \neg\phi$ can be expressed as a single L -sentence, hence it must hold in M as well.

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 by providing conditions on a structure such that if certain kinds of generic queries are definable with the help of symbols from L, they can be defined without them. The first such result is:

First Main Result :

Theorem 2.2 *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 (on tuples). 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 [3], [10], [18], [32] or [34] 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 complete 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 $\langle \mathbb{U}, <\rangle$. For any infinite linearly ordered structure O and abstract query Q, the query Q_O induced on O is locally order generic. Conversely, any locally order generic query naturally determines an abstract order query.

Let L be a language disjoint from S, let $L' = L \cup \{<\}$, 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. 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 $FO_{og}(M)$ of order-generic queries definable over 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.

Second Main Result :

Theorem 2.3 *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, a class does not have the independence property (has NIP) if it does not admit an infinite parametrically defined class of definable sets which can be satisfied independently; the formal definition is below. Theorem 2.3 generalizes *all* previous collapse results, since the classes involved all satisfied NIP. Classes without the independence property include all stable structures, all o-minimal and quasi-o-minimal structures (e.g. additive number theory), all C-minimal structures [26] (e.g. regular trees), and all linear orders (even with unary predicates) [32]. Particular examples of interest include the real field (with exponentiation) and the p-adic numbers, Q_p. Further, structures in NIP admit a Vapnik-Chervonenkis dimension, which implies that every parametrically definable family is PAC-learnable [25]. Some other studies are in [23] and [32]. In [34], Shelah proves a fundamental dichotomy : Every unstable theory either has the independence property or the strict order property (roughly, interprets a linear order). The protypical structure with the independence property is the random graph.

Definition 2.4 *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.*

It is shown in [8] that the parity query is definable in the random graph; this shows the necessity of the hypothesis in Theorem 2.3 that T does not have the independence property. Both instability and the independence property are preserved under taking reducts.

We also give a result comparing 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 model theory is that $FO_{GEN}(\langle \mathbb{U},=\rangle)$ is a proper subset of $FO_{GEN}(\langle \mathbb{U},<\rangle)$ for any linear order $<$ [1]. However, our next main result shows that adding extra structure beyond order does not yield new generic queries, when the model does not have the independence property.

We will establish the following variant of Theorem 2.3.

Theorem 2.5 *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 2.5 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.

Since it is known that any abstract order query definable in an infinite linear order is definable as a class of finite ordered structures (this follows, for example, from the natural/active collapse result for o-minimal structures in [6], but was probably known earlier), Theorem 2.5 yields:

[1] Here is a slight variant on his example. Let the query language contain the similarity type for Boolean algebras. Now in this language require that the structure is a Boolean algebra. Since each element of a finite Boolean algebra is a join of atoms, first order quantification over the entire structure codes monadic second order quantification over the set of atoms. Now using the order, one can assert there are an even number of atoms by saying there is a subset X of the atoms such that every other atom is in X.

Corollary 2.6 *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 order 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 2.7 *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 or deterministic transitive closure of a graph; maximal matching in a bipartite graph; Eulerian cycle.*

Proof: The examples are known not to be expressible over ordered structures [1]. $\square_{2.7}$

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

The study of expressibility over embedded finite structures began in connection with a particular database formalism, constraint databases, in [21] (there are other frameworks for studying "mixed" structures, particularly the one presented in [15]). [21] focused on the real ordered group and ordered field. Both Theorem 2.3 and Corollary 2.7 were proved in the special case of the real ordered group in [29]. This result was extended to o-minimal structures in [6] and extended to quasi-o-minimal structures in [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, $\langle N, +, *, < \rangle$); hence collapse of any sort fails for such structures.

Many of the results on collapse in the previous literature rely implicitly or explicitly on the use of indiscernibles (the first explicit use being [28]). The results on collapse theorems for o-minimal and quasi-o-minimal structures rely on the 'locality property' of o-minimal/quasi-o-minimal theories: this property is used to extend certain mappings on indiscernibles. The locality property does not apply to stable theories, so the approach here relies instead on an analysis of expansions of models by a set of indiscernibles.

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 [6] and [28], but the extension result we need (Theorem 7.3) requires quite a bit of machinery to prove. In this section we outline a set of results (proved in the ensuing 3 sections) about expansions of structures by a predicate for an indiscernible set or sequence. These results will be used in the Section 7 to prove the key result, Theorem 7.3, 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 the (range of)* a set (sequence) 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 [10]. A basic result of stability theory (e.g. [3] 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 the set I or, more precisely, the structure (M, I) *small* if further

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

If the L^+-structure (N, J) is elementarily equivalent to a small (M, I), J or (N, J) is 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.

Unless explicitly denied, henceforth T^+ is a pseudo-small theory.

Notes on Pseudo-smallness: Those interested in the connection of pseudo-small structures to collapse results may want to look ahead to Lemma 7.5. In the ordered case, the completeness of $\langle I, <\rangle$ ensures that for any model (M', I') of T^+, the reduct to L is definably complete: for any L-formula $\phi(x, \vec{y})$ and any $\vec{m}$ in M, if $\{c \in I : \phi(c, \vec{m})\}$ is a bounded subset of I, then it has a least upper bound. This holds in a small model (M, I) using completeness, and tranfers using elementary equivalence.

The main result stated in this section is a quantifier-reduction result for each of the 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. As a corollary to this we eventually show that T^+ depends only on T and the L-types realized in I.

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,*

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

3.2 Some Assumptions and Reductions.

The proof of Theorem 3.3 will take up the next three and a half sections.

The following key assumption amounts to assuming our reduction result holds for basic formulas. The definability of types over all sets is tantamount to the definition of stability, so the assumption is clear for stable theories. For theories without the independence property, we prove it in Section 5. Our argument will use this definability of L-types to establish Theorem 3.3 by induction on formulas.

Assumption 3.4 *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 model (M, I) of T^+ if I is a set of (Δ, L)-indiscernibles in M then for every $\vec{m} \in M$ there is a $\vec{c}_{\vec{m}} \in I$ such that*

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

The following lemma shows that P-bounded formulas have defining schema, which will be useful both in proving the reduction results and in applying them to query collapse.

Lemma 3.5 *Let $M^+ \models T^+$. For each P-bounded formula $\phi(\vec{x}, \vec{y})$ there is a P-bounded formula $\delta(\vec{w}, \vec{y})$ that is a defining schema for ϕ over I: That is, for each $\vec{m}$ in M there is $\vec{c}_{\vec{m}}$ in I such that*

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

Proof: Let $\phi(\vec{x}, \vec{y})$ be normalized as $(Q_1 z_1) \ldots (Q_k z_k)\ \eta(\vec{x}, \vec{y}, \vec{z})$ where the Q_i's are P-bounded quantifiers and η is a basic formula. By Assumption 3.4 there is a defining schema $\eta'(\vec{y}, \vec{z}, \vec{w})$ for $\eta(\vec{x}, \vec{y}, \vec{z})$ over I. That is,

$$M^+ \models \forall \vec{x}\ \exists \vec{w} \in P\ \forall \vec{y} \in P\ \forall \vec{z} \in P\ [\eta(\vec{x}, \vec{y}, \vec{z}) \leftrightarrow \eta'(\vec{y}, \vec{z}, \vec{w})].$$

Consider $\delta(\vec{w}, \vec{y})$ given by $(Q_1 z_1) \ldots (Q_K z_k)\ \eta'(\vec{y}, \vec{z}, \vec{w})$. It is easy to verify that $\delta(\vec{w}, \vec{y})$ is a P-bounded formula and a defining schema for ϕ: For each $\vec{m} \in M$ choose $\vec{c}_{\vec{m}} \in I$ so that $\forall \vec{y} \in P\ \forall \vec{z} \in P\ (\eta(\vec{m}, \vec{y}, \vec{z}) \leftrightarrow \eta'(\vec{y}, \vec{z}, \vec{c}_{\vec{m}}))$. Then $\delta(\vec{c}_{\vec{m}}, \vec{y})$ defines the ϕ-type of $\vec{m}$ over I. $\square_{3.5}$

Looking ahead a bit, the main application of the above result is the definability of all L^+-types over I.

Corollary 3.6 (Corollary to as-yet-unproved Theorem 3.3) *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 and Lemma 3.5 there is a P-bounded formula $\delta(\vec{w}, \vec{y})$ and 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 formulas 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. $\square_{3.6}$

Definition 3.7 *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. [34] 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 is 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 3.8 *Let I be a set of (Δ, L)-indiscernibles in L-structure M. Suppose that (M, I) satisfies Assumption 3.4 and 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 Assumption 3.4, 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. $\square_{3.8}$

Thus, to prove our main quantifier reduction result for a pseudo-small T^+, we need only show the hypothesis of Proposition 3.8 holds in an $|I|^+$-saturated model M of T. In the next sections we will give separate arguments for this in the stable and not independence property cases.

4 Quantifier Reduction in Expansions by Indiscernibles: Stable Case

Throughout this section we assume M is an L-structure for a stable theory T, I is a set of L-indiscernibles (i.e. we fix Δ as $\{x = y\}$), and M is $|I|^+$-saturated. Note that since equality is in L, for the purposes of this section a basic formula is just an L-formula.

We fix $\beta(x_1, \ldots, x_n, z, y_1, \ldots, y_R) \in L$; our goal will be to show the stable part of Theorem 3.3 by showing the reduction in the hypothesis of Proposition 3.8 holds for every equality-formula $\delta(\vec{w}, \vec{y})$. We begin by verifying Assumption 3.4, and fixing some notation in the process.

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, the trace of an L-formula on an arbitrary subset is uniformly defined with parameters in that set. If the set is a set of indiscernibles then the formula can be taken to involve only equality. In the next definition, we consider a sequence $\vec{x}, z$ in two parts only for easy reference in the application.

Lemma 4.1 *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*

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

Proof. By stability ([3] III.1.24), for any formula $\beta(\vec{x}, z, \vec{y})$ there is a formula $\delta(\vec{w}, \vec{y})$ such that for any $\vec{m}, b$, for any set F, there is a $\vec{c} \in F$ such that for all $\vec{f} \in F$, $\beta(\vec{m}, b, \vec{f}) \leftrightarrow \delta(\vec{c}, \vec{f})$. By indiscernibility, we can find an equality formula $\psi(\vec{w}, \vec{y})$ which is equivalent, for all $\vec{w}, \vec{y}$ in I, to δ. $\square_{4.1}$

The following definitions establish crucial notation for the remainder of the proof.

Notation 4.2 *($J = J_\beta$) We have fixed $\beta(x_1, \ldots, x_n, z, y_1, \ldots, y_R) \in L$. 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 4.3 1. *$\delta(\vec{c}, \vec{y})$ is* correct *for $\vec{m}, d$* on *X if for every L-tuple $\vec{e} \in X$, $\beta(\vec{m}, d, \vec{e}) \leftrightarrow \delta(\vec{c}, \vec{e})$.*

2. *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 3.8, 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 introduce an auxiliary formula $\psi_\delta(\vec{x}, \vec{w})$.

Definition 4.4 *(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 4.5 *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. $\square_{4.5}$

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 4.5, 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 3.8 and thus Theorem 3.3 for the stable case follow from

Lemma 4.6 (Goal 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:*

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

We will prove Lemma 4.6, by downward induction on the length of $\vec{w}$, starting at J_β. We first need the following observation.

Proposition 4.7 *If $\mathrm{tp}_\beta(\vec{m}, d/I)$ does not split over $X_0 \subset I$, $\delta(\vec{c}, \vec{y})$ is correct for $\vec{m}, d$ on $X \supset X_0$, and $|X - X_0| \geq R$ then $\delta(\vec{c}, \vec{y})$ defines $\mathrm{tp}_\beta(\vec{m}, d/I)$.*

Proof. If $\mathrm{tp}_\beta(\vec{m}, d/I)$ does not split over X_0 then $\mathrm{tp}_\beta(\vec{m}, d/I)$ is defined over X_0. This definition is equivalent to δ since every equality type of an R-tuple over X_0 is realized in X and $\mathrm{tp}_\beta(\vec{m}, d/I)$ does not split over X_0. $\square_{4.7}$

4.1 Inductive proof of Lemma 4.6

We now start the proof of Lemma 4.6. We have fixed β and $J = J_\beta$ in Notation 4.2. We begin with the base case, $n = J$. Let $J' = max\{J + R, 2JR\} + 1$.

Definition 4.8 *(ψ'_δ) Let $\psi'_\delta(\vec{x}, w_1, \ldots, w_J)$ hold of $\vec{m}, \vec{c}$ if and only if the c_i are distinct and there exist $E = \{e_1, \ldots, e_{J'}\}$ containing C, the range of $\vec{c}$, and a d such that*

1. *$\mathrm{tp}_\beta(\vec{m}, d/I)$ splits in E over every proper subset of C,*
2. *and $\delta(\vec{c}, \vec{y})$ is correct for $\vec{m}, d$ on E .*

Since the arbitrary d is chosen after the set E is selected from P and the conditions i) and ii) are expressible in L, ψ'_δ is a P-bounded formula.

Claim 4.9 *$\psi'_\delta(\vec{m}, \vec{c})$ holds if and only if $\psi_\delta(\vec{m}, \vec{c})$ holds.*

That is, ψ'_δ holds if and only if there is some d such that $\delta(\vec{c}, \vec{y})$ defines $\mathrm{tp}_\beta(\vec{m}, d/I)$ and $\mathrm{tp}_\beta(\vec{m}, d/I)$ splits over every proper subset of $\vec{c}$.

Proof of Claim 4.9. The implication from ψ_δ to ψ'_δ is immediate, since we can always witness all the necessary splittings with $2JR \leq J'$ elements.

Suppose the converse fails. Then we have a $C = \{c_1, \ldots c_J\}$, $E = \{e_1, \ldots e_{J'}\}$ containing C and $\vec{m}, d$ satisfying ψ'_δ for which the β-type of $\vec{m}, d$ over I is not given by δ. By the definition of J, we know that there is a set $S \subseteq I$ of size at most J such that the β-type of $\vec{m}, d$ over I is definable over S. If S is contained in E, then we are done by Proposition 4.7, since the size of E guarantees that it contains R elements not in S . So assume that $S - E$ is nonempty. Then $S - C$ is nonempty, and since $|S| \leq |C| = J$ we infer there is a c_0 in $C - S$.

Letting $C_0 = C - \{c_0\}$, we know, by assumption that there is a splitting over C_0 with witness tuples t_1 and t_2 lying completely in E. Let t'_1 and t'_2 be obtained from t_1 and t_2 by replacing any occurrence of c_0 with an element in $E - (C \cup S)$. Such elements exist since the size of E is bigger than $2J$. Since $c_0 \notin S$, we have $t_i \equiv_S t'_i$. As $\mathrm{tp}_\beta(\vec{m}, d/I)$ does not split over S, this implies that t'_1 and t'_2 still witness a split for $\mathrm{tp}_\beta(\vec{m}, d/I)$ over C_0.

But t'_1 and t'_2 are contained in E and have the same equality type over C (they have the same equality type over C_0 since the t_i's did, and neither of them contains c_0). Hence, $\delta(\vec{c}, \vec{y})$ must agree on t'_1 and t'_2. This and the last sentence in the previous paragraph contradict the fact that $\delta(\vec{c}, \vec{y})$ is correct on E. $\square_{4.9}$

For $n = J (= J_\beta)$ we have found an appropriate P-bounded formula ψ'_δ which is equivalent to ψ_δ. We turn to the induction step. Suppose we have Lemma 4.6 for each i with $n < i \leq J$; we will prove that $\psi_\delta(\vec{x}, w_1, \ldots w_n)$ is $\vec{w}, I$-equivalent to a P-bounded formula. The result follows easily from

Proposition 4.10 *There is an n' such that for any $\vec{m}$ and $C = \{c_1, \ldots c_n\}$, $\psi_\delta(\vec{m}, \vec{c})$ if and only if the c_i are distinct and for every set G with $|G| < n'$, there is a set E of size at most $J'' = max(n' + 2JR + 1, 3J + 1)$ containing G and a d such that $\mathrm{tp}_\beta(\vec{m}, d/I)$ equality splits in E over every proper subset of C, and $\delta(\vec{c}, \vec{y})$ is correct for $\vec{m}, d$ on E.*

Since the clauses after 'and a d such that' are definable in L this proposition provides a P-bounded formula equivalent to ψ_δ.

Proof of Proposition 4.10. The implication from $\psi_\delta(\vec{m}, \vec{c})$ is immediate for any n' taking E to be G plus $2JR$ witnesses for splitting. We will choose n' below so the converse holds.

We first need some further notation to employ the induction hypothesis. For $n < j \leq J$, let $DEF(j)$ denote the set of equality formulas $\delta'(w_1, \ldots, w_j, \vec{y})$ in $w_1, \ldots, w_j$ and $y_1, \ldots, y_R$. Note that up to equivalence $DEF(j)$ is finite.

Let $p(n,r)$ be a function such that for any finite subset A of I there is a set $A' \subseteq I$ with $|A'| < p(|A|, r)$ such that all $(\{=\}, r)$-types over A are realized in A'. It is easy to see that for equality one such function is $n + r$.

For each $\delta' \in DEF(j)$ we know that $\psi_{\delta'}(\vec{x}, \vec{w})$ is $\vec{w}, I$-equivalent to a P-bounded formula. By Lemma 3.5, each $\psi_{\delta'}(\vec{x}, \vec{w})$ type over I has a P-bounded defining schema. Let m' be the number of parameters which occur in any of these definitions, Now choose $n' = p(m' + J, R)$. We show that this n' suffices for Proposition 4.10.

Let $\vec{a} = \langle a_1, \ldots, a_{m'} \rangle$ (not necessarily distinct) from I be such that for each $\delta' \in DEF(j)$, the $\psi_{\delta'}(\vec{x}, \vec{w})$-type of $\vec{m}$ over I is definable (by an equality-formula) over $a_1, \ldots, a_{m'}$. Such an $\vec{a}$ exists because the trace on I of a P-bounded formula is definable in equality (as shown in the proof of Corollary 3.6).

Let G be a finite set of size at most n' such that every $(\{=\}, R + J)$-type over $\vec{a}\vec{c}$ is realised in G.

We now apply the hypothesis of Proposition 4.10 to this G. Then we have $\vec{a}\vec{c} \subset G \subset E \subset I$, and a d such that $\mathrm{tp}_\beta(\vec{m}, d/I)$ splits in E over every proper subset of $\vec{c}$ and $\delta(\vec{c}, \vec{y})$ is correct for $\vec{m}, d$ on E. We claim that $\delta(\vec{c}, \vec{y})$ is correct for $\vec{m}, d$ on I so $\psi_\delta(\vec{c}, m)$ holds. Suppose not. Then $\delta(\vec{c}, \vec{y})$ does not define $\mathrm{tp}_\beta(\vec{m}, d/I)$. By Proposition 4.7, since $\delta(\vec{c}, \vec{y})$ is correct on G and G contains $\vec{c}$ as well as R elements not in $\vec{c}$, there must be a β-split for $\vec{m}, d$ over $c_1, \ldots, c_n$.

We now note the following:

Claim 4.11 *There are $c_{n+1}, \ldots, c_H$, with $H \leq J$, such that $\mathrm{tp}_\beta(\vec{m}, d/E)$ splits over every proper subset of $c_1, \ldots, c_H$ but not over $c_1, \ldots, c_H$.*

Proof of claim: We know that there is $S = \{s_1, \ldots, s_H\}$ with $H \leq J$ such that $\mathrm{tp}_\beta(\vec{m}, d/E)$ is equality-defined over S. We will show that each c_i for $i \leq n$ must be in this set S so for $i > n$ we can take s_i for the required c_i. Suppose not and fix c in C but not in S. Let $\vec{e}_1, \vec{e}_2$ in E realize a splitting over $C - \{c\}$: that is, $\vec{e}_1$ and $\vec{e}_2$ have the same equality type over $C - \{c\}$ but $\neg[\beta(\vec{m}, d, \vec{e}_1) \leftrightarrow \beta(\vec{m}, d, \vec{e}_2)]$.

We first show that $\vec{e}_1, \vec{e}_2$ can be taken to exclude c. Since there is no splitting for $\mathrm{tp}_\beta(\vec{m}, d/E)$ over S and $c \notin S$, we can replace c in the $\vec{e}_i$'s by any other element of E distinct from c and the other elements of S without changing the truth value of $\beta(\vec{m}, d, \vec{e}_i)$. So we may assume that the original $\vec{e}_1, \vec{e}_2$ do not contain c. Hence we now know $\vec{e}_1, \vec{e}_2$ have the same equality type over C, since they agree on $C - \{c\}$ and neither contains c.

But there is no splitting for $\mathrm{tp}_\beta(\vec{m}, d/E)$ over C either, since $\delta(\vec{c}, \vec{y})$ is correct on E; hence we can uniformly replace any elements of $\vec{e}_1$ and $\vec{e}_2$ that are in $S - C$ by elements of $E - (C \cup S)$ without changing the truth value of $\beta(\vec{m}, d, \vec{e}_i)$ (note that there are at most $|S| \leq J$ elements to replace and $E - (C \cup S)$ contains more than J elements to replace these elements with). Thus, we can assume that $\vec{e}_1$ and $\vec{e}_2$ contain only elements of S that are in C. However, we now have that $\vec{e}_1$ and $\vec{e}_2$ agree on all elements of S, since they agree on C and have no elements from $S - C$. This contradicts the fact that $\mathrm{tp}_\beta(\vec{m}, d/E)$ does not equality-split over S. $\square_{4.11}$

Note that $\mathrm{tp}_\beta(\vec{m}, d/I)$ is defined by $\delta'(c_1, \ldots, c_H, \vec{y})$, where δ' is an equality formula.

Claim 4.12 *Fix any $\{b_{n+1}, \ldots b_H\} \subseteq I$ and d such that $\mathrm{tp}_\beta(\vec{m}, d/I)$ is defined over $C^+ = \{c_1, \ldots c_n, b_{n+1}, \ldots b_H\}$ by $\delta'(c_1, \ldots c_n, b_{n+1}, \ldots b_H, \vec{y})$. For any sequence of distinct elements $\vec{f}$ of length R in I, if $\vec{f} \cap \{b_{n+1}, \ldots b_H\} \subseteq G$ then $[\beta(\vec{m}, d, \vec{f}) \leftrightarrow \delta(c_1, \ldots c_n, \vec{f})]$.*

Proof. Consider a counterexample $\vec{f}$: $\vec{f} \cap \{b_{n+1}, \ldots, b_H\}$ is contained in G and $\neg[\beta(\vec{m}, d, \vec{f}) \leftrightarrow \delta(c_1, \ldots c_n, \vec{f})]$. Since there is no β-split for $\vec{m}, d$ over C^+, we know $\neg[\beta(\vec{m}, d, \vec{f'}) \leftrightarrow \delta(c_1, \ldots c_n, \vec{f'})]$ for any $\vec{f'} \in I$ with $\vec{f'} \equiv_{C^+} \vec{f}$. Since E realizes every $R + J$ type over C, we can choose an $\vec{f'} \in E$ with $\vec{f'} \equiv_C \vec{f}$ and $\vec{f} \cap G = \vec{f'} \cap G$. But $\vec{f} \cap \{b_{n+1}, \ldots, b_H\}$ is contained in G. So $\vec{f'} \equiv_{C^+} \vec{f}$. This contradicts the hypothesis that $\delta(\vec{c}, \vec{y})$ is correct for $\vec{m}, d$ on E. $\square_{4.12}$

Now we claim:

Claim 4.13 *Given any finite subset S of I, there is a d_S in M such that for all $\vec{f} \in S$, $\beta(\vec{m}, d_S, \vec{f}) \leftrightarrow \delta(c_1, \ldots, c_n, \vec{f})$ holds.*

Proof of claim 4.13. Fix any finite $S \subset I$. Replace all elements in $\{c_{n+1} \ldots c_H\} - G$ by distinct elements of I not in $\vec{a} \cup \{c_1 \ldots c_n\} \cup S$ to choose $h_{n+1}, \ldots, h_H$ such that $h_{n+1}, \ldots, h_H \equiv_{\vec{a}\vec{c}} c_{n+1}, \ldots c_H$ and $H = \{h_{n+1}, \ldots, h_H\} - G$ is disjoint from S. (This is easy since $\vec{a} \cup \{c_1 \ldots c_n\}$ is contained in G.) Then, $\psi_{\delta'}(\vec{m}, c_1, \ldots, c_n, h_{n+1}, \ldots h_H)$ holds, since the $\psi_{\delta'}$-type of $\vec{m}$ does not equality-split over $\vec{a}$. By the definition of $\psi_{\delta'}$, there is a d_S such that for any tuple of distinct elements $\vec{f} \in I$, $\models [\beta(\vec{m}, d_S, \vec{f}) \leftrightarrow \delta'(c_1, \ldots, c_n, h_{n+1}, \ldots h_H, \vec{f})]$. Claim 4.12 implies $\beta(\vec{m}, d_S, \vec{f}) \leftrightarrow \delta(c_1, \ldots, c_n, \vec{f})$ holds whenever $\vec{f} \cap H$ is empty. But $S \cap H = \emptyset$, so $\beta(\vec{m}, d_S, \vec{f}) \leftrightarrow \delta(c_1, \ldots, c_n, \vec{f})$ holds for all $\vec{f} \in S$ as required; this completes the proof of Claim 4.13. $\square_{4.13}$

Claim 4.13 and compactness and the saturation of M^+ provide a d to witness $\psi_\delta(\vec{m}, c_1, \ldots, c_n)$. Hence we now have a proof of Proposition 4.10, and hence a proof of Lemma 4.6 and Theorem 3.3 for the stable case. $\square_{4.10}$

5 Type definability in theories without the independence property

In this section we prove a definability result for types over sequences of order indiscernibles in models of a theory without the independence property. This result is a step to proving Theorem 3.3 for the ordered case. It is also of independent interest and does not depend on the hypotheses and conventions stated earlier; hence we suspend our various conventions and state explicitly the hypotheses of this section. We will apply the result in the next section to verify Assumption 3.4.

Definition 5.1 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 $X \subseteq m$, there is an $\vec{a}_X$ such $\phi(\vec{b}_i, \vec{a}_X)$ holds if and only if $i \in X$.*

2. *ϕ does not have the independence property in T (lacks IP, for short) 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 if no ϕ has the independence property in T. A model does not have the independence property if its theory does not have it.*

We work with T a complete countable L-theory which does not have the independence property. M is an arbitrary model of T. We have a set I from M linearly ordered by $<$, such that $\langle I, < \rangle$ is a sequence of order indiscernibles and is a complete dense linear order. Nothing is asserted about the properties of $<$ on the rest of M, and we do not assume that $<$ is in L.

The *trace* of a formula $\phi(\vec{m}, \vec{y})$ on I is the set of $\vec{c}$ from I satisfying $\phi(\vec{m}, \vec{c})$. We noted in Lemma 4.1 that if M is a model of a stable theory, the trace of an L-formula on a set of indiscernibles is uniformly defined by an equality formula with parameters in that set. We now prove the analogous result in a theory without the independence property: defining the trace of a formula on a sequence of order indiscernibles by an order formula with parameters from the indiscernibles. It appears that this definability result is new even in the extremely restricted case that one assumes T is a O-minimal theory and I is a sequence of indiscernibles that is complete for the given order. The definability result is:

Theorem 5.2 *If M lacks IP and I is order-indiscernible with order type a complete dense linear order then for every L-formula $\phi(\vec{x}, \vec{y})$ there is a $<$-formula $\psi(\vec{w}, \vec{y})$ such that for every $\vec{m}$ there is a $\vec{c}_{\vec{m}} \in I$ such that*

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

Note: We prove the lemma above assuming that I is ordered by a complete dense linear order. However, the above statement is a property of the theory of (M, I), and so it remains true in any model (M', I') elementary equivalent to such an (M, I). Thus, we can apply it to any pseudo-small structure.

Marker and Steinhorn [27] and Pillay [30] proved an analogous result: in an o-minimal theory, types over a model which is Dedekind complete are definable. We ask whether this result holds for any theory without the independence property where the models are linearly ordered.

The following definitions are needed before beginning the proof of Theorem 5.2. Let $S \subseteq I^n$ for some n. We say $\vec{e} \in I^n$ is a *boundary point* of S if there are points in S and S^c in every neighborhood of $\vec{e}$ in the product topology on I^n: i.e. if it is on the topological boundary of S in the usual sense. Given a finite set of points F in I^n, an *equivalence class* for F is the set of points in I^n satisfying some complete 1-type in $<$ over F; i.e. a maximal set of points in I^n satisfying all the same inequalities over F. If $\vec{c}_1, \ldots, \vec{c}_{i-1}$ is a sequence of vectors from I^R, we denote by C_i the union of the ranges of the $\vec{c_j}$, for $j < i$. Note that an equivalence class for F is definable in $<$ over F and is a connected (in the product topology) subset of I^n. This is because a complete dense linear order is connected (Completeness implies that every open set S is a union of disjoint intervals that are maximal in S. This plus density easily implies connectedness). Since every equivalence class is a product of intervals, and an interval within a complete order is complete, we get the above.

The key definition is:

A *coherent sequence of boundary points* for a set S is a (finite or infinite) sequence of vectors $\vec{c}_1, \ldots, \vec{c}_n \ldots$ from I^n such that any $\vec{c}_i$ is a boundary point of $S \cap E_i$ relative to the subspace topology in E_i, where E_i is the equivalence class of $\vec{c}_i$ over C_i. Put another way: a sequence is coherent if for any neighborhood O of $\vec{c}_i$ in I^n, $O \cap E_i$ intersects both S and S^c.

A *boundary point cover* for S is a maximal coherent sequence of boundary points for S.

We say the m-tuples $\langle p_i : i < \omega \rangle$ are *pairwise disjoint* if their ranges are disjoint.

We actually prove the following proposition, which entails Theorem 5.2.

Proposition 5.3 *Let S be an L-definable subset of I^R, where I is a sequence of indiscernibles in a model M without the independence property. Then the following hold.*

1. *Every finite coherent sequence of boundary points for S extends to a finite boundary point cover for S.*
2. *If C is a finite boundary point cover for S, then S is definable from $<$ over C, and hence S does not split over the union of all elements in C.*
3. *For any formula $\theta(\vec{x}, \vec{y})$ from L, there is an integer K_θ that bounds the size of a coherent set of boundary points for the set $\{\vec{y} : \theta(\vec{y}, \vec{m})\}$ uniformly for every $\vec{m}$ in M.*

The following (which does not require completeness of the order) is the key result.

Lemma 5.4 *Suppose $\gamma(x_1 \ldots x_R, \vec{y})$ does not have the independence property in T. There is K such that for any $\vec{a}$, there cannot be a coherent sequence $\langle \vec{p}_i : i < K \rangle$ of boundary points for $D = \{\vec{b} \in I^R : \gamma(\vec{b}, \vec{a})\}$.*

Proof of Lemma 5.4: Suppose γ does not have the m-independence property. Choose K large enough so that: for any sequence $\langle \vec{p}_i : i \in K \rangle$ we can choose a subsequence with length m whose elements all realize the same $<$-type over the empty set and such that the type of the ordered pair of vectors $\langle \vec{p}_i, \vec{p}_j \rangle$ is independent of i and j (subject to $i < j$). Such a K exists by Ramsey's theorem, and we claim that this K works. Suppose not, and consider a coherent sequence $\langle \vec{p}_i : i \in K \rangle$ for a set $D \subset I^R$ of the form $\{\vec{b} : \gamma(\vec{b}, \vec{a})\}$. Replace the sequence with a subsequence $\langle \vec{p}_i : i \in m \rangle$ satisfying the additional restrictions just described. Note that coherence implies that any neighborhood of $\vec{p}_i$ contains elements of D and D^c that are of the same type as $\vec{p}_i$ over P_{i-1}.

Let E_i be the equivalence class of $\vec{p}_i$ over P_{i-1}. Let s be the set of indices $k < R$ such that $\vec{p}_i(k) = \vec{p}_{i+1}(k)$ holds for some (hence all, by the property of the chosen subsequence) i. Let $t = s^c$, let $\vec{g}$ be the common value of the $\vec{p}_i|s$, and let $D_{\vec{g}}$ be $\{\vec{f} \in I^t : \vec{f} \cup \vec{g} \in D\}$. Since $\langle \vec{p}_i : i \in m \rangle$ is a coherent sequence of boundary points for D, we get $\vec{p}_i|t$ is a coherent sequence of boundary points for $D_{\vec{g}}$. Moreover, since every pair $\langle \vec{p}_i, \vec{p}_j \rangle$ for $i < j$ has the same type, the $\vec{p}_i|t$ are pairwise disjoint. (If $\vec{p}_i(k) = \vec{p}_j(n)$ for some $i \neq j, k \neq n$, then $\vec{p}_i(k) = \vec{p}_j(n)$ for all $i \neq j$, and hence $\vec{p}_i(k) = \vec{p}_j(k)$, which implies k is in s). The fact that $\vec{p}_i|t$ are a pairwise disjoint coherent sequence of boundary points for $D_{\vec{g}}$ allows us to easily construct for any subset X of m a sequence $\vec{q}_i^X$ of elements of I^t such that the order type of $\langle \vec{q}_i^X : i < m \rangle$ over the empty set does not depend on X but $\vec{q}_i^X \in D_{\vec{g}}$ if and only if $i \in X$.

But now, by concatenating the above $\vec{q}_i^X$ with $\vec{g}$, we get for any subset X of m a p_i^X in I^R such that the order type of $\langle \vec{p}_i^X : i < m \rangle$ over the empty set does not depend on X but $\gamma(\vec{p}_i^X, \vec{a})$ if and only if $i \in X$.

This observation and the order indiscernibility of I shows that for each $X \subset m$ there are $\vec{e}_1 \ldots \vec{e}_m$ and $\vec{a}^X$ such that $\gamma(\vec{e}_i, \vec{a}^X)$ if and only if $i \in X$, which contradicts the assumption that T does not have the m-independence property. $\square_{5.4}$

We use the following simple topological observation at several other stages in the main argument.

Observation 5.5 *For any R, if X is a connected (in the product topology) subset of I^R and Y is a subset of I^R such that both $X \cap Y$ and $X \cap Y^c$ are nonempty then X contains a boundary point of Y.*

Proof. If no point in X is a boundary-point of Y we could partition X into two open (in the product topology) sets: namely, points with an open set around them contained in Y, and points with an open set around them contained in the complement of Y. This contradicts the connectedness of X. $\square_{5.5}$

The following is an observation about complete dense orders:

Lemma 5.6 *Suppose that $\langle I, < \rangle$ is a complete dense linear order. If $B_i = \{\vec{b}_1, \ldots \vec{b}_i\}$ is a coherent sequence of boundary points of $D \subseteq I^R$ then either D is quantifier-free definable in the structure $\langle I, < \rangle$ from parameters in (the union of the ranges of) B_i, or we can add a new element $\vec{b}_{i+1}$ and extend the coherent sequence.*

Proof. Fix such a coherent sequence $B = \{\vec{b}_1 \ldots, \vec{b}_i\}$. Partition the tuples in I^R according to their type over $\{\vec{b}_1 \ldots, \vec{b}_i\}$, and let E be the corresponding equivalence relation. The equivalence classes are definable from B_i, so if D is not definable there must be some equivalence class C that contains both points in D and points not in D.

Define $\vec{g}$ as follows. Let $t \subseteq R$ be $\{j :$ for some $\vec{c} \in C$ there is $k \leq i$ with $\vec{b}_k(j) = \vec{c}(j)\}$. For $j \in t$, let $\vec{g}(j) = e$ if for some $k \leq i$ and some (equivalently all, since the elements of c are indiscernible over B_i) $\vec{c} \in C$, $\vec{b}_k(j) = \vec{c}(j) = e$. Since C has at least two points, we must have $t \neq R$. Let s denote the complement of t in R. Let $D_{\vec{g}}$ be the collection of functions $\vec{r} \in I^s$ such that $\vec{r} \cup \vec{g}$ is in D and let $C_{\vec{g}}$ be those $\vec{r}$ mapping s to I, such that $\vec{r} \cup \vec{g} \in C$.

Since $C_{\vec{g}}$ is all realizations of a type over $B_i \cup \vec{g}$, it is an equivalence class and is thus connected (using the completeness, as remarked in the beginning of this section); by Observation 5.5 there is a $D_{\vec{g}}$-boundary-point $\vec{f}$ in $C_{\vec{g}}$. Let $\vec{b}_{i+1} = \vec{g} \cup \vec{f}$ (considered as functions). Then $\vec{b}_{i+1}$ is a boundary-point of $D \cap C$ in C as required. $\square_{5.6}$

This establishes all three parts of Proposition 5.3. Since T does not have the independence property, for each formula $\theta(\vec{x},\vec{y})$ there is, by Lemma 5.4, a bound K_θ on the length of a coherent sequence of boundary points for $D = \{\vec{i} \in I : \theta(\vec{m},\vec{i})\}$. By Lemma 5.6 every coherent sequence of boundary points extends to a maximal coherent sequence, and D is definable by a quantifier free order formula over the elements in this sequence, since the sequence is a boundary point cover. $\Box_{5.3}$

We can remove the restriction that $(I,<)$ be a *complete* dense linear order if we weaken the definability conclusion. Recall that a linear order is *weakly o-minimal* if every definable subset is a disjoint union of convex sets.

Corollary 5.7 *If M lacks IP and I is a densely ordered sequence of order-indiscernibles with order type then for every L-formula $\phi(\vec{x},\vec{y})$ the trace of ϕ on $(I,<)$ is a disjoint union of convex sets. That is, the induced structure on $(I,<)$ is weakly o-minimal.*

Proof. Since I is a sequence of order indiscernibles we can find an elementary extension N of M and a sequence of order indiscernibles $J \subset N$ such that $(J,<)$ is the completion of $(I,<)$. Now applying Theorem 5.2, every definable subset of J is a union of intervals. Thus, the trace on I is a union of convex sets. $\Box_{5.7}$

6 Quantifier Reduction in Expansions by Indiscernibles: theories without the independence property

6.1 Preliminaries

Let T be an L-theory without the independence property. Fix an $|I|^+$-saturated model M of T and a complete densely ordered sequence I of order indiscernibles in M. Let $M^+ = (M, I, <)$ be the result of expanding M by interpreting $<$ as $<$ and P as I. There is no assumption that (M, I) does not have the independence property. Throughout this section T^+ is the theory of M^+.

Our goal is to show the 'order part' of Theorem 3.3. In the context of this section, a basic formula is a Boolean combination of formulas in $<$ and in L. Note that the trace on I of a basic formula is a Boolean combination of order formulas and the traces of L-formulas. Since the latter are order definable by Theorem 5.2, and the definition is quantifier free since the ordering is dense, we have Assumption 3.4. Hence it suffices to show that the hypothesis of Proposition 3.8 holds.

We fix a basic formula $\beta(\vec{x},d,\vec{c})$, with the goal of showing that the reduction in the hypothesis of Proposition 3.8 holds.

Definition 6.1 *For any $\vec{m},d$ in M let $B_{\vec{m},d}$ denote $\{\vec{c} \in I : \beta(\vec{m},d,\vec{c})\}$. Fix $K = K_\beta$ so that for any $\vec{m},d$, a coherent sequence of boundary points for $B_{\vec{m},d}$ has at most K elements.*

Definition 6.2 *(ν_δ) For any quantifier-free order formula $\delta(w_1,\ldots,w_n,y_1,\ldots,y_R)$, let $\nu_\delta(\vec{x},w_1,\ldots w_n)$ be the L^+-formula such that, for arbitrary $\vec{m}$ and $\vec{c} \in I$ we have $\nu_\delta(\vec{m},\vec{c})$ holds if and only if for some d, $\delta(\vec{c},\vec{y})$ defines $\mathrm{tp}_\beta(\vec{m},d/I)$.*

We say a sequence $\vec{b}_1\ldots\vec{b}_n$ is *σ-good* if the union of the ranges of the $\vec{b}_i$ is ordered so that $\vec{b}_1\ldots\vec{b}_n$ is a boundary point cover for the set of $\vec{c} \in I$ such that $\sigma(\vec{b}_1,\ldots,\vec{b}_n,\vec{c})$ holds.

To prove the hypothesis in Proposition 3.8, and hence to complete the proof of Theorem 3.3 it clearly suffices to show each ν_δ is equivalent to a P-bounded formula. We will actually deal not with ν_δ directly, but with a "boundary-point" version, that is easier to induct upon.

Definition 6.3 *(ρ_σ) Let $\sigma(\vec{w}_1 \dots \vec{w}_n, y_1, \dots y_R)$ be a quantifier-free order-formula in L with no parameters. Let $\rho_\sigma(\vec{x}, \vec{w}_1, \dots, \vec{w}_n)$ hold of $\vec{m}, \vec{b}_1 \dots \vec{b}_n$ if and only if there is d in M such that $\vec{b}_1 \dots \vec{b}_n$ form a boundary-point cover for $B_{\vec{m},d}$, and $\sigma(\vec{b}_1, \dots, \vec{b}_n, \vec{y})$ defines $\mathrm{tp}_\beta(\vec{m}, d/I)$.*

We will prove by downward induction on n that for any $<$-formula σ, ρ_σ is equivalent to a P-bounded formula. It follows from this that every ν_δ is equivalent to a P-bounded formula, since

Proposition 6.4 *$\nu_\delta(\vec{x}, w_1, \dots, w_n)$ holds of $\vec{m}, c_1, \dots c_n$ iff there are $i \leq K$, $\vec{b}_1 \dots \vec{b}_i$ from I and $\sigma(w_1, \dots, w_i, \vec{y})$ such that*

1. *$\delta(\vec{c}, \vec{y})$ and $\sigma(\vec{b}, \vec{y})$ are equivalent on I,*
2. *$\rho_\sigma(\vec{m}, \vec{b}_1 \dots \vec{b}_i)$ holds,*
3. *and $\rho_\sigma(\vec{m}, \vec{b}'_1 \dots \vec{b}'_i)$ holds for any $\vec{b}'_1 \dots \vec{b}'_i \equiv_{\vec{C}} \vec{b}_1 \dots \vec{b}_i$.*

Proof of Proposition: For one direction: suppose that $\nu_\delta(\vec{x}, w_1, \dots, w_n)$ holds of $\vec{m}, c_1, \dots c_n$, with d a witness. Proposition 5.3 implies that there is a boundary point cover $\vec{b}_1, \dots \vec{b}_i$ for the set $B_{\vec{m},d}$ with $i \leq K$ and a (quantifier-free) order formula $\sigma(\vec{b}, \vec{w})$ defining $B_{\vec{m},d}$. Suppose $\vec{b}'_1, \dots \vec{b}'_i$ is such that $\vec{b}'_1 \dots \vec{b}'_i \equiv_{\vec{C}} \vec{b}_1 \dots \vec{b}_i$. The statement that $\vec{b}_1, \dots \vec{b}_i$ is a boundary point cover for the set defined by $\delta(\vec{c}, \vec{y})$ is equivalent, by quantifier elimination in $\langle I, < \rangle$, to a quantifier-free order formula in $\vec{c}, \vec{b}$. As $\vec{b}'_1 \dots \vec{b}'_i, \vec{c}$ also satisfies this formula, $\vec{b}'_1 \dots \vec{b}'_i$ is also a boundary point cover for $B_{\vec{m},d}$.

For the other direction: if the conclusion holds, $\delta(\vec{c}, \vec{y})$ and $\sigma(\vec{b}, \vec{y})$ are equivalent on I and $\rho_\sigma(\vec{m}, \vec{b}_1 \dots \vec{b}_i)$ holds; thus $\nu_\delta(\vec{m}, \vec{c})$ holds. $\Box_{6.4}$

We prove the following lemma by reverse induction on n, the number of w variables.

Lemma 6.5 (Goal Lemma for Order Case) *Every $\rho_\sigma(\vec{x}, \vec{w})$ is $\vec{w}, I$ equivalent to a P-bounded formula.*

Since K is the limit on the size of a boundary point cover, ρ_σ is clearly equivalent to false when the number of variables n is greater than K. So we can use $n = K + 1$ as an easy base case of our induction: we proceed to the induction step.

We now assume that for each $\sigma(\vec{w}_1, \dots \vec{w}_i, \vec{y})$ with $i > n$, $\rho_\sigma(\vec{x}, \vec{w}_1, \dots \vec{w}_i)$ is $\vec{w}, I$ equivalent to a P-bounded formula. Here we show the same for $\sigma(\vec{w}_1, \dots \vec{w}_n, \vec{y})$. Letting E_i denote the equivalence class of $\vec{b}_i$ over $\vec{b}_1, \dots \vec{b}_{i-1}$, the proof of the induction step will follow from:

Proposition 6.6 *Suppose $\sigma(\vec{w}_1, \dots \vec{w}_n, \vec{y})$ is an order formula. Then there are n', p' such that for any $\vec{m}$ and $\vec{b} = \langle \vec{b}_1, \dots \vec{b}_n \rangle$, $\rho_\sigma(\vec{m}, \vec{b})$ holds if and only if $\vec{b}_1 \dots \vec{b}_n$ is σ-good and for every set G with $|G| < n'$, and for every set of neighborhoods N_i around $\vec{b}_i$, there exist*

1. *for $1 \leq i \leq n$, $\vec{f}_i, \vec{g}_i$ in $N_i \cap E_i$,*

2. *and a set E with $|E| < p'$ realizing all $(<, R)$-types over the union of G with the $\vec{f}_i, \vec{g}_i$ and the endpoints of N_i,*

3. *and a d in M*

such that $\sigma(\vec{b}_1, \ldots \vec{b}_n, \vec{y})$ is correct for $\vec{m}, d$ on E and $\neg[\beta(\vec{m}, d, \vec{f}_i) \equiv \beta(\vec{m}, d, \vec{g}_i)]$.

Using this proposition, it is clear that ρ_σ is equivalent to a P-bounded formula; since, fixing such an n' and p', the formula expressing the conclusion can be expressed as a P-bounded formula.

Let m' be the sum of the sizes of the defining schema for each ρ'_σ, as σ varies over $<$-formulas with variables $\vec{w}_1, \ldots \vec{w}_K$ and y. Such an m' exists by the induction hypothesis and Lemma 3.5.

Let $n' = (m' + n + 1)$. Let $p(n, r)$ be a function such that for any finite subset A of I there is a set $A' \subseteq I$ with $|A'| < p(|A|, r)$ such that all $(\{<\}, r)$-types over A are realized in A'. Let $p' = p(m', r)$. We will show that this n', p' witness Proposition 6.6. For this n' and p', denote the formula expressing the conclusion of Proposition 6.6 by $\rho'_\sigma(\vec{m}, \vec{b})$.

It is easy to see that ρ_σ implies ρ'_σ; the required witness d can always be taken to be the witness for ρ_σ and the existence of E and $\vec{f}, \vec{g}$ follow from the size of p' and the fact that $\vec{b}_1 \ldots \vec{b}_n$ is σ-good.

We now start the proof of the other implication. Suppose that ρ'_σ holds. Let $\eta(x)$ denote the L-type over $\vec{m} \cup I$ such that $\eta(d)$ asserts that d is such that $\sigma(b_1, \ldots b_n, y)$ defines the β type of $\vec{m}, d$ over I. It will take several pages to prove

Claim 6.7 *The type η is finitely satisfiable.*

Proof of Claim 6.7. Let η_0 be a finite subset of η, and let S be all the constants mentioned in η_0. Choose G so that $\vec{b} = \langle \vec{b}_1, \ldots \vec{b}_n \rangle$ is contained in G and for each i with $K \geq i > n$ and every i-variable $<$-formula $\sigma(\vec{w}_1, \ldots, \vec{w}_i, \vec{y})$, the $\rho_\sigma(\vec{x}, \vec{w}_1, \ldots, \vec{w}_i)$-type of $\vec{m}$ over I is definable over G.

Choose rectangular neighborhoods N_i of the $\vec{b}_i$ for $1 \leq i \leq n$ to separate the $\vec{b}_i$'s from the elements of $G \cup S$ that are distinct from $\vec{b}_i$: that is, choose both the upper $\vec{u}_i$ and lower endpoints $\vec{l}_i$ of N_i so that no element of $G \cup S$ lies in $(\vec{l}_i(m), \vec{b}_i(m))$ or $(\vec{b}_i(m), \vec{u}_i(m))$.

Applying the hypothesis of Proposition 6.6 to this G and $\langle N_i : i \leq n \rangle$, and by the choice of p', we get the corresponding $\vec{f}_i, \vec{g}_i$, E and d_0. That is, $\vec{f}_i, \vec{g}_i \in N_i \cap e_i$, $\sigma(\vec{b}_1, \ldots, \vec{b}_n, y)$ is correct for $B_{\vec{m}, d_0}$ on E and E realizes all $(<, R)$-types over $G^+ = G \cup \bigcup_i \{\vec{f}_i, \vec{g}_i, \vec{l}_i, \vec{u}_i\}$, where $N_i = (\vec{l}_i, \vec{u}_i)$ and $\neg\sigma(\vec{b}, f_i) \leftrightarrow \sigma(\vec{b}, g_i)$. Let N denote the collection of all components of the endpoint vectors of the N_i.

Note that since $\vec{b}$ is σ-good, it is a coherent sequence of boundary points for $\sigma(\vec{b}, \vec{y})$.

Lemma 6.8 *There is a coherent sequence of boundary points $\vec{b}'_i : i \leq n$ for $B_{\vec{m}, d_0}$ such that $\vec{b}'_i \in N_i$.*

Proof. We define the $\vec{b'}_i$ inductively. Since $B_{\vec{m}, d}$ changes sign on N_1, Observation 5.5 implies that N_1 contains a boundary point $\vec{b}'_1$. Let E'_{i+1} denote the equivalence class of $\vec{b}_{i+1}$ over $\vec{b'}_1, \ldots \vec{b'}_i$. Given $\vec{b'}_1, \ldots \vec{b'}_i$, we show that there is $\vec{b'}_{i+1} \in N_{i+1}$ which is a boundary point for $B_{\vec{m}, d} \cap E'_{i+1}$. To do this, it suffices (by Observation 5.5) to show that $N_{i+1} \cap E'_{i+1}$ contains two points which disagree on $B_{\vec{m}, d}$. We will do this by showing that $\vec{f}_{i+1}$ and $\vec{g}_{i+1}$ are in the same equivalence class for $\vec{b'}_1, \ldots \vec{b'}_i$. Let B_i denote the union of the ranges of the $\vec{b}_k$ for $k \leq i$. We know that they are in the same class E_{i+1} over $\vec{b}_1, \ldots \vec{b}_i$.

Let s be the set of indices j such that $\vec{b}_{i+1}(j) = \vec{b}_m(n)$, for some $m < i, n \leq R$ and let $\vec{e} = \vec{b}_{i+1}|s$. Then $\vec{b}_{i+1}$ is a boundary point of the set H of $\vec{h} \in B_{\vec{m},d}$ which extend $\vec{e}$ and $E_{i+1} \subseteq H$. In addition, $\vec{f}_{i+1}$,$\vec{g}_{i+1}$, and $\vec{b}_{i+1}$ all have the same restriction to s. Let $t = s^c$. Then the range $\vec{b}_{i+1}|t$ is disjoint from the ranges of the $\vec{b}_j$ for $j \leq i$.

Our choice of $\langle N_j : j \leq i \rangle$ guarantees that for each fixed $j \leq i$, the endpoints $\vec{l}_j$ and $\vec{u}_j$ of N_j satisfy the same inequalities over $\vec{b}_{i+1}|t$ that $\vec{b}_j$ satisfies. Hence all points in N_j have the same type as $\vec{b}_j$ over $\vec{b}_{i+1}|t$. So $\vec{b}'_j$, in particular, has the same type as $\vec{b}_j$ over $\vec{b}_{i+1}|t$. Thus B'_j is disjoint from $\vec{b}_{i+1}|t$.

Now, the definition of N_{i+1} guarantees that for all points $\vec{p}$ in N_{i+1}, $\vec{p}|t$ satisfies the same inequalities over B'_i that $\vec{b}_{i+1}|t$ does. In particular, B'_i is disjoint from $\vec{a}|t$ for every $\vec{a}$ in N_{i+1}. Further, $\vec{f}_{i+1}|t$ and $\vec{g}_{i+1}|t$ both satisfy the same inequalities over B'_i that $\vec{b}_{i+1}$ does. Thus, $\vec{f}_{i+1}$ and $\vec{g}_{i+1}$ are in the same equivalence class over B'_i as required. $\square_{6.8}$

Fix, by Lemma 6.8, $\vec{b}' = \langle \vec{b}'_i : i \leq n \rangle$ with $\vec{b}'_i \in N_i$ and $\vec{b}'$ a coherent sequence of boundary points for $B_{\vec{m},d_0}$.

We argue that if $\vec{b}'$ is a maximal coherent sequence of boundary points for $B_{\vec{m},d_0}$, then d_0 realizes η. For then, there is no splitting for $B_{\vec{m},d_0}$ over $b'_1, \ldots b'_n$ (i.e. $\mathrm{tp}_\beta(\vec{m}, d_0/I)$ does not split over $b'_1, \ldots b'_n$). This implies $\sigma(\vec{b}, y)$ is correct for $\vec{m}, d_0$ on S. To see this implication, choose for $\vec{q} \in S$ an $\vec{e} \in E$ which realizes the same $<$-type as $\vec{q}$ does over G^+. Note that, in particular, $\vec{e}$ and $\vec{q}$ have the same $<$-type over N and the same $<$-type over $\vec{b}$. But now the fact that the N_i separates elements of $S - B$ from B (where B is the range of $\vec{b}$) and the fact that $\vec{b}'_i$ lies in N_i gives us that $\mathrm{tp}_<(\vec{e}/\vec{b}') = \mathrm{tp}_<(\vec{q}/\vec{b}')$: e.g. if $\vec{e}(j) > \vec{b}'_i(k)$ then there is an endpoint $\vec{u}$ of N_i such that $\vec{e}(j) > \vec{u}(j) > \vec{b}'_i(k)$ and hence $\vec{q}(j) > \vec{u}(j) > \vec{b}'_i(k)$, since $\vec{e}$ and $\vec{q}$ have the same order-type over N. Now, $\mathrm{tp}_<(\vec{e}/\vec{b}) = \mathrm{tp}_<(\vec{q}/\vec{b})$ implies $\sigma(\vec{b}, \vec{q}) \leftrightarrow \sigma(\vec{b}, \vec{e})$ and $\sigma(\vec{b}, \vec{e})$ holds if and only if $\vec{e} \in B_{\vec{m},d_0}$ as $\sigma(\vec{b}, \vec{y})$ is correct for $\vec{m}, d_0$ on E. Since there is no splitting for $B_{\vec{m},d_0}$ over $b'_1, \ldots b'_n$, $\mathrm{tp}_<(\vec{e}/\vec{b}') = \mathrm{tp}_<(\vec{q}/\vec{b}')$ implies $\vec{e} \in B_{\vec{m},d_0}$ if and only $\vec{q} \in B_{\vec{m},d_0}$. Thus, $\sigma(\vec{b}, \vec{y})$ is correct for $\vec{m}, d_0$ on S and so d_0 realizes η as required.

Hence, we reduce to assuming $\vec{b}'_1, \ldots \vec{b}'_n$ are not a boundary point cover for $B_{\vec{m},d_0}$. By Proposition 5.3 1), these extend to a boundary point cover $B^+ = \{\vec{b}'_1 \ldots b'_H\}$ for $B_{\vec{m},d_0}$, with $n < H \leq K$. By the choice of $B' = \{b'_{n+1}, \ldots, b'_H\}$, some $\sigma'(b'_1, \ldots, \vec{b}'_H, \vec{y})$ defines $\mathrm{tp}_\beta(\vec{m}, d_0/I)$.

We now want to make use of our induction hypothesis on $\sigma'(b'_1, \ldots, \vec{b}'_H, \vec{y})$. Namely, we know that since $H > n$, the $\rho_{\sigma'}$ type of $\vec{m}$ over I is definable over G. This will be eventually applied in Claim 6.11. Before continuing the argument, we establish a canonical representation for equivalence classes with respect to the order. An equivalence class for some set W, i.e. a complete order type in variables $y_1 \ldots y_R$ over W, is specified by a list of intervals $\{I_k : k \leq R\}$ with endpoints in W where each I_k is either open or a singleton, plus an ordering among the variables $y_1 \ldots y_R$. This specifies that the class is the intersection of a product of the $\{I_k : k \leq R\}$, with the set of R-tuples satisfying the type over $\emptyset$ given by the ordering. Given an equivalence class t, we will call the sequence $\{I_k : k \leq R\}$ the *interval sequence* for t. Note that in the above definitions, we are talking about intervals I_k in the ordering on I, not subsets of I^R as in the previous paragraphs. We will refer to an interval with endpoints that are consecutive elements s_i, s_{i+1} from some set $S = \{s_1, \ldots < s_i < s_{i+1} \ldots < s_p\}$ as an *S-interval*. Given an S-interval I and a function f with domain S, we will talk about the *f-image* of the interval: namely, the interval whose endpoints are the f-images of the endpoints of I, and which is of the same type as I (half-open, open, etc.). Given a point x in I, the *S-interval of x* means the smallest S-interval that contains x.

We call a $G\vec{b}'$ interval which is either open or a singleton and that contains an element of E a *nice interval*. We now note the following

Lemma 6.9 *1. If e is an equivalence class for $G\vec{b}'$, and the interval sequence for e contains only nice intervals, then $\sigma'(\vec{b}', \vec{y})$ agrees with $\sigma(\vec{b}, \vec{y})$ on e.*

2. *Suppose we have a map f on $G\vec{b}'$ that is the identity on G and maps $\vec{b}'$ bijectively onto $\vec{b}''$, and such that $tp_<(\vec{b}'/G) = tp_<(\vec{b}''/G)$. If e is an equivalence class for $G\vec{b}''$, and the interval sequence for e contains only intervals that are the f-images of nice intervals, then $\sigma'(\vec{b}'', \vec{y})$ agrees with $\sigma(\vec{b}, \vec{y})$ on e.*

Proof of Lemma 6.9. To see part 1), notice that $\sigma'(\vec{b}', \vec{y})$ and $\sigma(\vec{b}, \vec{y})$ are constant on any $G\vec{b}'$ equivalence class, in particular on e. Since the interval sequence for e contains only nice intervals, e contains a tuple from E, and $\sigma'(\vec{b}', \vec{y})$ and $\sigma(\vec{b}, \vec{y})$ must agree on this tuple, hence they agree on the whole class.

Part 2) follows from 1) by transferring from $G\vec{b}'$ to $G\vec{b}''$, noting that the assertion that $\sigma'(\vec{b}', \vec{y})$ and $\sigma(\vec{b}, \vec{y})$ agree on specified intervals with endpoints in $G\vec{b}'$ is part of the type of $\vec{b}'$ over $G\vec{b}$. $\square_{6.9}$

Now we will construct a sequence $\vec{b}''$ and a function f so that we can apply part 2) to get the equivalence of $\sigma(\vec{b}'', \vec{y})$ and $\sigma(\vec{b}, \vec{y})$ on all R-tuples from S. We define a function f fixing G (and thus fixing $\vec{b}_i$ for $i \leq n$) and mapping $\vec{b}'_j$ to $f(\vec{b}'_j) = \vec{b}''_j$. Let $\{g_1, \ldots g_s\}$ enumerate G in increasing order, and let $g_0 = -\infty$, $g_{s+1} = \infty$. We will define f by cases, looking at the part of $G\vec{b}'$ in each interval (g_i, g_{i+1}). For $i \leq s+1$ do the following:

Find a nice $G\vec{b}'$ interval contained in the interval (g_i, g_{i+1}). Such an interval (l, u) exists, since E realizes all $(<, R)$-types over G. If l is not g_i, then let $f(l)$ be an element in I lying above g_i but below all elements of S in (g_i, g_{i+1}). Similarly, if u is not g_{i+1}, then let $f(u)$ be an element in I lying below g_{i+1} but above all elements of S in (g_i, g_{i+1}). Define f on the rest of the elements of $\vec{b}'$ in (g_i, g_{i+1}) so that the order type of $G\vec{b}'$ is preserved. That is, preserve the ordering among the elements of $\vec{b}'$ but move all of the elements of $\vec{b}' \in (g_i, g_{i+1})$ that were strictly below l to the interval $(g_i, f(l))$ and move all $\vec{b}' \in (g_i, g_{i+1})$ that were above u to the interval $(f(u), g_{i+1})$. Let f be the identity on the endpoints g_i, g_{i+1}.

This completes the construction of b''. The construction for (g_1, g_2) is illustrated in Figure 6.1.

It is clear that the resulting $\vec{b}''$ has the same type as $\vec{b}'$ over G. We now note the crucial property of this construction.

Observation 6.10 *For every R-tuple $\vec{h}$ from S, let the $G\vec{b}''$ equivalence class of $\vec{h}$ have interval sequence $J_1, \ldots, J_R$. Then each J_k is the image under the function f of a nice interval.*

Proof: Let $i \leq R$ be given. We will show that $\vec{h}(i)$ is in the image of a nice interval. If $\vec{h}(i)$ is in the interior of its G-interval then the construction guarantees that the $G\vec{b}''$-interval containing $\vec{h}(i)$ is the image of a nice interval. If $\vec{h}(i)$ is an endpoint of its G-interval, then $\vec{h}(i)$ is in G, hence in E, hence $\{\vec{h}(i)\}$ is itself a nice interval. The construction guarantees that f maps every element of $\vec{b}' \cap G$ to itself, hence $\vec{h}(i)$ is again in the f-image of a nice interval. $\square_{6.10}$

We now put these constructions together to show the finite satisfiability of η.

Claim 6.11 *For any finite subset S of I and set of neighborhoods N_i about $\vec{b}_i$ for $i \leq n$, there is a d_S in M and a sequence $\vec{f}_i, \vec{g}_i$ for $i \leq n$ in N_i such that for all $\vec{f} \in S$, $\beta(\vec{m}, d_S, \vec{f}) \leftrightarrow \sigma(b_1, \ldots, b_n, \vec{f})$ holds and $\neg(\sigma(\vec{b}, \vec{f}_i) \leftrightarrow \sigma(\vec{b}, \vec{g}_i))$ holds for each $i \leq n$.*

Proof of Claim 6.11: The $\vec{f}_i, \vec{g}_i$ for $i \leq n$ that are required in the conclusion of the claim are the $\vec{f}_i, \vec{g}_i : i \leq n$ chosen at the beginning of the proof of Claim 6.7. Let $\vec{b}''$ be as constructed just before Observation 6.10.

We argue that $\rho_{\sigma'}(\vec{m}, \vec{b}''_1, \ldots, \vec{b}''_n, \vec{b}''_{n+1}, \ldots \vec{b}''_H)$ holds. The $\rho_{\sigma'}$-type of $\vec{m}$ over I is definable by a P-bounded formula over G, by the defining property of G, and since $G \subset I$ this formula is equivalent over I to a $<$-formula. Hence the $\rho_{\sigma'}$-type of $\vec{m}$ over I does not $<$-split over G. But $\rho_{\sigma'}(\vec{m}, \vec{b}'_1, \ldots, \vec{b}'_H)$ holds, and hence $\rho_{\sigma'}(\vec{m}, \vec{b}''_1, \ldots, \vec{b}''_n, \vec{b}''_{n+1}, \ldots \vec{b}''_H)$ does also, since $\vec{b}''$ and $\vec{b}'$ have the same $<$-type over G.

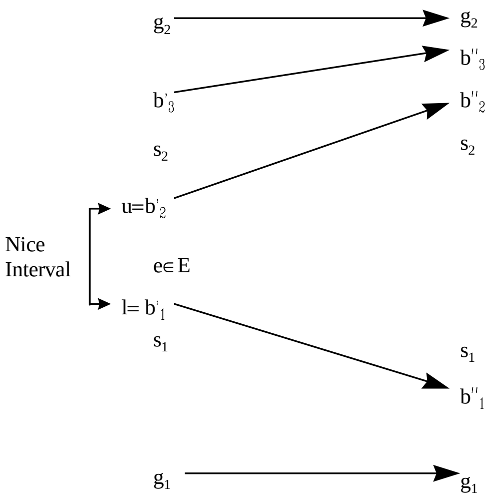

Figure 1: The Construction of the mapping f from $G\vec{b}'$ to $G\vec{b}''$.

Now, by the definition of $\rho_{\sigma'}$, we get a d_S such that for any tuple of distinct elements $\vec{f} \in I$, $\models [\beta(\vec{m}, d_S, \vec{f}) \leftrightarrow \sigma'(\vec{b}'', \vec{f})]$. Since $\vec{g}, \vec{b}'_1, \ldots \vec{b}'_H$ and $\vec{g}, \vec{b}''_1, \ldots \vec{b}''_H$ satisfy the same quantifier free order type, $\sigma'(\vec{b}'', \vec{y})$ agrees with $\sigma(\vec{b}, \vec{y})$ on the image of the $\vec{g}\vec{b}'$-equivalence classes on which $\sigma'(\vec{b}', \vec{y})$ agrees with $\sigma(\vec{b}, \vec{y})$. quantifier-free order type: By the construction of $\vec{b}''$, these classes include every R-tuple from S. So we have $\beta(\vec{m}, d_S, \vec{f}) \leftrightarrow \sigma(\vec{b}_1, \ldots, \vec{b}_n, \vec{f})$ holding for all $\vec{f} \in S$. This completes the proof of Claim 6.11. $\square_{6.11}$

Claim 6.11 implies Claim 6.7, the type η is finitely satisfiable, which completes the proof of Proposition 6.6. This finishes the proof that the hypothesis of Proposition 3.8 is satisfied. Hence Theorem 3.3 for theories without the independence property is proved. $\square_{6.7,\ 6.6,\ 3.3}$

7 Applications of the Quantifier Reduction Results to Query Collapse Theorems

We now apply the results on reductions of indiscernibles proved in the previous sections. We apply them to prove the expressivity results on embedded finite structures stated in Section 2, and also to prove a number of other interesting facts about expansions of models by indiscernibles. We start in Subsection 7.1 by stating a consequence of the reductions that captures all we actually need from the previous sections for the query collapse statements. Subsection 7.2 starts the applications to query collapse; we present a lemma allowing us to concentrate a number of finite partial counterexamples into a single infinite counterexample. Then we combine that result with the extension result of the previous section to prove the main collapse results stated at the beginning of the paper.

7.1 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 [12] and [10], which is also referred to as 'back-and-forth equivalence' [18].

Definition 7.1 *For any language L, two L-structures A and B are* partially L-isomorphic, *written $A \approx_{L,p} B$ 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 $A \approx_{L,p} B$, then A and B are elementarily equivalent in any sublanguage of L.

Notation 7.2 *Let I be a subset of N and let L^* be the expansion of L^+ by constants $\{c_\alpha : \alpha < |I|\}$. Fix a permutation g of I and an enumeration $\{a_\alpha : \alpha < |I|\}$ of I. Then N_1^g is the expansion of N where c_α is interpreted as a_α and N_2^g is the expansion of N where c_α is interpreted as $g(a_\alpha)$. We will omit the g if it is fixed in context.*

Note that $N_1 \approx_{L^*,p} N_2$ implies that the partial isomorphisms in the back-and-forth system can be taken to all extend g.

Here is the result that is relevant for collapse theorems:

Theorem 7.3 *Let N be a model that is stable (does not have the independence property) and I be a set (sequence) of (order) indiscernibles. Suppose that (N, I) is pseudo-small and that (N, I) is ω_1-saturated.*

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

$$N_1^g \approx_{L^*,p} N_2^g.$$

Proof: We use Δ to mark whether we are in the order or equality case. Note that in the order case, pseudo-small requires the (N, I) to be elementarily equivalent to (N', I'), where I' is an order-indiscernible sequence of complete dense order type. The quantifier-reduction results proved in Section 6 are all properties of the L^+-theory, and hence carry over from small (N', I') to pseudo-small (N, I).

It suffices to show the collection of pairs of finite sequences $\langle \vec{a}, \vec{b} \rangle$ such that $\langle N_1, \vec{a} \rangle$ is L^*-elementary equivalent to $\langle N_2, \vec{b} \rangle$ form a back-and-forth system. So suppose $\langle N_1, \vec{a} \rangle$ is L^*-elementarily equivalent to $\langle N_2, \vec{b} \rangle$, and a' is an element of N_1. We shall find b' such that $\langle N_1, \vec{a}, a' \rangle$ is L^*-elementarily equivalent to $\langle N_2, \vec{b}, b' \rangle$.

It suffices to find, a b' such that for each $\rho(x, \vec{w}, \vec{v})$ in L^+, the g-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 3.6, 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. $\square_{7.3}$

7.2 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 that applies in both the stable and not independence property case, and which connects the query collapse questions mentioned in the introduction to results about indiscernibles. This technique is implicit in [6] and is explicit — in a slightly different formalism – in [4]. The proofs presented here use elementary extension arguments patterned directly on the proof of Lindstrom's theorem reported in [13] and somewhat reminiscent of [4], rather than nonstandard analysis as in [6].

We then use the connection results (Lemma 7.5) and the results in the previous section on definability of types over I in L^+ to prove the expressivity results on embedded finite structures stated in Section 2, 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 stucture 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 7.4 *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 $M =< \mathbb{U}, \ldots\rangle$ exhibits *collapse to equality* if every Δ-generic query definable over M is definable in the structure $\langle \mathbb{U}, =\rangle$.

We will consider various 'collapse theorems'. In each case a failure to collapse will yield a family of finite structures as in the hypotheses of the next lemma. We will apply this lemma to produce a single infinite counterexample. Then we will apply the definability of types to eliminate that counterexample. In Lemma 7.5), when $<$ is in Δ , we know only that $<$ is a linear order of the indiscernibles; nothing is known about the behavior of $<$ on the rest of the model.

Lemma 7.5 (Concentration Lemma) *Let I be a set (Δ, L) indiscernibles in a structure M with universe $\mathbb{U}$. Suppose there are finite $\Delta\cup S$-interpretations $\vec{A}_n$, $\vec{B}_n$ with* $\operatorname{dom} A_n, \operatorname{dom} B_n \subseteq I$ *with $(I, \Delta, \vec{A}_n) =_{n,S} (I, \Delta, \vec{B}_n)$ and $(M, \vec{A}_n) \models \phi$ 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

$$\forall \vec{x} \in \mathbf{P}, \forall \vec{y} \in \mathbf{P}, \forall x \in \mathbf{P} \left(\exists y \in \mathbf{P} \left(\bigwedge_{0<i<r} (\chi_i(\vec{x}, x) \leftrightarrow \chi_i'(\vec{y}, y)) \right) \rightarrow \bigwedge_{0<i<r} (\chi_i(\vec{x}, x) \leftrightarrow \chi_i'(\vec{y}, \mathbf{f}_n(\vec{x}, \vec{y}, x))) \right)$$

and

$$\forall \vec{x} \in \mathbf{P}, \forall \vec{y} \in \mathbf{P}, \forall x \in \mathbf{P} \left(\exists y \in \mathbf{P} \left(\bigwedge_{0<i<r} (\chi_i(\vec{x}, x) \leftrightarrow \chi_i'(\vec{y}, y)) \right) \rightarrow \bigwedge_{0<i<r} (\chi_i(\vec{x}, \mathbf{g}_n(\vec{x}, \vec{y}, y)) \leftrightarrow \chi_i'(\vec{y}, y)) \right).$$

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. $\square_{7.5}$

We combine Theorem 7.3 with Lemma 7.5 to get the collapse results. We begin with stable T and generic queries.

Theorem 7.6 *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 $\operatorname{dom} A_n, \operatorname{dom} B_n \subseteq \mathbb{U}$ with $(\mathbb{U}, \vec{A}_n) =_{n,S} (\mathbb{U}, \vec{B}_n)$ and $M(A_n) \models \phi$ and $M(B_n) \models \neg\phi$. By Remark 7.4 we have (possibly rechoosing the A_n, B_n) that $A_n =_{n,S} B_n$. Fix a countable set of indiscernibles $I = \{i_0, i_1, \ldots\}$ in M. By Remark 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 7.4, we may assume, $(I, \vec{A}'_n) =_{n,S} (I, \vec{B}'_n)$.

We now have the hypotheses of Lemma 7.5 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 7.5. 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 7.2.

Now, we apply Theorem 7.3: 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. $\square_{7.6}$

We can now derive immediately the following:

Theorem 2.2: 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 now show the analagous results to Theorem 7.6 for collapse to order. We first have

Theorem 2.3 : Let $(\mathbb{U}, <)$ be an infinite linear order and M an $L \cup \{<\}$-structure with domain $\mathbb{U}$, 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.

Proof. Suppose not 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 7.6, we may assume $A_n =_{n,S,<} B_n$. And, again by Remark 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 7.4 again, we have the hypotheses of Lemma 7.5 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 7.6, using Theorem 7.3 for theories without the independence property. $\square_{2.3}$

Finally, we derive

Theorem 2.5 : 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 for any dense linear order (without endpoints) $\langle D, <\rangle$, Q is definable over $\langle D, <\rangle$.

Proof. 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 $\operatorname{dom} A_n, \operatorname{dom} B_n \subseteq I$ and $(I, <, \vec{A}_n) =_{n,S,<} (I, <, \vec{B}_n)$ 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 7.5 and we construct (M', I'). We now have a model (M', I') which is pseudo-small and in which I is linearly ordered by $<$ (although we know nothing of $<$ on the rest of M). 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 7.3, 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 7.6. $\square_{2.5}$

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.

The following result crystalizes the principle behind the arguments of this section and indicates a wider but difficult to specify collection of generalizations.

Theorem 7.7 *Let M be a structure which contains an infinite set I of (order)-indiscernibles such that any (order)-permutation of I can be extended to a back-and-forth system. Then M exhibits generic collapse to equality (order).*

Two interesting open problems are: does generic collapse to order hold in a model if and only if some definable set has NIP; does generic collapse to equality hold if and only if the model has nonempty stable part (defined in [23]).

8 The theory of the expanded structure

Let (M, I) be pseudo-small. We prove a number of results about T^+, the theory of (M, I). We have studied pseudo-small theories. A priori, each $|I|^+$-saturated model determines its own theory T^+ which is pseudo-small. Our first result is that for stable theories T and theories without the independence property, the complete theory of a pseudo-small expansion depends only 'type' (which we define next) of the indiscernible sequence.

If I is a set of (Δ, L)-indiscernibles in a model M the *type* Φ of I is the collection types $p_n(\vec{x})$ where p_n is realized by any n (properly ordered) elements of I.

Theorem 8.1 *Suppose the complete countable theory T is stable (does not have the independence property). Let M_1, M_2 be models of T and I_i is a set of (Δ, L)-indiscernibles (sequence of order-indiscernibles) in M_i such that I_1 and I_2 realize the same type. If (M_1, I_1) and (M_2, I_2) are pseudo-small then*

1. *For every L^+-formula ϕ, there is a P-bounded formula $\phi'(x)$ which is equivalent to $\phi(x)$ on both (M_1, I_1) and (M_2, I_2).*
2. *$Th(M_1, I_1, \Delta) = Th(M_2, I_2, \Delta)$.*

Proof. The first statement depends on the observation that the reductions in Sections 3, 4, or 6 never depended on the specific model M but only the types of the (Δ, L)-indiscernible set. The second then follows by proving by induction on the quantifier rank of the P-bounded formulas that for any P-bounded $\phi(\vec{x})$ and any (properly ordered) sequences of distinct elements $\vec{a} \in I_1$, $\vec{b} \in I_2$,

$$M_1 \models \phi(\vec{a}) \text{ if and only if } M_2 \models \phi(\vec{b}).$$

The base step of the induction holds as I_1 and I_2 realize the same types and the induction is trivial.

$\square_{8.1}$

We now show that T^+, the theory of pseudo-small expansions of models of T is stable and then that if (M, I) is any model of T^+ such that M is saturated, any permutation of I extends to an automorphism of M. We require one definition and a corollary to our elimination results.

Definition 8.2 *Let $p \in S(C)$ and f an elementary bijective map from C to D. Then fp denotes $\{\phi(x, f\vec{c}) : \phi(x, \vec{c}) \in p\}$.*

Lemma 8.3 *If $(N, P) \models T^+$, with $P(N) = I$, $C \subseteq N$ and α is an L-elementary permutation of IC which fixes C pointwise then for any $\vec{e}, \vec{e}' \in N$, if $\mathrm{tp}_L(\vec{e}'/IC) = \alpha\mathrm{tp}_L(\vec{e}/IC)$ then $\mathrm{tp}_{L^+}(\vec{e}'/C) = \mathrm{tp}_{L^+}(\vec{e}/C)$.*

Proof. We want to show that for any L^+-formula $\phi(\vec{x}, \vec{c})$, $M^+ \models [\phi(\vec{e}, \vec{c}) \leftrightarrow \phi(\vec{e'}, \vec{c})]$. By Theorem 3.3, we may assume ϕ is P-bounded. So we show by induction on the number of P-quantifiers that for all P-bounded ϕ, $M^+ \models [\phi(\vec{e}, \vec{c}) \leftrightarrow \phi(\vec{e'}, \vec{c})]$. Since $\mathrm{tp}_L(\vec{e'}/C) = \mathrm{tp}_L(\vec{e}/C)$, the initial stage holds. Now suppose that for L^+ formulas $\psi(\vec{w}, \vec{c})$ with less than n P-quantifiers, the lemma holds. Let $M^+ \models (\exists z \in P)\theta(\vec{e'}, z, \vec{c})$ where θ has fewer than n P-quantifiers. Then this is witnessed by some $e_0 \in I$: $M^+ \models \theta(\vec{e'}, e_0, \vec{c})$. Now $\mathrm{tp}_L(\vec{e'}/IC) = \alpha\mathrm{tp}_L(\vec{e}/IC)$ implies $\mathrm{tp}_L(\vec{e'}e_0/IC) = \alpha\mathrm{tp}_L(\vec{e}\alpha(e_0)/IC)$. So by induction $M^+ \models \theta(\vec{e}, \alpha(e_0), \vec{c})$ and the result follows. $\square_{8.3}$

Corollary 8.4 *If T is λ-stable then T^+ is λ-stable.*

Proof. We first show that T^+ is stable, i.e. stable in some cardinality κ. Since T is stable, T is stable in each κ with $\kappa^\omega = \kappa$. Fix such a κ and a model (N, I) of T^+ which is a κ^+ saturated L^+-structure. Let $C \subseteq N$ have power κ. We can expand C' to C with $|C| = |C'|$ such that $I - C$ is an indiscernible set over C. It suffices to show that only κ one types over C are realized in N. Let E be a complete set of representatives of realizations of the L^+-types over C. Define an equivalence relation on E by $e \sim e'$ if for some elementary permutation of CI which fixes C pointwise, $\mathrm{tp}_L(e'/IC) = \alpha\mathrm{tp}_L(e/IC)$.

For any e, there is a countable subset I_e of I such that $\mathrm{tp}_L(e/IC)$ does not fork over I_eC and $\mathrm{tp}_L(e/IC)$ is the unique nonforking extension of $\mathrm{stp}_L(e/I_eC)$. If $\mathrm{tp}_L(I_e/C) = \mathrm{tp}_L(I_{e'}/C)$, there is an elementary permutation α of IC which fixes C pointwise and maps $I_{e'}$ to I_e. If $\mathrm{stp}(e'/I_eC) = \alpha\mathrm{stp}(e/I_eC) = q$, then $\mathrm{tp}_L(e'/IC)$ and $\alpha\mathrm{tp}_L(e/IC)$ are both nonforking extensions of q. Thus, they are equal and $e \sim e'$. Thus the $\sim$ equivalence class of an element e depends only on i) $\mathrm{tp}_L(I_e/C)$ and ii) $\mathrm{stp}_L(\vec{e}/I_eC)$. Now there are κ^ω choices for i) and κ for ii) by stability so there are only κ-equivalence classes. But by Corollary 8.3, each equivalence class determines an L^+-type so we establish stability in L^+.

To extend this result to all κ (in which T is stable) we may first assume that T is superstable. This implies that in the second paragraph I_e can be chosen finite. So, κ^ω at the end of the paragraph becomes κ and the result follows. $\square_{8.4}$

Now we show that every permutation of a set of indiscernible of a saturated model M^+ of T^+ extends to an automorphism of M.

Theorem 8.5 *Let T be a stable theory, M a model of T, and I a set of indiscernibles in M with (M, I) saturated and (M, I) pseudo-small. Then every permutation of I extends to an automorphism of M.*

Proof. Since (M, I) is pseudo-small so by Corollary 3.6 we have every L^+-type over I is defined by an equality formula. Let g be a permutation of I. We proceed exactly as in the proof of Theorem 7.3. The only distinction is that the type $g_1(q)$ is over $J_0 \cup \{\vec{b}\}$ where $\vec{b}$ now can have any cardinality less than $|M|$; here we rely on full saturation of the model.

$\square_{8.5}$

In the language of [19] and [24] we have shown that if (M, I) is saturated and pseudo-small, M is *symmetric* over I; the last argument is essentially the same as Theorem 8 of [24].

Note: It is essential that we extend a permutation of a single set - not a partial isomorphism between two sets of indiscernibles. For example, consider an equivalence relation with two infinite classes. If one set of indiscernibles exhausts one class, while a second of the same size omits a few points of the other class, then no bijection of the two sets can be extended to make the associated M_1 and M_2 partially isomorphic, let alone to an automorphism.

9 Conclusions

There are several questions on the scope of the P-bounded reduction theorems proved in the paper. Can we characterize the assumptions on the model and predicate P that allow P-bounded reduction to hold for more general expansions by unary predicates, not just indiscernibles? The results in [6],[4] and [14] are relevant here, since they give sufficient conditions for such reductions over expansions by arbitrary pseudo-finite sets. There are a number of analogies with Poizat's theory of beautiful pairs [33]. Casanovas and Ziegler [9] have provided a generalization of our results here and of Poizat's theory to show that if (M, A) is small and A does not have the finite cover property, then (M, A) is stable.

The above work is entirely in a stable context. In the order context we have the following conjecture.

Conjecture 9.1 *If T does not have the independence property, then T^+ does not have the independence property.*

The proofs of the bounded-quantifier collapse result here rely on approximation arguments: we approximate a formula asserting the existence of a witness for many elements of I by a particularly nice conjunction of first-order L formulas. In this sense, they seem to be extensions of the results in [22]. We are also interested in seeing if the results here on theories without the independence property can be proved using the techniques of [23]. We are interested in using the techniques to expand a stable theory to get stronger collapse results, such as those given for more restricted theories in [6].

Belegradek, Stoubouskin, and Tsaitlin conjecture in [4] that if M is an expansion of $(Z, +, <)$ and locally order-generic queries on M do not collapse to order then the theory of M is undecidable. In view of Theorem 2.3, this conjecture can be reformulated as: any expansion of $(Z, +, <)$ that has the independence property is undecidable.

The structure M is called *symmetric* over I if every permutation of I extends to an automorphism of M. We have shown that if I is pseudo-small and (M, I) is saturated then M is *symmetric* over I. One naturally hopes

Conjecture 9.2 *If (M, I) is pseudo-small and M is an uncountable saturated model then M is* symmetric *over I.*

Even more strongly we speculate: If M is an uncountable saturated model of a stable theory and I is an indiscernible set, then M is *symmetric* over I.

This holds by standard stability arguments if the average type of the indiscernibles is regular, which implies the result for superstable T. The results here show that this holds provided $\langle M, I\rangle$ is elementary equivalent to $\langle M_0, I_0\rangle$, where M_0 is $|I_0|^+$-saturated.

Acknowledgements: We would like to acknowlege helpful conversations with a number of people. Marko Djordevich checked the proof of part 1 of Theorem 8.1 and suggested the simple inductive proof of the second. The remark in Theorem 8.5, that the Theorem also holds with the hypothesis of pseudo-small stable replaced by the hypothesis superstable was made by Hrushovski and verified with Pillay. The goal of extending the collapse-to-order results for o-minimal structures in [6] to a class of models closed under reduct was suggested by Jan Van den Bussche.

References

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

[2] A.K. Ailamazyan, M.M. Gilula, A.P. Stolboushkin and G.F. Shvarts. Reduction of a relational model with infinite domains to the finite-domain case.

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

[4] O.Belagradek, A. Stolboushkin, M. Tsaitlin. Extended order-generic queries. Manuscript. To appear.

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

[6] M. Benedikt, G. Dong, L. Libkin, L. Wong. Relational expressive power of constraint query languages. In *Proceedings of 15th ACM Symposium on Principles of Database Systems*, pages 5–17, Montreal Canada, June 1996.

[7] M. Benedikt and L. Libkin. On the structure of queries in constraint query languages. In *Proceedings of 11th IEEE Symposium on Logic in Computer Science, New Brunswick, New Jersey* 1996.

[8] M. Benedikt and L. Libkin. Languages for Relational Databases over Interpreted Structures. In *Proceedings of 16th ACM Symposium on Principles of Database Systems*, Tuscon Arizona, June 1996.

[9] E. Casanovas and M. Ziegler. Stable theories with a new predicate. preprint.

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

[11] A. Dawar, S. Lindell, S. Weinstein. First order logic, fixed point logic, and linear order. In *Computer Science Logic '95*, LNCS vol. 1092, 1996, pages 161–177.

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

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

[14] J.Flum and M. Ziegler. Pseudo-finite Homogeneity and Saturation. Preprint. 1997.

[15] E. Grädel and Y. Gurevich. Metafinite model theory. In *Proceedings of Logic and Comput. Complexity*, LNCS vol. 960, 1994, pages 313–366.

[16] S. Grumbach and J. Su. First-order definability over constraint databases. In *Proc. Conf. on Constr. Progr. 95.*

[17] L. Hella,P. Kolaitis, and K.Luosto. Almost Everywhere Equivalence of Logics *Bulletin of Symbolic Logic*, December 1996.

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

[19] W. Hodges and A. Pillay. Cohomology of Structures and some problems of Ahlbrandt and Ziegler *J. London Math. Soc.*, 50 (1994) 1-16.

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

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

[22] H.J. Keisler. Finite Approximation of Infinitely Long Formulas. In *The Theory of Models* Addison, Henkin, and Tarski,eds. North-Holland, 1965,pp 158-169.

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

[24] H. Kikyo and A. Tsuboi. On reduction properties. *The Journal of Symbolic Logic* (59) 1994 900–911.

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

[26] D. Macpherson and C. Steinhorn. On variants of O-minimality. preprint.

[27] C. Marker and C. Steinhorn. Types definable in O-minimal theories. *The Journal of Symbolic Logic* (59) 1994 185-198.

[28] 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.

[29] 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.

[30] A. Pillay. Definability of types and pairs of O-minimal structures. *The Journal of Symbolic Logic* (59) 1994 1400-1409.

[31] A. Pillay and S. Shelah. Classification theory over a predicate I. *Notre Dame Journal of Formal Logic* 26 (1985) 361-376.

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

[33] Bruno Poizat. Paires de structure stables. *The Journal of Symbolic Logic* (48) 1983 239–249.

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